Research in Commotion:

Measuring AI Research and Development through Conference Call Transcripts

Paul E. Soto
Federal Reserve Board

ABSTRACT

This paper introduces a novel measure of firm-level Artificial Intelligence (AI) Research & Development—the AIR Index—derived from the semantic similarity between earnings conference call transcripts and leading AI research papers. The AIR Index varies widely across industries, with sustained strength in computer and electronic manufacturing, and accelerating growth in computing infrastructure and educational services seen after the introduction of ChatGPT in November 2022. I find that the AIR Index is associated with an immediate increase in Tobin's Q and can help explain the cross-section of cumulative absolute returns following the conference call, suggestive of investors valuing substantive AI discussions in the near-term. A sharp rise in the AIR Index leads to persistent increases in year-over-year capex growth, lasting about a year before tapering off, indicative of the life cycle of AI-induced capital deepening. However, I find no significant effects of AI R&D on productivity or employment. Using industry level survey data from Census, I find that recent growth in the AIR Index correlates with broader AI adoption trends. The positive association of the AIR Index with capex and valuation holds across previous time periods, suggesting that Generative AI may be the latest form of an ongoing technical innovation process, albeit at an accelerated pace.

JEL CLASSIFICATION: O32, E22, C49
KEYWORDS: artificial intelligence, capital expenditure, corporate finance, natural language processing, productivity.

* This version is from January 2025. Paul E. Soto: Federal Reserve Board, Washington D.C., USA, paul.e.soto@frb.gov. I would like to thank David M. Byrne, Leland Crane, and Robert Kurtzman for helpful comments and suggestions. The analysis and conclusions set forth are those of the author and do not indicate concurrence by the Federal Reserve Board of Governors.

I. INTRODUCTION

Artificial intelligence (AI) is becoming a common feature in many aspects of the economy, yet understanding how firms invest in and implement these technologies is crucial for assessing its potential impact at the firm-level. Traditional measures of research and development (R&D) investment, such as capital expenditure or patent counts, often reflect only the outcomes of AI initiatives. This paper attempts to estimate firms' early-stage AI initiatives by proposing a novel higher-frequency measure quantifying AI R&D based on discussions in conference call transcripts. By leveraging natural language processing (NLP) techniques, this new firm-level measure—abbreviated as the AIR Index (from A.I. R&D)—allows for a close examination of the historical growth in AI integration across industries, as well as the effect of AI R&D on key firm variables.

The AIR Index is constructed by extracting and analyzing the content of conference call transcripts and comparing it to AI academic research. Using a sentence transformer model, I quantify both the conference call text and a corpus of academic papers from computer science (the broader field encompassing artificial intelligence and machine learning) as dense vectors that capture their semantic meaning. The AIR Index for a firm in a given quarter is the average similarity between the vectors of the conference call and contemporaneous computer science literature. I then assess how a firms' AI-related discussions align with key firm characteristics such as market valuation, productivity, and capital expenditure. The intuition behind this approach is that genuine engagement with AI, as reflected in substantive research-driven discussions, is likely to be valued by investors and may influence a firms input factors—including larger capital expenditures— or productivity.

I find that firms with higher AIR indices experience significant increases in Tobin's Q—a proxy for market valuation of intangible assets. To ensure that the AIR Index is not merely measuring the hype around AI, the analysis includes a control for the frequency of AI buzzwords—that is, the count of the phrases "artificial intelligence" and "machine learning", two expressions that have been tied to investor exuberance over the last business cycle.[1] The AIR

[1] Evidence from previous technological cycles, such as the dot-com bubble of the late 1990s, suggests that firms deliberately engage in superficial and strategic communication—such as name changes— to exploit investor sentiment around hyped technologies (Cooper et al. 2001). The phrases "artificial intelligence" and "machine learning" were chosen because of their colloquial and technological usage to describe recent technological innovation, and their potential association with investor sentiment. As mentioned in Brynjolfsson and Mcafee

Index remains positive and significant, while the measure of AI buzzwords is small in magnitude and insignificant. This suggests that investors place a premium on meaningful AI integration rather than superficial mentions of AI hype. However, this impact on Tobin's Q is short-lived—lasting no more than two years and highlighting the market's immediate but transient response to AI-related discussions.

While Tobin's Q captures lower-frequency valuation effects, I also assess the impact of the AIR index on cumulative abnormal returns (CAR) immediately following the conference calls. Using a five-factor Fama-French model (Fama and French, 2015), I find that the AIR Index helps explain the cross-section of returns, with a one standard deviation increase in the index leading to a 24 basis point rise in 1-day CAR. The effect persists over short horizons (3 and 30 days) but dissipates near 60 and 90 days, consistent with the effect of Tobin's Q that investor reactions to AI R&D discussions are transitory.

Traditional measures of investment often struggle to capture expenditure in new technologies, especially emerging fields like AI. Brynjolffson et al (2021) argue that investments in intangible assets, such as AI, may initially lead to an increase in capital expenditure without immediately reflecting in productivity gains. This aligns with this papers findings, as I show that the AIR Index is associated with higher growth in capex—lasting about a year before tapering off. However, the increased valuation and larger AI-driven investments do not translate into significant improvements in labor productivity nor changes in the firms workforce. This lag is consistent with the J-Curve effect (Brynjolffson et al 2021; Brynjolffson and McElheran 2016), where the benefits of technological investments become apparent only after a period of organizational adaptation and integration.

Lastly, I find that the impact of AI R&D on valuation and capital investments is not simply a recent phenomenon tied to the rise of LLMs or ChatGPT. Restricting the analysis to the years 2004-2015 and 2004-2020, the results continue to suggest that the AIR Index has a positive and significant relationship with Tobin's Q and capital expenditure growth well before the recent generative AI hype. Nonetheless, using the U.S. Census Bureau survey on AI adoption, I find that industries with higher AIR Index growth following the 2022 release of ChatGPT show greater adoption rates in 2024. The robustness of these results is indicative of Generative AI

(2017), "The most important general-purpose technology of our era is artificial intelligence, particularly machine learning (ML)."

being the latest form of an ongoing technical innovation process, albeit at an accelerated pace since 2020 with noteworthy AIR index increases in high-tech industries such as computer and electronic manufacturing, as well as educational services.

This paper contributes to three strands of literature. First, this paper provides early firm-level evidence of the potential impact of AI on the economy. Historically, IT investments of the 1980s and 1990s led to IT-induced productivity gains, particularly in industries related to IT-producers (Stiroh 2000). More recently, in the context of artificial intelligence, Autor et al. (2003) and Acemoglu and Restrepo (2018) suggest AI will affect the economy via the labor force, arguing that while automation—a potential side effect of AI technologies—may replace some jobs, it simultaneously creates new tasks where labor maintains a comparative advantage. Eloundou et al. (2023) and Eisfeldt et al. (2023) characterize the labor market impacts of generative AI, showing that the release of ChatGPT significantly impacted firm valuations and the ability to scale non-routine tasks. Baily et al. (2023) suggest that AI-driven productivity gains may manifest more rapidly than past technological cycles due to the ease of AI integration into existing frameworks. While recent surveys indicate widespread and growing adoption of GenAI tools (Bick et al. 2024 and Humlum and Vestergaard 2024), the impact on firm activities remains to be seen. The analysis of the AIR Index on key firm metrics suggests that while AI initiatives drive short-term boosts in firm valuations, productivity effects remain delayed.

Second, this paper contributes to the literature on the potential challenges of measuring new technologies, particularly in the context of growth accounting. Corrado et al. (2005) propose an expanded framework for measuring technology-driven capital. Byrne et al. (2016) and Byrne et al. (2018) argue that standard measures of labor productivity and total factor productivity may underestimate the contributions of modern technologies like IT due to accounting discrepancies. The AIR Index offers not only a new way of measuring firm-level exposure to new technologies, but also the extent of capital deepening following a sharp rise in AI R&D.

Lastly, this paper contributes to the usage of textual data and natural language processing (NLP) in economics and finance. Ash and Hansen (2023) emphasize how unstructured data from sources textual sources can reveal important firm characteristics not encoded in quantitative variables. Prior studies using firm level documents, such as earnings conference calls and regulatory filings, have explored the use of NLP techniques to measure firm level exposure to pandemics (Hassan et al. 2023; Davis et al. 2020), geopolitical risk (Caldara and Iacoviello

2022) and bank-level uncertainty (Soto 2021). Overall, this paper demonstrates the value of unstructured data for tracking firm- and sector-specific AI R&D trends and their potential economic impact.

Most similar to this paper, Babina et al. (2024) measure firm-level AI investments by analyzing text from resumes and job postings over 2010–2018, finding that AI investment leads to higher valuations and larger employment and sales growth. Their results largely mirror the valuation effects reported here, yet they also observe a positive employment impact—absent in this paper—and no significant effect on productivity (measured by sales per worker). The divergence in employment outcomes may stem from Babina et al.'s (2024) focus on AI-skilled human capital, as their text corpus reflects largely labor demand for AI roles, whereas this paper's AIR Index captures a broader measure of R&D investment.

The paper is structured as follows: Section II outlines the data used in the analysis. Section III summarizes the evolution of AI research through the lens of computer science academic papers and illustrates how this corpus is used to develop the AIR Index at the firm-level. Section IV presents the impact of the AIR Index on valuation, productivity and capital expenditure. Finally, Section V concludes.

II. DATA

First, I download data on public firms from Compustat to examine firm performance and market valuation. Specifically, I collect the following variables: at (total assets), capx (capital expenditures), csho (common shares outstanding), ceq (common equity), dt (total debt), dp (depreciation), emp (number of employees), oibdp (operating income before depreciation), ppent (property, plant, and equipment—net), sale (net sales) and prcc (closing stock price). These variables allow me to compute Tobin's Q, a proxy for market valuation of intangibles, defined as [(prcc × csho)+at–ceq]/at. Additionally, I estimate labor productivity as sales per employee.

I merge the firm data from Compustat with quarterly earnings conference call transcripts obtained from S&P Global. Since S&P Global uses its own identifier system (companyid), I employ a crosswalk within WRDS to link the companyid to the gvkey used in Compustat. For stock price information, I further merge the dataset with CRSP using the CUSIP associated with each companyid. This enables me to estimate cumulative abnormal returns (CAR) for each

firm following its earnings call. Specifically, for each call and firm, I regress daily returns from the year prior to the call on the Fama-French five-factor model (Fama and French, 2015) to estimate predicted returns for the day of the call and subsequent days. CAR is calculated as the summed difference between actual returns and these predicted returns, aggregated over various horizons such as 1 day, 3 days, 30 days, 60 days, and 90 days. I merge firm tickers with data from the I/B/E/S (Institutional Brokers Estimate System) database to obtain earnings surprises, defined as the difference between the actual and forecasted EPS divided by the standard deviation of the forecasts. The resulting dataset includes 4,589 firms with available earnings call data and firm-level variables including capital expenditure, Tobin's Q, and market reactions following the call.

Lastly, as a basis for AI research, I construct a corpus of influential computer science papers from 2001 to 2024. Using the Semantic Scholar API, I retrieve titles and abstracts for the top 100 cited journal articles and conference papers (as of October 2024) each year within the "computer science" field.[2] This selection process prioritizes papers with substantial academic influence, ensuring a focus on the most impactful AI research.

III. METHODOLOGY

Calculating firm-level AI R&D discussions involves comparing the text of conference call transcripts with a relevant corpus on artificial intelligence. This section begins with a discussion of trends in AI research over the past twenty years, using the corpus of influential "computer science" papers as a basis for AI research. I then explain how I measure the AI R&D (AIR) index using these academic papers.

AI RESEARCH OVER TIME

I begin by documenting the dynamic themes in computer science and artificial intelligence research. Figure 1 presents a series of word clouds that visualize the changing focus in AI research abstracts.[3] The word clouds highlight a progressive shift in terminology, supporting

[2] This classification may result in innovations in computer science unrelated to AI—such as cryptography, systems theory, and hardware-specific research—to be treated as "AI." However, as the subsequent analysis of these papers shows, these types of breakthroughs are not the most prevalent in highest cited computer science research.

[3] Each year, I filtered stopwords and commonly recurring scientific words, such as "technique", "system", and "results", to remove uninformative and commonly recurring terms.

the methodology proposed in this paper of using a dynamic approach to measure AI innovation. Words like “parameter,” “efficient,” and “measurement” are prominent in the early 2010s, reflecting a period where research largely centered on optimizing traditional machine learning models.[4] Much of the research around this time focused on refining machine learning models for increased accuracy and computational efficiency.

In 2014, however, the word “architecture” begins to surface, signaling the beginning of a shift toward more complex models like neural networks. By 2017, “deep learning” dominates the abstracts in the field, marking a period where the research interests had moved decisively towards scalable, data-intensive models. A particularly notable milestone in this change was the publication of “Attention is All You Need” in 2017, introducing the transformer architecture. This work, now recognized as a crucial precursor to large language models (LLMs), redefined the way models can handle vast and complex datasets, paving the way for today’s most advanced AI tools.

The evolution accelerates further in 2022, as terms like “transformer,” “benchmark,” “diffusion,” and “generate” become the focus of academic AI papers. This vocabulary shift reflects the rising significance of generative AI and the use of standardized benchmarks to evaluate large models. In 2023, the year following ChatGPT’s release in 2022, “LLM” and “ChatGPT” emerge as the most frequent terms, illustrating the significant impact of conversational AI and large language models on the research landscape.

Overall, Figure 1 captures the dynamic nature of AI research and the movement of research from traditional machine learning approaches to architectures defined by transformers and LLMs. These evolving trends underscore the need for a dynamic semantic similarity approach to track AI developments, as the conceptual focus in AI has continuously shifted, reflecting a rapid pace of technological advancement. In the next section, I describe how the abstracts are used to measure firm-level exposure to AI R&D.

[4] Traditional machine learning can be considered a subfield of AI, encompassing algorithms that learn patterns and make decisions based on data, of without explicit functional forms. Given their reliance on large datasets, these models are sometimes referred as "data-driven" or "big data" approaches. While they automate specific tasks like clustering and classification, machine learning is often considered a simpler form of AI as it focuses on narrow tasks rather than full cognitive abilities.

MEASURING FIRM-LEVEL EXPOSURE TO AI R&D

The AIR Index compares the content of earnings conference call transcripts to leading academic papers on AI and ML. For each conference call and each response *r* within a firm f's call at quarter q, I convert the text into a 384-dimensional dense vector representation—$Text_Vector_{r,f,q}$—using the sentence transformer model (Reimers 2019).[5] This model translates a given text into an embedding that captures the semantic meaning of the text. I then compute the average embedding of all academic paper abstracts related to AI and ML over the last three years of the conference call, resulting in another 384-dimensional vector $AI_Papers_Vector_{past\ 3\ years}$. To measure AI R&D exposure of a given response in an earnings call, I calculate the cosine similarity between the sentence transformer representation of the response and the average embedding of the AI/ML papers. Higher cosine similarity scores indicate that the response is more aligned with AI/ML research, while lower scores suggest less relevance. I compute the firm's AIR Index at the conference call level by averaging the cosine similarity scores across all responses in a given call. Thus the AIR Index for a firm f in quarter q is calculated as follows:

$$AIR_Index_{f,q} = \frac{\sum_{r=1}^{R} cos_sim(Text_Vector_{r,f,q}, AI_Papers_Vector_{past\ 3\ years})}{R}$$

This results in a firm-quarter level dataset that reflects the firm's AI R&D discussion intensity based on the content of its earnings conference calls.

Table 1 presents a selection of the highest-ranked responses based on the AIR Index similarity score. The first example highlights Cerence in 2023Q3, an AI software company specializing in the automotive industry. During this call, Cerence discussed the integration of generative AI and large language models (LLMs) to enhance user experiences, explicitly referencing the transformer architecture—a core component of recent advances in AI. This alignment with a popular contemporaneous academic theme underlines the relevance of the semantic-based methodology. Other examples from 2023 include C3, an enterprise AI software firm; TaskUs, an outsourcing company focused on customer support; and Cineverse, an entertainment company that mentioned training LLMs on their own film-based data. Notably, NVIDIA's 2019 call demonstrates an early commitment to previous cutting-edge AI models such as BERT and its successors RoBERTa and GPT-2. Collectively, Table 1 underscores that firms

[5] https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2

scoring high with this measure are not only engaged with cutting-edge AI research but are incorporating detailed, research-driven terminology into their discussions.

Figure 2 shows the average AIR Index over time. In the early 2010s, we observe a gradual upward trend, reaching a notable peak toward the end of 2015 before tapering slightly. The increase from 2013 to 2015 is primarily driven by companies in fields such as biological sciences, chemical manufacturing, and computer and electronic product manufacturing. For instance, in Q1 2015, Mobileye, a driver assistance technology firm, devoted part of its conference call to discussing "optical-object detection" alongside "deep learning," emphasizing that "we've worked with deep learning technology for a couple of years already." Similarly, Merrimack Pharmaceuticals, a biopharmaceutical firm, highlighted its systems biology approach, describing it as an "integrated engineering, computing, and biology big data approach to understanding the complex network interactions that drive cancer."

Between 2018 and 2023, a steady rise in the AIR Index emerges, reflecting the increased alignment of conference call content with research on deep learning and advanced model architectures. This period captures firms' efforts to leverage deep learning models and data-driven approaches. The figure also shows that in any given quarter, the AIR Index distribution is relatively symmetric, with the mean and median closely aligned. Additionally, the rise in the interquartile range over time indicates that firms more broadly are engaging with detailed AI research topics in their conference calls. Nonetheless, this aggregated trend conceals important industry-specific differences, which are explored in the following figure.

Figure 3 presents the AIR Index over time across key industries, highlighting variation in the distribution of AI R&D exposure. Each subplot includes a dotted black line representing the overall average across industries, providing a benchmark for each sector's performance. Within each industry, the distribution of firms is shown quarterly, with the mean of that industry in a given quarter indicated by a black diamond.

The first three subplots focus on manufacturing sectors with traditionally higher-than-average AI R&D exposure: Electronic/Computer Manufacturing, Chemicals Manufacturing, and Machinery Manufacturing. Notably, Electronic and Computer Manufacturing has maintained above-average AI R&D exposure, particularly in recent years, while AI-focused discussions in Chemicals Manufacturing have tapered off. Other manufacturing sectors generally fall below the average AI R&D exposure.

In recent years, sectors outside manufacturing have shown strong AI engagement, particularly Computing Infrastructure; Professional, Scientific, and Technical Services; and Educational Services. Computing Infrastructure and Educational Services display significant right-skewed distributions. This tail skewness suggests that a subset of firms within these sectors are deeply engaged in AI research, driving much of the industry's conversation.

Figure 3 also includes vertical red lines representing structural break points identified using the Bai and Perron (1998) method to detect shifts in the average AIR Index within each industry over time. Nearly all industries experience a structural break sometime between 2019 and 2022. Educational Services and Computing Infrastructure experienced another ramp up in their discussions of AI R&D, with structural breaks observed again in 2022Q4 and 2023Q1, respectively, coinciding with the release of ChatGPT in November 2022. Indeed, both of these industries have seen stark transformations since then. For instance, education platforms, e.g. Chegg, which traditionally relied on subscription models for textbooks, tutoring, and study guides, have faced steep competition from cheaper AI tools like ChatGPT.[6] Contrastingly, high requirements to run large scale AI operations have raised computing infrastructure. Overall, the presence of detected breaks in the AIR Index, which was already trending upward across industries, points to an acceleration in AI-related initiatives and a growing focus on LLMs and generative AI more broadly.

IV. Impact of the AIR Index on Firm Characteristics

First, I examine the impact of the AIR Index on Tobin's Q, the ratio of a firm's market value to the replacement cost of its assets—a widely used measure of firm valuation that proxies for intangible assets (Hayashi 1982). Firms with high Tobin's Q are typically associated with a greater stock of intangible assets, including intellectual property, human capital, and technological capabilities that are not fully captured on balance sheets. As a result, I estimate the following baseline regression:

$$TobinsQ_{f,y} = \beta_0 + \beta_1 AIR_Index_{f,y}, + X_{f,y} + \epsilon_{f,y}$$

[6] Chegg has faced significant competition from free and low-cost AI alternatives, prompting the launch of CheggMate in 2023, an AI-powered learning tool built with GPT-4. Nonetheless, its stock price has declined by 99% since 2021 (WSJ, November 2024).

where $TobinsQ_{f,y}$ is the year-over-year growth in Tobin's Q for firm f in year y, $AIR_Index_{f,y}$ is firm f's yearly average of the AIR Index for year y, and $X_{f,y}$ is a vector of firm-level controls, starting off with the log of total assets. Standard errors are clustered at the firm level. The AIR Index has been standard normalized, allowing the coefficient β_1 to be interpreted as the effect of a one standard deviation change in the AIR Index on Tobin's Q.

Table 3 shows the results. In column 1, I note a positive and statistically significant relationship between the AIR Index and Tobin's Q growth. In Column 2, I control for industry dynamics by including 4-digit NAICS industry-year fixed effects, further ensuring that the AIR Index is not simply capturing broad industry trends. The results remain robust and significant. To address concerns that the AIR Index may be conflating substantive AI research with superficial mentions of AI-related buzzwords, in Column 3 I replace the AIR Index with a simple count of mentions of "artificial intelligence" and "machine learning" in conference calls. The coefficient is insignificant, suggesting that superficial mentions of AI buzzwords have little influence on market valuation. Finally, in Column 4, I include both the AIR Index and the buzzword count as explanatory variables, along with all controls. The coefficient on the AIR Index remains positive and statistically significant. This result implies that investors place greater value on firms with substantive AI R&D discussions, as reflected in the AIR Index, rather than firms that merely mention AI-related buzzwords.

Figure 4 shows the time dynamics of the AIR Index's effect on Tobin's Q, using local projections to trace the impact of AI R&D discussions on intangible valuation over time (Jordà 2005). The horizontal axis represents the time horizon H, with each point corresponding to the estimated coefficient from a regression of year-over-year Tobin's Q growth at y+H on the AIR Index at year y, controlling for firm size, AI buzzword frequency, firm fixed effects, and 4-digit NAICS industry-year fixed effects. Prior to the fiscal year in which AI R&D discussions begin, there is no significant relationship between the AIR Index and Tobin's Q, as reflected in the insignificance of the coefficients for H equal to -3, -2 and -1. This finding suggests that AI-related market reactions are not driven by pre-existing firm trends or market expectations that precede the discussions. The effect of the AIR Index begins to materialize in the fiscal year when AI R&D discussions are first introduced (i.e. H=0), with a positive and statistically significant impact on Tobin's Q. The effect remains positive and significant at H=1, albeit with a slight reduction in magnitude, suggesting that the initial market enthusiasm carries over into the

following year but begins to decline. Beyond this short-term horizon, the effect of AI R&D discussions on Tobin's Q diminishes and becomes statistically insignificant for future time horizons. This chart is suggestive of the transitory nature of the market's response, highlighting that while AI R&D may generate immediate positive expectations, these effects are short-lived.

I next examine whether AI R&D discussions affect valuation at a higher frequency. While Tobin's Q captures long-term valuation changes, cumulative abnormal returns (CAR) show how markets react immediately following earnings conference calls. Recent extensions to the original three-factor Fama-French framework, such as Fama and French (2015) and Hou, Xue, and Zhang (2015), motivate my use of a five-factor Fama-French model to estimate predicted returns. I compute CAR over 1-day, 3-day, 30-day, 60-day, and 90-day horizons. Table 4 shows that a one standard deviation increase in the AIR Index leads to a 24 basis point increase in 1-day CAR (column 1), indicating an immediate positive market reaction. Including firm-year fixed effects (column 2) significantly increases the model's explanatory power while leaving the AIR Index coefficient robust and significant. Adding a control for earnings surprise—typically associated with post-call price changes—confirms the AIR Index remains a strong predictor of CAR. Notably, AI buzzword mentions are insignificant across all specifications, consistent with the market's focus on substantive AI R&D discussions.

Over longer horizons, the effect of the AIR Index diminishes. While significant over 3-day and 30-day CAR windows, the association dissipates, becoming insignificant at the 60-day and 90-day marks (columns 4–7). This suggests that the market reaction to AI R&D discussions is transitory, concentrated in the short term. These findings align with the increases in Tobin's Q. Collectively, the results highlight the role of AI R&D discussions in shaping investor expectations, with meaningful but time-limited effects on abnormal returns.

Table 5 presents the broader implications of AI R&D discussions across a range of dependent variables. Column 1 reports the effect of the AIR Index on year-over-year labor productivity growth as measured by sales per employee This specification suggests no contemporaneous impacts of AI discussions on productivity. While AI R&D may contribute to firm value (as seen in Tobin's Q), these discussions do not appear to immediately translate into productivity gains. In fact, in columns 2 through 4, I check whether productivity effects appear in the years following a sharp increase in the AIR index. In all cases, I find no growth in labor productivity. In columns 5 through 8, I examine the effect of the AIR Index on labor, measured

by the year-over-year growth in the log of employees. The AIR Index also shows no significant impact, indicating that firms engaging in AI R&D discussions do not necessarily experience immediate changes in their labor force, consistent with AI-driven innovations immediately being linked to skill-based rather than labor-intensive operations. Similar to productivity, the effects are muted over the following three-year time horizons. Overall, the results point to no productivity or labor effects from AI R&D intensity.

Given that R&D is often associated with substantial investments, in Figure 5, I examine quarterly data to assess how quickly increases in AI R&D translates into capital investment. Using local projections, I regress year-over-year quarterly capital expenditures growth on the AIR Index. Six quarters prior to a one standard deviation increase in the AIR Index, there is no significant impact on capital expenditures. However, beginning just one quarter after a sharp rise in AI R&D discussions, year-over-year capex growth jumps by a statistically significant 1%, a trend that persists for about four quarters before diminishing and becoming insignificant by the fifth quarter. This pattern aligns with capital deepening observed following major technological advances, as firms ramp up spending on fixed investments to harness the potential of new technologies. It also underscores the cyclical nature of AI-driven investments, documenting that firm investment growth lasts about year as they operationalize and integrate emerging AI applications into the firms operations.

Table 6 further demonstrates that these effects are not solely driven by recent trends in large language models (LLMs) and ChatGPT post-2022. Even when limiting the sample to 2004-2015 (columns 1 and 2) or 2004-2020 (columns 3 and 4)—well before generative AI gained widespread attention— the AIR Index is positively and significantly related to both Tobin's Q and subsequent quarter capital expenditure growth. During these earlier periods, capex growth increased by 1.5% and 1.3% year-over-year, respectively, following a one standard deviation rise in AI R&D. This consistent relationship across different time frames highlights the persistent impact of AI R&D on valuation and capital investments.

Finally, I examine the extent to which AI R&D is linked to AI adoption across firms. I leverage data from the U.S. Census Bureau's Business Trends and Outlook Survey (BTOS), which includes firms' responses to the question: "During the next six months, do you think this business will be using AI in producing goods or services?" According to the BTOS, AI adoption rates remain relatively low, around 5-10% as of 2024Q4 on average. However, I find a positive

association between the industry-level growth in the AIR Index and AI adoption, particularly after the release of ChatGPT. In other words, industries experiencing a sharper rise in AI-related discussions tend to report higher adoption rates. For instance, Educational Services (NAICS 61), Information (NAICS 51), and Professional, Scientific, and Technical Services (NAICS 54) experienced a 45%, 35%, and 27% average growth in AI R&D discussions, respectively. These industries now report adoption rates between 10%-35%, suggesting that recent growth in AI R&D is translating into more robust adoption of AI tools to be used in production.

V. CONCLUSION

This paper introduces a novel measure, the AIR Index, to capture the extent of AI R&D by firms through the similarity of their earnings conference call discussions to academic AI research. As AI becomes increasingly integrated into the economy, understanding how firms are investing in AI-intensive capital will be crucial to ultimately determine the productivity effects of this emerging technology. The AIR Index provides insight into this process, showing that more substantive discussions of AI R&D are valued by investors, leading to immediate price reactions following the call and higher Tobin's Q the following year, independent of mere mentions of AI-related buzzwords. Importantly, substantive AI research discussions are linked to capital deepening, with sustained year-over-year capex growth lasting one year on average. However, these discussions are not immediately reflected in productivity or labor force changes, consistent with the notion that AI, like other transformative technologies, requires upfront capital investments before yielding tangible productivity gains.

REFERENCES

Acemoglu, Daron, and Pascual Restrepo. "Artificial intelligence, automation, and work." *The economics of artificial intelligence: An agenda*. University of Chicago Press, 2018. 197-236.

Autor, David H., Frank Levy, and Richard J. Murnane. "The skill content of recent technological change: An empirical exploration." *The Quarterly Journal of Economics* 118.4 (2003): 1279-1333.

Ash, Elliott, and Stephen Hansen. "Text algorithms in economics." *Annual Review of Economics* 15.1 (2023): 659-688.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. "Artificial intelligence, firm growth, and product innovation." *Journal of Financial Economics* 151 (2024): 103745.

Bai, Jushan, and Pierre Perron. "Estimating and testing linear models with multiple structural changes." *Econometrica* (1998): 47-78.

Baily, Martin, Erik Brynjolfsson, and Anton Korinek. "Machines of mind: The case for an AI-powered productivity boom." (2023).

Bick, Alexander, Adam Blandin, and David J. Deming. *The Rapid Adoption of Generative AI*. No. w32966. National Bureau of Economic Research, 2024.

Brynjolfsson, Erik, and Andrew Mcafee. "The business of artificial intelligence: What it can—And cannot—Do for your organization." *Harvard Business Review Digital Articles* 7 (2017): 3-11.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. "The productivity J-curve: How intangibles complement general purpose technologies." *American Economic Journal: Macroeconomics* 13.1 (2021): 333-372.

Brynjolfsson, Erik, and Kristina McElheran. "The rapid adoption of data-driven decision-making." *American Economic Review* 106.5 (2016): 133-139.

Byrne, David M., Carol Corrado, and Daniel E. Sichel. *The rise of cloud computing: minding your P's, Q's and K's*. No. w25188. National Bureau of Economic Research, 2018.

Byrne, David M., John G. Fernald, and Marshall B. Reinsdorf. "Does the United States have a productivity slowdown or a measurement problem?." *Brookings Papers on Economic Activity* 2016.1 (2016): 109-182.

Caldara, Dario, and Matteo Iacoviello. "Measuring geopolitical risk." *American Economic Review* 112.4 (2022): 1194-1225.

Cooper, Michael J., Orlin Dimitrov, and P. Raghavendra Rau. "A rose. com by any other name." *The Journal of Finance* 56.6 (2001): 2371-2388.

Corrado, Carol, Charles Hulten, and Daniel Sichel. "Measuring capital and technology: an expanded framework." *Measuring capital in the new economy*. University of Chicago Press, 2005. 11-46.

Davis, Steven J., Stephen Hansen, and Cristhian Seminario-Amez. *Firm-level risk exposures and stock returns in the wake of COVID-19*. No. w27867. National Bureau of Economic Research, 2020.

Eisfeldt, Andrea L., Gregor Schubert, and Miao Ben Zhang. *Generative AI and firm values*. No. w31222. National Bureau of Economic Research, 2023.

Eloundou, Tyna, Sam Manning, Pamela Mishkin and Daniel Rock. "Gpts are gpts: An early look at the labor market impact potential of large language models." *arXiv preprint arXiv:2303.10130* (2023).

Fama, Eugene F., and Kenneth R. French. "A five-factor asset pricing model." *Journal of Financial Economics* 116.1 (2015): 1-22.

Hassan, Tarek A., et al. "Firm-level exposure to epidemic diseases: Covid-19, SARS, and H1N1." *The Review of Financial Studies* 36.12 (2023): 4919-4964.

Hayashi, Fumio. "Tobin's marginal q and average q: A neoclassical interpretation." *Econometrica* (1982): 213-224.

Hou, Kewei, Chen Xue, and Lu Zhang. "Digesting anomalies: An investment approach." *The Review of Financial Studies* 28.3 (2015): 650-705.

Humlum, Anders, and Emilie Vestergaard. "The Adoption of ChatGPT." *University of Chicago, Becker Friedman Institute for Economics Working Paper* 2024-50 (2024).

Jordà, Òscar. "Estimation and inference of impulse responses by local projections." *American Economic Review* 95.1 (2005): 161-182.

Reimers, N. "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks." arXiv preprint arXiv:1908.10084 (2019).

Soto, Paul E. "Breaking the Word Bank: measurement and effects of bank level uncertainty." *Journal of Financial Services Research* 59.1 (2021): 1-45.

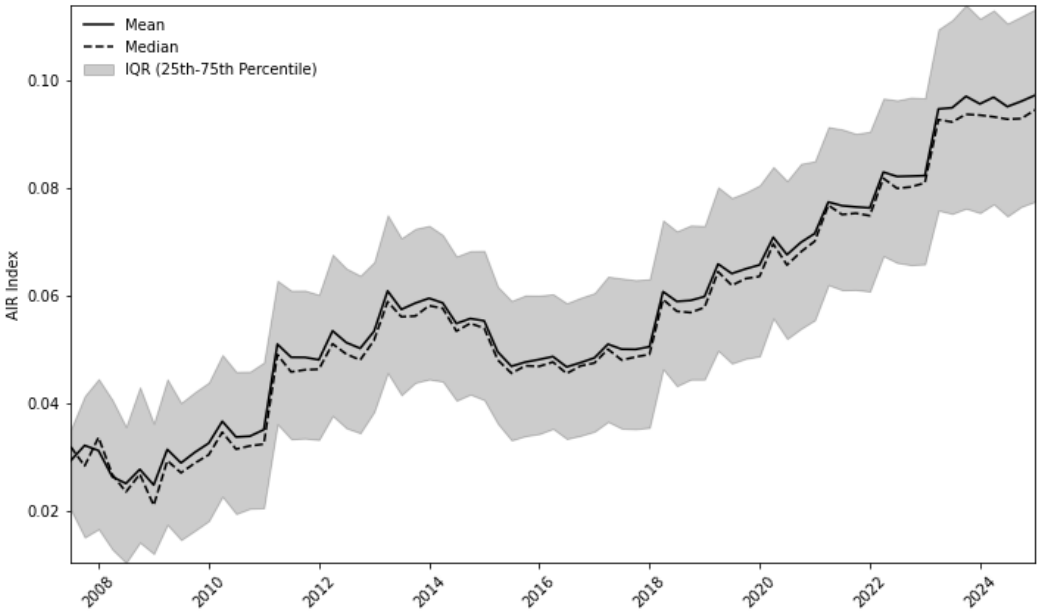
Stiroh, Kevin J. "Information technology and the US productivity revival: what do the industry data say?." *American Economic Review* 92.5 (2002): 1559-1576.

Figure 1: Highly Cited Computer Science Abstracts over Time



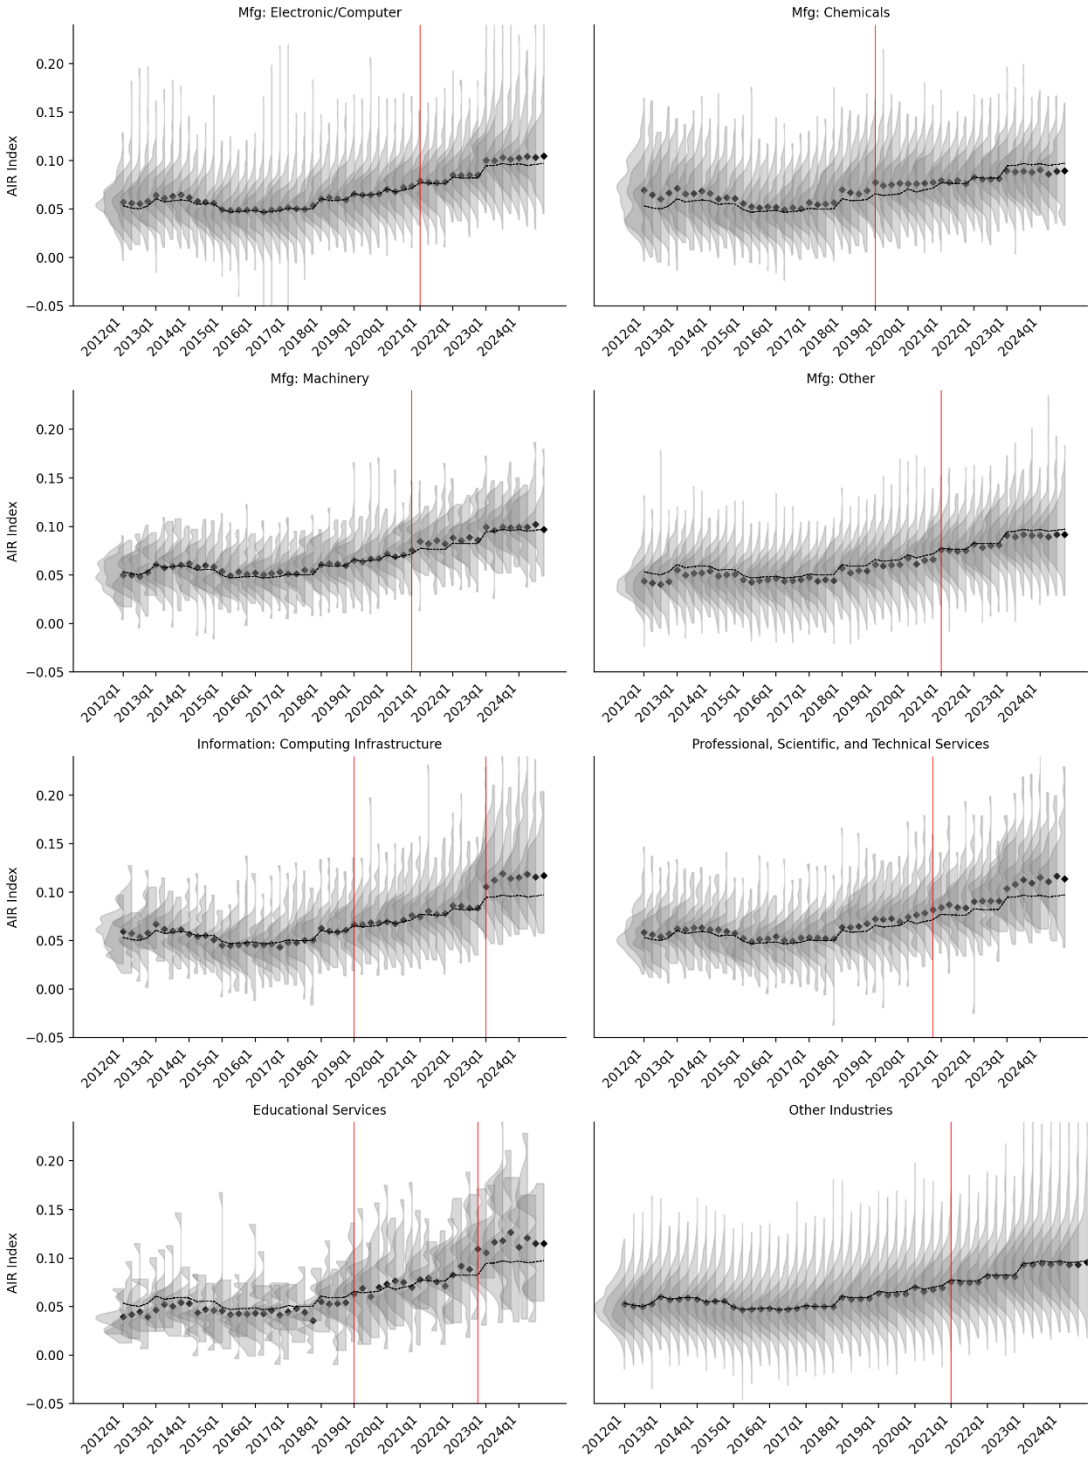
Note: This figure presents word clouds of the abstracts for leading academic research papers in the computer science field. Each row shows a word cloud of the top 100 papers released that year, measured via citation counts as of October 2024, with the text filtered for stopwords and commonly recurring scientific words, such as "technique", "system", and "results".

Figure 2: AIR Index- AI R&D Discussions over Time



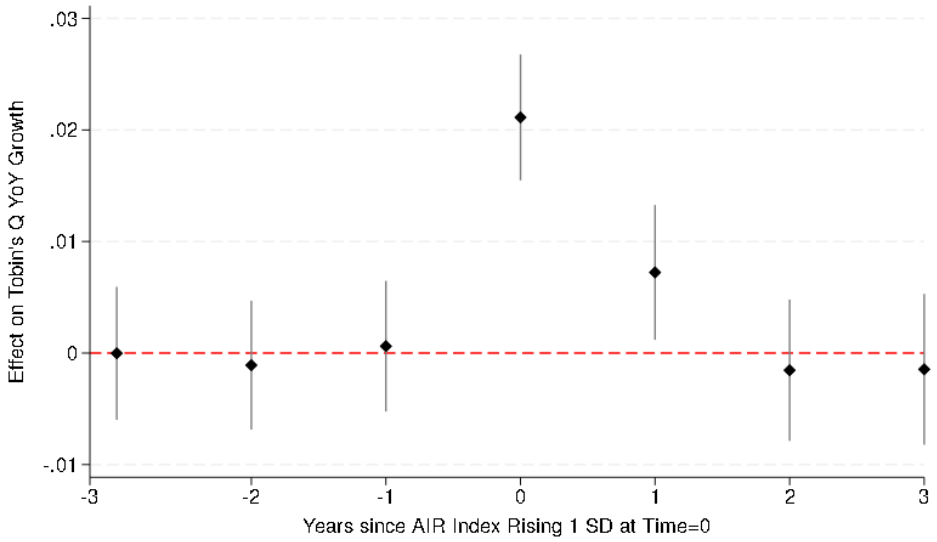
Note: This figure presents the time series of the new AI R&D measure developed in this paper. The measure is calculated by converting firm responses from quarterly earnings calls into sentence embeddings using a sentence transformer model and computing the cosine similarity between these embeddings and the average embedding of well-cited computer science academic papers. The measure is aggregated to the firm-quarter level by averaging cosine similarity scores across responses, and the figure shows the average (solid line), median (dotted line), and interquartile range (shaded region) across all firms.

Figure 3: AIR Index by Industry



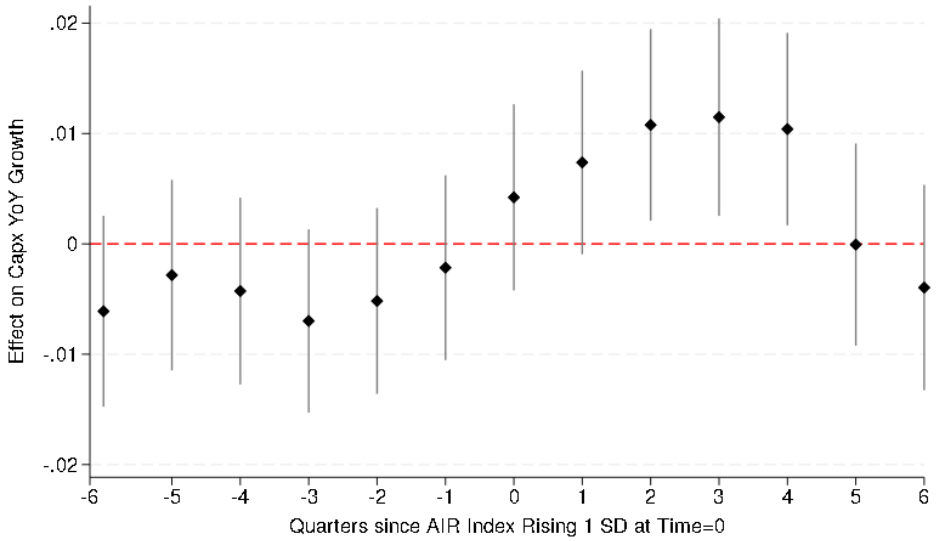
Note: This figure presents the AIR Index over time across key industries. Each subplot represents a distinct industry, with a dotted black line denoting the overall average across all industries, serving as a benchmark for comparison. Industries shown include Electronics/Computer Manufacturing (NAICS 335, 334), Chemicals Manufacturing (NAICS 325), Machinery Manufacturing (NAICS 333), Other Manufacturing (all other NAICS within 31-33), Computing Infrastructure (NAICS 518), Professional, Scientific, and Technical Services (NAICS 54), Educational Services (NAICS 61), and Other Industries (all remaining NAICS codes). Within each subplot, the quarterly distribution of firms' AIR Index scores is depicted vertically, with the mean for that industry in each quarter marked by a black diamond. Bai and Perron (1998) structural breaks are shown by a vertical red line.

Figure 4: Impact of AIR Index on Valuation



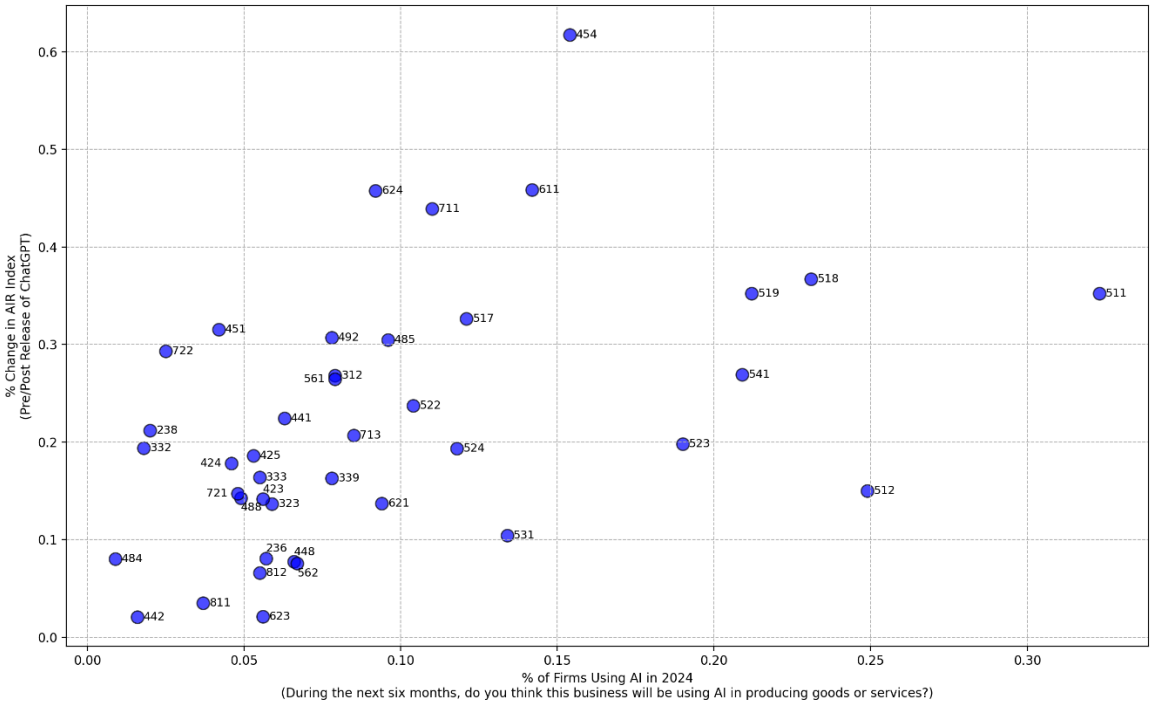
Note: This figure shows the local projections of the impact of AI R&D discussions (AIR Index) on year-over-year Tobin's Q growth. The horizontal axis represents the time horizon in years (H) and each point corresponds to the estimated coefficient from a regression of $TobinsQ_{t+H}$ on the AIR Index at time t, controlling for the log of total assets, AI buzzwords, firm-level fixed effects, and 4-digit NAICS industry-year fixed effects. Standard errors, clustered at the firm level, are shown with 90% confidence intervals.

Figure 5: AIR Index and Capital Deepening



Note: This figure shows the local projections of the impact of AI R&D discussions (AIR Index) on year-over-year capital expenditure growth. The horizontal axis represents the time horizon in quarters (H) and each point corresponds to the estimated coefficient from a regression of $Capex_{t+H}$ on the AIR Index at time t, controlling for the log of total assets, AI buzzwords, firm-level fixed effects, and 4-digit NAICS industry-year fixed effects. Standard errors, clustered at the firm level, are shown with 90% confidence intervals.

Figure 6: AIR Index and AI Adoption



Note: This figure shows a scatter plot of the relationship between industry-level AI adoption rates and growth in AI R&D discussions in conference calls. AI adoption rates are based on responses to the November 2024 BTOS question, "During the next six months, do you think this business will be using AI in producing goods or services?" Each data point represents a 3-digit NAICS industry, with the x-axis displaying the corresponding BTOS adoption response and the x-axis showing the average growth in the AIR Index within the industry from 2020Q4- 2022Q3 (pre ChatGPT release) to 2022Q4-2024Q3 (post ChatGPT release). The industries represented include 11 (Agriculture, Forestry, Fishing, and Hunting); 21 (Mining, Quarrying, and Oil and Gas Extraction); 22 (Utilities); 23 (Construction); 31-33 (Manufacturing); 42 (Wholesale Trade); 44-45 (Retail Trade); 48-49 (Transportation and Warehousing); 51 (Information); 52 (Finance and Insurance); 53 (Real Estate and Rental and Leasing); 54 (Professional, Scientific, and Technical Services); 56 (Administrative and Support and Waste Management and Remediation Services); 61 (Educational Services); 62 (Health Care and Social Assistance); 71 (Arts, Entertainment, and Recreation); 72 (Accommodation and Food Services); and 81 (Other Services, except Public Administration).

Table 1: Examples of Responses with High AIR Index Scores

Company	Date	Text	AIR
Cerence	2023Q3	With decades of extensive vertical expertise in the automotive industry and a rich history of leading AI innovation, Cerence is uniquely positioned to bring the latest advances in AI into the car. We bring unmatched experience and knowledge to the application of generative AI and large language models in transportation as well as a strategic methodical focus on creating groundbreaking user experiences. As we envision the future of in-car experiences, we are keenly focused on solving user problems by harnessing transformer-based foundational AI models. These models enable us to develop intuitive voice and multimodal user interfaces as well as generative AI applications that empower our customers to deliver high-value user experiences.	0.681
C3.ai	2023Q3	So you are commenting on the fact that these large language models tend to be almost exclusively limited to, okay, text, HTML and code. So other sorts of data, they don't know how to ingest Okay. Good. Good. Okay. Now we -- so let's talk about this. We are the masters of the universe, ingesting what you call multimodal data. images, okay, images from space, trajectories of hypersonics, high-speed telemetry, trading volume, the rate at which electrons are going across the grid, enterprise data, free text. And so we're using our standard architecture to ingest those data, okay? We're using one of our standard deep learning models to basically parse out this data and store all the relationships in a vector data store.	0.664
TaskUs	2023Q2	The work we're doing for generative AI companies is incredibly exciting. We spun up adversarial testing teams. These are teams of people who are provoking large language models and image generation models to produce offensive content or content that violates terms of service. We're then documenting those violations, so that generative AI engineering teams can further refine their models to protect their users. In addition to this, we're actually helping to train many of these models answering complex questions, recruiting experts in particular areas of subject matter expertise to help build the future of large language models. Today, we work for 2 of the 3 leading large language models.	0.634
Cineverse	2024Q3	These large language models, or LLMs, that are powering major AI company products require exceedingly large volumes of video to teach them everything about the world around us. These LLMs need to be trained on everything from how a horse runs through the woods to a flow of pedestrians crossing a busy street intersection or the movement and sounds of how an ocean wave crashes on a sandy beach and so on. The most effective way of doing so requires movie and television content, which, by its own nature, encompasses the full human experience in extremely high quality and consistency. By combining our vast independent film library, proprietary content distribution technology and extensive experience as a content aggregator, we find ourselves uniquely positioned to provide these leading AI developers with the most expansive and high-quality video training data sets available without the legal encumbrances hindering the major Hollywood studios.	0.612
Alphabet	2022Q3	This is our seventh year as an AI-first company, and we intuitively know how to incorporate AI into our products. Large language models make them even more helpful models like PaLM 2 and soon Gemini, which we are building to be multimodal. These advances provide an opportunity to reimagine many of our products, including our most important product, Search. We are in a period of incredible innovation for Search, which has continuously evolved over the years. This quarter saw our next major evolution with the launch of the Search Generative Experience, or SGE, which uses the power of generative AI to make Search even more natural and intuitive.	0.600
NVIDIA	2019Q3	In the area of training, the thing that's really exciting everybody, and everybody is racing towards, is training these large gigantic natural language understanding models, language models. The transformer model that was introduced by Google, called BERT, has since been enhanced into XLned and RoBERTa and, gosh, so many different, GP2, and Microsoft's MASS. And there's so many different versions of these language models. And in the AI, NLU, natural language understanding, is one of the most important areas that everybody's racing to go to. And so these models are really, really large. It's over 1,000x larger than image models that we're training just a few years ago, and they're just gigantic models. It's one of the reasons why we built the DGX SuperPOD so that we could train these gigantic models in a reasonable amount of time.	0.596
Sound Group	2023Q4	With respect to technological advancement, we continue to strengthen our R&D capabilities to provide more customized product support. This, in turn, empowers product innovation and drives progress across our global business. We have also consistently deepened the integration and development of multi-model AIGC technologies into our innovative business framework. Through motion-enhanced model training and algorithm optimization, we are fortifying our AI assets, unlocking vast opportunities for ongoing product innovation. On voice technologies, we have been constantly improving the stability and performance of automatic speech recognition and text to speech, or ASR and TTS, for both Chinese and English languages, adapting them for diverse scenarios.	0.589
Palantir Technologies	2023Q3	At AIPCon this past June, we introduced Palantir's AI platform, a core set of technologies designed to bring LLMs to your enterprise to supercharge and accelerate your experiences. From integrating data and hydrating your ontology to building AI-enabled applications and human agent teams with copilots. AIP enables you to deploy LLMs anchored in your data on your private network and to safely orchestrate your enterprise with tools, actions and other AI models. All of this in a controlled, governed and trusted AI operating system. The accelerating pace of AI developments continues to be all inspiring. The key to capturing value is a fundamental recognition that we are dealing with something new and different that demand solving new integration and engineering challenges.	0.578

Note: This table presents selected responses from earnings conference call transcripts that exhibit high AIR Index scores. The AIR Index is calculated by converting firm responses from quarterly earnings calls into sentence embeddings using a sentence transformer model and computing the cosine similarity between these embeddings and the average embedding of well-cited computer science academic papers.

Table 2: Summary Statistics

	Mean	SD	P25	P50	P75
Yearly Variables					
TobinsQ	2.20	2.05	1.15	1.58	2.46
TobinsQ YoY	-0.01	0.21	-0.08	0.00	0.08
Sales/Emp	920.83	9636.02	197.97	321.38	580.04
Sales/Emp YoY	0.03	0.37	-0.06	0.02	0.10
Employees	19.97	70.07	0.65	3.30	13.50
Employees YoY	0.04	0.27	-0.03	0.03	0.11
Log(Assets)	7.45	2.10	6.05	7.45	8.86
Quarterly Variables					
CapEx	2.80	2.03	1.14	2.68	4.17
CapEx YoY	0.07	0.95	-0.28	0.03	0.42
AI R&D	0.06	0.03	0.04	0.06	0.08
AI Buzzwords	0.00	0.04	0.00	0.00	0.00
$\beta_{mkt_premium}$	1.00	0.42	0.77	0.99	1.23
β_{SMB}	0.65	0.70	0.17	0.59	1.07
β_{HML}	-0.01	0.78	-0.39	-0.01	0.36
β_{RMW}	-0.18	1.01	-0.62	-0.03	0.40
β_{CMA}	-0.01	1.12	-0.51	0.02	0.52
$CAR_{f,0:1\ Day}$	0.01	9.27	-4.00	0.00	4.07
$CAR_{f,0:3\ Days}$	-0.08	10.26	-4.56	0.00	4.40
$CAR_{f,0:30\ Days}$	-0.48	17.52	-8.00	0.00	6.77
$CAR_{f,0:60\ Days}$	-0.71	23.95	-10.86	0.00	8.84
$CAR_{f,0:90\ Days}$	-1.09	30.94	-14.14	0.00	11.48
SUE	1.56	4.05	-0.31	1.01	2.89

Note: This table presents summary statistics for the primary variables used in the analysis, categorized by their measurement frequency. Yearly variables include Tobin's Q, labor productivity (measured as sales per employee), total employees, and the logarithm of assets. Quarterly variables consist of capital expenditure, the AIR Index, AI Buzzwords, betas from a Fama-French 5-Factor model estimated over the year prior to the call, the cumulative abnormal returns over various horizons using the 5-factor model, and the standardized unexpected earnings (SUE). The AIR Index is estimated by converting conference call responses into sentence embeddings, which are then compared to embeddings of well-cited computer science research papers. AI Buzzwords captures the frequency of "artificial intelligence" and "machine learning" mentions in conference calls. Year-over-year (YoY) growth rates for each variable are also shown, calculated as the logarithmic difference from the previous year.

Table 3: Impact of AIR Index on Tobins Q

	(1)	(2)	(3)	(4)
Dependent Variable	$TobinsQ_{f,y}$			
$AIR_Index_{f,y}$	0.0278***	0.0213***		0.0211***
	(0.00186)	(0.00343)		(0.00343)
$AI\ Buzzwords_{f,y}$			0.00188	0.000936
			(0.00173)	(0.00184)
N	33381	32465	32465	32465
R-Squared	0.176	0.322	0.320	0.322
Year*4-Digit Industry FE	N	Y	Y	Y
Firm FE	Y	Y	Y	Y

Note: This table shows estimates of regressions of year-over-year Tobin's Q growth on AI R&D discussions (AIR Index) at the yearly level. Standard errors are clustered at the firm level. ***, **, * denote significance at the 99%, 95%, and 90% confidence levels, respectively.

Table 4: Impact of AIR Index on Cumulative Abnormal Returns

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent Variable	$CAR_{f,\ 0:1\ days}$			$CAR_{f,\ 0:3\ days}$	$CAR_{f,\ 0:30\ days}$	$CAR_{f,\ 0:60\ days}$	$CAR_{f,\ 0:90\ days}$
$AIR_Index_{f,yq}$	0.235***	0.920***	0.836***	0.910***	0.923***	0.513*	0.234
	(0.0526)	(0.109)	(0.104)	(0.117)	(0.198)	(0.270)	(0.323)
$AI\ Buzzwords_{f,yq}$	0.0563	0.0166	0.0347	0.0362	0.149	0.148	-0.0534
	(0.0509)	(0.0801)	(0.0768)	(0.0787)	(0.153)	(0.185)	(0.236)
$SUE_{f,yq}$			3.145***	3.164***	3.163***	2.616***	1.552***
			(0.0704)	(0.0773)	(0.110)	(0.144)	(0.172)
N	62909	62909	62909	62909	62909	62909	62909
R-squared	0.003	0.309	0.364	0.355	0.305	0.313	0.385
Firm*Year FE	N	Y	Y	Y	Y	Y	Y
Quarter FE	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y

Note: This table shows estimates of regressions of cumulative abnormal returns (CAR), denoted as $CAR_{f,\ 0:D\ days}$, for D equal 1 day (columns 1-3), 3 days (column 4), 30 days (column 5), 60 days (column 6), and 90 days (column 7), on R&D discussions (AIR Index) at the firm-quarter level. $CAR_{f,\ 0:D\ days}$ is measured as the cumulative excess returns over D days, as forecasted by a 5-factor Fama-French model estimated using data over the previous year of the conference call. Columns 3 through 7 include a control for earnings surprise (SUE, standardized unexpected earnings). Standard errors are clustered at the firm level. ***, **, * denote significance at the 99%, 95%, and 90% confidence levels, respectively.

Table 5: Impact of AIR Index on Labor Productivity and Employment

Dependent Variable	(1) $LaborProd_{f,y}$	(2) $LaborProd_{f,y+1}$	(3) $LaborProd_{f,y+2}$	(4) $LaborProd_{f,y+3}$	(5) $Emp_{f,y}$	(6) $Emp_{f,y+1}$	(7) $Emp_{f,y+2}$	(8) $Emp_{f,y+3}$
$AIR_Index_{f,y}$	-0.00301	-0.00318	-0.00985	0.00170	-0.000810	0.00170	0.00422	0.00129
	(0.00591)	(0.00667)	(0.00736)	(0.00788)	(0.00413)	(0.00436)	(0.00460)	(0.00511)
AI $Buzzwords_{f,y}$	-0.00208	0.000533	0.00241	0.000270	0.00171	0.00157	-0.000576	0.000143
	(0.00168)	(0.00140)	(0.00222)	(0.00201)	(0.00158)	(0.00126)	(0.00174)	(0.00210)
N	31344	28793	25165	21753	31750	29132	25420	21934
R-squared	0.261	0.237	0.234	0.254	0.370	0.352	0.363	0.365
Year*4-Digit Industry FE	Y	Y	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y

Note: This table shows estimates of regressions of year-over-year labor productivity and employment growth on AI R&D discussions (AIR Index) at the yearly level. Standard errors are clustered at the firm level. ***, **, * denote significance at the 99%, 95%, and 90% confidence levels, respectively.

Table 6: AIR Index prior to the GenAI Era

	(1)	(2)	(3)	(4)
	2004-2015		2004-2020	
Dependent Variable	$TobinsQ_{f,y}$	$Capx_{f,yq+1}$	$TobinsQ_{f,y}$	$Capx_{f,yq+1}$
$AIR_Index_{f,t}$	0.0250***	0.0146*	0.0184***	0.0121**
	(0.00678)	(0.00826)	(0.00421)	(0.00588)
AI $Buzzwords_{f,t}$	0.0260*	-0.0202	-0.00156	-0.000163
	(0.0145)	(0.0129)	(0.00236)	(0.00422)
N	13299	49334	25828	94668
R-Squared	0.379	0.238	0.302	0.258
Year*4-Digit Industry FE	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y

Note: This table shows estimates of regressions of year-over-year Tobin's Q and Capx growth on AI R&D discussions (AIR Index). Columns 1 and 3 (2 and 4) use yearly (quarterly) data. The sample is restricted to 2004-2015 (2004-2020) in columns 1 and 2 (3 and 4). Standard errors are clustered at the firm level. ***, **, * denote significance at the 99%, 95%, and 90% confidence levels, respectively.