Artificial Intelligence, Opportunity, and Regulatory Uncertainty: Implications for Asset Pricing

Kris Shen*

April 10th, 2025

ABSTRACT

This paper studies two channels—opportunity and regulatory uncertainty—through which Artificial Intelligence (AI) affects the stock prices and risk premia. On the one hand, advances in AI present firms with opportunities, leading them to exhibit characteristics of growth firms and earn lower expected returns. On the other hand, firms face increased regulatory uncertainty in AI development, increasing their political risk exposure and resulting in higher expected returns. Using conference call transcripts, I construct a firm-level measure of *AI Exposure* that captures the level of attention analysts and managers devote to AI-related topics at specific points in time. Empirically and theoretically, I show that these two channels exert opposing effects: firms focused on opportunity earn a negative AI risk premium, while those more affected by regulatory uncertainty earn a positive AI risk premium.

JEL Classification: G10; G11; G12; G18; O33

Keywords: Artificial intelligence; Asset pricing; Risk premium; Opportunity; Regulation

---

*Fisher College of Business at The Ohio State University. I thank Andrei Gonçalves, Sergey Sarkisyan, René Stulz, Daisy Wang, Yufeng Wu, Cynthia Yin, and Shaojun Zhang for their helpful discussions and comments. I am also grateful to the OSU Finance Student Seminar participants for their valuable feedback.

I. Introduction

Artificial Intelligence (AI) has emerged as a transformative force, driving both unprecedented opportunities and regulatory uncertainties for firms across industries. By 2030, AI is projected to contribute up to $15.7 trillion to the global economy, with 40% of productivity gains stemming from AI-driven automation and innovation (PwC, 2020). Firms utilizing or investing AI technologies for productivity, product innovation, and market expansion have seen revenue increases of up to 20% in some sectors (McKinsey, 2023). However, the rapid adoption of AI also exposes firms to regulatory uncertainty, as governments worldwide try to come up with frameworks to address ethical concerns, data privacy issues, and algorithmic accountability worries on AI. For example, the European Union's AI Act, proposed in 2021, aims to classify AI systems by risk levels, potentially imposing stringent compliance costs on firms. Meanwhile, 47% of executives cite regulatory uncertainty as a top barrier to AI adoption (Deloitte, 2024), highlighting the tension between the potential opportunities AI brings and the regulation risks AI poses. Consistent with these estimates and surveys, examples from 10-K annual filings reveal that firms explicitly discuss both the opportunities and regulatory uncertainties surrounding AI (see Appendix B for details). This duality, where AI simultaneously creates opportunities for growth and increasing productivity while introducing regulatory uncertainties, has profound implications for the asset prices. This paper studies how these two channels affect stock prices and risk premia.

AI, which Deming, Ong, and Summers (2025) argue to be the next General-Purpose Technologies (GPTs) like electricity and steam power, will bring firms a great amount of opportunities that tend to be widespread and long-spreading. This differentiates AI from other recent technology breakthroughs documented in the literature. For example, in the labor market, studies show that generative AI models can increase workers' productivity in writing, customer service, and programming tasks (Brynjolfsson, Li, and Raymond, 2023; Peng et al., 2023; Noy and Zhang, 2023). Firms investing or taking advantage of these tools will increase their productivity. However, firms starting to invest or adapt to using AI in their operations might see negative profitability in the beginning due to long-term R&D investments or payment for third-party AI use like ChatGPT4. These firms are primarily concerned with rising opportunities, balancing entry costs against the potential for future gains as they explore AI investment or adoption. I call these firms *AI Starters*. There are also other firms that have invested in AI or used AI technologies in their daily operations for quite some time and they see big profitability growth when there are positive AI opportunity shocks like the launch of ChatGPT3. However, their advanced integration of AI exposes them disproportionately to regulatory risks, as policymakers increasingly target established

AI creators or adopters with compliance requirements, ethical scrutiny, or restrictions on market dominance. I call these firms *AI Elites*.

This paper hence focuses on testing the null hypothesis that these two forces—opportunity and regulatory uncertainty—exert opposing effects on stock prices and risk premia for *AI Starters* and *AI Elites*. Specifically, I hypothesize that for *AI Starters*, the opportunity channel dominates, as markets price their potential for future gains despite near-term costs, while for *AI Elites*, the regulation channel dominates, as their exposure to AI amplifies risks from policy changes. These two channels lead to contrasting asset pricing implications for the two groups.

In the first part of this paper, I document the empirical results regarding asset prices and risk premia associated with AI. The initial step is to construct a firm-level *AI Exposure* measure to proxy for firms' exposure to AI advances. I begin by creating an AI Vocabulary (AIV) through textual analysis of 159,444 academic paper titles from the Web of Science, specifically related to *artificial intelligence*. These papers span from 2010 to 2024 and are accessible via ProQuest. I extract unigrams and bigrams, compute their TF-IDF scores, and compile the AIV with the most frequent unigrams and bigrams. I utilize both unigrams and bigrams to balance the tradeoff between using purely unigrams, which can introduce excessive noise, and purely bigrams, which may overly restrict the scope of what can be captured in a document. Figure 1 displays the AIV WordCloud, where a larger font size indicates higher TF-IDF scores. The WordCloud reveals that the main terms align well with those identified by Babina et al. (2024a), who analyzed AI topics using resumes and job posting data. After establishing an AIV, I then construct an *AI News Index* through textual analysis of the RavenPack newspaper database. I analyze approximately 170 million news headlines from 2010 to 2024, counting each month how many headlines contain any of the AIV bigrams and defining these as AI-related news. Figure 2 depicts the *RavenPack AI News Index*, showing the number of AI-related news on a monthly basis. The figure illustrates that the *RavenPack AI News Index* captures some of the main AI-related events since 2010, with the ChatGPT launch at the end of 2023 marking the most significant rise in AI-related news. Prior to 2017, the development of AI was focused on the implications and advancements in machine learning and deep learning.

With the AIV in place, I construct a firm-level *AI Exposure* measure derived from firms' conference call transcripts. Following Sautner et al. (2023), I conduct textual analysis on these transcripts and compute the fraction of AI-related unigrams and bigrams, presented in the previously defined AIV, relative to the total number of unigrams and bigrams in any given transcripts. This method provides a firm-level *AI Exposure* measure that captures the attention that financial analysts and managers pay to AI topics at specific times. To further study what this measure mainly captures, I decompose this measure into different versions that speak to whether this measure is more about opportunity, regulation, or risk. Furthermore, I create two versions of this measure that speak to the tone on whether it is more positive or negative.

The second step involves conducting a long set of empirical asset pricing tests to see how the two channels, opportunity and regulatory uncertainty, affect stock prices and risk premia. To evaluate the opposing effects of these two channels—where *AI Starters* are more opportunity-driven and *AI Elites* are more concerned about regulatory risks—I conduct several complementary empirical tests. First, typical empirical asset pricing tests including portfolio sorts on *AI Exposure* measures and factor regressions reveal stark contrasts in risk premia: value-weighted portfolios (dominated by *AI Elites*) exhibit statistically significant positive return spreads for high AI Exposure firms, consistent with regulatory risks driving higher expected returns. Conversely, equal-weighted portfolios (tilted toward *AI Starters*) show negative spreads, aligning with the opportunity channel's dominance and lower expected returns for *AI Starters*. These results hold robustly across different factor models, with significant unexplained alphas for *AI Elites*. More specifically, I divide firms into portfolios based on their firm-level *AI Exposure* and run time-series regressions of the monthly return spread between the 1st decile and the 10th decile on different sets of factors. These factor models include the CAPM model, the ICAPM model (Chabi-Yo, Gonçalves, and Loudis, 2025), the Fama and French three-factor (FF3) model (Fama and French, 1996), the Fama and French three-factor model plus the momentum factor (FF4) by Carhart (1997), the Fama French five-factor (FF5) models (Fama and French, 2015), and the Hou, Xue, and Zhang q and q5 factor models (Hou, Xue, and Zhang, 2015; Hou et al., 2020). I find there is a significantly positive α in all factor models, with the maximum monthly $\alpha = 99.1$ basis points (t=3.85) using q factors and the minimum monthly $\alpha = 66.6$ basis points (t=3.24) using FF4. These leading factor models cannot explain this risk premium associated with *AI Exposure* in the full sample from 2009 to 2024. In subsequent subsample tests using recent years, the factor alphas increase in both quantitative and statistical magnitudes. Furthermore, consistent with Babina et al. (2024a), I do find that growth options help explain the risk premium associated with AI, which validates the fact that opportunity channel does exist among *AI Elites* but does not dominate the price impact.

Second, Fama-MacBeth regressions further validate the results from portfolio sorts and factor regressions. I run Fama-MacBeth regressions in Ordinary Least Squares (OLS) and Weighted Least Squares (WLS) with market value as the weights. The overall AI Exposure-return relationship is negative in both full-sample and subsample (more recent years) OLS regressions. The coefficients on the overall *AI Exposure* are -0.0775 (t-statistics: -2.19) and

-0.0867 (t-statistics: -1.52), respectively. However, the topic-specific AI exposure measure, $AIExposure^{Opp}$ for AI opportunity-focused firms, reveal significantly negative coefficients, with a coefficient on full-sample and subsample of -0.104 (t-statistics: -4.41) and -0.109 (t-statistics: -2.47), respectively. However, the overall AI Exposure-return relationship is positive in both full-sample and subsample (more recent years) WLS regressions, which are in line with the fact that *AI Starters* dominate the OLS regression results, while *AI Elites* dominate the WLS regression results. The sign change in the overall AI Exposure-return relationship from OLS to WLS validates the opposing effects of opportunity and regulatory uncertainty on expected returns.

Third, to mitigate the concern that ex post realized returns might be a bad proxy for ex ante expected returns, I follow Hou, van Dijk, and Zhang (2012) in constructing a few implied cost of capital (ICC) measures and rerun the Fama-MacBeth regressions. These ICC measures include Gebhardt, Lee, and Swaminathan (2001), Claus and Thomas (2001), Ohlson and Juettner-Nauroth (2005), Easton and Monahan (2005), and a composite measure which is the average of the previous four measures. I denote them as GLS, CT, OJ, MPEG, and Composite, respectively. I run Fama-MacBeth regressions using the two main measures: the overall $AIExposure$ measure and opportunity specific measure, $AIExposure^{Opp}$. There is a mix of results for the overall $AIExposure$ measure. GLS and OJ show negative while statistically insignificant coefficients on *AI Exposure*, while OJ, MPEG, and Composite show positive and statistically significant coefficients. However, all coefficients, except for the MPEG measure, are negative and statistically significant for $AIExposure^{Opp}$. Since $AIExposure^{Opp}$ should mainly capture the opportunity channel, the consistent negative relationship between ICC measures and $AIExposure^{Opp}$ validates the null hypothesis that the opportunity channel dominates *AI Starters* and there should be a negative relationship between $AIExposure^{Opp}$ and expected returns.

Fourth, in addition to using ICC measures as an alternative proxy for ex ante expected returns, I also apply an option-implied expected return measure following Martin and Wagner (2019). One limitation of the option-implied measure is that it only applies to S&P 500 firms, which in my sample tend to be *AI Elites*. Hence, I would expect to see a significant positive expected return spread between a value-weighted decile portfolio with the highest *AI Exposure* stocks and a value-weighted decile portfolio with the lowest *AI Exposure* stocks. That is exactly what I see in the real data. The return spreads, H-L, are statistically positive. The magnitude of the monthly return spread ranges from 0.45% to 0.73%, which is relatively smaller in magnitude but comparable to the VW return spread using realized returns, 0.96%.

Fifth, I re-estimate the risk premia associated with *AI Exposure* using a three-pass procedure following Giglio and Xiu (2021), to take into account the omitted variable bias and measurement error. The three-pass approach is, in essence, a combination of Principal Component Analysis (PCA) with two-pass cross-sectional regressions, to generate a consistent estimate of risk premia. Assume that returns follow a linear factor model with p factors, and my goal is to estimate the risk premium of one of them, g_t, which is the AI risk factor associated with *AI Exposure*. The three-pass approach consists of the following steps. First, I use PCA to recover factors and their loadings using a total number of 3,127 equity portfolios. Second, I run cross-sectional regressions using only the principal components, without the factor of interest, g_t, to estimate their risk premia. Third, I run time-series regressions of g_t onto the principal components to estimate the loadings on the principal components. The risk premium of g_t is calculated as the product of the loadings estimated in step two and their risk premia estimated in step three. I showcase that although VW AI factors tend to be more positive than EW AI factors, which is consistent with my hypothesis that *AI Starters* tend to earn a negative risk premium while *AI Elites* tend to earn a positive risk premium, most of the three-pass risk premia tend to be statistically insignificant. However, Giglio and Xiu (2021) argue that the three-pass procedure performs better as $T \to \infty$ (their sample spans 1080 months), whereas my sample covers only 168 months; the limited time-series length may diminish statistical power and lead to most insignificance.

Finally, I conduct several event studies that utilize plausibly exogenous shocks to disentangle these two channels. I include several plausibly exogenous shocks: (1) the launch of ChatGPT (November 30, 2022) as an opportunity shock, (2) Biden's AI Executive Order (October 30, 2023) as an AI-specific regulation shock, and (3) Trump's and Biden's election wins (2016 and 2020) as general regulation shocks. For each event, I compute cumulative abnormal returns (CARs) over a 10-day post-event window (CAR[1,10]), adjusting for market movements using CAPM betas estimated from a pre-event window. Portfolios are sorted into quintiles by *AI Exposure* and weighted equally. The results are striking: after Biden's AI Executive Order, high-exposure firms underperform low-exposure firms by -1.84% (t = -2.49), but this gap diminishes when excluding large firms (*AI Elites*). Conversely, ChatGPT's launch generates a +1.49% spread (t = 2.72) favoring high *AI Exposure* firms, which dimishes when excluding small firms (*AI Starters*). General political shocks (Trump's and Biden's election wins) show no significant differential effects. These results validate that the regulatory channel predominantly affects *AI Elites*, while the opportunity channel primarily influences *AI Starters*.

After documenting the empirical findings, I construct a competitive equilibrium model to show how the two channels, opportunity and regulation uncertainty, affect the stock prices and risk premia theoretically, building on Pástor and Veronesi (2012, 2013) and Hsu, Li, and Tsou (2023). The main model incorporates two types of regime changes that directly impact firms' profitability: regulation regime change and opportunity regime change. The regulation regime change occurs between favorable and unfavorable policies, where an unfavorable shift negatively affects firms' profitability. A favorable policy indicates a pro-AI government stance, including subsidies for AI research, tax incentives, or relaxed data privacy laws. In contrast, an unfavorable policy imposes stricter regulations and legal burdens, such as Biden's 2023 AI Executive Order emphasizing "*safe, secure, and trustworthy development and use of artificial intelligence*," which was widely viewed as restrictive to AI development. A regulatory shift occurred when Trump revoked this order in early 2025, with the White House stating that it had imposed "*unnecessarily burdensome requirements*" that could "*stifle private sector innovation and threaten American technological leadership.*" The model does not quantify how favorable or unfavorable different administrations are but highlights how these shifts influence stock prices. The opportunity regime change reflects AI's aggregate development stage, transitioning from early to mature. In the early stage, AI development involves high entry costs due to long-run R&D expenses and frictions such as the unavailability of large language models like ChatGPT. In contrast, the mature stage sees lower entry costs, aided by available open-source codes. For instance, in early 2025, the Chinese tech company DeepSeek launched a large AI model, building on existing LLMs and reducing training costs to under $6 million—whereas similar developments in the early 2010s could have cost up to $1 billion. This model illustrates how these regime changes influence firms' expected profitability and stock prices over time. In the model, the government makes decisions on regulation regime change with an endogenous political cost, while I assume the opportunity regime change is exogenously given. Households maximize a CRRA utility and liquidate firms' value at terminal date. All agents observe signals—such as political news—and learn about the political cost following a Bayesian Learning process.

The main intuition of the model is that for *AI Elites*, which are more concerned with regulatory uncertainty, firms with high *AI Exposure* are more sensitive to unfavorable regulatory shifts and face greater exposure to regulatory regime change risk. Consequently, these firms experience larger stock price declines following a materialized regulatory shock (e.g., Biden's AI Executive Order), implying that *AI Elites* should earn higher ex-ante expected returns. In contrast, for *AI Starters*, which are more focused on opportunity, firms with high *AI Exposure* are more sensitive to AI's transition from early-stage to mature-stage development. As a result, they see larger stock price gains when an opportunity shock materializes (e.g., the launch of ChatGPT), suggesting that *AI Starters* should earn lower ex-ante expected returns.

The last part of this paper provides validation of the main measures used in the empirical analysis, namely, *AIV* and firm-level *AI Exposure* measure. Following the mimicking portfolio approach of Engle et al. (2020), I show that using my *AIV*, investors can create a time series of *AI News Index* to capture the innovation in AI development. In addition, investors can use my firm-level *AI Exposure* measures to construct a mimicking portfolio to hedge the innovation in the *AI News Index*.

The contributions of this paper are twofold. First, it is the first study to examine how the two channels, opportunity and regulatory uncertainty, simultaneously affect stock prices and risk premia for firms with AI exposure. The closest study to mine is Babina et al. (2024a). They focus more on the systematic risk measured by CAPM market risk and find that most AI-invested firms bear systematic risk resulting from growth options. However, my measure differs from their labor-channel measure and captures more about "soft information" from conference call transcripts. My measure aims to capture investors' and managers' views on AI topics through conference call transcripts. Conference calls represent an important corporate event where managers and analysts have the opportunity to sit together to review the firm's financial results for a given period and discuss their initiatives and perspectives on various corporate topics. I have developed different versions of this measure to capture overall AI exposure and topic-specific aspects such as opportunities and regulations. The literature has shown that using conference calls is beneficial in capturing broader channels or perspectives of investors' perception of firm risk (Sautner et al., 2023). My focus is also on studying whether and how the opportunity and regulation uncertainty affect the AI risk premia. Through a series of empirical tests, including Fama-MacBeth and portfolio regressions, I demonstrate that growth options cannot fully account for the AI risk premium.

Second, this paper adds to the existing body of research on the impact of AI on firm outcomes. Using the launch of Google's TensorFlow as an exogenous shock, Rock (2019) demonstrates that increases in AI investment lead to significant gains in market valuation. Acemoglu et al. (2022) explores the effects of AI exposure on labor demand using data from Burning Glass job postings. Both Cockburn, Henderson, and Stern (2018) and Babina et al. (2024b) find that AI investments enhance firm-level product innovation. The studies most similar to mine are by Eisfeldt, Schubert, and Zhang (2023) and Babina et al. (2024a). Eisfeldt, Schubert, and Zhang (2023) investigates the impact of Generative AI—specifically the public release of ChatGPT—on equity returns at the firm level. Babina et al. (2024a) shows that firms increasing their AI investments, particularly through the growth of AI-related workers, experience an increase in their systematic risk as measured by the CAPM. Unlike these studies which focus on the labor channel, this paper uses a measure based on the general perceptions of investors and managers as expressed in conference call transcripts, capturing not only the labor channel but potentially the physical capital and business idea channels, which are crucial for evaluating asset prices in existing literature.

The rest of the paper is organized as follows. Section II describes the data, and how I construct the main empirical measures. In Section III, I present the empirical results. Section IV shows the general equilibrium model. In Section V, I validate the main measures used in this paper. In Section VI, I conclude the paper.

II. Data and Measures

This section describes the data and measures. In Section II.A, I first construct an *Artificial Intelligence Vocabulary (AIV)* using AI-related academic papers from Web of Science through ProQuest. This *AIV* allows me to capture the most frequent AI-related terms that will be used to conduct textual analysis in firms' conference call transcripts. In Section II.B, I construct the firm-level *AI Exposure* measure by studying how frequent *AIV* terms are discussed in firms' quarterly call transcripts. This measure is to capture how much attention investors and managers have devoted to AI-related topics.

A. Artificial Intelligence Vocabulary (AIV)

To construct an *Artificial Intelligence Vocabulary (AIV)*, I study a corpus of Web of Science texts on the subject of AI. More specifically, I collect a total number of 159,444 academic paper titles from Web of Science via ProQuest. Following the literature (see, e.g., Engle et al., 2020), I extract the unigrams and bigrams that are directly related to *artificial intelligence*, and compute their TF-IDF scores. Figure 1 presents word clouds that summarize the terms extracted from academic paper titles in the Web of Science. The larger the term size, the more frequent the term appears in the corpus (measured in TF-IDF score). The years range from 2010 to 2024. Panel (a) shows most frequent bigrams. Panel (b) shows most frequent unigrams. Panel (a) shows that most terms are largely related to machine learning and deep learning topics, which is consistent with the fact that these topics are at the core of AI. Other than those, I also see a lot of data-related terms like data mining and big data. Panel (b) is largely consistent with Panel (a) but extracts the unigrams that are also focused on the similar topics in machine learning and deep learning. Overall, these terms, including both unigrams and bigrams, largely capture the main intuitive terms that we could think of in the topics of AI. Here, I am not claiming these are the best keywords that are representative to AI. However, these represents well the terms that are most studied in academic research for this given period from 2010 to 2024.

For robustness, I construct the *AIV* using a larger and more general sample of 744,044 Web of Science academic articles or books that are directly related to: *artificial intelligence, machine learning, natural language processing, and computer vision*. These are the terms used in Babina et al. (2024a) when they study the systematic risk of AI through the labor channel. Appendix Figure D.1 show the word cloud of bigrams. The top bigrams are consistent with the Figure 1 Panel (a) in capturing mostly machine learning and deep learning topics.

I select the top 50 unigrams and top 50 bigrams as my final *AIV* in Figure 1 for simplicity. Specifying a different number, such as 25 or 100, would alter the quantitative outcomes, though the main qualitative results would remain consistent. Determining the optimal number of unigrams and bigrams to include in the final AIV is beyond the scope of this paper. Appendix Table D.2 and Appendix Table D.3 list the top 50 bigrams and top 50 unigrams by their TF-IDF scores.

B. *Firm-level AI Exposure*

B.1. AI Exposure Measure

After having an AIV, I follow Sautner et al. (2023) in constructing a firm-level *AI Exposure* measure using firms' call transcripts. As argued in the literature, exposure measures studying call transcripts well capture the attention devoted to a topic by the call participants, such as executives, analysts, investors and media. I use around 53 million quarterly call transcripts from 2009 to 2024 via Refinitiv. The overall firm-level *AI Exposure* measure is defined as the fraction of unigrams and bigrams from AIV over all unigrams and bigrams of a transcript,

$$AIExposure_{i,t} = \frac{1}{U_{i,t} + B_{i,t}} \left(\sum_{u=1}^{U_{i,t}} \mathbb{1}[u \in AIV] + \sum_{b=1}^{B_{i,t}} \mathbb{1}[b \in AIV] \right), \tag{1}$$

where $u = 0, 1, ..., U_{i,t}$ and $b = 0, 1, ..., B_{i,t}$ are unigrams and bigrams in the call transcripts of firm i in quarter t, and $\mathbb{1}[\cdot]$ is an indicator function. The annual measure is the average value aggregated from quarter measures.

To study the specific context of each sentence containing AI-related terms, I include different versions of the *AI Exposure* measures based on different topics. These topics include opportunity, regulation, and risk. I also study the tone of the transcript and define two versions of the *AI Exposure* measures as either positive or negative (Loughran and McDonald, 2011). The definition for these topic-specific *AI Exposure* measures, for example regarding opportunity, is defined as the fraction of unigrams and bigrams from AIV over all unigrams and bigrams of a transcript, but conditional on opportunity topic bigrams existing in the same sentence that contains AIV terms,

$$AIExposure^{Opp}_{i,t} = \frac{1}{U_{i,t} + B_{i,t}} \left(\sum_{u=1}^{U_{i,t}} (\mathbb{1}[u \in AIV] * \mathbb{1}[u, o \in S]) + \sum_{b=1}^{B_{i,t}} (\mathbb{1}[b \in AIV] * \mathbb{1}[b, o \in S]) \right), \tag{2}$$

where $\mathbb{1}[u/b, o \in S]$ denotes that a sentence with AI-related bigrams or unigrams also include opportunity-related bigrams or unigrams in the same sentence. Other topic-specific *AI Exposure* measures are defined in a similar way. The topic bigrams for opportunity, regulation, and risk are from Sautner et al. (2023) and extended by ChatGPT4. The positive and negative bigrams are from Loughran and McDonald (2011) and extended by ChatGPT4. Throughout the texts, I denote the overall *AI Exposure* measure, the *AI Exposure* regarding opportunity, the *AI Exposure* regarding regulation, the *AI Exposure* regarding risk, the *AI Exposure* with positive tones, and the *AI Exposure* with negative tones as $AIExposure$, $AIExposure^{Opp}$, $AIExposure^{Reg}$, $AIExposure^{Risk}$, $AIExposure^{Pos}$, and $AIExposure^{Neg}$, respectively.

Table I presents the summary statistics of different versions of *AI Exposure* and their correlations at the firm-quarter level. For the ease of exposition, I multiply by 1,000 for each measure. Panel A shows that the mean of *AI Exposure* is 4.4 while $AIExposure^{Opp}$ has the mean of 2.44 which is more than half the size of the mean of *AI Exposure*. Hence, the overall *AI Exposure* mostly captures topics related to opportunity. The mean of $AIExposure^{Reg}$ and $AIExposure^{Risk}$ are similar in magnitudes, 0.78 and 0.90 respectively, which are around 20% of the mean of the overall $AIExposure$. Interestingly, the mean for $AIExposure^{Pos}$ is 0.86 while the mean for $AIExposure^{Neg}$ is only 0.17, which indicates that any sentences that mention AI-related topics with a non-neutral tone are more possible to have a positive tone rather than a negative tone. Panel B shows the correlation among different *AI Exposure* measures. Similar to the mean, $AIExposure^{Opp}$ has the highest correlation with the overall $AIExposure$, and then it is $AIExposure^{Pos}$ and $AIExposure^{Reg}$.

Figure 3 shows the summary statistics of the overall *AI Exposure*, $AIExposure^{Opp}$, and $AIExposure^{Reg}$. Panel (a) presents the monthly mean of *AI Exposure*, $AIExposure^{Opp}$, and $AIExposure^{Reg}$, and Panel (b) presents the unique number of firms over time. The full dataset spans from July 2009 to June 2023. There are several observations that are worth noticing. First, the average means of all three measures have increased since 2009 and there is a boost near 2023. Second, the unique number of firms included in the full sample has been increasing since 2009. This is in line with the fact that over years, the transcript database covers more and more firms. The spikes shown in the figure are due to the quarterly basis measurement and there are clustering on transcripts on specific quarters.

B.2. Summary Statistics of AI Exposure Sorted Portfolios

After getting the overall firm-level *AI Exposure* measures, I present the summary statistics (mean) of the overall *AI Exposure* sorted portfolios. Table II presents the mean of the firm characteristics by quintiles. At the end of June of each year, stocks are ranked by their *AI Exposure* and sorted into quintiles. *AI Exposure* is defined as the fraction of AI-related unigrams and bigrams in a firm's call transcripts, and here it is multiplied by 1000 for ease of exposition. Firm size, *ln(ME)*, is measured as the log of market equity in June of year t. *BE/ME* is measured as the ratio between book equity at the end of June of year t and market equity at the end of December of year $t-1$. *ROE* is profitability, measured as income before extraordinary items divided by book equity. *ROA* is another profitability measure, defined as income before extraordinary items divided by total assets. *Asset Growth* is defined as the ratio between the change of total assets from year $t-2$ to year $t-1$ and total assets in year $t-2$. I follow Fahlenbrach, Rageth, and Stulz (2021) to construct the following three financial flexibility measures. *Cash/Assets* is the ratio of cash to total assets. *St Debt/Assets* is the ratio of debt in current liabilities to total assets. *Lt Debt/Assets* is the ratio of long-term debt to total assets. *Tangibility* is property, plant, and equipment divided by total assets. *Book Leverage* is the sum of current liabilities and long-term debt divided by total assets. *WW Index* is the Whited and Wu index used to measure financial constraint (see, Whited and Wu, 2006). I include all topic-specific *AI Exposure* measures as well. After winsorizing at the 1st and 99th percentiles to limit the influence of outliers, all firm characteristics variables are standardized to have a mean of zero and a standard deviation of one. The sample period is 2009 to 2024 at an annual frequency.

As shown in Table II, firms with higher *AI Exposure* tend to be smaller firms, growth firms (lower BE/ME and higher Asset Growth), with lower profitability, higher financial flexibility (higher cash ratio and lower debt ratio/leverage), and lower tangibility. With these characteristics pattern, I link to the anomaly literature in thinking ex ante whether *AI Exposure* could be positively priced or negatively priced in the stock market.

On the one hand, firms with smaller size should earn higher expected returns (Banz, 1981), and firms with lower tangibility or higher intangibility should earn higher expected returns (Eisfeldt and Papanikolaou, 2013). On the other hand, firms that behave like growth firms (Chan, Hamao, and Lakonishok, 1991), with lower profitability (Hou, Xue, and Zhang, 2015) and higher financial flexibility (Gamba and Triantis, 2008) should earn lower expected returns. Overall, firms with higher *AI Exposure* could be positively or negatively priced in the stock market. In the later empirical tests, I will dig more into this question.

B.3. AI Exposure by Industry

After having the overall firm-level *AI Exposure* measure, I aggregate to the industry level by firms' Standard Industrial Classification (SIC) code. Table III presents the mean of *AI Exposure* measures of top ten industries classified by the SIC 2-digit code. Panel A shows the top ten industries by the overall *AI Exposure* measure. The top four industries with the highest average *AI Exposure* are Educational Services, Insurance Agents, Brokers, & Service, Local & Interurban Passenger Transit, and Business Services. In general, Services face the highest AI exposure, potentially due to large impacts in areas such as customer support, where tasks traditionally performed by humans could be automated through technologies like chatbots. In Transportation, the advancements in autonomous vehicle technologies and AI-powered driving assistance could increasingly replace human-operated processes and hence have a high AI exposure. Appendix Table D.4 shows a more detailed decomposition by industries at the SIC 2-digit code level. Panel B shows the top ten industries by the overall $AIExposure^{Opp}$ measure. Seven out of ten match with the ranking by the overall *AI Exposure* measure in Panel A. This verifies the fact that the overall *AI Exposure* measure mainly captures opportunity.

III. Empirical Results

This section focuses on testing the null hypothesis in this paper that the two forces, regulatory uncertainty and opportunity, have opposing effects on firms' prices and expected stock returns. More specifically, the opportunity channel dominates *AI Starters*, while the regulation channel dominates *AI Elites*. As discussed in the introduction, I define the *AI Starters* as the small firms that might just start AI projects that suffer from profitability loss and more prone to opportunity shocks, and the *AI Elites* as the big firms that are well-developed and mature in AI projects and more prune to regulation shocks. To test this null hypothesis, I perform four sets of empirical tests. The following empirical tests will be around the prediction that the opportunity channel dominates *AI Starters* and *AI Starters* with higher $AIExposure$ should earn lower expected returns, while the regulation channel dominates *AI Elites* and *AI Elites* with higher $AIExposure$ should earn higher expected returns. Section III.A follows the typical asset pricing in portfolio sorts and factor regressions. Section III.B conducts multivariate Fama-MacBeth regressions with one-year forwarded realized returns. Section III.C also conducts multivariate Fama-MacBeth regressions, but with implied cost of capital estimates, instead of realized returns. Section III.D estimates the risk premium with option-implied expected return measure. Section III.E estimates risk premia associated with *AI Exposure* following a three-pass procedure, in order to take into account the omitted variable bias and measurement error. Section III.F conducts events studies to show the price reaction heterogeneity upon plausibly exogenous opportunity (innovation) or regulation (political) shocks.

A. Portfolio Sorts and Factor Regressions

First, I follow the typical asset pricing approaches in conducting portfolio sorts on $AIExposure$ measures (see, e.g., Davis, Fama, and French, 2000). In the first test, I calculate the return spread of value-weighted (VW) portfolios sorted on $AIExposure$ and $AIExposure^{Opp}$ between the highest decile (H) and lowest decile (L). If the null hypothesis is correct that the regulation channel dominates the *AI Elites*, then I should see a statistically positive return spread between H and L. However, if I calculate the return spread of equal-weighted (EW) portfolios sorted on $AIExposure$ and $AIExposure^{Opp}$, I should see a less positive or even negative return spread since the EW portfolios would be dominated by *AI Starters* that tend to be smaller in size, and hence more prune to the opportunity.

Table IV presents average monthly value-weighted (VW) or equal-weighted (EW) returns (in percentage) of portfolios sorted by overall firm-level *AI Exposure* and $AIExposure^{Opp}$. Here, I do not include other versions of measures due to the fact that other measures have too many zeros. The t-statistics are computed using heteroscedasticity and autocorrelation consistent Newey-West (1987) standard error estimates with a lag length of 12 months. Panel A presents VW returns of decile portfolios of different versions of *AI Exposure* measures. H-L is the return spread between the highest decile and the lowest decile. There is an overall increase in VW returns from portfolio L to portfolio H for both *AI Exposure* and $AIExposure^{Opp}$ sorted portfolios, although not strictly monotonic. The mean of monthly return spread, H-L, reaches 96 basis points for *AI Exposure* and 60 basis points for $AIExposure^{Opp}$, with a t-statistics of 2.65 and 1.67, respectively. Panel B presents VW returns of quintile portfolios of different versions of *AI Exposure* measures. The mean of monthly return spread, H-L, reaches 77 basis points for *AI Exposure* and 56 basis points for $AIExposure^{Opp}$, with a t-statistics of 2.56 and 2.08, respectively. Both Panel A and Panel B show that the return spread, H-L, is positive and (almost always) statistically significant for both *AI Exposure* and $AIExposure^{Opp}$ sorted portfolios. This is consistent with the fact that VW returns are usually dominated by big firms and big firms in my sample tend to be *AI Elites* that are predicted to be dominated by the regulation channel and hence firms with higher *AI Exposure* should earn positive expected returns. Panel C presents EW returns of decile portfolios of different versions of *AI Exposure* measures. The mean of monthly return spread, H-L, reaches -32 basis points for *AI Exposure* and -19 basis points for $AIExposure^{Opp}$, with a t-statistics of -0.73 and -0.83, respectively. Panel D presents EW returns of quintile portfolios of different versions of *AI Exposure* measures. The mean of monthly return spread, H-L, reaches -19 basis points for *AI Exposure* and -24 basis points for $AIExposure^{Opp}$, with a t-statistics of -0.53 and -1.03, respectively. Except for the case of $AIExposure^{Opp}$ in EW returns by quintiles, all return spread, H-L, is negative but statistically insignificant. This is consistent with the fact that EW returns are usually dominated by small firms and small firms in my sample tend to be *AI Starters* that are predicted to be dominated by the opportunity channel and hence firms with higher *AI Exposure* should earn lower expected returns.

There is a concern that the *AI Exposure* measures are correlated with common anomalies and once these anomalies are taken into account, the return spread will become statistically insignificant. Hence, I include sequential sorts. Table V present average monthly VW returns (in percentage) of sequentially sorted portfolios, first sorted by anomalies, and then sorted by overall firm-level *AI Exposure*. In each panel, I report sequential sort returns using full sample, and subsample that includes the second half of the full sample, i.e., recent years. This table shows that even after controlling for leading anomalies, the return spread, H-L, is overall positive and statistically significant, except for size, R&D, and cash holdings in the full sample. In the subsample, all return spread, H-L, is positive and statistically significant. This is in line with the fact that AI has become more prominent in recent years.

Then, I follow the standard procedure in factor regression tests to see if there is a risk premium associated with the *AI Exposure* measures that could not be explained by the leading risk factors. I first form ten portfolios based on *AI Exposure*. Specifically, at the end of June of each year, I allocate stocks into ten *AI Exposure*-sorted portfolios. I compute the portfolio returns (value weighted) difference between highest *AI Exposure* and lowest *AI Exposure*, and run time-series regressions of the returns spread on leading factor models. Here I include seven factor models. The first one is the CAPM model; The second one is the Chabi-Yo, Gonçalves, and Loudis (2025) ICAPM model; The third is the Fama and French (1996) three factor model; The fourth is the Fama and French (1996) three factor model plus Carhart (1997) momentum factor; The fifth is the Fama and French (2015) five factor model; the sixth is the Hou, Xue, and Zhang (2015) four q factors; the seventh is the Hou et al. (2020) five q5 factors. Data on the Fama-French factors and the Carhart factor are downloaded from Kenneth French's website. Data on q factors are from Lu Zhang's website.

Table VI presents the regressions estimation results using the full sample from July 2009 to June 2023. Overall, the results show that the risk premium associated with *AI Exposure* cannot be fully explained by leading factor models. There is a significantly positive α in all factor models at a monthly basis, with the maximum $\alpha = 99.1$ basis points (t=3.85) using q factors and the minimum $\alpha = 66.6$ (t=3.24) using Fama-French four factors. Since the portfolio returns are value weighted and more dominated by big stocks which in my sample tend to be *AI Elites*, the statistically positive alphas in these factor regressions indicate that *AI Elites* are dominated by the regulation channel and firms with higher *AI Exposure* earn a positive premium.

The comparison between Column (5) and Column (6) shows that expected growth does help explain the risk premium associated with *AI Exposure*. This is consistent with the mechanism as explained in Babina et al. (2024a). They find that growth options best explain the higher systematic risk associated with higher AI-investment firms in the labor channel.

Figure 4 plots the cumulative returns and twelve-month rolling alpha. Panel (a) presents the cumulative returns of the return spread between high- and low- *AIExposure* or $AIExposure^{Opp}$ portfolios, either VW or EW. The VW cumulative returns sorted by *AIExposure* quickly become positive and reaches a nearly 200% cumulative returns by the end of the sample. The VW cumulative returns sorted by $AIExposure^{Opp}$ experiences a negative value before 2018 but becomes positive shortly after 2018, and reaches more than 100% cumulative returns by the end of the sample. Panel (b)-(e) show the twelve-month rolling factor alphas of a strategy that longs the high *AIExposure* or $AIExposure^{Opp}$ portfolio and short the low *AIExposure* or $AIExposure^{Opp}$ portfolio. Panel (b) and (c) show that VW factor alphas are largely positive for both measures, but the *AIExposure* sorted factor alphas are far more positive than the $AIExposure^{Opp}$ sorted factor alphas. Panel (d) and (e) show that EW factor alphas are more negative in general and both experience a big negative alpha post 2021.

Overall, the results from portfolio sorts and factor regressions are consistent with the null hypothesis that *AI Elites* are dominated by the regulation channel and *AI Starters* are dominated by the opportunity channel.

B. *Fama-MacBeth Regressions*

Here, I study the *AI Exposure*-returns relationship by running Fama-MacBeth regressions with a long list of control firm-level characteristics. More specifically, I run cross-sectional regressions on a monthly basis from July of year t to June of year $t+1$. Each month, individual stock returns are regressed on *AI Exposure* from year $t-1$ and various control variables available by the end of June of year t. The definition for the control variables are the same as Section II.B.2. The t-statistics are based on standard errors adjusted using the Newey-West method. The full sample is from July 2009 to June 2023. I run Fama-MacBeth regressions in Ordinary Least Squares (OLS) and Weighted Least Squares (WLS) with market value as the weights.

Table VII presents the Fama-MacBeth regressions using the full sample. Model (1) shows that there is a significantly negative relationship between the overall *AI Exposure* measure and excess returns. Furthermore, the coefficients are negative and statistically significant for model (2) and (3) for $AIExposure^{Opp}$ and $AIExposure^{Reg}$, while the coefficient is statistically insignificant for $AIExposure^{Risk}$. ROE and R&D/Asset ratio positively predict the excess returns. While financial flexibility also helps explains the excess returns, size, BE/ME, and Tangibility do not have statistically significant coefficients. The OLS regressions are often dominated by small firms, which in my context tend to be *AI Starters*. I further include WLS regressions with the market value as weights and I see positive (while not significant) coefficient between excess returns and *AI Exposure* ($AIExposure^{Opp}$).

One concern is that AI is not prevalent in the early 2010s and investors might not care. I hence do subsample regression and ex ante I expect to see a more pronounced negative relationship between *AI Exposure* measures and excess returns. Table VIII presents the results on a subsample from July 2016 to June 2023, and the overall *AI Exposure*-returns relationship becomes less statistically significant in model (1), while the magnitude of the negative relationship becomes greater. The coefficient is still significantly negative in model (2) for $AIExposure^{Opp}$ and with a bigger magnitude. The WLS model (5) and model (6) show greater positive coefficients between excess returns and *AI Exposure* ($AIExposure^{Opp}$).

C. *Implied Cost of Capital*

There is a continuing debate in the literature on what the best proxies for expected returns are. To mitigate the concern that ex-post returns might be a bad proxy for expected returns, I estimate a few Implied Cost of Capital (ICC) measures and run Fama-MacBeth regressions to compare with ex-post returns.

Table IX presents the results. I follow Hou, van Dijk, and Zhang (2012) in constructing four main ICC measures. Column (1) reports regressions of ex-post realized returns on *AI Exposure*; column (2) reports results of ICC estimates based on Gebhardt, Lee, and Swaminathan (2001); column (3) reports results of ICC estimates based on Claus and Thomas (2001); column (4) reports results of ICC estimates based on Ohlson and Juettner-Nauroth (2005); column (5) reports results of ICC estimates based on Easton and Monahan (2005); column (6) reports results of ICC estimates that are computed as the average of the previous four ICC estimates. I denote these ICC measures as GLS, CT, OJ, MPEG, and Composite, respectively. Panel A regresses on the overall $AIExposure$ measure, while Panel B regresses on $AIExposure^{Opp}$.

There is a mixed result. In Panel A, Column (1) shows that *AI Exposure* negatively predict future ex-post realized turns as discussed before. Column (2) shows that *AI Exposure* negatively predicts GLS, but not statistically significant. Column (3) to (6) show that *AI Exposure* positively predicts CT, OJ, MPEG and the Composite ICC measure. Here, *AI Exposure* is adjusted to mean zero and one standard deviation. Hence to interpret the coefficient, for example in Column (6), one standard deviation increase in *AI Exposure* is associated with 0.649 percentage point increase in the Composite ICC measure.

In Panel B, Column (1) to (6) all show a negative coefficient between $AIExposure^{Opp}$ and realized returns/ICC measures. All coefficients are statistically significant except for

MPEG measure. Since $AIExposure^{Opp}$ should mainly capture the opportunity channel, the consistent negative relationship between ICC measures and $AIExposure^{Opp}$ validate the null hypothesis that the opportunity channel dominates *AI Starters* and there should be a negative $AIExposure^{Opp}$-returns relationship.

D. Option-Implied Returns

Apart from the ICC measures, I apply another option-implied ex ante expected returns measure following Martin and Wagner (2019). The monthly expected stock returns for firms in the S&P 500 using the 1-month, 3-month, 6-month, 9-month, and 12-month option-implied measures are obtained from Vilkov (2023). One limitation of this measure is that it only applies to S&P 500 firms, which in my sample tend to be *AI Elites*. Hence, I would expect to see a positive expected return spread between a value-weighted decile portfolio with the highest *AI Exposure* stocks and a value-weighted decile portfolio with the lowest *AI Exposure* stocks. This is exactly what I see in the real data.

Table X presents the average monthly value-weighted (VW) option-implied returns (in percentage) of portfolios sorted by overall firm-level *AI Exposure*. Panel A presents VW average monthly 1-month returns of decile portfolios. H-L is the return spread between the highest decile and the lowest decile. Panel B presents VW average monthly 3-month returns of decile portfolios. Panel C presents VW average monthly 6-month returns of decile portfolios. Panel D presents VW average monthly 9-month returns of decile portfolios. Panel E presents VW average monthly 12-month returns of decile portfolios. The return spreads, H-L, are statistically positive across all panels. The magnitude of the monthly return spread ranges from 0.45% to 0.73%. This is comparable to the VW return spread using realized returns in Table IV, which is 0.96%.

E. Three-Pass Procedure with Latent Factors

One big critique regarding the traditional two-pass cross-sectional regressions, like Fama-MacBeth regressions and factor regressions, is that they are susceptible to omitted variable bias (see, e.g., Burmeister and McElroy, 1988; Jagannathan and Wang, 1998). Estimating risk premia using the traditional regressions cannot guarantee resolving the omitted variable bias since we cannot possibly account for all priced risks in the market. This section follows the three-pass approach proposed by Giglio and Xiu (2021) to take omitted variable bias and measurement error into account and to estimate the risk premia associated with *AI Exposure*. The three-pass approach is, in essence, a combination of Principal Component Analysis (PCA) with two-pass cross-sectional regressions, to generate a consistent estimate of AI risk premium. Assume that returns follow a linear factor model with p factors, and my goal is to estimate the risk premium of one of them, g_t, which is the AI risk factor associated with *AI Exposure*.

The three-pass approach consists of the following steps. First, I use PCA to recover factors and their loadings using a large set of portfolios. Second, I run cross-sectional regressions using only the principal components, without the factor of interest, g_t, to estimate their risk premia. Third, I run time-series regressions of g_t onto the principal components to estimate the loadings on the principal components. The risk premium of g_t is calculated as the product of the loadings estimated in step two and their risk premia estimated in step three.

In the first step, I include a wide range set of testing portfolios. First, I include 202 standard equity portfolios available on Kenneth French's website as Giglio and Xiu (2021) do, covering the most well-known dimensions of risk. These include: 25 portfolios sorted by size and book-to-market ratio, 17 industry portfolios, 25 portfolios sorted by operating profitability and investment, 25 portfolios sorted by size and variance, 35 portfolios sorted by size and net issuance, 25 portfolios sorted by size and accruals, 25 portfolios sorted by size and beta, and 25 portfolios sorted by size and momentum. Second, I include 2925 equity portfolios available on Lu Zhang's website with monthly two-way sorted (3×5) portfolios based on 195 anomalies including momentum (42), value-versus-growth (32), investment (32), profitability (48), intangibles (31), and frictions (10).

One concern regarding this approach is on selecting parameters to display best results. To mitigate this concern, I follow the three-pass procedure to estimate the risk premium of eight AI factors constructed either using H-L with decile or quintile sorting and value-weighted or equal-weighted for *AI Exposure* or $AIExposure^{Opp}$. I include a range of one to ten for Lags when computing Newey-West standard errors. I also include a range of one to ten for the number of PCs.

Figure 5 plots the estimated risk premium for the eight AI factors. The four circle plots represent value-weighted AI factors, while the four triangle plots represent equal-weighted AI factors. "AIE" denotes the return spread between the highest and lowest deciles sorted by $AIExposure$, and "AIE Opp" denotes the return spread between the highest and lowest deciles sorted by $AIExposure^{Opp}$. When the notation "5" is used, it indicates the return spread between the highest and lowest quintiles. For each factor, I estimate the risk premia using different sets of testing portfolios, a range of one to ten lags for computing Newey-West standard errors and a range of one to ten for principal components. The dataset spans from July 2009 to June 2023.

There are several observations. First, overall, the three-pass risk premia for both value-weighted (VW) and equal-weighted (EW) factors are mostly insignificant at the 95% confidence level. However, there are exceptions: the EW factors exhibit negative significant risk premia, with the minimum monthly risk premia around -100 basis points, and no factors show positive significant risk premia. This contradicts my null hypothesis, as the significance of the risk premia essentially disappears. Nonetheless, Giglio and Xiu (2021) argue that the three-pass procedure performs better as $T \to \infty$ (their sample spans 1080 months), whereas my sample covers only 168 months; the limited time-series length may diminish statistical power.

Second, VW factors are, on average, more positive than EW factors, while EW factors are more negative than VW factors across all setups. The maximum monthly risk premia for VW factors reach almost 40 basis points, compared to less than 10 basis points for EW factors. This finding aligns with the null hypothesis that *AI Starters* earn a more negative risk premium, while *AI Elites* earn a more positive risk premium.

F. Event Study

In this section, I conduct two sets of event studies to test the null hypothesis that the opportunity channel dominates *AI Starters*, while the regulation channel dominates *AI Elites*. The event studies are to show whether there is price reactions heterogeneity for firms with different *AI Exposure* upon plausibly exogenous opportunity or regulation shocks. For opportunity shocks to AI, I include the launch of ChatGPT on Nov 30th, 2022. The launch of ChatGPT is often treated as a plausibly exogenous shock to AI innovations and bring opportunities for firms to boost their productivity across all industries. For example, Eisfeldt, Schubert, and Zhang (2023) show a positive "Artificial Minus Human" return spread after the launch of ChatGPT. For regulation shocks to AI, I include Trump's U.S. presidential election win on Nov 8th, 2016, Biden's U.S. presidential election win on Nov 3rd, 2020, and Biden's AI Executive Order on Oct 30th, 2023. Trump's election is widely assumed to be favorable to AI related regulations, while Biden's election is relatively restrictive. However, Biden signed an AI Executive Order to further restrict AI development and uses in the US, which is more restrictive.

I follow the literature (see, e.g., Brown and Huang, 2020; Fahlenbrach, Ko, and Stulz, 2024) in estimating average cumulative abnormal returns (CAR) in percentage after opportunity or regulation shocks in each quintile portfolio sorted on *AI Exposure*. More specifically, I compute the CARs over a 10-day window from day $t + 1$ to day $t + 10$ after the event day t by equally weighted across *AI Exposure*-sorted portfolios, which I denote as CAR[1,10]. I follow the typical way by using a 250-trading day window that ends 25 days before the event day to estimate CAPM beta (daily market returns and the risk-free rate obtained from the French website).

Table XI presents the results. Panel A computes the CAR[1,10] upon the Trump election win on Nov 8th, 2016 by quintile and high minus low (H-L). It shows that there is no significant CAR difference since the CAR[1,10] spread between high and low quintile is only -0.75% with a t-statistics of -0.89. The same observation holds for the Biden election win on Nov 3rd, 2020 in Panel B that the CAR[1,10] spread between high and low quintile is only -0.27% with a t-statistics of -0.45. These two events are general political events which are not targeting AI. Hence, it is not surprising to see no heterogeneity in price reactions across different *AI Exposure* quintiles. However, in Panel C, the CAR[1,10] spread between high and low quintile is -1.84% with a t-statistics of -2.49. This is consistent with the fact that Biden's AI Executive Order on Oct 30th, 2023 is largely restricting AI developments and uses in the US so that we see firms with higher *AI Exposure* experience a more negative price reactions. To test whether regulation channel dominates *AI Elites* that tend to be bigger firms, in Panel D, I exclude big stocks (top 20 percentile in size) and show that price reactions diminish in all quintiles and the CAR[1,10] spread between high and low quintile. In Panel E, I show that the launch of ChatGPT3 as a plausibly exogenous AI innovation shock leads to a statistically positive CAR[1,10] spread between high and low quintile, with a 1.49% return spread and a t-statistics of 2.72. There is a monotonic increase in CAR[1,10] from the lowest *AI Exposure*-sorted quintile to the highest *AI Exposure*-sorted quintile. To test whether opportunity channel dominates *AI Starters* that tend to be smaller firms, in Panel F, I exclude tiny stocks (bottom 20 percentile in size) and show that price reactions diminish in all quintiles and the CAR[1,10] spread between high and low quintile.

The results from these event studies are consistent with the null hypothesis that both regulation and opportunity shocks lead to significant price reactions, and regulation channel dominates the *AI Elites* and opportunity channel dominates the *AI Starters*.

IV. Theory

After documenting the empirical findings, I construct a competitive equilibrium model to show how the two forces, regulation uncertainty and opportunity, affect the stock prices and risk premium, building on Pástor and Veronesi (2012, 2013) and Hsu, Li, and Tsou (2023). In the base model, I assume there are two types of regime changes that directly affect firms' profitability process: regulation regime change and opportunity regime change.

The regulation regime change is between Favorable (F) and Unfavorable (U) policy. I assume the change from Favorable (F) and Unfavorable (U) policy would negatively affect the profitability process which will be incorporated in the profitability process. Here, the Favorable policy means the government's attitude towards AI development, investments, and uses are generally favorable, which could be related to government subsidies for AI research, positive tax incentives, or relaxed data privacy laws. However, the Unfavorable policy means that government's attitude is more restrictive and legally burdensome. As previously discussed, Biden's AI Executive Order in 2023 emphasized, "safe, secure, and trustworthy development and use of artificial intelligence," which is usually deemed as unfavorable to AI development and uses in the US. For example, Trump revoked Biden's EO in 2025. On January 23, 2025, the White House stated that "*The Biden AI Executive Order established unnecessarily burdensome requirements for companies developing and deploying AI that would stifle private sector innovation and threaten American technological leadership.*" This model's main purpose is not to quantify how unfavorable or favorable different presidential administrations are towards the AI development and uses, but to showcase the regulation regime change could affect prices.

The opportunity regime change is regarding the overall AI development stages between Early (E) and Mature (M). I assume the change of overall AI development from Early (E) and Mature (M) would positively affect the profitability process. In the Early stages of AI development or uses, there are high entry costs since most of the AI development or uses incur long-run R&D expenses and high costs due to other frictions like unavailability of large language models (ChatGPT). However, in the Mature stages of AI development or uses, the entry costs are lower due to less frictions. For example, building on open resources of existing LLMs, the Chinese tech company, DeepSeek, launched a large AI model in early 2025 and said that training one model cost less than $6 million, which would have cost up to $1 billion back in early 2010s.

A. *Firm Profitability*

I assume there is a continuum of firms of $i \in [0, 1]$ and a finite horizon $t \in [0, T]$. The stage of AI development and uses for firm i at time t is denoted as $s_t^i \in [0, 1]$. Specifically, when $s_t^i = 0$, firms are classified as *AI Starters*, where profitability is more sensitive to innovation or opportunity shocks. Conversely, when $s_t^i = 1$, firms are classified as *AI Elites*, where profitability is more sensitive to regulatory changes. The profitability process for firm i at time t is modeled as:

$$d\Pi_t^i = \left(\mu + \xi_t^i \left[s_t^i \cdot g_t + (1 - s_t^i) \cdot \eta_t\right]\right) dt + \sigma dZ_t + \sigma_1 dZ_t^i, \tag{3}$$

where μ represents the baseline profitability growth rate, and ξ_t^i captures firm i's exposure to AI, which could be deemed as the *AI Exposure* measure in the empirical setting and captures how much attention investors and managers devote to AI developments, investments and uses. The term g_t reflects regulatory uncertainty, where $g_t^F > 0$ denotes a Favorable regulatory environment and $g_t^U < 0$ denotes an Unfavorable one. Additionally, η_t captures the opportunity environment, representing the overall stage of AI development and uses, with $\eta_t^E < 0$ characterizing an Early-stage environment and $\eta_t^M > 0$ characterizing a Mature-stage environment. The terms σdZ_t and $\sigma_1 dZ_t^i$ correspond to aggregate and idiosyncratic shocks, respectively.

The profitability process can be interpreted differently depending on the stage of AI development and uses. For *AI Elites* ($s_t^i \approx 1$), profitability is predominantly driven by regulatory uncertainty, resulting in the simplified profitability process:

$$d\Pi_t^i \approx \left(\mu + \xi_t^i g_t\right) dt + \sigma dZ_t + \sigma_1 dZ_t^i. \tag{4}$$

In contrast, for *AI Starters* ($s_t^i \approx 0$), profitability is more influenced by the overall AI development stage, with the profitability process expressed as:

$$d\Pi_t^i \approx \left(\mu + \xi_t^i \eta_t\right) dt + \sigma dZ_t + \sigma_1 dZ_t^i. \tag{5}$$

In this base model, I consider two types of firms with either high or low *AI Exposure*, i.e., ξ_t^H or ξ_t^L. g_t captures the impact of AI-related regulation. For simplicity, I consider two cases for regulation regimes, Favorable (F) or Unfavorable (U). I further assume that,

$$g = \begin{cases} g_t^F > 0 & \text{means Favorable regulation on AI carries a positive effect} \\ g_t^U < 0 & \text{means Unfavorable regulation on AI carries a negative effect,} \end{cases} \tag{6}$$

because under Early AI-stage, firms bears negative cash flow effects for high entry costs. For example,

I also consider two cases for opportunity regimes, Early (E) or Mature (M) for the aggregate AI development and uses in the US. I further assume that,

$$\eta = \begin{cases} \eta_t^M > 0 & \text{means aggregate Mature AI-stage carries a positive effect} \\ \eta_t^E < 0 & \text{means aggregate Early AI-stage carries a negative effect,} \end{cases} \tag{7}$$

If we assume the current AI regulation in the U.S. is more Unfavorable during Biden's administration, then we should expect a negative relationship between profitability, i.e., ROE, and firm-level *AI Exposure* since,

$$g_t^U(\xi_t^H - \xi_t^L) < 0, \tag{8}$$

where $\xi_t^H > \xi_t^L > 0$. Likewise, in the Early-AI stages such as early 2010s, the entry costs are high and impose a negative effect on profitability, and so we should expect a negative relationship between profitability, i.e., ROE, and firm-level *AI Exposure* since,

$$\eta_t^E(\xi_t^H - \xi_t^L) < 0, \tag{9}$$

where $\xi_t^H > \xi_t^L > 0$. With these assumptions, I have two corresponding predictions that I can test in the real data:

Prediction 1. *Firms with higher AI Exposure should have lower ROE under the overall Unfavorable policy environment or the Early-stage AI environment.*

Prediction 2. *Firms with higher AI Exposure are more sensitive to policy changes.*

I show in the empirical section that this assumption is consistent with what I find in the data. The prediction cannot be directly tested regarding the Early-stage part since AI is still at the early stage as widely assumed. However, I have discussed the assumption is decent using the example of Deepseek. The following test will be focused on the Unfavorable part. The following panel regression is proposed to test these two predictions,

$$\text{ROE}_{i,t} = \beta_0 + \beta_1 AIExposure_{i,t} + \beta_2 \mathbb{1}_{\{g_t = g_t^U\}} + \beta_3 (AIExposure_{i,t} \times \mathbb{1}_{\{g_t = g_t^U\}}) + \gamma X_{i,t} + \varepsilon_{i,t}, \tag{10}$$

where $\text{ROE}_{i,t}$ represents the return on equity for firm i at time t. The variable $AIExposure_{i,t}$ is defined as the firm-level AI exposure as before. The dummy variable $\mathbb{1}_{\{g_t = g_t^U\}}$ captures period with Unfavorable AI regulatory, where $g_t^U < 0$. The vector $X_{i,t}$ includes control variables such as firm size, leverage, and industry fixed effects..

Table XII reports panel regressions of ROE on *AI Exposure* and other firm characteristics. *isUnfav* is a dummy variable if year is between 2021 and 2023 during Biden's administration, when the AI policy is relatively unfavorable. I include the interaction term between *AI Exposure* and *isUnfav*. *Current ROE* is defined as the ROE at year t, *Lagged ROE* at year $t-1$, and *Future ROE* at year $t+1$. Column (1) and (2) have *Current ROE* as the dependent variable and include *Lagged ROE* as one of the independent variables. Column (3) and (4) have *Future ROE* as the dependent variable and include *Current ROE* as one of the independent variables. The coefficients of interest are β_2 (coefficient on *isUnfav*) and β_3 (coefficient on *AI Exposure* $\times$ *isUnfav*). The Prediction 1 and Prediction 2 basically test whether β_2 and β_3 are statistically negative. First, the results show that the firm-level *AI Exposure* does negatively comove with current ROE and negatively predict future ROE, i.e., β_2 is statistically negative. The only exception is Column (2) when clustering SE by Industry $\times$ Year. Second, the results show that β_3 is statistically negative across all regressions regarding the SE clustering level.

B. Household and Government Maximization Problems

Households/investors maximize a CRRA utility function,

$$U(W_t) = \frac{W_t^{1-\gamma}}{1-\gamma}, \quad \gamma > 1, \tag{11}$$

where γ is the risk aversion and W_t is the wealth of households at time t. To keep the model trackable, I assume that there is no dividend payment before T and households liquidate and consume all wealth at time T.

Government decides on whether to have a policy change at time τ, and households observe the decision at τ. Government maximizes a similar problem as households, but face a non-pecuniary cost (or benefit) associated with an AI-policy change from Favorable to Unfavorable,

$$\max_{\tau > t} \left\{ \mathbb{E}_t \left[\frac{\phi(c) W_\tau^{1-\gamma}}{1-\gamma} \mid U \right], \mathbb{E}_t \left[\frac{W_\tau^{1-\gamma}}{1-\gamma} \mid F \right] \right\}, \tag{12}$$

where

$$\phi(c) = 1 + e^c, \tag{13}$$

which could be interpreted as the shadow cost of imposing an Unfavorable AI policy from the government's perspective. For example, Biden Administration's Executive Orders on AI Safety, Security, and Privacy could be treated as an Unfavorable AI policy, which imposes costs for the government. And the prior distribution of c is assumed to be drawn from a normal distribution,

$$c \sim N\left(-\frac{\sigma_c^2}{2}, \sigma_c^2\right), \tag{14}$$

where σ_c captures the policy uncertainty. Using a Moment Generating Function shows that, $\mathbb{E}[e^c] = 1$. The market clears at time T,

$$W_T = B_T = \int_0^1 B_T^i di, \tag{15}$$

where B_T^i denotes firm i's capital at time T.

C. *Bayesian Learning*

The political cost, c, is unknown to all agents in this economy and all agents learn about c for any time before the policy decision, i.e., $t < \tau$. The Bayesian learning process involves updating signals, s, with the mean and an independent noise,

$$ds_t = cdt + hdZ_t^c, \tag{16}$$

where h governs the magnitude of noise, and dZ_t^c is an independent Brownian motion from aggregate and idiosyncratic shocks in Equation (3). We can think of these ds_t as a steady flow of political news related to AI policy.

Lemma 1. *Combining the signals, Equation (16), and prior distribution of the political cost, Equation (14), we can get the posterior distribution of the political cost,*

$$c \,|\, \mathcal{F}_t^s \sim N(\hat{c}_t, \hat{\sigma}_{c,t}^2) \tag{17}$$

where

$$d\hat{c}_t = \hat{\sigma}_{c,t}^2 h^{-1} d\hat{Z}_t^c, \tag{18}$$

and

$$\hat{\sigma}_{c,t}^2 = \frac{1}{\frac{1}{\sigma_c^2} + \frac{1}{h^2}t} \tag{19}$$

Proof. See the Proof in Appendix E.A. □

The intuition behind the posterior distribution of the political cost is straightforward: If the signal is noisy (h is large), the update is small. If the signal is precise (h is small), the update is large. The uncertainty about cost decreases over time as more signals are observed. The rate of learning depends on the noise level (h) and prior political uncertainty (σ_c).

D. *Government Optimal Choice*

Government learns about c and makes a policy-change decision at τ, i.e., from Favorable (F) to Unfavorable (U) iff:

$$\mathbb{E}_\tau \left[\frac{W_T^{1-\gamma}}{1-\gamma} \mid F \right] < \mathbb{E}_\tau \left[\phi(c) \frac{W_T^{1-\gamma}}{1-\gamma} \mid U \right] \tag{20}$$

As assumed, the government AI policy change will directly affect profitability, and the two expectations in equation (20) should be calculated under different stochastic processes for the aggregate capital $B_t = \int_0^1 B_t^i di$, combined with market clear condition equation (15). Without loss of generality, I further assume that, $\xi^i \sim \text{Uniform}(0,2)$ with $\mathbb{E}[\xi^i] = 1$, and each firm's AI stage, $s^i \sim \text{Bernoulli}(p)$, where $s^i = 1$ (AI Elites) with probability p, and $s^i = 0$ (AI Starters) with probability $1-p$, where p captures the proportion of AI Elite firms in the continuum of firms.

Lemma 2. *The aggregate capital at time T, $B_T = \int_0^1 B_T^i di$, is given by*

$$B_T = B_\tau e^{\left(\mu + pg + (1-p)\eta - \frac{1}{2}\sigma^2\right)(T-\tau) + \sigma(Z_T - Z_\tau)},$$

where $g \equiv g^F$ under Favorable regulation, $g \equiv g^U$ under Unfavorable regulation, $\eta \equiv \eta^E$ in Early-stage AI environment, and $\eta \equiv \eta^M$ in Mature-stage AI environment.

Proof. See the Proof in Appendix E.B. □

It is straightforward to see that the aggregate capital at terminal time T depends on the regulation regime and opportunity regime. With Lemma 2, we can further simplify the condition on which the government will choose to switch from a Favorable (F) AI policy to an Unfavorable (U) AI policy. The following proposition shows the condition.

Proposition 1. *The government will switch from a Favorable (F) to an Unfavorable (U) AI policy at time τ if and only if the realized political cost c exceeds the threshold:*

$$c > \underline{c}(\tau) \equiv \log\left(e^{(\gamma-1)\left(g^F - g^U\right)(T-\tau)} - 1\right),$$

where $g^F > 0$ and $g^U < 0$ represent the regulatory impacts on profitability under Favorable and Unfavorable regimes, respectively. The threshold $\underline{c}(\tau)$ increases with risk aversion ($\gamma > 1$) and the magnitude of the regulatory gap $|g^F - g^U|$.

Proof. See the Proof in Appendix E.C. □

Corollary 1. *The time-t perceived probability of a government AI policy shift from Favorable (F) to Unfavorable (U) at time τ ($t < \tau$) is:*

$$p_{\tau|t} = 1 - \Phi\left(\underline{c}(\tau); \hat{c}_t, \hat{\sigma}_{c,t}^2\right),$$

where $\Phi(\cdot; \hat{c}_t, \hat{\sigma}_{c,t}^2)$ is the CDF of the posterior normal distribution $c \sim N(\hat{c}_t, \hat{\sigma}_{c,t}^2)$, and $\underline{c}(\tau)$ is defined in Proposition 1.

Proof. See the Proof in Appendix E.D. □

The above Proposition 1 and Corollary 1 lead to several testable predictions. I will discuss them one by one below.

Prediction 3. *When investors have higher risk aversion, i.e., $\gamma > 1$, it becomes less likely to have an AI-related policy change from Favorable to Unfavorable.*

Prediction 4. *When the impact of policy change from Favorable to Unfavorable ($g^F - g^U$) becomes larger, it becomes less likely to have a policy change from Favorable to Unfavorable.*

Prediction 5. *The Trump's US presidential victory at the end of 2016 decreases (or neutral) the probability of adopting Unfavorable policy change on AI development. He actually enhanced AI development, e.g., 2019 American AI Initiative. Firms with higher AI Exposure should experience more positive stock prices reactions upon these events.*

Prediction 6. *The 2019 Trump's American AI Initiative marks a more Favorable AI political environment. The impact of AI Favorable policy, i.e., g^F in the model, should become more positive. This could be tested to see if more labor displacement threats happen, more business ideas replaced by AI, and more physical capital becomes obsolete.*

Prediction 7. *Biden signed Unfavorable (or less Favorable) Executive Order on Oct, 2023. In the model, it means, the probability of change from Favorable to Unfavorable increases, or the impact of AI Favorable policy, i.e., g^F in the model, should become less positive or even become negative as g^U (Unfavorable) in the model. AI labor displacement effect reduced. Firms with higher AI Exposure react less positively or even negatively in the stock prices upon these events.*

The event studies in the previous empirical section speak partially to the above predictions (5-7). I defer additional empirical tests to future work to further examine these predictions.

E. SDF, Stock Returns, and Risk Premia

Now we turn to the stock market implications, we can derive the SDF and see how this policy shock from Favorable to Unfavorable is priced in the stock market. We first derive the state price density in the following proposition:

Proposition 2. *Before a regulatory or opportunity regime shift ($t < \tau$), the state price density is:*

$$\pi_t = B_t^{-\gamma}\Omega_t,$$

where

$$\Omega_t = e^{\left(-\gamma\mu+\frac{1}{2}\gamma(\gamma+1)\sigma^2\right)(T-t)-\gamma\left(pg^F+(1-p)\eta\right)(\tau-t)}\left[p_{\tau|t}e^{-\gamma\left(pg^U+(1-p)\eta\right)(T-\tau)}+(1-p_{\tau|t})e^{-\gamma\left(pg^F+(1-p)\eta\right)(T-\tau)}\right]$$

and $p_{\tau|t}$ is the probability of a shift to Unfavorable AI policy (Corollary 1).

Proof. See the Proof in Appendix E.E. □

After deriving the state price density, we can apply Itô's Lemma and get the process for SDF in the following proposition:

Proposition 3. *The SDF is,*

$$\frac{d\pi_t}{\pi_t} = \mathbb{E}_t\left[\frac{d\pi_T}{\pi_T}\right] - \lambda dZ_t - \lambda_{c,t}d\hat{Z}_t^C$$

where,

$$\lambda = \gamma\sigma > 0$$

and

$$\lambda_{c,t} = \frac{1}{\Omega_t}\frac{\partial\Omega_t}{\partial\hat{C}_t}\hat{\sigma}_{c,t}^2\eta^{-1} > 0$$

Proof. See the Proof in Appendix E.F. □

Prediction 8. *AI policy change shock risk (from Favorable to Unfavorable) is negatively priced in the stock market, e.g., Biden's EO should decrease prices for high AI-exposure firms.*

With the derived SDF process, we can further write down the stock returns and risk premia assuming zero risk-free rate as the continuous time model typically does. The following proposition states the process for stock realized returns and risk premia. Before we jump to the proposition, we need a lemma to help us prove the proposition.

Lemma 3. *For $t < \tau$, the stock price for firm i is given by:*

$$M_t^i = B_t^i\Theta_t^i,$$

where

$$\Theta_t^i = e^{(\mu-\gamma\sigma^2)(T-t)+\xi^i\left(pg^F+(1-p)\eta\right)(\tau-t)}\left[\phi_t e^{\xi^i\left(pg^U+(1-p)\eta\right)(T-\tau)} + (1-\phi_t)e^{\xi^i\left(pg^F+(1-p)\eta\right)(T-\tau)}\right],$$

and

$$\phi_t = \frac{p_{\tau|t}}{p_{\tau|t} + (1-p_{\tau|t})e^{-\gamma p(g^F-g^U)(T-\tau)}}.$$

Proof. See the Proof in Appendix E.G. □

Proposition 4. *Firm i's stock realized returns:*

$$\frac{dM_t^i}{M_t^i} = \mathbb{E}_t\left[\frac{dM_T^i}{M_T^i}\right] + \sigma dZ_t + \sigma_1 dZ_t^i + \beta_{M,t}^i d\hat{Z}_t^c,$$

where

$$\beta_{M,t}^i = \frac{1}{\Theta^i}\frac{\partial\Theta_t^i}{\partial\hat{c}_t}\hat{\sigma}_{c,t}^2 < 0,$$

and risk premia can be expressed as:

$$\mathbb{E}_t\left[\frac{dM_T^i}{M_T^i}\right] = \sigma\lambda dt + \beta_{M,t}^i\lambda_{c,t}dt,$$

where

$$\frac{\partial\beta_{M,t}^i}{\partial\xi^i} = \begin{cases} > 0 & \text{if } p = 0 \text{ (Firms are all AI Starters)}, \\ < 0 & \text{if } p = 1 \text{ (Firms are all AI Elites)}, \end{cases} \qquad \text{and} \quad \lambda_{c,t} < 0.$$

Proof. See the Proof in Appendix E.H. □

Lemma 4. *For any two firms, i and j, with AI Exposure $\xi^i > \xi^j$, $\forall i \neq j$, return premia spread between i and j is,*

$$\mathbb{E}_t\left[\frac{dM_T^i}{M_T^i}\right] - \mathbb{E}_t\left[\frac{dM_T^j}{M_T^j}\right] = (\beta_{M,t}^i - \beta_{M,t}^j)\lambda_{c,t}dt \begin{cases} > 0 & \text{if } p = 0 \text{ (Firms are all AI Starters)}, \\ < 0 & \text{if } p = 1 \text{ (Firms are all AI Elites)}, \end{cases}$$

Proof. See the Proof in Appendix E.I. □

Prediction 9. *There should be a statistically significant positive H-L return spread among AI Exposure-sorted portfolios if value-weighted (AI Elites dominate), and a statistically significant negative H-L return spread among AI Exposure-sorted portfolios if equal-weighted (AI Starters dominate).*

The results from the previous empirical section (Table IV) show largely consistency with this prediction. Panel A presents VW returns of decile portfolios of different versions of *AI Exposure* measures. H-L is the return spread between the highest decile and the lowest decile. Panel B presents VW returns of quintile portfolios of different versions of *AI Exposure* measures. Both Panel A and Panel B show that the return spread, H-L, is positive and statistically significant for both *AI Exposure* and $AIExposure^{Opp}$ sorted portfolios. This is consistent with the fact that VW returns are usually dominated by big firms and big firms in my sample tend to be *AI Elites* that are predicted to be dominated by the regulation channel and hence firms with higher *AI Exposure* should earn positive expected returns. Panel C presents EW returns of decile portfolios of different versions of *AI Exposure* measures. Panel D presents EW returns of quintile portfolios of different versions of *AI Exposure* measures. Except for the case of $AIExposure^{Opp}$ in EW returns by quintiles, all return spread, H-L, is significantly negative. This is consistent with the fact that EW returns are usually dominated by small firms and small firms in my sample tend to be *AI Starters* that are predicted to be dominated by the opportunity channel and hence firms with higher *AI Exposure* should earn lower expected returns.

V. Validation: Hedging AI Innovations

This section validates the main measures used in the empirical analysis, namely, *AIV* and firm-level *AI Exposure*. Following the mimicking portfolio approach of Engle et al. (2020), I show that using my *AIV*, investors can create a time series of *AI News Index* to capture the innovation in AI development. In addition, investors can use my firm-level *AI Exposure* measures to construct a mimicking portfolio to hedge the innovation in the *AI News Index*.

A. RavenPack AI News Index

After getting the *AIV*, I construct the *RavenPack AI News Index*. I use approximately 170 million news headlines from the Dow Jones and Press Release Edition of the RavenPack News database from 2020 to 2024. I compute a monthly *RavenPack AI News Index* in two version, *Number of Headlines* and *Fraction of AI-Related Headlines*. *Number of Headlines* is defined as the total number of headlines that contain unigrams or bigrams in the previously defined *AIV* on a monthly basis. *Fraction of AI-Related Headlines* is the fraction of total number of headlines that contain unigrams or bigrams in the previously defined *AIV* over the total number of headlines on a monthly basis.

Figure 2 presents the RavenPack AI News Index from 2010 to 2024, along with key AI-related news announcements highlighted in the plot. The plot displays the *Number of Headlines*, representing the monthly count of RavenPack news headlines that include *AIV* terms. The *RavenPack AI News Index* effectively captures the primary trends in AI-related news. Notably, significant events include the launch of ChatGPT3 at the end of 2022 and ChatGPT4 in March 2023, which correspond to notable increases in the *Number of Headlines*. Prior to 2017, AI-related news was sparse. Key developments in AI-related topics during this period were focused on machine learning and deep learning applications. For instance, Google's Neural Networks team succeeded in identifying cats from millions of YouTube videos, DeepMind excelled at Atari games, Facebook introduced DeepFace, Google launched TensorFlow, and AlphaGo triumphed over Lee Sedol in the Go match. After 2017, applications of machine learning and deep learning expanded into everyday home and customer services, exemplified by products like Google Home, Amazon Alexa, and Waymo's taxi service initiation. This period also saw AI applications in biotech advancing the development of vaccines and other technologies during the COVID pandemic.

After constructing *AI News Index*, I define the series of *AI Innovations*. *AI Innovations* is defined as residuals, ε_t, from an AR(1) process:

$$AINewsIndex_t = \beta * AINewsIndex_{t-1} + \varepsilon_t \tag{21}$$

where t is on a monthly basis. Running this regression gives me a series of 168 full months *AI Innovations* that will be included in my final analysis, from July 2009 to June 2023.

B. Constructing Hedging Portfolio

In this part, I apply the mimicking portfolio approach and construct hedging portfolio for innovations in *RavenPack AI News Index*. To disentangle the AI risk factor, the one condition for mimicking portfolio approach is to spanning all risk factors in the projection portfolios. Obviously this is impossible to do. To mitigate the concern, I not only include a long-short portfolio sorted on *AI Exposure* but also include Fama and French (1993) three factors in constructing the projection portfolio. I explain this in a regression,

$$\begin{aligned} AIInnovations_t = \xi &+ w_{AIExp} Z_{t-1}^{AIExp'} r_t + w_{SIZE} Z_{t-1}^{SIZE'} r_t \\ &+ w_{HML} Z_{t-1}^{HML'} r_t + w_{MKT} Z_{t-1}^{MKT'} r_t + e_t, \end{aligned} \tag{22}$$

where *AI Innovations* are the residual series from an AR(1) regressions on *AI News Index*. For ease of exposition, I follow the notation of Engle et al. (2020). I set $Z_{t-1}^{AIExp'} r_t$ to represent the returns of a long-short portfolio, which longs the top half of firms with higher *AI Exposure* and shorts the bottom half with lower *AI Exposure*. Similarly, I set size (using cross-sectionally standardized market value to create $Z_{t-1}^{SIZE'}$, so that half the firms, sorted by market value, have positive weight, and half have negative weight; note that this portfolio will be long large firms and short small firms), value (using cross-sectionally standardized values of book-to-market to create $Z_{t-1}^{HML'}$), and the market (setting $Z_{t-1}^{MKT'}$ to equal the share of total market value). w_{AIExp}, w_{SIZE}, w_{HML}, and w_{MKT} capture the weights for each factor for the final mimicking portfolio.

C. In-Sample Fit

I first run Equation (22) for the in-sample fit using the full sample from July 2009 to June 2023. Table XIII presents the in-sample regression results. Column (1) runs the *AI Innovations* on the mimicking portfolio only with *AI Exposure* sorted portfolio; Column (2) adds MKT portfolio; Column (3) adds SIZE portfolio; Column (4) adds HML portfolio; Column (5) is the full regression.

Compare these five results, we can observe several things. First, the full regression in Column (5) gives the highest in-sample regression R-squared, 8.19%. Second, full regression in Column (5) gives the *AI Exposure* sorted portfolio highest weights. Third, the only statistically significant weights are the *AI Exposure* sorted portfolio. However, from Column (5), we see other factors help boost the R-squared.

A possible explanation for this relatively low R-squared could result from the fact that the full sample starts from July 2009, when AI was not a big thing. This could also be seen in the Figure 2. Intuitively, we would see a higher weight on the *AI Exposure* sorted portfolio and higher overall R-squared if we run the in-sample fit on the later years from July 2016. Table XIV presents the subsample in-sample regression results from July 2016 to June 2023. The regression results are consistent with the intuition, that both the weight on the *AI Exposure* sorted portfolio and overall R-squared are higher now. The overall R-squared is 10.4%.

In the above regressions, the results show that the *AI Exposure* sorted portfolio really matters. Hence, if we construct the *AI Exposure* sorted portfolio by quartile sorting and long the top quartile and short the bottom quartile, instead of long the top half and short the bottom half, we might get more variations and hence better in-sample fitness. Table XV presents the subsample in-sample regression results from July 2016 to June 2023 and quartile sorting on *AI Exposure* portfolio. In Column (5), both the weight on the *AI Exposure* sorted portfolio and overall R-squared are higher now. The overall R-squared is 19.6%.

These in-sample tests show that investors can always figure out a better subsample to boost the in-sample fitness. While the most important thing an investor should care is the out-of-sample performance. In the next part, I show the out-of-sample performance of the above in-sample models.

D. Out-of-Sample Performance

To test the out-of-sample performance of the above in-sample models, I use the estimated in-sample weights and run regressions with these weights for a later sample period in rolling window style. More specifically, I run regression using data from the initial month, t_{min}, for which I observe $AIExposure$ and $AIInnovations_t$ series, to month $t-1$, for any month t, but making sure each regression has at least 30 months.

Figure 6 presents the out-of-sample performance for each of the in-sample estimated models. The left-hand-side graphs (a)(c)(e) show binscatter plots between the returns of the hedge portfolios and the AI News Index Innovations, while the right-hand-side graphs (b)(d)(f) depict the time series of the two. Panel (a) and (b) corresponds to the in-sample results in Table XIII with full sample. There is a positive, 18.57%, correlation between the constructed mimicking portfolio returns and *AI News Index Innovations*. This means that the AI mimicking portfolio does earn higher returns when the positive AI news innovations materialize. Panel (c) and (d) corresponds to the in-sample results in Table XIV with recent years from July 2016 to June 2023. There is a positive, 20.24%, correlation between the constructed mimicking portfolio returns and *AI News Index Innovations*. Panel (e) and (f) corresponds to the in-sample results in Table XV with recent years from July 2016 to June 2023, and sorting is based on the top quartile minus the bottom quartile. There is a positive, 40.20%, correlation between the constructed mimicking portfolio returns and *AI News Index Innovations*.

The last subsample out-of-sample correlation reaches 40.20%, which is great to investors. However, here I am not claiming this is the optimal hedging subsample or best hedging strategy. My goal is to show it is valid to apply the mimicking portfolio approach in hedging long-run AI risk.

VI. Conclusion

AI has emerged as a transformative force reshaping firms' risk profiles and opportunities. This paper studies two channels, opportunity and regulatory uncertainty, through which AI affects the stock prices and risk premia. I study this question between firms that I call *AI Starters* and *AI Elites*. The *AI Starters* are firms that start to invest or adapt to using AI in their operations, while the *AI Elites* are firms that have invested in AI or used AI technologies in their daily operations for quite some time. The main finding in this paper is that *AI Starters* are more opportunity-driven and *AI Elites* are more concerned about regulatory risks, which leads to *AI Starters* showing a negative AI Exposure-returns relationship while *AI Elites* showing a positive AI Exposure-returns relationship.

The findings are largely robust across several complementary empirical tests. First, standard asset pricing tests, including portfolio sorts on *AI Exposure* measures and factor regressions, show a positive return spread and significant positive alpha for *AI Elites*, while *AI Starters* exhibit a negative return spread. Second, Fama-MacBeth regressions reveal a negative relationship between *AI Exposure* and returns for *AI Starters*, particularly for the opportunity-related $AIExposure^{Opp}$ measure. Third, this negative *AI Exposure*-returns relationship also holds for Implied Cost of Capital estimates. Fourth, the result is robust when using option-implied expected return measures. Fifth, the risk premia mostly disappear following a three-pass procedure, which accounts for omitted variable bias and measurement error, although the heterogeneous effects from the two channels remain. Lastly, using plausibly exogenous regulation and opportunity shocks, I show that the real price reactions are consistent with the main findings.

These results highlight a growing need for investors to account for AI-driven opportunities and regulatory uncertainties in their asset allocation strategies. In the last part of the of the paper, I implement a mimicking portfolio approach to showcase that AI innovations are hedgeable. Future research could further investigate whether there could be an optimal approach in hedging this AI innovations by leveraging the opposing effects of the two channels discussed in this paper. As firms increasingly integrate AI technologies, understanding and mitigating associated risks will be critical for achieving sustainable financial performance. Future research could extend the theoretical framework to explore cross-industry variations in AI risk exposure and assess its impact on asset prices.

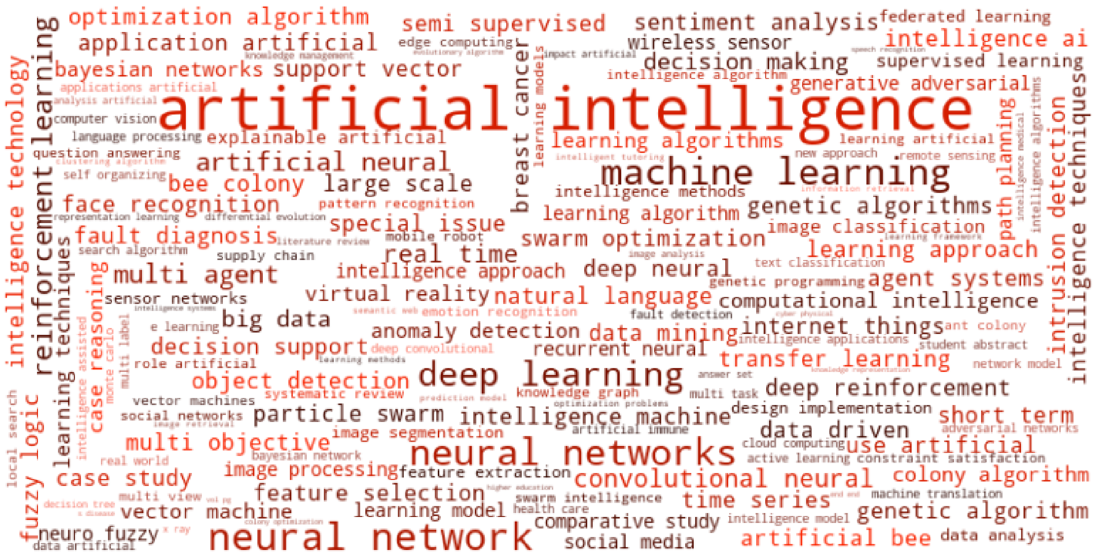

(a) **Artificial Intelligence Bigrams**

(b) **Artificial Intelligence Unigrams**

Figure 1. Artificial Intelligence Vocabulary WordCloud. The figures present word clouds that summarize the *Artificial Intelligence Vocabulary (AIV)* extracted from academic paper titles in the Web of Science. Following the literature, I focus on unigrams and bigrams. The larger the term size, the more frequent the terms appear in the corpus (measured in TF-IDF score). The Web of Science paper titles are downloaded from ProQuest, and cover 159,444 papers specifically related to *artificial intelligence*. The years range from 2010 to 2024. Panel (a) shows most frequent bigrams. Panel (b) shows most frequent unigrams.

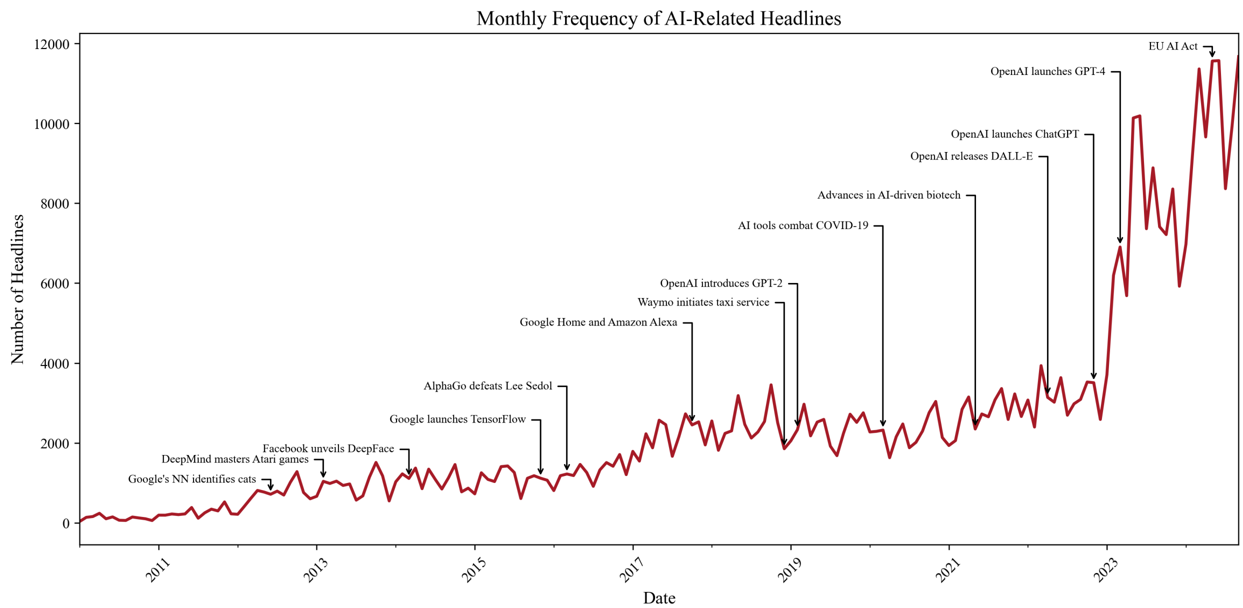

Figure 2. RavenPack AI News Index. The figures present the RavenPack AI News Index from 2010 to 2024, together with main AI-relevant news announcements marked in the plot. The y-axis variable, *Number of Headlines*, represents the monthly count of RavenPack news headlines containing AIV terms. The plot utilizes the Dow Jones and Press Release Edition of the RavenPack News database, which includes approximately 170 million news headlines from 2010 to 2024.

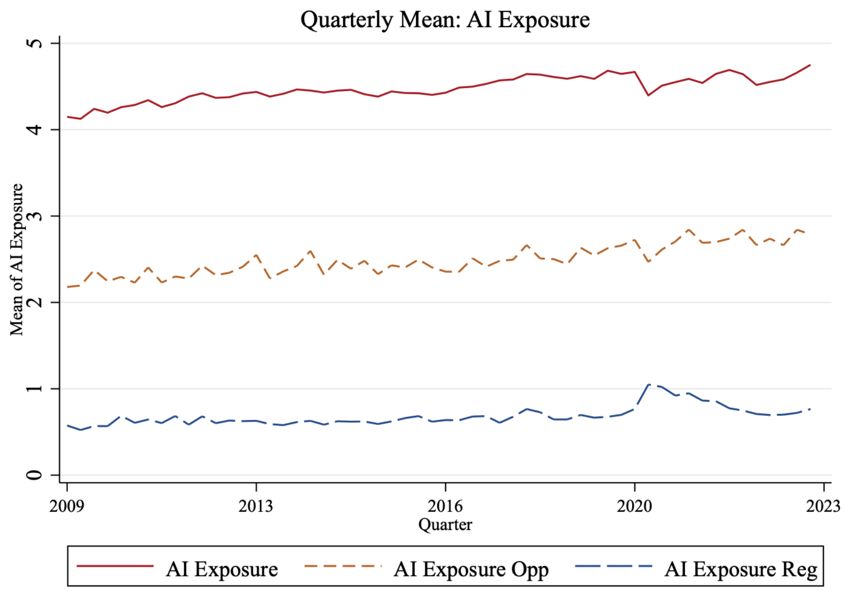

(a) **Monthly Mean of AI Exposure**

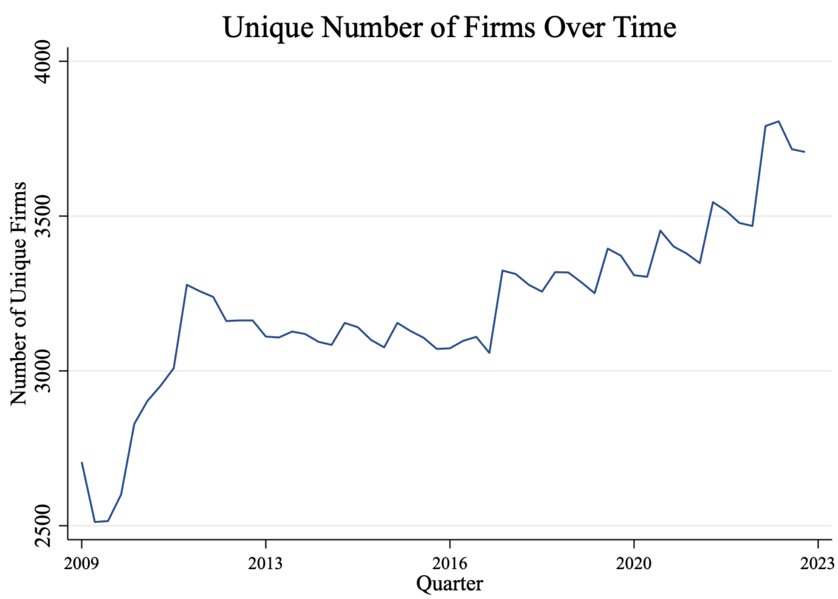

(b) **Unique Number of Firms Over Time**

Figure 3. Summary Statistics of AI Exposure. This figure presents summary statistics of firm-level *AI Exposure* measures on a monthly basis. Panel (a) presents the monthly mean of *AI Exposure*, $AIExposure^{Opp}$, and $AIExposure^{Reg}$, and Panel (b) presents the unique number of firms over time. The dataset spans from July 2009 to June 2023.

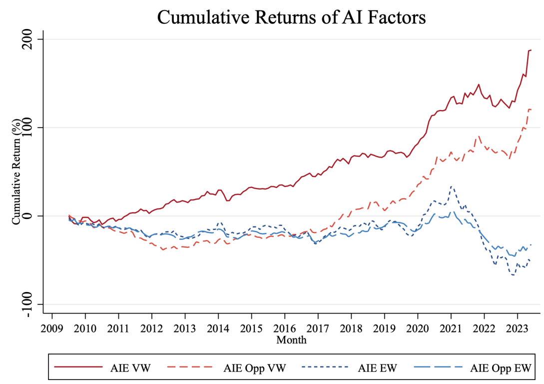

(a) **Cumulative Returns**

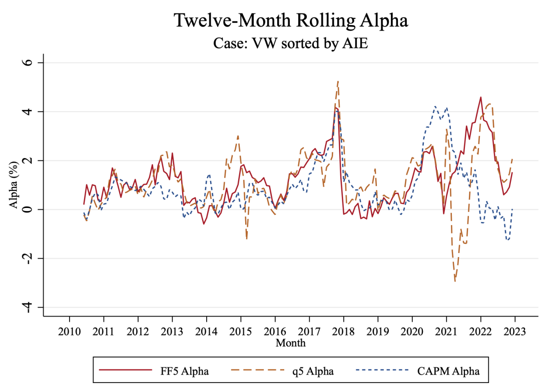

(b) **VW sorted by AIE**

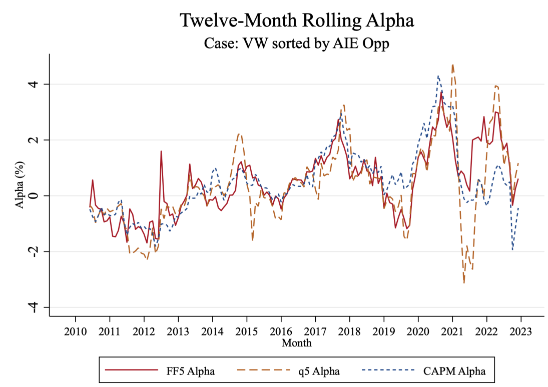

(c) **VW sorted by AIE Opp**

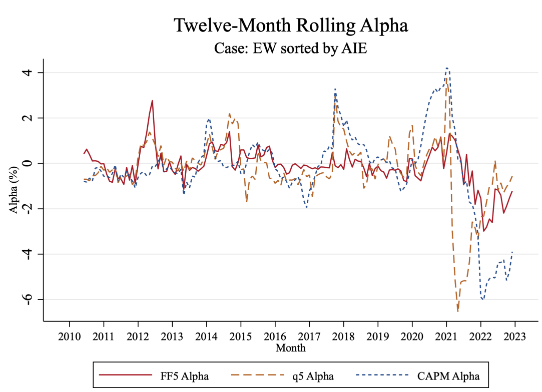

(d) **EW sorted by AIE**

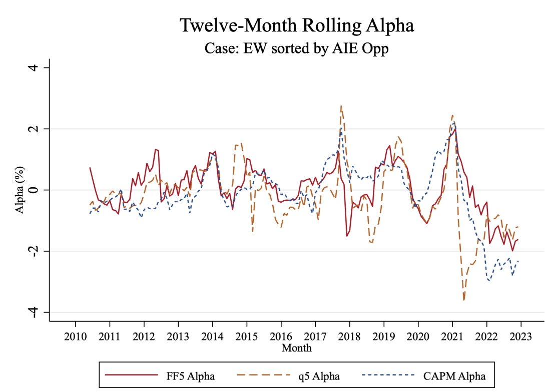

(e) **EW sorted by AIE Opp**

Figure 4. Cumulative Returns and Twelve-Month Rolling Alpha. Panel (a) presents the cumulative returns of the return spread between high- and low- $AIExposure$ or $AIExposure^{Opp}$ portfolios, either VW or EW. Panel (b)-(e) show the twelve-month rolling factor alphas of a strategy that longs the high $AIExposure$ or $AIExposure^{Opp}$ portfolio and short the low $AIExposure$ or $AIExposure^{Opp}$ portfolio.

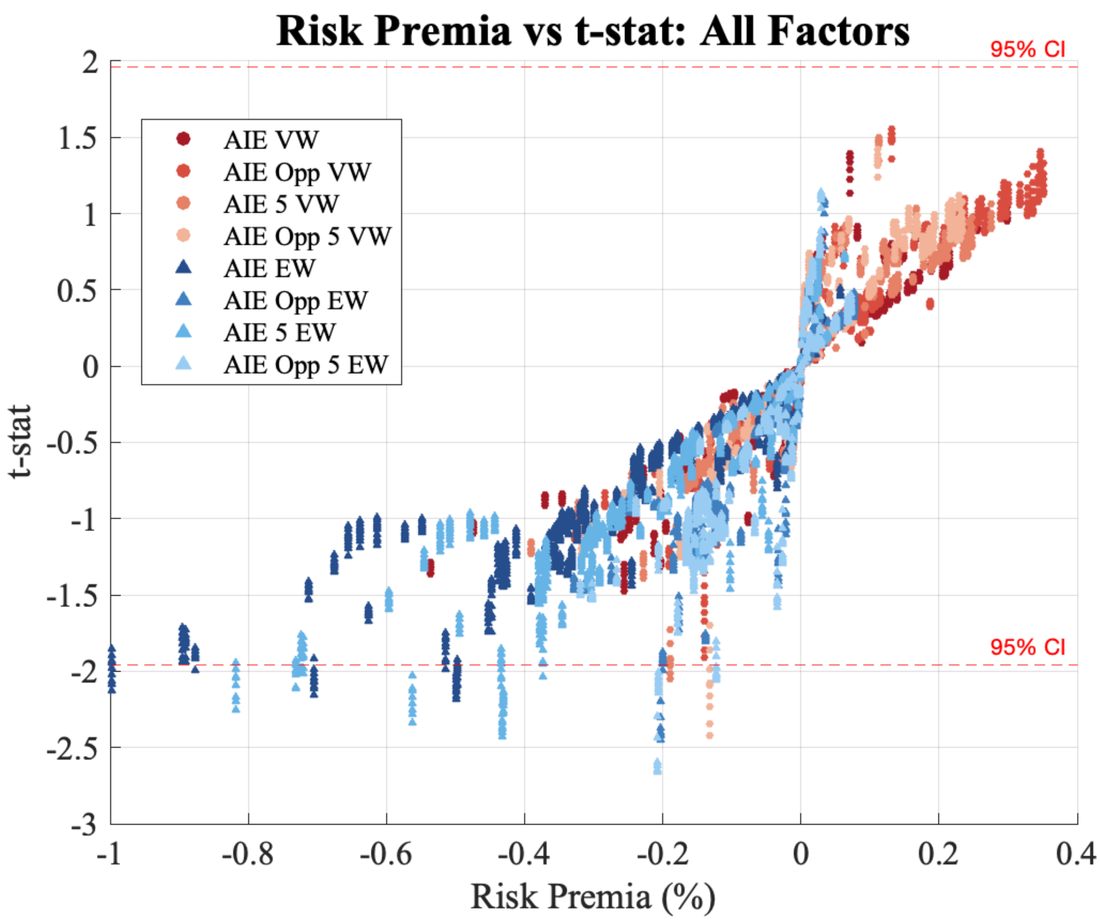

Figure 5. Three-Pass Risk Premia of AI Factors. This figure shows the risk premia estimates from the three-pass procedure for AI factors constructed using firm-level $AIExposure$ measures on a monthly basis, following Giglio and Xiu (2021). The four circle plots represent value-weighted AI factors, while the four triangle plots represent equal-weighted AI factors. "AIE" denotes the return spread between the highest and lowest deciles sorted by $AIExposure$, and "AIE Opp" denotes the return spread between the highest and lowest deciles sorted by $AIExposure^{Opp}$. When the notation "5" is used, it indicates the return spread between the highest and lowest quintiles. For each factor, I estimate the risk premia using different sets of testing portfolios, a range of one to ten lags for computing Newey-West standard errors and a range of one to ten for principal components. The dataset spans from July 2009 to June 2023.

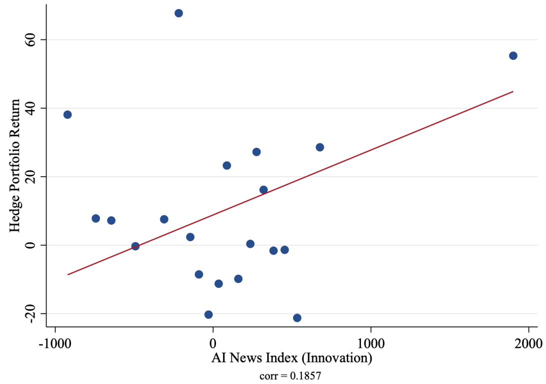

(a) **Binscatter (Full Sample)**

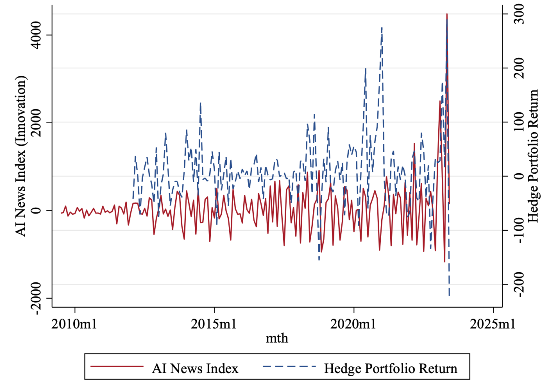

(b) **Time Series: Full Sample**

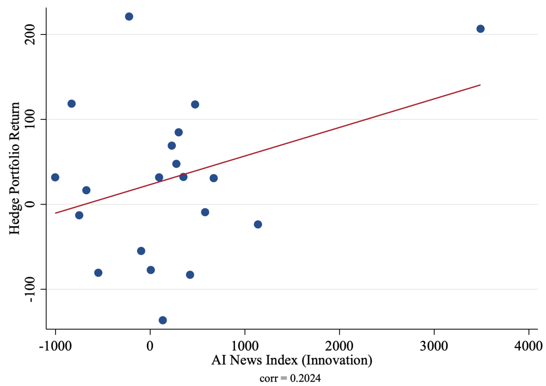

(c) **Binscatter (Subsample: Recent Years)**

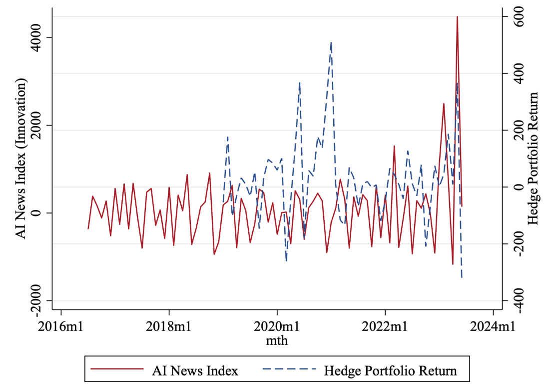

(d) **Time Series (Subsample: Recent Years)**

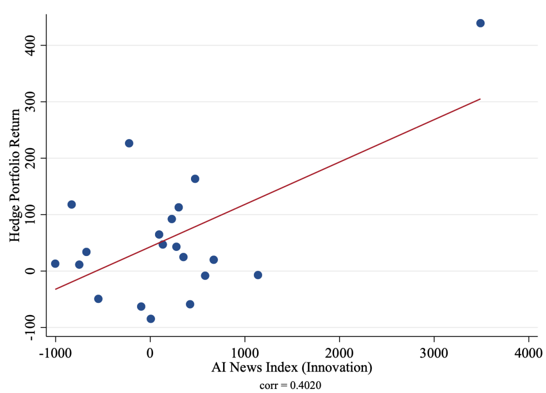

(e) **Binscatter (Subsample: Recent Years + Quartile Mimicking Portfolio)**

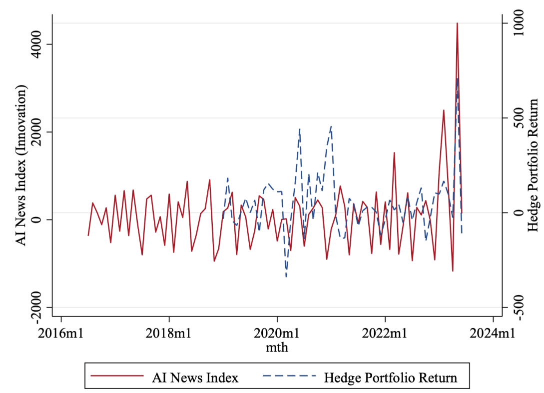

(f) **Time Series (Subsample: Recent Years + Quartile Mimicking Portfolio)**

Figure 6. Out-of-sample Performance. This figure presents out-of-sample performance of hedge portfolios constructed to hedge the RavenPack AI News Index. The left-hand graphs (a)(c)(e) show binscatter plots between the returns of the hedge portfolios and the AI News Index Innovations, while the right-hand graphs (b)(d)(f) depict the time series of the two. The full sample analysis covers months from July 2009 to June 2023, with a subsample with recent years from July 2016 to June 2023. For the full sample, the mimicking portfolios are sorted by the top half minus the bottom half; for the subsample, sorting is based on the top quartile minus the bottom quartile.

Table I
Summary Statistics and Correlation: AI Exposure Measures

This table presents the summary statistics and correlation of the overall firm-level *AI Exposure* and different versions of topic-specific or tone-specific *AI Exposure* measures at the firm-quarter level. Panel A presents the summary statistics of different versions of *AI Exposure* measures. Panel B presents the correlation among different versions of *AI Exposure* measures. I denote the overall *AI Exposure* measure, the *AI Exposure* regarding opportunity, the *AI Exposure* regarding regulation, the *AI Exposure* regarding risk, the *AI Exposure* with positive tones, and the *AI Exposure* with negative tones as $AIExposure$, $AIExposure^{Opp}$, $AIExposure^{Reg}$, $AIExposure^{Risk}$, $AIExposure^{Pos}$, and $AIExposure^{Neg}$, respectively. For the ease of exposition, I multiply 1,000 for each measure.

Panel A: Summary Statistics

Variable	Observation	Mean	SD	10%	25%	50%	75%	90%	Max
$AIExposure_{i,t}$	391,232	4.40	1.73	2.55	3.22	4.11	5.27	6.60	24.83
$AIExposure_{i,t}^{Opp}$	391,232	2.44	2.64	0.00	0.00	2.41	4.44	5.91	24.83
$AIExposure_{i,t}^{Reg}$	391,232	0.78	1.87	0.00	0.00	0.00	0.00	3.99	24.83
$AIExposure_{i,t}^{Risk}$	391,232	0.90	1.88	0.00	0.00	0.00	0.00	4.14	22.42
$AIExposure_{i,t}^{Pos}$	391,232	0.86	1.97	0.00	0.00	0.00	0.00	4.22	23.71
$AIExposure_{i,t}^{Neg}$	391,232	0.17	0.88	0.00	0.00	0.00	0.00	0.00	19.63

Panel B: Correlation

	$AIExposure_{i,t}$	$AIExposure_{i,t}^{Opp}$	$AIExposure_{i,t}^{Reg}$	$AIExposure_{i,t}^{Risk}$	$AIExposure_{i,t}^{Pos}$	$AIExposure_{i,t}^{Neg}$
$AIExposure_{i,t}$	1.000					
$AIExposure_{i,t}^{Opp}$	0.478	1.000				
$AIExposure_{i,t}^{Reg}$	0.203	0.240	1.000			
$AIExposure_{i,t}^{Risk}$	0.119	0.192	0.355	1.000		
$AIExposure_{i,t}^{Pos}$	0.241	0.176	0.081	0.052	1.000	
$AIExposure_{i,t}^{Neg}$	0.033	0.031	0.026	0.028	0.016	1.000

Table II
Characteristics of AI Exposure-Sorted Portfolios (Quintile)

This table reports the summary statistics for characteristics of *AI Exposure*-sorted quintile portfolios. At the end of June of each year, stocks are ranked by their *AI Exposure* and sorted into quintiles. *AI Exposure* is defined as the fraction of AI-related unigrams and bigrams in a firm's call transcripts, and here it is multiplied by 1000 for ease of exposition. Firm size, *ln(ME)*, is measured as the log of market equity in June of year t. *BE/ME* is measured as the ratio between book equity at the end of June of year t and market equity at the end of December of year $t-1$. *ROE* is profitability, measured as income before extraordinary items divided by book equity. *ROA* is another profitability measure, defined as income before extraordinary items divided by total assets. *R&D/Assets* is defined as the ratio of R&D to lagged total assets. *Firm Age* measures the age of a firm starting from its initial listing in the CRSP database. *Asset Growth* is defined as the ratio between the change of total assets from year $t-2$ to year $t-1$ and total assets in year $t-2$. I follow Fahlenbrach, Rageth, and Stulz (2021) to construct the following three financial flexibility measures. *Cash/Assets* is the ratio of cash to total assets. *St Debt/Assets* is the ratio of debt in current liabilities to total assets. *Lt Debt/Assets* is the ratio of long-term debt to total assets. *Tangibility* is property, plant, and equipment divided by total assets. *Book Leverage* is the sum of current liabilities and long-term debt divided by total assets. *WW Index* is the Whited and Wu index used to measure financial constraint (see, Whited and Wu, 2006). After winsorizing at the 1st and 99th percentiles to limit the influence of outliers, all firm characteristics variables are standardized to have a mean of zero and a standard deviation of one. The sample period is 2009 to 2024 at an annual frequency.

Quintile	1	2	3	4	5
	mean	mean	mean	mean	mean
$AIExposure$	-1.14	-0.62	-0.18	0.37	1.57
$AIExposure^{Opp}$	-0.76	-0.37	-0.08	0.30	0.92
$AIExposure^{Reg}$	-0.32	-0.14	-0.01	0.11	0.36
$AIExposure^{Risk}$	-0.23	-0.03	0.04	0.09	0.13
$AIExposure^{Pos}$	-0.45	-0.22	-0.04	0.19	0.53
$AIExposure^{Neg}$	-0.07	-0.01	0.04	0.05	-0.00
Size	0.05	0.10	0.06	-0.02	-0.19
BE/ME	0.05	0.03	-0.01	-0.04	-0.02
ROA	0.25	0.23	0.15	-0.03	-0.60
ROE	0.19	0.15	0.08	-0.05	-0.38
R&D/Assets	-0.39	-0.33	-0.19	0.13	0.78
Firm Age	0.21	0.15	0.06	-0.10	-0.31
Asset Growth	-0.09	-0.07	-0.02	0.04	0.14
Cash/Assets	-0.46	-0.36	-0.18	0.16	0.83
St Debt/Assets	0.10	0.01	-0.01	-0.03	-0.07
Lt Debt/Assets	0.17	0.18	0.08	-0.07	-0.37
Tangibility	0.21	0.17	0.07	-0.10	-0.32
Book Leverage	0.20	0.17	0.07	-0.08	-0.37
WW Index	-0.00	0.01	0.01	0.02	-0.04

Table III
AI Exposure by Industry (SIC2)
Full Sample: July 2009 to June 2023

This table shows the average *AI Exposure* for the top ten industries based on their SIC 2-digit codes. Panel A lists the top ten industries with the highest overall *AI Exposure*, while Panel B shows the top ten industries ranked by the *AI Exposure* measure based on opportunities ($AIExposure^{Opp}$).

Panel A: $AIExposure_{i,t}$

Industry (SIC2)	N	Mean	SD	Median
Educational Services	1,582	6.29	2.40	5.82
Insurance Agents, Brokers, & Service	887	5.56	3.00	4.59
Local & Interurban Passenger Transit	176	5.50	2.33	4.98
Business Services	29,793	5.46	1.91	5.25
Chemical & Allied Products	24,532	5.40	2.12	5.14
Instruments & Related Products	12,424	5.38	2.05	5.03
Electronic & Other Electric Equipment	15,788	5.05	1.91	4.82
Engineering & Management Services	3,582	4.99	1.79	4.77
Eating & Drinking Places	2,925	4.85	1.43	4.72
Communications	7,630	4.80	1.91	4.57

Panel B: $AIExposure^{Opp}_{i,t}$

Industry (SIC2)	N	Mean	SD	Median
Insurance Agents, Brokers, & Service	887	3.89	3.84	3.75
Business Services	29,793	3.78	3.08	4.37
Educational Services	1,582	3.75	3.88	4.04
Instruments & Related Products	12,424	3.62	3.07	3.99
Local & Interurban Passenger Transit	176	3.36	3.41	3.72
Engineering & Management Services	3,582	3.30	2.89	3.72
Miscellaneous Manufacturing Industries	1,369	3.30	2.41	3.70
Security & Commodity Brokers	5,094	3.21	2.55	3.72
Electronic & Other Electric Equipment	15,788	3.07	2.99	3.43
Industrial Machinery & Equipment	9,478	2.95	2.74	3.15

Table IV
Portfolio Sorts and Returns

This table presents average monthly value-weighted (VW) or equal-weighted (EW) returns (in percentage) of portfolios sorted by overall firm-level *AI Exposure* and $AIExposure^{Opp}$. Panel A presents VW returns of decile portfolios of different versions of *AI Exposure* measures. H-L is the return spread between the highest decile and the lowest decile. Panel B presents VW returns of quintile portfolios of different versions of *AI Exposure* measures. Panel C presents EW returns of decile portfolios of different versions of *AI Exposure* measures. Panel D presents EW returns of quintile portfolios of different versions of *AI Exposure* measures. The t-statistics are computed using heteroscedasticity and autocorrelation consistent Newey-West (1987) standard error estimates with a lag length of 12 months.

Panel A. VW Returns by Deciles

$AIExposure$	L	2	3	4	5	6	7	8	9	H	H-L
Mean	1.47	1.50	1.59	1.37	1.51	1.43	1.69	1.80	2.07	2.43	0.96
[t]	5.54	4.94	6.65	5.62	5.21	5.44	5.65	6.06	5.83	5.06	2.65
$AIExposure^{Opp}$	L	2	3	4	5	6	7	8	9	H	H-L
Mean	1.74	1.21	1.46	1.49	1.58	1.38	1.51	1.77	2.03	2.34	0.60
[t]	6.05	4.32	5.13	5.47	6.56	4.89	6.09	6.01	6.83	5.00	1.67

Panel B. VW Returns by Quintiles

$AIExposure$	L	2	3	4	H	H-L
Mean	1.49	1.49	1.50	1.77	2.26	0.77
[t]	5.27	6.27	5.67	6.13	5.77	2.56
$AIExposure^{Opp}$	L	2	3	4	H	H-L
Mean	1.63	1.47	1.48	1.63	2.19	0.56
[t]	6.07	5.39	5.85	6.15	6.09	2.08

Panel C. EW Returns by Deciles

$AIExposure$	L	2	3	4	5	6	7	8	9	H	H-L
Mean	1.23	1.10	1.26	1.19	1.13	1.04	1.06	0.92	1.04	0.91	-0.32
[t]	2.69	2.41	2.69	2.74	2.57	2.31	2.19	1.87	2.00	1.43	-0.73
$AIExposure^{Opp}$	L	2	3	4	5	6	7	8	9	H	H-L
Mean	1.07	1.08	1.19	1.10	1.22	1.08	1.06	1.15	0.93	0.88	-0.19
[t]	2.17	2.15	2.64	2.46	2.78	2.45	2.42	2.31	1.75	1.50	-0.83

Panel D. EW Returns by Quintiles

$AIExposure$	L	2	3	4	H	H-L
Mean	1.16	1.22	1.08	0.99	0.97	-0.19
[t]	2.56	2.72	2.45	2.05	1.70	-0.53
$AIExposure^{Opp}$	L	2	3	4	H	H-L
Mean	1.14	1.14	1.15	1.11	0.90	-0.24
[t]	2.40	2.55	2.62	2.37	1.63	-1.03

Table V
Sequential Sorts and Returns

This table presents average monthly value-weighted (VW) returns (in percentage) of sequentially sorted portfolios, first sorted by anomalies, and then sorted by overall firm-level *AI Exposure*. In each panel, I report sequential sort returns using full sample, and subsample that includes the second half of the full sample, i.e., recent years. The t-statistics are computed using heteroscedasticity and autocorrelation consistent Newey-West (1987) standard error estimates with a lag length of 12 months.

Panel A. Size × AIE Sequential Sort

Full Sample	L	2	3	4	H	H-L
Mean	1.52	1.55	1.45	1.64	1.91	0.40
[t]	5.39	6.02	5.43	6.50	5.71	1.58
Recent Years	L	2	3	4	H	H-L
Mean	1.45	1.49	1.40	1.63	2.25	0.80
[t]	3.37	3.83	3.11	4.01	3.81	1.90

Panel B. B/M × AIE Sequential Sort

Full Sample	L	2	3	4	H	H-L
Mean	1.54	1.50	1.54	1.81	2.18	0.64
[t]	5.81	6.15	5.55	6.37	5.67	2.60
Recent Years	L	2	3	4	H	H-L
Mean	1.46	1.50	1.58	1.86	2.61	1.16
[t]	3.64	3.69	3.63	3.88	3.87	2.94

Panel C. R&D × AIE Sequential Sort

Full Sample	L	2	3	4	H	H-L
Mean	1.63	1.56	1.59	1.55	1.91	0.28
[t]	6.28	5.58	6.12	5.98	5.59	1.62
Recent Years	L	2	3	4	H	H-L
Mean	1.66	1.52	1.50	1.68	2.22	0.57
[t]	4.01	3.39	3.44	4.16	3.74	2.03

Panel D. Age × AIE Sequential Sort

Full Sample	L	2	3	4	H	H-L
Mean	1.53	1.58	1.50	1.61	1.98	0.45
[t]	5.50	5.69	5.97	5.87	6.24	1.96
Recent Years	L	2	3	4	H	H-L
Mean	1.44	1.57	1.47	1.61	2.33	0.89
[t]	3.39	3.57	3.50	3.60	4.26	2.49

Panel E. ROE × AIE Sequential Sort

Full Sample	L	2	3	4	H	H-L
Mean	1.51	1.51	1.52	1.63	1.98	0.47
[t]	5.13	6.24	5.47	6.41	6.08	2.07
Recent Years	L	2	3	4	H	H-L
Mean	1.43	1.50	1.45	1.72	2.25	0.82
[t]	3.14	3.89	3.02	4.45	3.89	2.10

Panel F. AG × AIE Sequential Sort

Full Sample	L	2	3	4	H	H-L
Mean	1.51	1.48	1.57	1.63	2.16	0.65
[t]	5.66	5.71	5.76	6.09	6.13	2.49
Recent Years	L	2	3	4	H	H-L
Mean	1.43	1.45	1.57	1.65	2.57	1.14
[t]	3.66	3.35	3.41	4.02	4.13	2.69

Panel G. CA × AIE Sequential Sort

Full Sample	L	2	3	4	H	H-L
Mean	1.58	1.67	1.54	1.54	1.83	0.25
[t]	5.89	5.88	5.87	5.67	6.06	1.66
Recent Years	L	2	3	4	H	H-L
Mean	1.54	1.63	1.56	1.67	2.03	0.49
[t]	3.64	3.58	3.53	3.74	3.79	2.13

Panel H. BL × AIE Sequential Sort

Full Sample	L	2	3	4	H	H-L
Mean	1.49	1.58	1.43	1.69	2.07	0.58
[t]	5.54	6.00	5.40	6.01	6.16	2.27
Recent Years	L	2	3	4	H	H-L
Mean	1.40	1.54	1.40	1.80	2.40	1.00
[t]	3.47	3.59	3.21	3.85	4.10	2.31

Panel I. TANT × AIE Sequential Sort

Full Sample	L	2	3	4	H	H-L
Mean	1.49	1.59	1.52	1.74	1.98	0.49
[t]	5.52	6.22	5.82	6.18	5.33	1.96
Recent Years	L	2	3	4	H	H-L
Mean	1.43	1.47	1.61	1.82	2.32	0.89
[t]	3.43	3.64	3.76	3.98	3.49	2.02

Table VI
Regressions for H-L Portfolios Sorted on AI Exposure

This table reports the estimated monthly α (basis points) and coefficients of regression of dependent variable *H-L* (based on VW returns) on a sets of leading factors (see, Chabi-Yo, Gonçalves, and Loudis, 2025; Fama and French, 1996; Fama and French, 2015; Hou, Xue, and Zhang, 2015; Hou et al., 2020). *H-L* is the monthly return spread between firms with High *AI Exposure* (10th decile) and Low *AI Exposure* (1st decile). Column (1) reports the CAPM; Column (2) reports the ICAPM; Column (3) reports Fama-French 3 factors; Column (4) reports Fama-French 3 factors plus Carhart (1997) momentum factor-UMD; Column (5) reports Fama-French 5 factors; Column (6) reports q factors; Column (7) reports q5 factors. The sample period is from July 2009 to June 2023 at an monthly frequency. The t-statistics are based on standard errors estimated using the Newey-West correction for 12 lags. ***, **, and * correspond to statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1) CAPM	(2) ICAPM	(3) FF3	(4) FF4	(5) FF5	(6) q	(7) q5
α	66.88**	72.92**	67.90***	66.62***	81.19***	99.09***	83.58***
	(2.25)	(2.58)	(3.22)	(3.24)	(3.65)	(3.85)	(3.61)
MktRF	11.80	5.726	10.63	11.22	8.800	0.871	4.244
	(1.55)	(0.73)	(1.51)	(1.51)	(1.19)	(0.10)	(0.55)
$r_{\mathbb{E}}$		-14.96					
		(-0.67)					
$r_{\mathbb{V}}$		6.986					
		(0.34)					
SMB			30.47*	30.71*	23.82		
			(1.94)	(1.93)	(1.61)		
HML			-63.20***	-62.22***	-47.44***		
			(-10.91)	(-11.72)	(-5.21)		
UMD				2.767			
				(0.32)			
RMW					-13.22		
					(-0.76)		
CMA					-34.07**		
					(-2.09)		
q^{ME}						14.68	26.79*
						(0.96)	(1.66)
q^{IA}						-79.78***	-59.95***
						(-8.28)	(-5.75)
q^{ROE}						-14.49	-40.18***
						(-1.59)	(-3.55)
q^{EG}							51.92***
							(4.22)
R-squared	0.023	0.035	0.359	0.360	0.386	0.304	0.351
Observations	162	162	162	162	162	162	162

Table VII
Fama-MacBeth Regressions on AI Exposure
Full Sample: July 2009 to June 2023

This table reports Fama-MacBeth regressions of individual stock excess returns (in percentage) on *AI Exposure* and other firm characteristics. Cross-sectional regressions are conducted monthly from July of year t to June of year $t+1$. Each month, individual stock returns are regressed on *AI Exposure* from year $t-1$ and various control variables available by the end of June of year t. The list of control variables are defined as Table II. Column (1)-(4) report Ordinary Least Square (OLS) regression results. Column (5)-(8) report Weighted Least Square (WLS) regression results using market equity as the weights. The t-statistics are based on standard errors adjusted using the Newey-West method. The full sample is from July 2009 to June 2023.

	OLS				WLS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI Exposure	-0.0775**				0.0806			
	(-2.19)				(1.32)			
AI Exposure Opp		-0.104***				0.102		
		(-4.41)				(1.39)		
AI Exposure Reg			-0.0566**				-0.0593**	
			(-2.06)				(-1.86)	
AI Exposure Risk				-0.0283				0.0113
				(-0.91)				(0.38)
Size	-0.0887	-0.0809	-0.0817	-0.0838	-0.0764	-0.0771	-0.0756	-0.0768
	(-1.62)	(-1.50)	(-1.51)	(-1.57)	(-1.40)	(-1.41)	(-1.39)	(-1.41)
BE/ME	0.139	0.140	0.143	0.138	0.132	0.131	0.129	0.130
	(0.96)	(0.96)	(0.98)	(0.94)	(0.91)	(0.90)	(0.88)	(0.89)
ROE	0.540***	0.543***	0.545***	0.546***	-0.151	-0.155	-0.153	-0.153
	(4.11)	(4.13)	(4.13)	(4.14)	(-1.00)	(-1.03)	(-1.03)	(-1.03)
R&D/Assets	0.154**	0.147**	0.139**	0.137**	0.295**	0.324***	0.332***	0.323***
	(2.61)	(2.50)	(2.38)	(2.35)	(2.59)	(2.94)	(2.92)	(2.83)
Firm Age	-0.0682	-0.0686	-0.0685	-0.0682	-0.0749	-0.0752	-0.0750	-0.0747
	(-1.50)	(-1.48)	(-1.46)	(-1.45)	(-1.60)	(-1.61)	(-1.59)	(-1.58)
Asset Growth	-0.0693	-0.0707	-0.0662	-0.0662	-0.0317	-0.0344	-0.0287	-0.0306
	(-1.23)	(-1.25)	(-1.18)	(-1.15)	(-0.43)	(-0.47)	(-0.40)	(-0.41)
Cash/Assets	0.349***	0.350***	0.374***	0.374***	0.526***	0.545***	0.582***	0.577***
	(4.29)	(4.33)	(4.13)	(4.35)	(4.28)	(4.39)	(4.26)	(4.31)
St Debt/Assets	-0.0638*	-0.0652*	-0.0599*	-0.0597*	-0.0917*	-0.0964*	-0.0902*	-0.0898*
	(-1.43)	(-1.47)	(-1.36)	(-1.36)	(-1.80)	(-1.88)	(-1.79)	(-1.76)
Lt Debt/Assets	0.0829	0.0865	0.0863	0.0898	0.0508	0.0544	0.0542	0.0595
	(1.43)	(1.48)	(1.46)	(1.51)	(0.91)	(0.97)	(0.96)	(1.05)
Tangibility	0.00989	0.0103	0.00700	0.00810	0.0493	0.0498	0.0451	0.0456
	(0.12)	(0.12)	(0.09)	(0.10)	(0.50)	(0.51)	(0.47)	(0.48)
Constant	1.724***	1.452**	1.437**	1.827***	2.265***	2.223**	2.136**	2.138**
	(3.16)	(2.57)	(2.53)	(3.08)	(2.64)	(2.57)	(2.48)	(2.34)
Industry (SIC2) FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	528170	528170	528170	528170	528170	528170	528170	528170
R-squared	0.286	0.286	0.284	0.285	0.286	0.285	0.284	0.285

Table VIII
Fama-MacBeth Regressions on AI Exposure
Subsample: July 2016 to June 2023

This table reports Fama-MacBeth regressions of individual stock excess returns (in percentage) on *AI Exposure* and other firm characteristics. Cross-sectional regressions are conducted monthly from July of year t to June of year $t+1$. Each month, individual stock returns are regressed on *AI Exposure* from year $t-1$ and various control variables available by the end of June of year t. The list of control variables are defined as Table II. Column (1)-(4) report Ordinary Least Square (OLS) regression results. Column (5)-(8) report Weighted Least Square (WLS) regression results using market equity as the weights. The t-statistics are based on standard errors adjusted using the Newey-West method. The subsample is from July 2016 to June 2023.

	OLS				WLS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI Exposure	-0.0867				0.184*			
	(-1.52)				(1.96)			
AI Exposure Opp		-0.109*				0.216**		
		(-2.47)				(2.48)		
AI Exposure Reg			-0.0420				0.000160	
			(-1.01)				(0.01)	
AI Exposure Risk				0.0121				0.0570
				(0.24)				(1.57)
Size	-0.0146	-0.00903	-0.0128	-0.0158	-0.0314	-0.0340	-0.0266	-0.0265
	(-0.22)	(-0.14)	(-0.19)	(-0.24)	(-0.98)	(-1.06)	(-0.81)	(-0.81)
BE/ME	0.0342	0.0342	0.0327	0.0282	-0.00613	0.00511	0.0129	0.0129
	(0.65)	(0.65)	(0.62)	(0.54)	(-0.06)	(0.05)	(0.13)	(0.13)
ROE	0.483*	0.486*	0.489*	0.487*	0.00345	0.000945	0.0165	0.00237
	(2.58)	(2.57)	(2.58)	(2.58)	(0.02)	(0.00)	(0.09)	(0.01)
R&D/Assets	0.130	0.122	0.116	0.117	0.414**	0.453***	0.484***	0.473***
	(1.39)	(1.27)	(1.22)	(1.24)	(2.61)	(2.84)	(3.00)	(2.89)
Firm Age	-0.00410	-0.00291	-0.00206	-0.00366	-0.133*	-0.137*	-0.137*	-0.136*
	(-0.04)	(-0.03)	(-0.02)	(-0.03)	(-1.72)	(-1.74)	(-1.73)	(-1.70)
Asset Growth	-0.122	-0.121	-0.121	-0.123	-0.105	-0.104	-0.0995	-0.0988
	(-2.07)	(-2.06)	(-2.07)	(-2.07)	(-1.19)	(-1.17)	(-1.15)	(-1.07)
Cash/Assets	0.259*	0.244	0.238	0.232	0.454***	0.460***	0.508***	0.512***
	(2.16)	(1.95)	(1.88)	(1.82)	(3.18)	(3.24)	(3.19)	(3.50)
St Debt/Assets	-0.0906	-0.0917	-0.0880	-0.0911	0.00387	0.000215	0.00889	0.00969
	(-1.19)	(-1.22)	(-1.17)	(-1.21)	(0.07)	(0.00)	(0.16)	(0.18)
Lt Debt/Assets	0.0805	0.0879	0.0915	0.0917	0.134**	0.142***	0.145***	0.147***
	(1.24)	(1.29)	(1.33)	(1.34)	(2.58)	(2.81)	(2.80)	(2.79)
Tangibility	0.150	0.144	0.152	0.150	0.146	0.148	0.137	0.137
	(1.32)	(1.28)	(1.33)	(1.34)	(1.38)	(1.41)	(1.36)	(1.39)
Constant	0.731	1.120	0.469	0.875	1.091	0.804	0.948	1.181
	(0.60)	(1.34)	(0.42)	(0.95)	(1.61)	(1.09)	(1.31)	(1.54)
Industry (SIC2) FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	275346	275346	275346	275346	275343	275343	275343	275343
R-squared	0.107	0.107	0.106	0.106	0.308	0.307	0.306	0.306

Table IX
Multivariate Fama-MacBeth Regressions of Realized Returns and ICC Measures on AI Exposure

This table reports mutivariate Fama-MacBeth regressions of ex-post realized returns (in percentage) and Implied Cost of Capital (ICC) measures (in percentage) on *AI Exposure*. Column (1) reports regressions of realized returns on *AI Exposure*; column (2) reports results of ICC estimates based on Gebhardt, Lee, and Swaminathan (2001); column (3) reports results of ICC estimates based on Claus and Thomas (2001); column (4) reports results of ICC estimates based on Ohlson and Juettner-Nauroth (2005); column (5) reports results of ICC estimates based on Easton and Monahan (2005); column (6) reports results of ICC estimates that are computed as the average of the previous four ICC estimates. The detailed ICC estimates construction could be found in Hou, van Dijk, and Zhang (2012). The list of control variables are the same as Table VII. For the ease of exposition, their estimated coefficients are not shown here.

Panel A: $AIExposure$

	(1) Realized Ret	(2) GLS	(3) CT	(4) OJ	(5) MPEG	(6) Composite
AI Exposure	-0.0775**	-0.0128	0.0443	0.332***	1.348***	0.649***
	(-2.19)	(-0.17)	(1.02)	(5.51)	(4.80)	(3.32)
Industry (SIC2) FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	528174	302464	255532	208890	233935	335901
R-squared	0.106	0.332	0.325	0.255	0.388	0.339

Panel B: $AIExposure^{Opp}$

	(1) Realized Ret	(2) GLS	(3) CT	(4) OJ	(5) MPEG	(6) Composite
AI Exposure Opp	-0.104***	-0.0869*	-0.113***	-0.0821*	-0.0139	-0.0813*
	(-4.41)	(-1.69)	(-4.32)	(-1.72)	(-0.15)	(-1.83)
Industry (SIC2) FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	528174	302464	255532	208890	233935	335901
R-squared	0.106	0.332	0.325	0.254	0.383	0.336

Table X
Portfolio Sorts and Option-Implied Returns

This table presents average monthly value-weighted (VW) option-implied returns (in percentage) of portfolios sorted by overall firm-level *AI Exposure*. The option-implied expected returns are estimated following Martin and Wagner (2019). Panel A presents VW average monthly 1-month returns of decile portfolios. H-L is the return spread between the highest decile and the lowest decile. Panel B presents VW average monthly 3-month returns of decile portfolios. Panel C presents VW average monthly 6-month returns of decile portfolios. Panel D presents VW average monthly 9-month returns of decile portfolios. Panel E presents VW average monthly 12-month returns of decile portfolios.

Panel A. Average Monthly 1-month Expected Returns

$AIExposure$	L	2	3	4	5	6	7	8	9	H	H-L
Mean	3.79	4.21	4.10	3.81	3.61	3.95	3.81	3.85	4.11	4.23	0.45
[t]	6.98	6.80	6.99	7.82	7.11	7.44	6.83	7.75	8.24	8.35	1.79

Panel B. Average Monthly 3-month Expected Returns

$AIExposure$	L	2	3	4	5	6	7	8	9	H	H-L
Mean	3.86	4.19	4.10	3.90	3.68	4.01	3.91	3.97	4.20	4.43	0.57
[t]	7.77	7.61	8.02	8.86	8.22	8.22	7.91	8.94	9.37	8.80	1.88

Panel C. Average Monthly 6-month Expected Returns

$AIExposure$	L	2	3	4	5	6	7	8	9	H	H-L
Mean	4.00	4.33	4.21	4.04	3.82	4.13	4.05	4.12	4.34	4.59	0.59
[t]	8.57	8.43	9.14	10.09	9.71	9.39	8.95	10.13	10.66	9.91	1.91

Panel D. Average Monthly 9-month Expected Returns

$AIExposure$	L	2	3	4	5	6	7	8	9	H	H-L
Mean	3.51	3.79	3.67	3.62	3.38	3.67	3.64	3.71	3.95	4.20	0.69
[t]	8.30	8.74	10.32	10.14	10.57	9.84	9.17	10.72	11.85	9.59	1.77

Panel E. Average Monthly 12-month Expected Returns

$AIExposure$	L	2	3	4	5	6	7	8	9	H	H-L
Mean	3.47	3.72	3.63	3.59	3.35	3.62	3.63	3.68	3.92	4.20	0.73
[t]	8.46	9.13	11.10	10.68	11.35	10.48	9.48	11.28	12.57	10.05	1.85

Table XI
Event Studies: Cumulative Abnormal Returns (CAR)
Political and Innovation Shocks

This table presents cumulative abnormal returns (in percentage) post several AI related political and innovation stocks sorted into *AI Exposure*-sorted portfolios. The table reports daily cumulative abnormal returns over a 10-day window from day 0 (event day) to day 11 after the election. These cumulative abnormal returns are equally weighted across *AI Exposure*-sorted portfolios. I follow the typical way by using a 250-trading day window that ends 25 days before the event day to estimate CAPM beta. I focus on the period from the day of the U.S. presidential election to ten days post-election. I adjust the Cumulative Abnormal Return (CAR) of each stock for market trends using daily market returns and the risk-free rate obtained from the French website.

Panel A: Trump Election Win on Nov 8, 2016

CAR [1,10]	1	2	3	4	5	H - L
Daily Return	4.18	3.00	2.42	2.64	3.43	-0.75
[t]	5.96	6.07	6.42	6.44	7.19	-0.89

Panel B: Biden Election Win on Nov 3, 2020

CAR [1,10]	1	2	3	4	5	H - L
Daily Return	1.84	2.45	1.58	1.55	1.57	-0.27
[t]	5.19	4.91	4.20	3.17	3.26	-0.45

Panel C: Biden AI EO on Oct 30, 2023

CAR [1,10]	1	2	3	4	5	H - L
Daily Return	-2.22	-3.31	-2.59	-3.63	-4.06	-1.84
[t]	-5.60	-8.38	-5.01	-6.68	-6.51	-2.49

Panel D: Biden AI EO on Oct 30, 2023
Exclude Big Stocks

CAR [1,10]	1	2	3	4	5	H - L
Daily Return	-1.64	-2.78	-1.52	-2.66	-2.75	-1.11
[t]	-4.66	-7.76	-2.93	-5.11	-5.11	-1.73

Panel E: The launch of ChatGPT3 on Nov 30, 2022

CAR [1,10]	1	2	3	4	5	H - L
Daily Return	-0.35	-0.14	0.36	0.49	1.14	1.49
[t]	-1.61	-0.46	0.89	1.31	2.27	2.72

Panel F: The launch of ChatGPT3 on Nov 30, 2022
Exclude Tiny Stocks

CAR [1,10]	1	2	3	4	5	H - L
Daily Return	-0.23	0.12	0.40	0.12	0.61	0.83
[t]	-1.22	0.60	0.95	0.49	1.60	1.98

Table XII
Profitability Regressions on AI Exposure

This table reports panel regressions of ROE on *AI Exposure* and other firm characteristics. *isUnfav* is a dummy variable if year is between 2021 and 2023 during Biden's administration, when the AI policy is relatively unfavorable. I include the interaction term between *AI Exposure* and *isUnfav*. All other control variables are defined the same as Table II. The analysis spans from July 2009 to June 2023. *Current ROE* is defined as the ROE at year t, *Lagged ROE* at year $t-1$, and *Future ROE* at year $t+1$. Column (1) and (2) have *Current ROE* as the dependent variable and include *Lagged ROE* as one of the independent variables. Column (3) and (4) have *Future ROE* as the dependent variable and include *Current ROE* as one of the independent variables. Industry FE is included. Column (1) and (3) cluster standard errors by Firm, while Column (2) and (4) cluster standard errors by Industry × Year.

	(1) Current ROE	(2) Current ROE	(3) Future ROE	(4) Future ROE
AI Exposure	-0.0415***	-0.0415***	-0.0250***	-0.0250
	(-6.08)	(-3.11)	(-3.87)	(-1.76)
isUnfav	-0.0304***	-0.0304	-0.127***	-0.127***
	(-2.80)	(-0.47)	(-8.52)	(-3.54)
AI Exposure × isUnfav	-0.0506***	-0.0506***	-0.0618***	-0.0618*
	(-4.33)	(-3.75)	(-3.50)	(-1.96)
Lagged ROE	0.468***	0.468***		
	(21.93)	(6.96)		
Current ROE			0.496***	0.496***
			(22.58)	(7.01)
Size	0.219***	0.219***	0.115***	0.115***
	(21.98)	(5.43)	(15.97)	(6.57)
log(BE/ME)	0.186***	0.186***	-0.0103	-0.0103
	(15.20)	(4.26)	(-1.60)	(-1.07)
Asset Growth	-0.0523***	-0.0523***	-0.0518***	-0.0518**
	(-4.85)	(-3.58)	(-5.14)	(-2.54)
Cash/Assets	-0.0345***	-0.0345	-0.0834***	-0.0834**
	(-3.70)	(-0.99)	(-8.20)	(-2.18)
St Debt/Assets	-0.144***	-0.144***	-0.0252	-0.0252
	(-3.59)	(-4.02)	(-0.62)	(-0.99)
Lt Debt/Assets	-0.218**	-0.218**	-0.0455	-0.0455
	(-2.20)	(-2.74)	(-0.43)	(-0.89)
Tangibility	-0.00749	-0.00749	0.0223***	0.0223
	(-0.89)	(-0.58)	(3.02)	(1.41)
Book Leverage	0.113	0.113	0.00127	0.00127
	(1.05)	(1.41)	(0.01)	(0.02)
WW Index	0.00479	0.00479	0.00661***	0.00661*
	(1.59)	(1.27)	(2.84)	(2.04)
Constant	-0.0230***	-0.0230	0.00237	0.00237
	(-3.86)	(-1.55)	(0.44)	(0.09)
Industry FE	Yes	Yes	Yes	Yes
Cluster SE by Firm	Yes	No	Yes	No
Cluster SE by Industry × Year	No	Yes	No	Yes
Observations	35242	35242	32429	32429
R-squared	0.340	0.340	0.285	0.285

Table XIII
In-Sample Regressions for AI Innovations on Mimicking Portfolio Returns
Full Sample: July 2009 to June 2023

This table presents results from in-sample regression (1). The dependent variable captures innovations in the Ravenpack AI News Index measure. Observations are monthly, covering the period from July 2009 to June 2023. Following Engle et al. (2020), I set $Z_{t-1}^{AIExp'} r_t$ to represent the returns of a mimicking portfolio, which longs the top half of firms with higher *AI Exposure* and shorts the bottom half with lower *AI Exposure*. Similarly, I set size (using cross-sectionally standardized market value to create $Z_{t-1}^{SIZE'}$, so that half the firms, sorted by market value, have positive weight, and half have negative weight; note that this portfolio will be long large firms and short small firms), value (using cross-sectionally standardized values of book-to-market to create $Z_{t-1}^{HML'}$), and the market (setting $Z_{t-1}^{MKT'}$ to equal the share of total market value). The t-statistics are provided in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

$$\begin{aligned} AIInnovations_t = \xi + w_{AIExp} Z_{t-1}^{AIExp'} r_t + w_{SIZE} Z_{t-1}^{SIZE'} r_t \\ + w_{HML} Z_{t-1}^{HML'} r_t + w_{MKT} Z_{t-1}^{MKT'} r_t + e_t, \end{aligned} \tag{22}$$

	(1)	(2)	(3)	(4)	(5)
$Z_{t-1}^{AIExp'} r_t$	82.24***	82.25***	86.42***	92.40***	93.25***
	(3.60)	(3.59)	(3.73)	(2.87)	(2.87)
$Z_{t-1}^{MKT'} r_t$		-2.717			-7.683
		(-0.27)			(-0.72)
$Z_{t-1}^{SIZE'} r_t$			17.87		21.57
			(1.07)		(1.17)
$Z_{t-1}^{HML'} r_t$				8.226	4.805
				(0.45)	(0.25)
Constant	11.89	14.97	10.18	9.012	16.87
	(0.27)	(0.33)	(0.23)	(0.20)	(0.36)
Observations	168	168	168	168	168
R-squared	0.0724	0.0728	0.0787	0.0735	0.0819

t statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table XIV
In-Sample Regressions for AI Innovations on Mimicking Portfolio Returns
Subsample: July 2016 to June 2023

This table presents results from regression (1). The dependent variable captures innovations in the Ravenpack AI News Index measure. Observations are monthly, covering the period from July 2016 to June 2023. Following Engle et al. (2020), I set $Z_{t-1}^{AIExp'} r_t$ to represent the returns of a mimicking portfolio, which longs the top half of firms with higher *AI Exposure* and shorts the bottom half with lower *AI Exposure*. Similarly, I set size (using cross-sectionally standardized market value to create $Z_{t-1}^{SIZE'}$, so that half the firms, sorted by market value, have positive weight, and half have negative weight; note that this portfolio will be long large firms and short small firms), value (using cross-sectionally standardized values of book-to-market to create $Z_{t-1}^{HML'}$), and the market (setting $Z_{t-1}^{MKT'}$ to equal the share of total market value). The t-statistics are provided in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

$$AIInnovations_t = \xi + w_{AIExp} Z_{t-1}^{AIExp'} r_t + w_{SIZE} Z_{t-1}^{SIZE'} r_t + w_{HML} Z_{t-1}^{HML'} r_t + w_{MKT} Z_{t-1}^{MKT'} r_t + e_t, \quad (22)$$

	(1)	(2)	(3)	(4)	(5)
$Z_{t-1}^{AIExp'} r_t$	93.26***	94.39***	102.1***	107.1**	115.6**
	(2.71)	(2.73)	(2.88)	(2.10)	(2.24)
$Z_{t-1}^{MKT'} r_t$		-9.021			-17.29
		(-0.53)			(-0.94)
$Z_{t-1}^{SMB'} r_t$			29.10		38.23
			(1.00)		(1.20)
$Z_{t-1}^{HML'} r_t$				10.76	6.706
				(0.37)	(0.22)
Constant	35.26	43.90	30.30	27.78	40.65
	(0.41)	(0.50)	(0.35)	(0.31)	(0.45)
Observations	84	84	84	84	84
R-squared	0.0824	0.0855	0.0935	0.0839	0.104

t statistics in parentheses

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

Table XV
In-Sample Regressions for AI Innovations on Mimicking Portfolio Returns
Subsample: July 2016 to June 2023 + Quartile Sorting

This table presents results from regression (1). The dependent variable captures innovations in the Ravenpack AI News Index measure. Observations are monthly, covering the period from July 2016 to June 2023. Following Engle et al. (2020), I set $Z_{t-1}^{AIExp'} r_t$ to represent the returns of a mimicking portfolio, which longs the top quartile of firms with higher *AI Exposure* and shorts the bottom quartile with lower *AI Exposure*. Similarly, I set size (using cross-sectionally standardized market value to create $Z_{t-1}^{SIZE'}$, so that half the firms, sorted by market value, have positive weight, and half have negative weight; note that this portfolio will be long large firms and short small firms), value (using cross-sectionally standardized values of book-to-market to create $Z_{t-1}^{HML'}$), and the market (setting $Z_{t-1}^{MKT'}$ to equal the share of total market value). The t-statistics are provided in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

$$
\begin{aligned}
AIInnovations_t = \xi + w_{AIExp} Z_{t-1}^{AIExp'} r_t + w_{SIZE} Z_{t-1}^{SIZE'} r_t \\
+ w_{HML} Z_{t-1}^{HML'} r_t + w_{MKT} Z_{t-1}^{MKT'} r_t + e_t,
\end{aligned}
\tag{22}
$$

	(1)	(2)	(3)	(4)	(5)
$Z_{t-1}^{AIExp'} r_t$	77.87***	81.19***	82.04***	101.4***	112.8***
	(3.78)	(3.89)	(3.91)	(3.53)	(3.82)
$Z_{t-1}^{MKT'} r_t$		-16.45			-28.75
		(-0.99)			(-1.62)
$Z_{t-1}^{SMB'} r_t$			28.92		38.41
			(1.04)		(1.28)
$Z_{t-1}^{HML'} r_t$				30.95	31.01
				(1.17)	(1.13)
Constant	-2.915	9.852	-7.205	-27.65	-11.07
	(-0.03)	(0.12)	(-0.09)	(-0.32)	(-0.13)
Observations	84	84	84	84	84
R-squared	0.148	0.159	0.160	0.163	0.196

t statistics in parentheses

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

References

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo, 2022, Artificial Intelligence and jobs: Evidence from online vacancies, *Journal of Labor Economics* 40, S293–S340.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, 2024a, Artificial Intelligence and firms' systematic risk, Working Paper.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, 2024b, Artificial Intelligence, firm growth, and product innovation, *Journal of Financial Economics* 151, 103745.

Banz, Rolf W., 1981, The relationship between return and market value of common stocks, *Journal of Financial Economics* 9, 3–18.

Brown, Jeffrey R., and Jiekun Huang, 2020, All the president's friends: Political access and firm value, *Journal of Financial Economics* 138, 415–431.

Brynjolfsson, Erik, Danielle Li, and Lindsey R Raymond, 2023, Generative AI at work, Working Paper 31161, National Bureau of Economic Research.

Burmeister, Edwin, and Marjorie B. McElroy, 1988, Joint estimation of factor sensitivities and risk premia for the arbitrage pricing theory, *The Journal of Finance* 43, 721–733.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *The Journal of Finance* 52, 57–82.

Chabi-Yo, Fousseni, Andrei S. Gonçalves, and Johnathan A. Loudis, 2025, An intertemporal risk factor model, *Management Science, forthcoming* .

Chan, Louis K. C., Yasushi Hamao, and Josef Lakonishok, 1991, Fundamentals and stock returns in japan, *The Journal of Finance* 46, 1739–1764.

Claus, James, and Jacob Thomas, 2001, Equity premia as low as three percent? evidence from analysts' Earnings forecasts for domestic and international stock markets, *The Journal of Finance* 56, 1629–1666.

Cockburn, Iain M, Rebecca Henderson, and Scott Stern, 2018, The impact of Artificial Intelligence on innovation, Working Paper 24449, National Bureau of Economic Research.

Davis, James L., Eugene F. Fama, and Kenneth R. French, 2000, Characteristics, covariances, and average returns: 1929 to 1997, *The Journal of Finance* 55, 389–406.

Deloitte, 2024, State of Generative AI in the enterprise: Now decides next.

Deming, David J, Christopher Ong, and Lawrence H Summers, 2025, Technological disruption in the labor market, Working Paper 33323, National Bureau of Economic Research.

Easton, Peter D., and Steven J. Monahan, 2005, An evaluation of accounting-based measures of expected returns, *The Accounting Review* 80, 501–538.

Eisfeldt, Andrea L., and Dimitris Papanikolaou, 2013, Organization capital and the cross-section of expected returns, *The Journal of Finance* 68, 1365–1406.

Eisfeldt, Andrea L, Gregor Schubert, and Miao Ben Zhang, 2023, Generative AI and firm values, Working Paper 31222, National Bureau of Economic Research.

Engle, Robert F, Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel, 2020, Hedging climate change news, *The Review of Financial Studies* 33, 1184–1216.

Fahlenbrach, Rüdiger, Minsu Ko, and Rene M. Stulz, 2024, Bank payout policy, regulation, and politics, *Fisher College of Business Working Paper* European Corporate Governance Institute – Finance Working Paper No. 1003/2024, Swiss Finance Institute Research Paper No. 24-86.

Fahlenbrach, Rüdiger, Kevin Rageth, and René M Stulz, 2021, How valuable is financial flexibility when revenue stops? Evidence from the COVID-19 crisis, *The Review of Financial Studies* 34, 5474–5521.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Fama, Eugene F., and Kenneth R. French, 1996, Multifactor explanations of asset pricing anomalies, *The Journal of Finance* 51, 55–84.

Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.

Gamba, Andrea, and Alexander Triantis, 2008, The value of financial flexibility, *The Journal of Finance* 63, 2263–2296.

Gebhardt, William R., Charles M. C. Lee, and Bhaskaran Swaminathan, 2001, Toward an implied cost of capital, *Journal of Accounting Research* 39, 135–176.

Giglio, Stefano, and Dacheng Xiu, 2021, Asset pricing with omitted factors, *Journal of Political Economy* 129, 1947–1990.

Hou, Kewei, Haitao Mo, Chen Xue, and Lu Zhang, 2020, An augmented q-factor model with expected growth, *Review of Finance* 25, 1–41.

Hou, Kewei, Mathijs A. van Dijk, and Yinglei Zhang, 2012, The implied cost of capital: A new approach, *Journal of Accounting and Economics* 53, 504–526.

Hou, Kewei, Chen Xue, and Lu Zhang, 2015, Digesting anomalies: An investment approach, *The Review of Financial Studies* 28, 650–705.

Hsu, Po-Hsuan, Kai Li, and Chi-Yang Tsou, 2023, The pollution premium, *The Journal of Finance* 78, 1343–1392.

Jagannathan, Ravi, and Zhenyu Wang, 1998, An asymptotic theory for estimating beta-pricing models using cross-sectional regression, *The Journal of Finance* 53, 1285–1309.

Loughran, Tim, and Bill McDonald, 2011, When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks, *The Journal of Finance* 66, 35–65.

Martin, Ian W. R., and Christian Wagner, 2019, What is the expected return on a stock?, *The Journal of Finance* 74, 1887–1929.

McKinsey, 2023, AI-powered marketing and sales: Reach new heights with Generative AI.

Noy, Shakked, and Whitney Zhang, 2023, Experimental evidence on the productivity effects of Generative Artificial Intelligence, *Science* 381, 187–192.

Ohlson, James A., and Beate E. Juettner-Nauroth, 2005, Expected EPS and EPS growth as determinants of value, *Review of Accounting Studies* 10, 349–365.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer, 2023, The impact of AI on developer productivity: Evidence from GitHub Copilot .

PwC, 2020, Global Artificial Intelligence study: Sizing the prize.

Pástor, Ľuboš, and Pietro Veronesi, 2012, Uncertainty about government policy and stock prices, *The Journal of Finance* 67, 1219–1264.

Pástor, Ľuboš, and Pietro Veronesi, 2013, Political uncertainty and risk premia, *Journal of Financial Economics* 110, 520–545.

Rock, Daniel, 2019, Engineering value: The returns to technological talent and investments in Artificial Intelligence, Working Paper.

Sautner, Zacharias, Laurence Van Lent, Grigory Vilkov, and Ruishen Zhang, 2023, Firm-level climate change exposure, *The Journal of Finance* 78, 1449–1498.

Vilkov, Grigory, 2023, Option-implied data and analysis.

Whited, Toni M., and Guojun Wu, 2006, Financial constraints risk, *The Review of Financial Studies* 19, 531–559.

Appendix A 10K Annual Filing Examples

Here are some examples from 10-K annual filings reveal that firms explicitly discuss both the opportunities and regulatory uncertainties surrounding AI:

- Firms face **regulatory uncertainty** on AI:

  > *The Company's global operations are subject to complex and changing laws and regulations on subjects, including antitrust; privacy, data security and data localization, ..., machine learning and artificial intelligence* - **Apple, 2023**

  > "*Additionally, a number of states have recently introduced or passed legislation as it relates to disclosures of the use of artificial intelligence ("AI") in political advertising... which may impact the sale of political advertising.*" - **Fox, 2024**

- Firms face **opportunity** on AI:

  > "*Building on decades of analytics and data science expertise, the company accelerated its application of artificial intelligence in 2023 to drive innovation, increase employee productivity and deliver business outcomes.*" - **Chevron, 2024**

  > "*Therefore, one of our top priorities is to digitalize the Coca-Cola system by, among other things, ..., digitalizing operations through the use of data, artificial intelligence, automation, robotics and digital devices to increase efficiency and productivity.*" - **Coca-Cola, 2024**

Appendix B Conference Call Transcripts Examples

Here are some examples from conference call transcripts reveal that firms explicitly use either unigrams or bigrams about AI, opportunity, or regulatory uncertainty:

- Firms face **regulatory uncertainty** on AI:

  "*Indeed, the generative AI space remains very active with regulatory oversight, a top priority to continue development and adoption in a responsible manner... For example, we have already established governance processes for new technologies, including AI, to continuously assess the compliance of our revolving solutions with regulatory requirements and industry standards.* " - **TELUS International, Earnings Call 2023**

  "*In the U.S., even as support continues to build for a federal right of publicity, several states are taking action. The State of Tennessee recently enacted the Ensuring Likeness Voice and Image Security Act, known as the ELVIS Act that provides strong protections against generative AI voice cloning. We expect further action on these issues as there are ongoing legislative debates in jurisdictions around the world, but we are not waiting for these processes to complete...* " - **Universal Music Group, 2024**

- Firms face **opportunity** on AI:

  "*As we enter 2024, we are seeing much more optimism as growth prospects driven by new artificial intelligence capabilities start to emerge... Finally, we get a lot of questions about how MST is related to the fast-evolving developments in artificial intelligence. And I can tell you it's extensive.*" - **Atomera Incorporated, Earnings Call 2023**

  "*So again, artificial intelligence can be applied at any of these different stages of the digital thread. So you have discover, create, make and sell. When people typically think of us, they see us in this manufacturing this make area, where we talk about areas like manufacturing operations, but really everything from product ideation influences manufacturing costs and other areas upstream. So when we talk about the digital thread, again, it's artificial intelligence applied throughout.*" - **Rockwell Automation, Special Call 2023**

Appendix C Variable Definitions

Variable	Definition
Book Leverage	The sum of current liabilities (DLC) and long-term debt (DLTT) divided by total assets (AT).
Cash/Assets	The ratio of cash (CHE) to total assets (AT) following Fahlenbrach, Rageth, and Stulz (2021).
St Debt/Assets	The ratio of debt in current liabilities (DLC) to total assets (AT) following Fahlenbrach, Rageth, and Stulz (2021).
Lt Debt/Assets	The ratio of total long-term debt (DLTT) to total assets (AT) following Fahlenbrach, Rageth, and Stulz (2021).
ROE	The measure of profitability. It is calculated as the income before extraordinary items (IB) divided by book equity.
ROA	Profitability measure, defined as income before extraordinary items divided by total assets.
R&D/Assets	The ratio of R&D to lagged total assets.
Firm Age	The age of a firm starting from its initial listing in the CRSP database.
Asset Growth	The ratio of the change in total assets from year $t-2$ to year $t-1$ to total assets in year $t-2$.
Sales	Total sales, calculated as the total income generated by an organization through the sale of products and services. Preferably, consolidated sales data is used when available.
Firm Size	*ln(ME)*, measured as the log of market equity in June of year t.
BE/ME	The ratio between book equity at the end of June of year t and market equity at the end of December of year $t-1$.
SIC4	Four-digit SIC Code, categorizing businesses based on their primary activity per the US 1987 SIC classification.
Tangibility	Property, plant, and equipment (PPENT) divided by total assets (AT).
WW Index	The Whited-Wu index, computed as per Whited and Wu (2006). Higher values indicate greater financial constraint.
ICC: GLS	$M_t = B_t + \sum_{k=1}^{11} \left(\frac{E_t[(ROE_{t+k}-R)\times B_{t+k-1}]}{(1+R)^k}\right) + \left(\frac{E_t[(ROE_{t+12}-R)\times B_{t+11}]}{R(1+R)^{11}}\right)$, following Gebhardt, Lee, and Swaminathan (2001).
ICC: CT	$M_t = B_t + \sum_{k=1}^{5} \left(\frac{E[(ROE_{t+k}-R)\times B_{t+k-1}]}{(1+R)^k}\right) + \frac{E[(ROE_{t+5}-R)\times B_{t+4}](1+g)}{(R-g)\times(1+R)^5}$, following Claus and Thomas (2001).

Variable	Definition
ICC: OJ	$R = A + \sqrt{A^2 + \frac{E_t[E_{t+1}]}{M_t} \times (g - (\gamma - 1))}$ where $A = 0.5(\gamma - 1) + \frac{E_t[D_{t+1}]}{M_t}$ and $g = 0.5\left(\frac{E_t[E_{t+3}] - E_t[E_{t+2}]}{E_t[E_{t+2}]} + \frac{E_t[E_{t+5}] - E_t[E_{t+4}]}{E_t[E_{t+4}]}\right)$, following Ohlson and Juettner-Nauroth (2005).
ICC: MPEG	$M_t = \frac{E_t[E_{t+2}] + R \times (E_t[D_{t+1}] - E_t[E_{t+1}])}{R^2}$, following Easton and Monahan (2005).
TF-IDF	TF-IDF score is the Term Frequency (TF) times Inverse Document Frequency (IDF). Term Frequency (TF) measures how frequently a term occurs in a document. $$\mathrm{TF}(t, d) = \frac{\text{Number of times term } t \text{ appears in document } d}{\text{Total number of terms in document } d}$$ Inverse Document Frequency (IDF) measures how important a term is. It downweights terms that appear more frequently across multiple documents. $$\mathrm{IDF}(t, D) = \log\left(\frac{\text{Total number of documents } D}{\text{Number of documents with term } t}\right)$$

Appendix D Additional Figures and Tables

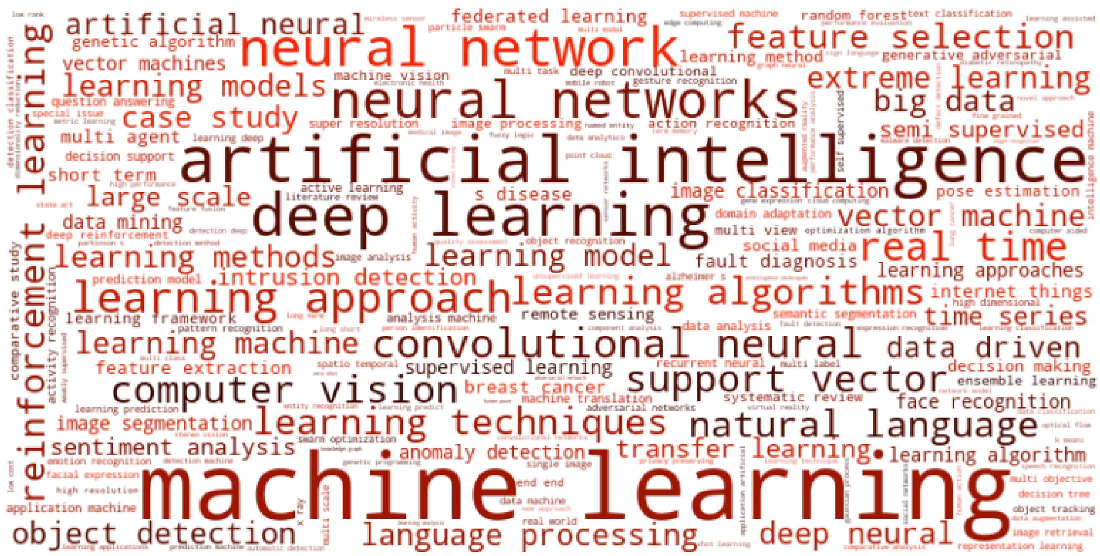

Figure D.1. Alternative Artificial Intelligence Vocabulary WordCloud. The figures present word clouds that summarize the bigrams extracted from a total number of 744,044 academic paper titles in the Web of Science that are directly related to *artificial intelligence, machine learning, natural language processing, and computer vision.* These are the keywords used in Babina et al. (2024a) when they study the systematic risk of AI through the labor channel. The larger the term size, the more frequent the terms appear in the corpus (measured in TF-IDF score). The Web of Science paper titles are downloaded from ProQuest. The years range from 2010 to 2024.

Table D.2
Top 50 Artificial Intelligence Bigrams

This table presents the initial AI bigrams derived from 159,444 titles of academic papers directly related to AI from the Web of Science. The bigrams are selected by computing the TF-IDF scores and these are the top 50 with highest scores.

artificial intelligence	neural network	machine learning
neural networks	deep learning	reinforcement learning
artificial neural	multi agent	real time
convolutional neural	big data	decision making
genetic algorithm	case study	application artificial
time series	support vector	data mining
decision support	deep neural	feature selection
learning approach	intelligence techniques	fuzzy logic
natural language	internet things	intelligence ai
particle swarm	multi objective	use artificial
special issue	swarm optimization	bee colony
deep reinforcement	fault diagnosis	artificial bee
large scale	intelligence machine	semi supervised
transfer learning	object detection	genetic algorithms
intrusion detection	intelligence technology	face recognition
data driven	agent systems	short term
breast cancer	optimization algorithm	

Table D.3
Top 50 Artificial Intelligence Unigrams

This table presents the initial AI unigrams derived from 159,444 titles of academic papers directly related to AI from the Web of Science. The unigrams are selected by computing the TF-IDF scores and these are the top 50 with highest scores.

artificial	base	intelligence
system	network	learning
model	algorithm	neural
approach	analysis	application
datum	learn	detection
machine	multi	method
deep	image	design
ai	classification	agent
intelligent	optimization	prediction
recognition	control	study
fuzzy	knowledge	research
time	problem	technology
human	feature	decision
information	management	technique
framework	robot	new
support	dynamic	process
development	logic	

Table D.4
AI Exposure by Industry (SIC2)
Full Sample: July 2009 to June 2023

Industry (SIC2)	AI Exposure	Industry (SIC1)
Educational Services	6.285266	Services
Insurance Agents, Brokers, & Service	5.561357	Finance, Insurance, & Real Estate
Local & Interurban Passenger Transit	5.495105	Transportation & Public Utilities
Business Services	5.464625	Services
Chemical & Allied Products	5.396434	Manufacturing
Instruments & Related Products	5.375045	Manufacturing
Electronic & Other Electric Equipment	5.047649	Manufacturing
Engineering & Management Services	4.989501	Services
Eating & Drinking Places	4.852603	Retail Trade
Communications	4.803536	Transportation & Public Utilities
Real Estate	4.79584	Finance, Insurance, & Real Estate
Furniture & Homefurnishings Stores	4.609057	Retail Trade
Industrial Machinery & Equipment	4.607254	Manufacturing
Health Services	4.544601	Services
Printing & Publishing	4.542002	Manufacturing
Hotels & Other Lodging Places	4.534035	Services
Legal Services	4.505624	Services
Security & Commodity Brokers	4.504373	Finance, Insurance, & Real Estate
Miscellaneous Retail	4.495735	Retail Trade
Automotive Dealers & Service Stations	4.447903	Retail Trade
Miscellaneous Manufacturing Industries	4.394613	Manufacturing
Personal Services	4.361629	Services
Electric, Gas, & Sanitary Services	4.355988	Transportation & Public Utilities
Transportation Services	4.337441	Transportation & Public Utilities
Metal Mining	4.291618	Mining
Auto Repair, Services, & Parking	4.2776	Services
Building Materials & Gardening Supplies	4.218395	Retail Trade
Transportation Equipment	4.200852	Manufacturing
Amusement & Recreation Services	4.153239	Services
Nondepository Institutions	4.148221	Finance, Insurance, & Real Estate

Continued on next page

Table D.4 Continued from previous page

Industry (SIC2)	AI Exposure	Industry (SIC1)
Non-Classifiable Establishments	4.127732	Nonclassifiable Establishments
Furniture & Fixtures	4.074484	Manufacturing
Motion Pictures	4.060033	Services
Food Stores	4.036707	Retail Trade
Insurance Carriers	4.028654	Finance, Insurance, & Real Estate
General Merchandise Stores	3.994649	Retail Trade
Transportation by Air	3.958159	Transportation & Public Utilities
Social Services	3.9469	Services
Trucking & Warehousing	3.927254	Transportation & Public Utilities
Holding & Other Investment Offices	3.916995	Finance, Insurance, & Real Estate
Pipelines, Except Natural Gas	3.909467	Transportation & Public Utilities
Agricultural Production – Crops	3.903013	Agriculture, Forestry, & Fishing
Rubber & Miscellaneous Plastics Products	3.870281	Manufacturing
Wholesale Trade – Durable Goods	3.854764	Wholesale Trade
Apparel & Accessory Stores	3.794587	Retail Trade
Oil & Gas Extraction	3.788054	Mining
Apparel & Other Textile Products	3.761449	Manufacturing
Textile Mill Products	3.726901	Manufacturing
Agricultural Services	3.714297	Agriculture, Forestry, & Fishing
Fabricated Metal Products	3.710879	Manufacturing
General Building Contractors	3.696674	Construction
Wholesale Trade – Nondurable Goods	3.69314	Wholesale Trade
Leather & Leather Products	3.685508	Manufacturing
Food & Kindred Products	3.648712	Manufacturing
Tobacco Products	3.604894	Manufacturing
Heavy Construction, Except Building	3.603374	Construction
Depository Institutions	3.542971	Finance, Insurance, & Real Estate
Railroad Transportation	3.538	Transportation & Public Utilities
Petroleum & Coal Products	3.512356	Manufacturing
Water Transportation	3.417509	Transportation & Public Utilities
Nonmetallic Minerals, Except Fuels	3.302386	Mining
Services, Not Elsewhere Classified	3.285075	Services

Continued on next page

Table D.4 Continued from previous page

Industry (SIC2)	AI Exposure	Industry (SIC1)
Stone, Clay, & Glass Products	3.240903	Manufacturing
Primary Metal Industries	3.235793	Manufacturing
Lumber & Wood Products	3.230363	Manufacturing
Paper & Allied Products	3.219572	Manufacturing
Special Trade Contractors	3.182942	Construction
Coal Mining	3.137699	Mining

Appendix E Theoretical Derivations and Proofs

A Proof of Lemma 1 (using Kalman-Bucy filter)

Lemma 1. *Combining the signals, Equation (16), and prior distribution of the political cost, Equation (14), we can get the posterior distribution of the political cost,*

$$c \sim N(\hat{c}_t, \hat{\sigma}^2_{c,t})$$

where

$$d\hat{c}_t = \hat{\sigma}^2_{c,t} h^{-1} d\hat{Z}^c_t,$$

and

$$\hat{\sigma}^2_{c,t} = \frac{1}{\frac{1}{\sigma_c^2} + \frac{1}{h^2}t}$$

Proof. I use a Kalman-Bucy filter to prove this Lemma. As agents observe the signals (ds_t), they update their beliefs about political cost using Bayesian updating. The Kalman-Bucy filter provides a way to compute the posterior distribution of the political cost given the observed signals. Assume the political cost, c, is a constant parameter (state) with prior distribution:

$$c \sim N\left(-\frac{\sigma_c^2}{2}, \sigma_c^2\right).$$

Agents observe a signal process s_t governed by:

$$ds_t = c\,dt + h\,dZ^c_t, \quad \text{where } dZ^c_t \text{ is a Brownian motion.}$$

This corresponds to a continuous-time filtering problem with:

- State equation: $dc = 0$ (since c is constant),
- Observation equation: $ds_t = c\,dt + h\,dZ^c_t$.

Applying a Kalman-Bucy filter, the posterior distribution, $c \,|\, \mathcal{F}^s_t$, is Gaussian, with mean $\hat{c}_t$ and variance $\hat{\sigma}^2_{c,t}$. The Kalman-Bucy filter gives the dynamics for $\hat{c}_t$ and $\hat{\sigma}^2_{c,t}$:

$$d\hat{c}_t = \frac{\hat{\sigma}^2_{c,t}}{h^2}\left(ds_t - \hat{c}_t\,dt\right), \tag{E1}$$

$$\frac{d\hat{\sigma}^2_{c,t}}{dt} = -\frac{\left(\hat{\sigma}^2_{c,t}\right)^2}{h^2}. \tag{E2}$$

Now we can solve the variance equation. The equation (E2) is solved as follows:

$$\frac{d}{dt}\left(\frac{1}{\hat{\sigma}^2_{c,t}}\right) = \frac{1}{h^2}.$$

Integrating from 0 to t, with initial condition $\hat{\sigma}^2_{c,0} = \sigma^2_c$:

$$\frac{1}{\hat{\sigma}^2_{c,t}} = \frac{1}{\sigma^2_c} + \frac{t}{h^2} \quad \Rightarrow \quad \hat{\sigma}^2_{c,t} = \frac{1}{\frac{1}{\sigma^2_c} + \frac{t}{h^2}}.$$

Then we update the posterior mean: Define the innovation process $d\hat{Z}^c_t$, a $\mathcal{F}^s_t$-Brownian motion:

$$d\hat{Z}^c_t = \frac{1}{h}\left(ds_t - \hat{c}_t\, dt\right).$$

Substitute $ds_t = \hat{c}_t\, dt + h\, d\hat{Z}^c_t$ into (E1) to obtain:

$$d\hat{c}_t = \hat{\sigma}^2_{c,t}\, h^{-1} d\hat{Z}^c_t.$$

The posterior distribution of c given observations up to time t is:

$$c \,|\, \mathcal{F}^s_t \sim N\left(\hat{c}_t, \hat{\sigma}^2_{c,t}\right),$$

with dynamics:

$$d\hat{c}_t = \hat{\sigma}^2_{c,t} h^{-1} d\hat{Z}^c_t, \quad \hat{\sigma}^2_{c,t} = \frac{1}{\frac{1}{\sigma^2_c} + \frac{t}{h^2}},$$

which is what we need to prove. □

B Proof of Lemma 2 (Aggregate Capital)

Lemma 2. *The aggregate capital at time T, $B_T = \int_0^1 B^i_T di$, is given by*

$$B_T = B_\tau e^{\left(\mu + pg + (1-p)\eta - \frac{1}{2}\sigma^2\right)(T-\tau) + \sigma(Z_T - Z_\tau)},$$

where $g \equiv g^F$ under Favorable regulation, $g \equiv g^U$ under Unfavorable regulation, $\eta \equiv \eta^E$ in Early-stage AI environment, and $\eta \equiv \eta^M$ in Mature-stage AI environment.

Proof. The capital growth follows $dB^i_t = B^i_t d\Pi^i_t$ with the profitability process, equation (3), we derive:

$$B_T^i = B_\tau^i \exp\left\{\left(\mu + \xi^i\left[s^i g + (1-s^i)\eta\right] - \frac{1}{2}\sigma^2 - \frac{1}{2}\sigma_1^2\right)(T-\tau) + \sigma(Z_T - Z_\tau) + \sigma_1(Z_T^i - Z_\tau^i)\right\}$$

The above equation is derived by applying Itô's Lemma to $ln(B_t^i)$ and taking integral from τ to T.
Aggregating across firms:

$$B_T = \int_0^1 B_T^i di = e^{\left(\mu - \frac{1}{2}\sigma^2 - \frac{1}{2}\sigma_1^2\right)(T-\tau)+\sigma(Z_T - Z_\tau)} \int_0^1 B_\tau^i e^{\xi^i[s^i g+(1-s^i)\eta](T-\tau)+\sigma_1(Z_T^i - Z_\tau^i)} di$$

Applying the Law of Large Numbers:

$$\int_0^1 B_\tau^i e^{\xi^i[s^i g+(1-s^i)\eta](T-\tau)+\sigma_1(Z_T^i - Z_\tau^i)} di \to \mathbb{E}^i\left[B_\tau^i\right]\mathbb{E}^i\left[e^{\xi^i[s^i g+(1-s^i)\eta](T-\tau)}\right] e^{\frac{1}{2}\sigma_1^2(T-\tau)}$$

Using cross-sectional independence, and the expectation value of Uniform distribution and Bernoulli distribution, $\xi^i \sim \text{Uniform}(0,2)$ with $\mathbb{E}[\xi^i] = 1$, and each firm's AI stage, $s^i \sim \text{Bernoulli}(p)$, where $s^i = 1$ (AI Elites) with probability p, and $s^i = 0$ (AI Starters) with probability $1-p$:

$$\mathbb{E}^i\left[e^{\xi^i[s^i g+(1-s^i)\eta](T-\tau)}\right] = e^{(pg+(1-p)\eta)(T-\tau)}$$

Given $\mathbb{E}^i[B_\tau^i] = B_\tau$, we obtain:

$$B_T = B_\tau e^{\left(\mu+pg+(1-p)\eta-\frac{1}{2}\sigma^2\right)(T-\tau)+\sigma(Z_T - Z_\tau)},$$

which is what we need to prove. □

C Proof of Proposition 1 (Optimal AI Cost Threshold)

Proposition 1. *The government will switch from a Favorable (F) to an Unfavorable (U) AI policy at time τ if and only if the realized political cost c exceeds the threshold:*

$$c > \underline{c}(\tau) \equiv \log\left(e^{(\gamma-1)\left(g^F - g^U\right)(T-\tau)} - 1\right),$$

where $g^F > 0$ and $g^U < 0$ represent the regulatory impacts on profitability under Favorable and Unfavorable regimes, respectively. The threshold $\underline{c}(\tau)$ increases with risk aversion ($\gamma > 1$)

and the magnitude of the regulatory gap $|g^F - g^U|$.

Proof. Using Lemma 2, aggregate capital under each regime is:

$$B_T^{\mathrm{F}} = B_\tau e^{\left(\mu + pg^F + (1-p)\eta - \frac{1}{2}\sigma^2\right)(T-\tau) + \sigma(Z_T - Z_\tau)},$$

$$B_T^{\mathrm{U}} = B_\tau e^{\left(\mu + pg^U + (1-p)\eta - \frac{1}{2}\sigma^2\right)(T-\tau) + \sigma(Z_T - Z_\tau)}.$$

The government compares expected utilities under both regimes:

$$\mathbb{E}_\tau\left[\frac{W_T^{1-\gamma}}{1-\gamma} \mid F\right] = \frac{B_\tau^{1-\gamma}}{1-\gamma} e^{(1-\gamma)\left(\mu + pg^F + (1-p)\eta - \frac{1}{2}\sigma^2\right)(T-\tau) + \frac{1}{2}(1-\gamma)^2\sigma^2(T-\tau)},$$

$$\mathbb{E}_\tau\left[\phi(c)\frac{W_T^{1-\gamma}}{1-\gamma} \mid U\right] = \frac{\phi(c)B_\tau^{1-\gamma}}{1-\gamma} e^{(1-\gamma)\left(\mu + pg^U + (1-p)\eta - \frac{1}{2}\sigma^2\right)(T-\tau) + \frac{1}{2}(1-\gamma)^2\sigma^2(T-\tau)}.$$

Substitute $\phi(c) = 1 + e^c$ into the inequality $\mathbb{E}_\tau[U(W_T) \mid F] < \mathbb{E}_\tau[\phi(c)U(W_T) \mid U]$:

$$e^{(1-\gamma)(pg^F + (1-p)\eta)(T-\tau)} < (1 + e^C)e^{(1-\gamma)(pg^U + (1-p)\eta)(T-\tau)}.$$

Cancel common terms and simplify:

$$e^{(1-\gamma)p(g^F - g^U)(T-\tau)} < 1 + e^c.$$

Take logarithms and rearrange:

$$c > \log\left(e^{(1-\gamma)p(g^F - g^U)(T-\tau)} - 1\right) \equiv \underline{c}(\tau),$$

which is what we need to prove. $\square$

D Proof of Corollary 1

Corollary 1. *The time-t perceived probability of a government AI policy shift from Favorable (F) to Unfavorable (U) at time* τ $(t < \tau)$ *is:*

$$p_{\tau|t} = 1 - \Phi\left(\underline{c}(\tau); \hat{c}_t, \hat{\sigma}_{c,t}^2\right),$$

where $\Phi(\cdot;\hat{c}_t,\hat{\sigma}^2_{c,t})$ is the CDF of the posterior normal distribution $c \sim N(\hat{c}_t,\hat{\sigma}^2_{c,t})$, and $\underline{c}(\tau)$ is defined in Proposition 1.

Proof. From Lemma 1, the posterior distribution of the political cost c, given information and signals $\mathcal{F}^s_t$, is:

$$c \,|\, \mathcal{F}^s_t \sim N\left(\hat{c}_t, \hat{\sigma}^2_{c,t}\right).$$

The government switches policies at τ iff $c > \underline{c}(\tau)$. The perceived probability of this event is:

$$p_{\tau|t} = \mathbb{P}\left(c > \underline{c}(\tau) \,|\, \mathcal{F}^s_t\right) = 1 - \Phi\left(\underline{c}(\tau); \hat{c}_t, \hat{\sigma}^2_{c,t}\right),$$

where $\Phi(\cdot;\hat{c}_t,\hat{\sigma}^2_{c,t})$ is the CDF of $N(\hat{c}_t,\hat{\sigma}^2_{c,t})$. More specifical:

$$\mathbb{P}\left(c > \underline{c}(\tau) \,|\, \mathcal{F}^s_t\right) = \mathbb{P}\left(\frac{c-\hat{c}_t}{\hat{\sigma}_{c,t}} > \frac{\underline{c}(\tau)-\hat{c}_t}{\hat{\sigma}_{c,t}} \,|\, \mathcal{F}^s_t\right) = 1 - \Phi\left(\frac{\underline{c}(\tau)-\hat{c}_t}{\hat{\sigma}_{c,t}}\right),$$

which is what we need to prove. □

E Proof of Proposition 2 (State Price Density)

Proposition 2 (State Price Density)**.** *Before a regulatory or opportunity regime shift ($t < \tau$), the state price density is:*

$$\pi_t = B_t^{-\gamma}\Omega_t,$$

where

$$\Omega_t = e^{\left(-\gamma\mu+\frac{1}{2}\gamma(\gamma+1)\sigma^2\right)(T-t)-\gamma\left(pg^F+(1-p)\eta\right)(\tau-t)}\left[p_{\tau|t}e^{-\gamma\left(pg^U+(1-p)\eta\right)(T-\tau)}+(1-p_{\tau|t})e^{-\gamma\left(pg^F+(1-p)\eta\right)(T-\tau)}\right]$$

and $p_{\tau|t}$ is the probability of a shift to Unfavorable AI policy (Corollary 1).

Proof. Households maximize CRRA utility $U(W_T) = \frac{W_T^{1-\gamma}}{1-\gamma}$. The state price density is derived from marginal utility:

$$\pi_t = \mathbb{E}_t\left[U'(W_T)\right] = \mathbb{E}_t\left[W_T^{-\gamma}\right].$$

By market clearing, aggregate wealth equals aggregate capital: $W_T = B_T = \int_0^1 B_T^i di$, where B_T^i follows the profitability process in Equation (3). From Lemma 2, aggregate capital evolves as:

$$B_T = B_t e^{\left(\mu + pg_T + (1-p)\eta_T - \frac{1}{2}\sigma^2\right)(T-t) + \sigma(Z_T - Z_t)},$$

with $g_T \in \{g^U, g^F\}$ (regulatory regimes) and $\eta_T \in \{\eta^E, \eta^M\}$ (opportunity regimes). Substituting B_T into π_t:

$$\pi_t = B_t^{-\gamma} \mathbb{E}_t \left[e^{-\gamma\left(\mu(T-t) + pg_T(T-t) + (1-p)\eta_T(T-t) - \frac{1}{2}\sigma^2(T-t) + \sigma(Z_T - Z_t)\right)} \right].$$

The stochastic term $-\gamma\sigma(Z_T - Z_t)$ is Gaussian, with expectation:

$$\mathbb{E}_t \left[e^{-\gamma\sigma(Z_T - Z_t)} \right] = e^{\frac{1}{2}\gamma^2\sigma^2(T-t)}.$$

Substitute this result and decompose the expectation over regimes:

$$\pi_t = B_t^{-\gamma} e^{\left(-\gamma\mu + \frac{1}{2}\gamma(\gamma+1)\sigma^2\right)(T-t)} \mathbb{E}_t \left[e^{-\gamma(pg_T + (1-p)\eta_T)(T-t)} \right].$$

Conditional on regulatory shifts ($g_T = g^U$ with probability $p_{\tau|t}$, or g^F otherwise), and opportunity regimes ($\eta_T = \eta^E$ or η^M), the expectation becomes:

$$\mathbb{E}_t \left[e^{-\gamma(pg_T + (1-p)\eta_T)(T-t)} \right] = e^{-\gamma\left(pg^F + (1-p)\eta\right)(\tau - t)} \left[p_{\tau|t} e^{-\gamma\left(pg^U + (1-p)\eta\right)(T-\tau)} + (1 - p_{\tau|t}) e^{-\gamma\left(pg^F + (1-p)\eta\right)(T-\tau)} \right].$$

Combining terms yields Ω_t, completing the proof. □

F Proof of Proposition 3 (SDF)

Proposition 3. *The SDF is,*

$$\frac{d\pi_t}{\pi_t} = \mathbb{E}_t \left[\frac{d\pi_T}{\pi_T} \right] - \lambda dZ_t - \lambda_{c,t} d\hat{Z}_t^C$$

where,

$$\lambda = \gamma\sigma > 0$$

and

$$\lambda_{c,t} = \frac{1}{\Omega_t} \frac{\partial \Omega_t}{\partial \hat{C}_t} \hat{\sigma}_{c,t}^2 \eta^{-1} > 0$$

Proof. The uncertainty risk price is given by:

$$
\begin{aligned}
\lambda_{c,t} &= -\frac{1}{\Omega_t}\frac{\partial \Omega_t}{\partial p_{\tau|t}}\frac{\partial p_{\tau|t}}{\partial \hat{c}_t}\hat{\sigma}^2_{c,t}\eta^{-1} \\
&= -\frac{e^{\left(-\gamma\mu+\frac{1}{2}\gamma(\gamma+1)\sigma^2\right)(T-t)-\gamma\left(pg^F+(1-p)\eta\right)(\tau-t)}\left[e^{-\gamma\left(pg^U+(1-p)\eta\right)(T-\tau)}-e^{-\gamma\left(pg^F+(1-p)\eta\right)(T-\tau)}\right]}{\Omega_t} \\
&\quad\times n(\underline{c}(\tau);\hat{c}_t,\hat{\sigma}^2_{c,t})\hat{\sigma}^2_{c,t}\eta^{-1} \\
&= -\left[\frac{(1-p_{\tau|t})(1-F_\tau)}{p_{\tau|t}+(1-p_{\tau|t})F_\tau}\right]n(\underline{c}(\tau);\hat{c}_t,\hat{\sigma}^2_{c,t})\hat{\sigma}^2_{c,t}\eta^{-1},
\end{aligned}
\tag{E3}
$$

where:

$$
F_\tau \equiv \frac{e^{-\gamma\left(pg^F+(1-p)\eta\right)(T-\tau)}}{e^{-\gamma\left(pg^U+(1-p)\eta\right)(T-\tau)}} = e^{-\gamma p(g^F-g^U)(T-\tau)} < 1. \tag{E4}
$$

The term $\frac{(1-p_{\tau|t})(1-F_\tau)}{p_{\tau|t}+(1-p_{\tau|t})F_\tau}$ is positive because $F_\tau < 1$. Since all other terms in (E3) are positive, $\lambda_{c,t} < 0$. This confirms that regulatory uncertainty shocks are *negatively priced* in equilibrium. □

G Proof of Lemma 3

Lemma 3. *For $t < \tau$, the stock price for firm i is given by:*

$$
M^i_t = B^i_t\Theta^i_t,
$$

where

$$
\Theta^i_t = e^{(\mu-\gamma\sigma^2)(T-t)+\xi^i\left(pg^F+(1-p)\eta\right)(\tau-t)}\left[\phi_t e^{\xi^i\left(pg^U+(1-p)\eta\right)(T-\tau)}+(1-\phi_t)e^{\xi^i\left(pg^F+(1-p)\eta\right)(T-\tau)}\right],
$$

and

$$
\phi_t = \frac{p_{\tau|t}}{p_{\tau|t}+(1-p_{\tau|t})e^{-\gamma p(g^F-g^U)(T-\tau)}}.
$$

Proof. The stock price M^i_t is derived using the law of iterated expectations and the state price density π_t from Proposition 2. For $t < \tau$, the market value satisfies:

$$
M^i_t = \mathbb{E}_t\left[\frac{\pi_T}{\pi_t}M^i_T\right].
$$

Conditional on a regime shift at τ, the stock price under the two regimes is:

$$M_t^{U,i} = B_t^i e^{(\mu-\gamma\sigma^2)(T-t)+\xi^i\left(pg^F+(1-p)\eta\right)(\tau-t)+\xi^i\left(pg^U+(1-p)\eta\right)(T-\tau)},$$

$$M_t^{F,i} = B_t^i e^{(\mu-\gamma\sigma^2+\xi^i\left(pg^F+(1-p)\eta\right))(T-t)}.$$

The unconditional stock price is a weighted average of the two regimes:

$$M_t^i = \phi_t M_t^{U,i} + (1-\phi_t)M_t^{F,i},$$

where ϕ_t is the probability-weighted adjustment factor:

$$\phi_t = \frac{p_{\tau|t}}{p_{\tau|t} + (1-p_{\tau|t})e^{-\gamma p(g^F-g^U)(T-\tau)}}.$$

Substituting $M_t^{U,i}$ and $M_t^{F,i}$ into the expression for M_t^i yields:

$$M_t^i = B_t^i \Theta_t^i,$$

where Θ_t^i is the same as defined in the proposition. □

H Proof of Proposition 4

Proposition 4. *Firm i's stock realized returns:*

$$\frac{dM_t^i}{M_t^i} = \mathbb{E}_t\left[\frac{dM_T^i}{M_T^i}\right] + \sigma dZ_t + \sigma_1 dZ_t^i + \beta_{M,t}^i d\hat{Z}_t^c,$$

where

$$\beta_{M,t}^i = \frac{1}{\Theta^i}\frac{\partial\Theta_t^i}{\partial\hat{c}_t}\hat{\sigma}_{c,t}^2 < 0,$$

and risk premia can be expressed as:

$$\mathbb{E}_t\left[\frac{dM_T^i}{M_T^i}\right] = \sigma\lambda dt + \beta_{M,t}^i\lambda_{c,t}dt,$$

where

$$\frac{\partial \beta^i_{M,t}}{\partial \xi^i} = \begin{cases} > 0 & \textit{if } p = 0 \textit{ (Firms are all AI Starters)}, \\ < 0 & \textit{if } p = 1 \textit{ (Firms are all AI Elites)}, \end{cases} \qquad \textit{and} \quad \lambda_{c,t} < 0.$$

Proof. For the ease of exposition, I assume at every point of time, there is an exogenous increase of possibility of opportunity regime change from Early to Mature. From Lemma 3, the stock price can be written as,

$$M^i_t = B^i_t \Theta^i_t,$$

where

$$\Theta^i_t = e^{(\mu - \gamma\sigma^2)(T-t) + \xi^i \left(pg^F + (1-p)\eta^E\right)(\tau - t)} \left[\phi_t \, e^{\xi^i \left(pg^U + (1-p)\eta^M\right)(T-\tau)} + (1-\phi_t)\, e^{\xi^i \left(pg^F + (1-p)\eta^E\right)(T-\tau)}\right].$$

Applying Itô's Lemma to $M^i_t = B^i_t \Theta^i_t$, the return dynamics are given by

$$\frac{dM^i_t}{M^i_t} = \mathbb{E}_t\left[\frac{dM^i_T}{M^i_T}\right] + \sigma\, dZ_t + \sigma_1\, dZ^i_t + \beta^i_{M,t}\, d\hat{Z}^c_t,$$

with $\beta^i_{M,t}$ derived as follows:

$$\beta^i_{M,t} = \frac{1}{\Theta^i_t}\frac{\partial \Theta^i_t}{\partial \phi_t}\frac{\partial \phi_t}{\partial p_{\tau|t}}\frac{\partial p_{\tau|t}}{\partial \hat{c}_t}\,\hat{\sigma}^2_{c,t},$$

$$\times \left\{p_{\tau|t} + (1-p_{\tau|t})e^{-\gamma p(g^F - g^U)(T-\tau)} - p_{\tau|t}\Big(1 - e^{-\gamma p(g^F-g^U)(T-\tau)}\Big) n(c(\tau); \hat{c}_t, \hat{\sigma}^2_{c,t})\right\}\, \hat{\sigma}^2_{c,t},$$

which can be rewritten as

$$\begin{aligned}\beta^i_{M,t} = & \left[\frac{1 - e^{\xi^i(p(g^F-g^U)+(1-p)(\eta^E-\eta^M))(T-\tau)}}{\phi_t + (1-\phi_t)e^{\xi^i(p(g^F-g^U)+(1-p)(\eta^E-\eta^M))(T-\tau)}}\right] \\ & \times \left[\frac{e^{-\gamma(p(g^F-g^U)+(1-p)(\eta^E-\eta^M))(T-\tau)}}{\left(p_{\tau|t} + (1-p_{\tau|t})e^{-\gamma(p(g^F-g^U)+(1-p)(\eta^E-\eta^M))(T-\tau)}\right)^2}\right] n(c(\tau); \hat{c}_t, \hat{\sigma}^2_{c,t})\, \hat{\sigma}^2_{c,t},\end{aligned}$$

or more compactly,

$$\beta^i_{M,t} = \left[\frac{1 - G^i_\tau}{\phi_t + (1-\phi_t)G^i_\tau}\right]\left[\frac{F_\tau}{\left(p_{\tau|t} + (1-p_{\tau|t})F_\tau\right)^2}\right] n(c(\tau); \hat{c}_t, \hat{\sigma}^2_{c,t})\, \hat{\sigma}^2_{c,t} < 0,$$

where

$$G_\tau^i = e^{\xi^i(p(g^F - g^U)+(1-p)(\eta^E - \eta^M))(T-\tau)} > 1, \quad F_\tau = e^{-\gamma(p(g^F - g^U)+(1-p)(\eta^E - \eta^M))(T-\tau)} < 1.$$

Since only G_τ^i depends on ξ^i, the sign of $\frac{\partial \beta_{M,t}^i}{\partial \xi^i}$ will be decided by G_τ^i. Specifically,

$$\frac{\partial}{\partial \xi^i}\left[\frac{1-G_\tau^i}{\phi_t + (1-\phi_t)\, G_\tau^i}\right] = \frac{\left(\frac{\partial G_\tau^i}{\partial \xi^i}\right)\left[\phi_t + (1-\phi_t)\, G_\tau^i\right] - \left(1-G_\tau^i\right)\frac{\partial}{\partial \xi^i}\left[\phi_t + (1-\phi_t)\, G_\tau^i\right]}{\left[\phi_t + (1-\phi_t)\, G_\tau^i\right]^2} < 0,$$

where $G_\tau^i > 1$ and the sign of $\frac{\partial G_\tau^i}{\partial \xi^i}$ depends on p, the proportion of AI Elites in the firms cross-section:

$$\frac{\partial G_\tau^i}{\partial \xi^i} = \begin{cases} > 0 & \text{if } p = 0 \text{ (Firms are all AI Starters)}, \\ < 0 & \text{if } p = 1 \text{ (Firms are all AI Elites)}, \end{cases} \qquad \text{and} \quad \lambda_{c,t} < 0.$$

This can be translated to:

$$\frac{\partial \beta_{M,t}^i}{\partial \xi^i} = \begin{cases} > 0 & \text{if } p = 0 \text{ (Firms are all AI Starters)}, \\ < 0 & \text{if } p = 1 \text{ (Firms are all AI Elites)}, \end{cases} \qquad \text{and} \quad \lambda_{c,t} < 0.$$

By definition, the risk premia are:

$$\mathbb{E}_t\left[\frac{dM_T^i}{M_T^i}\right] = \sigma\lambda dt + \beta_{M,t}^i \lambda_{c,t} dt,$$

where $\lambda = \gamma\sigma$ and $\lambda_{c,t} < 0$ as shown in Proposition 3. □

I *Proof of Lemma 4*

Lemma 4. *For any two firms, i and j, with AI Exposure $\xi^i > \xi^j$, $\forall i \neq j$, return premia spread between i and j is,*

$$\mathbb{E}_t\left[\frac{dM_T^i}{M_T^i}\right] - \mathbb{E}_t\left[\frac{dM_T^j}{M_T^j}\right] = (\beta_{M,t}^i - \beta_{M,t}^j)\lambda_{c,t} dt \begin{cases} > 0 & \textit{if } p = 0 \textit{ (Firms are all AI Starters)}, \\ < 0 & \textit{if } p = 1 \textit{ (Firms are all AI Elites)}, \end{cases}$$

Proof. From Proposition 4, we know

$$\frac{\partial \beta^i_{M,t}}{\partial \xi^i} = \begin{cases} < 0 & \text{if } p = 0 \text{ (Firms are all AI Starters)}, \\ > 0 & \text{if } p = 1 \text{ (Firms are all AI Elites)}, \end{cases} \quad \text{and} \quad \lambda_{c,t} < 0.$$

The lemma is then straightforward. $\square$