MARKET-PERCEIVED ARTIFICIAL INTELLIGENCE

INSIGHTS FROM THEMATIC ETF PORTFOLIO HOLDINGS

PHILIP HEGARDT

ERIK MILBURNE

Bachelor Thesis

Stockholm School of Economics

2026

Market-Perceived Artificial Intelligence: Insights from Thematic ETF Portfolio Holdings

Abstract:

How do financial markets identify and price firms' exposure to AI? We construct a novel market-perceived measure of U.S. firms' AI exposure using the holdings of AI-themed ETFs. We develop three measures that capture distinct dimensions of firm-level AI exposure and validate them using bidirectional evidence. AI-exposed firms underperformed non-AI peers by 1.9% around the European Parliament's formal approval of the EU AI Act and outperformed by 1.8% around the launch of ChatGPT. These effects survive 3-digit NAICS sector fixed effects and are directionally consistent with an existing fundamentals-based measure of AI exposure. Thematic ETF holdings offer a non-proprietary and easily replicable first-pass proxy for identifying AI-exposed firms, while providing a scalable framework that can be extended to future emerging market themes.

Keywords:

AI, Artificial Intelligence, Thematic ETFs, Event Study, EU AI Act, ChatGPT, Stock Returns

Authors:

Philip Hegardt (26039)
Erik Milburne (25782)

Tutor:

Alvin Chen, Assistant Professor, Department of Finance

Examiner:

Ramin Baghai, Professor, Department of Finance

Bachelor Thesis
Bachelor Program in Business & Economics
Stockholm School of Economics

I. Introduction

Recent regulatory developments around artificial intelligence are reshaping the operating environment for AI-exposed firms worldwide, with the European Union's AI Act standing as the most comprehensive legislative response to date. We construct the first market-perceived measure of U.S. publicly traded firms' exposure to AI, classifying firms based on their inclusion and weighting in AI-themed exchange-traded funds. Using this measure, we show that AI-related shocks generated significant and substantial cross-sectional differences in the valuation of U.S. firms both negatively and positively. Firms classified as AI-exposed underperformed their non-AI peers by 1.9% around the European Parliament's formal adoption of the Act on March 13, 2024, and outperformed their non-AI peers by 1.8% around the launch of ChatGPT on November 30, 2022. Both reactions survive 3-digit NAICS sector fixed effects, indicating that the documented effects are not entirely driven by AI firms' concentration in tech-adjacent sectors.

The interaction between rapid AI adoption and an emerging regulatory framework raises a fundamental empirical question: how does the stock market price firms' exposure to AI when confronted with both technological breakthroughs and regulatory shocks? Answering this question requires a credible measure of what constitutes an "AI stock", which remains a central challenge in the literature. Existing approaches generally rely on internal firm characteristics. Babina, Fedyk, He, and Hodson (2024) construct a measure based on the share of AI-skilled workers in a firm's labor force, using large-scale resume and job-posting data. Eisfeldt, Schubert, Zhang, and Taska (2023) classify firms based on the susceptibility of their occupational task composition to generative AI substitution. Ante and Saggu (2024) apply natural language processing to 10-K filings to quantify the prominence of AI in corporate disclosures. Each approach captures a meaningful dimension of how firms relate to AI internally, but they share a common limitation: for AI exposure to be reflected in stock prices, it is not enough for a firm to be AI-intensive, investors must also recognize it as such. None of the existing measures captures how the market itself classifies AI exposure.

We address this gap by constructing the first market-perceived measures of AI exposure, derived from the holdings and weightings of ten AI-thematic ETFs. We focus on ETFs that provide exposure to AI through their underlying holdings, excluding funds that use AI as an investment-selection tool. We validate this measure by examining how AI-classified firms react to two contrasting AI-related shocks: the EU AI Act and the launch of ChatGPT. This setting motivates two research questions.

1. Do AI-themed ETF holdings and weightings provide a meaningful market-perceived proxy for firm-level AI exposure?

2. Do AI-exposed firms earn systematically different abnormal returns around AI-

relevant regulatory and technological shocks?

To answer these questions, we construct three measures of firm-level AI exposure: a binary indicator capturing whether a firm appears in any AI ETF (*AI ETF*), a frequency-based measure reflecting how consistently a firm is held across the ten funds (*ETF Share*), and a value-weighted measure capturing how much of the aggregate AI ETF portfolio is invested in the firm (*ETF Weight*). We deliberately fix our classification to the pre-ChatGPT, pre-regulation period to ensure that ETF composition reflects underlying AI exposure rather than reactive portfolio rebalancing in response to the events we study.

These three measures capture how ETF providers, index constructors, and the broader market classify firms as AI-related; they do not claim to identify the firms with the greatest operational AI intensity in an objective sense, a question better addressed by fundamentals-based measures such as Babina et al. (2024). Our approach requires only that ETF inclusion reflects the prevailing market classification of AI exposure, since this classification should determine which firms experience the largest valuation effects when AI-related news arrives.

To test these measures, we conduct an event study around two contrasting shocks to the AI sector. First, we examine five legislative milestones of the EU AI Act between April 2021 and March 2024, which represent negative shocks to AI firms through the introduction of compliance costs and regulatory risk. Crucially, the EU AI Act has broad extraterritorial reach: it applies to U.S.-based firms either when their AI systems are placed on the EU market or when their AI outputs are used within the EU, regardless of whether the firm has a physical presence in the Union. Second, we examine the launch of ChatGPT on November 30, 2022, which constitutes an unambiguously positive shock to the commercial prospects of AI. The use of shocks with opposite predicted signs allows us to test whether our exposure measures identify firms whose valuations respond to AI-related news in both directions. This bidirectional validation provides a stronger validation exercise than one-sided tests based only on positive technological shocks, which are common in the literature (Ante and Saggu, 2024).

Our results support the validity of the proposed measures. Around the two cleanest legislative events, the Commission's proposal in April 2021 and the Parliament's formal approval in March 2024, AI-exposed firms earn significantly negative cumulative abnormal returns of 2.0% and 1.9%, respectively, relative to non-AI peers over a $[-1, +10]$ event window. Around the ChatGPT launch, AI-exposed firms outperform by 1.8%, with the value-weighted measure achieving particularly strong significance. When compared to Babina et al.'s (2024) measure, our measures (*AI ETF dummy* and *ETF Share*) are directionally consistent across all events, with *ETF Share* having a moderate correlation of 0.55 with Babina et al.'s (2024) AI Investment score.

Our findings make three contributions to the literature. First, we add to the growing body of work on measuring firm-level AI exposure (Babina et al., 2024; Eisfeldt et al., 2023; Ante and Saggu, 2024) by introducing a market-perceived alternative that is inexpensive to construct, transparent, and replicable using publicly available data. Second, we contribute to the literature on thematic investing and ETFs (Ben-David, Franzoni, Kim, and Moussawi, 2023; Blitz, 2021; Wu and Chen, 2022) by providing evidence that AI-themed ETF holdings, despite the often-vague selection criteria documented in prior work, identify firms whose valuations meaningfully respond to AI-related shocks. Finally, we contribute to literature on the financial market consequences of AI regulation by documenting cross-sectional return reactions to a major regulatory shock (the EU AI Act) which has received comparatively little empirical attention relative to the much-studied technological shocks such as the ChatGPT launch.

The thesis proceeds as follows. Section II reviews the related literature. Section III describes the ChatGPT launch and the EU AI Act. Section IV presents the data and sample construction. Section V constructs the ETF-based AI exposure measures, studies their pricing around AI-related shocks, and validates them against the labor-based measure of Babina et al. (2024). Section VI discusses implications and limitations. Section VII concludes.

II. Theoretical Framework and Literature

A. Classifying an "AI Stock"

Recent literature has proposed several approaches to measuring firm-level AI exposure, each capturing a different dimension of how firms relate to AI, yet none directly reflects how financial markets themselves identify AI-relevant firms.

Babina et al. (2024) measure AI exposure through a firm's active investment in AI-skilled human capital. Using a data-driven algorithm applied to 180 million job postings and resume histories for 535 million individuals, they classify workers as AI-related based on the co-occurrence of their skills with core AI competencies such as machine learning and natural language processing. Aggregating to the firm level, their measure captures the intensity with which a firm is building AI capabilities through hiring. They find that firms investing more heavily in AI talent experience faster growth through increased product innovation, suggesting that AI's primary effect to date has been in reducing the costs of product development.

Eisfeldt et al. (2023) take a different approach. Rather than measuring firms' AI investment, they measure firms' exposure to AI through their existing workforces. They classify roughly 19,000 occupational tasks according to whether ChatGPT can perform them more effectively, then map these tasks to firms using occupational employment shares. This yields a firm-level measure of the share of labor input exposed to potential substitution by generative AI. Around the release of ChatGPT, firms in the highest exposure quintile earned daily returns 0.44% higher than firms in the lowest quintile. This evidence is consistent with investors pricing the productivity effects of generative AI.

Ante and Saggu (2024) instead infer AI engagement from what firms say about themselves, applying natural language processing to 10-K filings for 3,395 NASDAQ-listed firms to quantify how prominently AI features in corporate disclosures. Their measure rests on the intuition that firms genuinely committed to AI will discuss it more extensively in their filings. Consistent with this, firms with higher AI disclosure scores earned significantly greater abnormal returns following the ChatGPT release.

These three approaches ask how deeply a firm is involved in AI operationally, but they share a common limitation: they measure AI exposure through internal firm characteristics that are often observed with a lag, require proprietary data, and may not reflect how investors actually perceive and price AI relevance. Our measures ask a different and more directly relevant question for asset pricing: which firms does the market classify as AI-exposed?

B. Thematic Investing and Exchange Traded Funds

Thematic investing provides exposure to long-run macroeconomic, technological, and geopolitical trends (Claus and Krippner, 2018). Thematic ETFs are a prominent implementation of this strategy: rather than requiring investors to select individual firms, these funds package securities associated with a specific theme into tradable portfolios (Somefun, Perchet, Yin, and Leote de Carvalho, 2023). Common themes include artificial intelligence, robotics, and ESG. This structure makes thematic investing accessible to investors who lack the resources or expertise to evaluate individual firms. Thematic ETF holdings therefore provide a useful setting for studying how market participants classify firms as exposed to emerging economic themes.

ETFs differ in the strategies they use to select holdings. Using FactSet ETF Classification System (2018), Vanilla and Equal strategies are rule-based: a Vanilla ETF tracks a broad market index such as the S&P 500, while an Equal ETF assigns the same weight to each constituent regardless of size. Active strategies, by contrast, are discretionary, with fund managers selecting holdings based on proprietary research, forecasts, and strategic judgement. Each of these strategies can be applied to a thematic mandate; an AI-themed ETF with a vanilla strategy, for example, simply tracks an AI-focused index such as the "NYSE FactSet Global Robotics and Artificial Intelligence Index". Most thematic ETFs follow an external index rather than constructing their own. According to Vanguard (2026), index providers build these indices using a combination of quantitative screens and qualitative criteria, and ETF issuers pay a licensing fee for the right to track them.

AI ETFs reflect the combined judgment of three groups: ETF issuers, index providers, and the investors who buy these products. A fund whose holdings strayed too far from the market's view of AI would struggle to attract capital, since it would no longer credibly track the theme. Competition therefore keeps inclusion criteria reasonably consistent across funds, and ETF holdings end up reflecting a consensus about which firms the market treats as AI-relevant.

C. AI's Effect on Firm Performance

A growing body of research documents that AI exposure has material consequences for firm performance, and therefore for firm valuations. Babina et al. (2024) find that a one-standard-deviation increase in firm-level AI investments over the 2010-2018 period is associated with a 19.5% increase in sales, an 18.1% increase in employment, and a 22.3% increase in market valuation, with the effects driven primarily by product innovation and concentrated among larger firms. Eisfeldt et al. (2023) find that firms with high workforce exposure to Generative AI experienced market value increases of approximately 5% relative to low-exposure firms in the two weeks following the launch of ChatGPT, accompanied by declines in wages and job postings for the most exposed occupations.

These findings provide the economic foundation for our event study. If AI exposure materially affects firm value, then regulatory shocks that alter the cost or feasibility of AI deployment should generate measurable cross-sectional differences in returns between AI-exposed and non-AI-exposed firms.

III. EU AI Act Legislative Events and the ChatGPT Release

We examine two contrasting shocks to the AI market: the EU AI Act, which likely weighs negatively on AI-exposed firms through compliance costs and regulatory risk, and the launch of ChatGPT, which represents an unambiguously positive shock to the commercial prospects of AI. Studying events of opposite sign allows us to test whether our exposure measures identify firms whose valuations are sensitive to AI-related news in either direction.

A. The EU AI Act

According to the European Parliament (2024), the EU AI Act is the European Union's comprehensive legal framework for regulating artificial intelligence. It establishes a risk-based classification system in which AI applications are categorised as posing unacceptable, high, limited, or minimal risk, with compliance obligations scaling accordingly. Unacceptable applications, such as social scoring, are prohibited outright, while high-risk applications, including AI used in employment, credit scoring, and critical infrastructure, face extensive requirements covering data governance, transparency, and human oversight. Non-compliance can result in fines of up to €35 million or 7% of global annual turnover. Crucially, the Act has explicit extraterritorial reach, applying to non-EU providers whenever their systems' outputs are used within the Union, meaning that U.S. AI firms are directly affected regardless of where they are headquartered.

The Act's economic consequences fall asymmetrically across firms: companies with significant AI exposure face direct compliance costs and potential liability, while firms with limited AI exposure are largely unaffected. This asymmetry is what makes the Act useful from an event study perspective, as it provides exactly the differential exposure needed to identify cross-sectional market reactions. The Act is also well-suited to event study methodology because its legislative milestones are discrete, officially announced events with exact dates. Moreover, the Act's legislative process unfolded over several years through clearly separated institutional steps, allowing us to study multiple distinct events and examine whether market reactions are consistent across these milestones.

From the full timeline we selected five events based on their 1) potential to materially shift investor expectations, focusing on events that either introduced new regulatory information or 2) substantially reduced uncertainty about the eventual scope and likelihood of implementation. Although the EU AI Act should affect AI firms negatively overall, the market reaction to each individual milestone may be positive or negative depending on how it compares to investor expectations. Table I summarises the scope of the events and their timing. We identified the legislative milestones using two sources: four events were obtained directly from the The EU Artificial Intelligence Act's (2026) official timeline, and the fifth, the European Parliament's plenary vote of 13 March 2024, was identified through the European Parliament's (2024) website.

Table I describes the events of the EU AI Act that were selected:

Table I

EU AI Act Legislative Events

This table lists the five EU AI Act legislative events used in the event study. Events are selected based on their significance in the legislative process.

No.	Date	Event	Description
1	21 Apr 2021	Commission Proposal	The European Commission published its proposal to regulate artificial intelligence in the European Union, marking the formal start of the legislative process.
2	6 Dec 2022	Council General Approach	The EU Council adopted its common position on the AI Act, establishing the Council's negotiating mandate ahead of trilogue discussions.
3	14 Jun 2023	Parliament Negotiating Position	The European Parliament adopted its negotiating position on the AI Act with 499 votes in favour, 28 against, and 93 abstentions, setting the Parliament's mandate for trilogue negotiations.
4	9 Dec 2023	Provisional Agreement	The European Parliament and the Council reached a provisional political agreement on the final text of the EU AI Act.
5	13 Mar 2024	Formal Parliament Approval	The European Parliament formally approved the EU AI Act, completing the legislative process.

Event 1: Commission Proposal (21 April 2021). This event provided the market with the first concrete regulatory framework for AI in Europe. While regulation had been anticipated, its scope and structure were uncertain, and the proposal both reduced this uncertainty and introduced the possibility of future compliance costs and differentiated effects across firms depending on their AI exposure. It marked the first major signal of how seriously the EU intended to regulate the AI sector (EU Artificial Intelligence Act, 2026).

Event 2: Council General Approach (6 December 2022). The Council's adoption of a general approach signalled that the legislation had gained substantial support among EU member states and was progressing beyond the Commission stage. By increasing the perceived likelihood of eventual implementation, the event reduced investor beliefs that the proposal would be significantly delayed, weakened, or abandoned (European Commission, 2023).

Event 3: Parliament Negotiating Position (14 June 2023). The European Parliament's formal adoption of its negotiating position moved the regulation closer to the trilogue stage and provided additional information about the direction and likely strictness of the final law. This influenced expectations not only about whether regulation would pass, but also about its eventual form, including the future compliance burden and strategic constraints faced by AI-related firms (European Parliament, 2023).

Event 4: Provisional Agreement (9 December 2023). A provisional agreement between Parliament and Council indicated that the main political negotiations had concluded and that the overall shape of the regulation was largely settled. This event substantially reduced legislative uncertainty, raising the probability of implementation sharply and allowing the market to begin pricing in the expected economic consequences of the regulation with greater confidence (European Council, 2023).

Event 5: Formal Parliament Approval (13 March 2024). The European Parliament's plenary vote formally adopted the EU AI Act, passing with a large majority of 523 votes in favour and 46 against. The vote confirmed that the regulation would proceed to final enactment without material changes, removing any remaining legislative risk. While much of the uncertainty had already been resolved at Event 4, this vote represented the final confirmation that the Act would enter into force (European Parliament, 2024).

B. ChatGPT Release

As a further use-case of our exposure measures, we examine the market reaction to the launch of ChatGPT on November 30, 2022. This event is the standard reference point in studies of AI-related firms, as it represents an unambiguous positive shock to the commercial prospects of artificial intelligence. If our ETF-based measures successfully identify firms with genuine AI exposure, they should exhibit significantly positive abnormal returns around this event relative to non-AI ETF firms.

C. Unexpected Events: Google Trends Data

The purpose of this section is to verify that our selected events were genuinely surprising to investors rather than already priced in by the market. If a legislative milestone was widely anticipated, the observed abnormal return around the event date will understate its true informational effect, potentially producing a null result even when the event was substantively important. To address this anticipation problem, we use Google Trends search volume as a proxy for investor attention, following Da, Engelberg, and Gao (2011), who establish search volume as a reliable indicator of attention and informational surprise. The intuition is straightforward: a genuinely unexpected event should generate a sharp spike in public search activity, whereas an anticipated event should produce little movement, since the information was already known.

We download weekly Google Trends indices for three search term combinations: *eu ai act + the eu ai act*, *"ai act" + "the ai act"*, and *eu ai regulation + eu ai law* (Google Trends, 2026). They were queried simultaneously over a single session to ensure all three series are normalised to the same 0-100 scale. Figure 1 shows search volume from January 2021 through December 2024, with our five event dates marked.

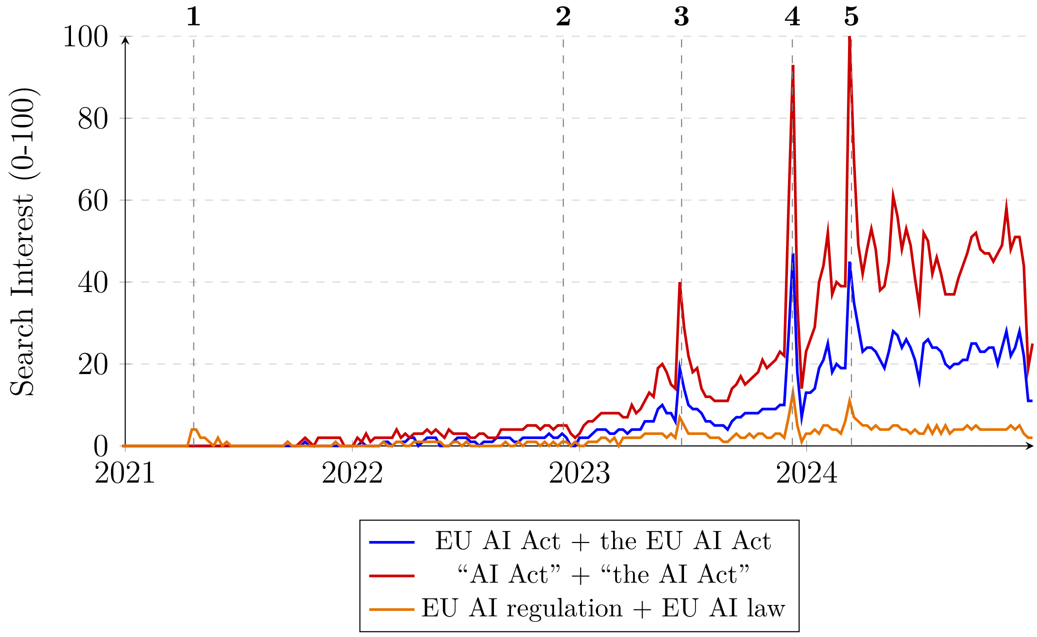

Figure 1. Google Trends Data: EU AI Act. This figure plots weekly Google Trends search interest for three groups of EU AI Act-related search terms worldwide from December 2020 to December 2024. Search interest is measured on a 0-100 scale and normalised jointly, so that 100 corresponds to the peak week of March 13, 2024, when the European Parliament formally adopted the EU AI Act. The dashed vertical lines mark the five legislative events used in the event study: Commission Proposal, April 21, 2021; Council General Approach, December 6, 2022; Parliament Position, June 14, 2023; Provisional Agreement, December 9, 2023; and Final Adoption, March 13, 2024.

Discrete spikes are visible in the weeks surrounding four of our five event dates, indicating that these milestones generated genuine investor attention rather than passing unnoticed. The one event without a clear spike, Event 2, was likely largely anticipated by the market, a possibility we discuss further in the empirical results.

Figure 2 plots search activity around the launch of ChatGPT, with a sharp spike visible immediately after the 30 November 2022 release. The pattern is consistent with the launch being largely unanticipated and triggering a sharp rise in public interest in AI.

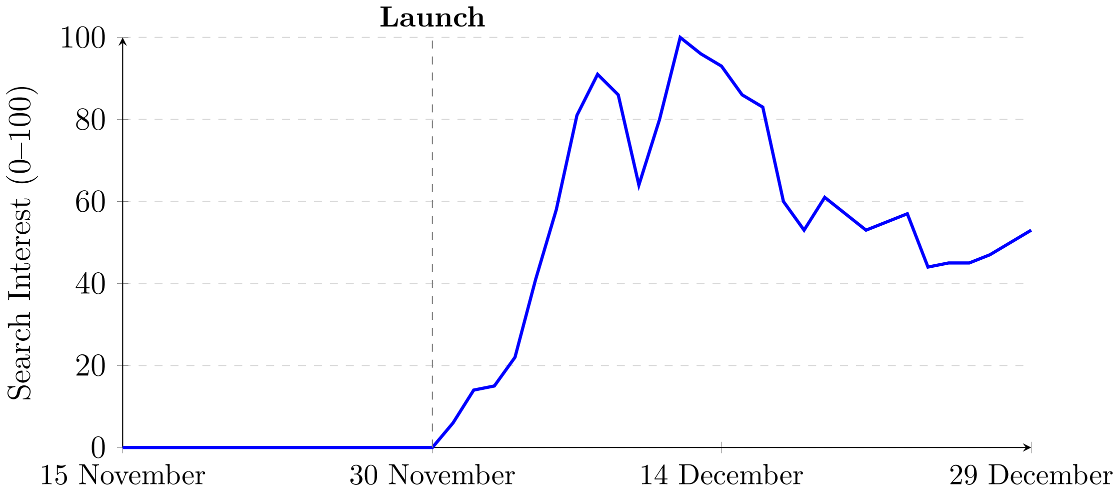

Figure 2. Google Trends Data: ChatGPT Launch. This figure plots daily Google Trends search interest for "ChatGPT" worldwide from November 15 to December 29, 2022. Search interest is measured on a 0-100 scale, where 100 corresponds to the peak day of December 12, 2022. The dashed vertical line marks the public launch of ChatGPT on November 30, 2022.

The combination of unanticipated events (visible as search-volume spikes) and clean event windows (free of confounding news) is what makes a particular milestone informative for our exposure measure.

IV. Data

This study draws on data from two main sources: Wharton Research Data Services (WRDS) and LSEG Workspace. We retrieve stock returns through WRDS, which provides access to the Centre for Research in Security Prices (CRSP) database, while ETF holdings, firm identifiers, 3-digit NAICS-codes and sample characteristics are obtained from LSEG Workspace.

We focus on firms listed in the U.S. because restricting to a single jurisdiction holds the domestic regulatory environment constant across firms. Our sample consists of the constituents of the Russell 3000 Index as of December 31, 2020, retrieved through the Vanguard Russell 3000 ETF (VTHR) (LSEG Workspace, 2026). The Russell 3000 measures the performance of the largest 3,000 U.S. firms by market capitalization and represents approximately 98% of the investable U.S. equity market (FTSE Russell, 2020), making it a natural benchmark universe for studying U.S. public firms.

To identify AI-related ETFs, we search the complete universe of ETFs available in LSEG Workspace for trading between 1 January 2016 and 31 December 2020 for funds containing either "AI" or "Artificial Intelligence" in their name. The period is restricted to 2016-2020 to ensure that ETF composition reflects pre-regulatory AI investment themes rather than investor responses to the legislative events studied. The end date of 2020 is chosen because the EU AI Act's legislative process begins in early 2021 and continues until 2024. The screening initially yields 12 ETFs. We exclude 2 that use AI as a portfolio selection tool rather than investing in AI-related firms, leaving 10 qualifying ETFs with a combined 438 unique constituent firms and valid CUSIP identifiers. Tables II list these ETFs alongside their descriptive statistics.

Table II
AI ETFs: Selection Methodology and Asset Selection Criteria

This table summarises the selection methodology of the ten AI-themed ETFs in our sample. For each fund, we report the ticker, full name, replication strategy, top three portfolio holdings with their respective weights as of December 31, 2020, and a representative excerpt of the asset selection criteria as stated in the fund's official documentation. Strategy types are classified as *Active* for discretionary stock selection by the fund manager, *Vanilla* for passive replication of an underlying index using market-capitalisation weighting, *Equal* for passive replication with equal weighting across constituents, and *ESG* for index-based strategies with an additional environmental, social, and governance screen.

Ticker	**Name**	**Strategy**	**Top 3 Holdings**	**Asset Selection Criteria Excerpt**
AIFD	TCW Artificial Intelligence ETF	Active	Alphabet (3.96%) Amazon (3.91%) Meta (3.89%)	"…companies that are positioned to benefit from the long-term adoption of artificial intelligence across industries, sectors, and geographies…" (TCW, 2020)
BOTZ	Global X Robotics & Artificial Intelligence ETF	Vanilla	Intuitive (7.97%) Nvidia (7.86%) Fanuc (7.72%)	"…companies that potentially stand to benefit from increased adoption and utilization of robotics and artificial intelligence…" (Global X, 2020)
AIVC	Amplify Bloomberg AI Value Chain ETF	Equal	Nice (4.84%) Elastic (4.48%) Open Text (4.17%)	"…a global mix of semiconductor, cloud/software and hardware companies that form the foundation of artificial intelligence (AI) technologies…" (Amplify, 2020)
RBOT	Global X Robotics & AI Index ETF	Vanilla	Intuitive (7.97%) Nvidia (7.88%) Fanuc (7.77%)	"…equity securities of companies that are involved in the development of robotics and/or artificial intelligence…" (Global X, 2020)
ROBT	First Trust Nasdaq AI and Robotics ETF	Equal	Blue Prism (2.20%) Pros (2.12%) Ciena (2.11%)	"…companies engaged in Artificial Intelligence (AI), robotics and automation […] as determined by the Consumer Technology Association (CTA)…" (First Trust, 2020)
AIQ	Global X Artificial Intelligence & Technology ETF	Vanilla	Qualcomm (4.09%) Samsung (3.97%) Meituan (3.56%)	"…companies that potentially stand to benefit from the further development and utilization of artificial intelligence (AI)…" (Global X, 2020)

Continued on next page

Table II continued

Ticker	Name	Strategy	Top 3 Holdings	Asset Selection Criteria Excerpt
ARTY	iShares Future AI & Tech ETF	Vanilla	MTS Systems (1.28%) Baidu (1.28%) LG (1.28%)	"... U.S. and non-U.S. companies that provide products and services that are expected to contribute to artificial intelligence (AI) technologies in areas including generative AI, AI data and infrastructure, AI software, and AI services... " (iShares, 2020)
GOAI	Amundi MSCI Robotics & AI UCITS ETF Acc	ESG	Bayer (9.19%) Repsol (7.52%) Santander (6.42%)	"... companies associated with the increased adoption and utilization of artificial intelligence, robots and automation... " (Amundi, 2020)
XAID	Xtrackers AI and Big Data UCITS ETF 1C	Vanilla	Strategy (2.51%) NCR Voyix (1.66%) 8x8 (1.64%)	"... companies that are most active in filing patents related to Deep Learning, Natural Language Processing, Image Recognition, Speech Recognition & Chatbots, Cloud Computing, Cybersecurity, and Big Data... " (Xtrackers, 2020)
AIAG	L&G Artificial Intelligence UCITS ETF USD Acc	Vanilla	Baidu (2.38%) Proofpoint (1.85%) Varonis (1.81%)	"... companies developing the technology and infrastructure enabling AI... " (L&G, 2020)

After cross-referencing against the Russell 3000, 226 firms appear in at least one AI-related ETF and form our AI-exposed sample, while the remaining 2,136 Russell 3000 firms form the benchmark sample. 638 firms are excluded because they were delisted after 31 December 2020 and therefore lack a valid CUSIP identifier. Table III reports descriptive statistics for both groups.

Table III

AI ETFs: Inception, AUM, and Russell 3000 Coverage

This table reports descriptive statistics for the ten AI-themed ETFs in our sample as of December 31, 2020. For each fund, we report the ticker, inception date, total assets under management (AUM, in millions of USD), and the total number of holdings. The final two columns restrict attention to the subset of each fund's holdings that are constituents of the Russell 3000 index, reporting the corresponding AUM and number of holdings. The bottom row reports the aggregate values across all ten ETFs. The Russell 3000 columns form the basis of our AI-exposed sample construction described in Section IV.

		Total		Russell 3000	
Ticker	Inception	AUM	# Stocks	AUM	# Stocks
AIFD	2017-08-31	24	55	24	49
BOTZ	2016-12-09	2,357	30	2,272	11
AIVC	2016-03-08	60	66	59	42
RBOT	2016-12-09	60	29	58	11
ROBT	2018-02-21	185	107	183	56
AIQ	2018-05-11	147	84	146	58
ARTY	2018-06-26	294	112	292	53
GOAI	2019-09-04	79	75	78	40
XAID	2019-01-29	81	97	81	65
AIAG	2019-07-02	277	69	271	53
Total		**3,566**	**724**	**3,464**	**438**

Notes: AUM figures are in millions of USD. The Russell 3000 columns report only those holdings that are constituents of the Russell 3000 at the same date.

Table IV reports summary statistics for the different groups.

Table IV

Summary Statistics

This table reports summary statistics for the main firm-level variables used in the analysis. Panel A reports statistics for the full sample of U.S. publicly traded firms. Panels B and C split the sample into firms classified as AI-exposed through AI-themed ETF inclusion and firms not included in AI-themed ETFs. MktCap is measured in billions of U.S. dollars, Leverage is measured in percent, Beta captures firms' market risk, and Volatility captures return variability. ETF Share and ETF Weight capture the breadth and value-weighted intensity of firms' exposure across AI-themed ETFs.

Variable	N	Mean	Median	Std. Dev.	Min	Max
Panel A: Full Sample (N = 2,362)						
MktCap ($B)	2,324	16.864	2.602	80.162	0.062	2,232.280
Leverage (%)	2,318	39.929	36.650	47.551	−347.560	1,207.600
Beta	2,314	1.319	1.240	0.851	−9.450	7.098
Volatility	2,318	0.466	0.405	0.238	0.012	4.995
ETF Share	226	0.184	0.100	0.132	0.100	0.800
ETF Weight	226	0.0044	0.0009	0.0152	0.0000	0.1325
Panel B: AI ETF Firms (N = 226)						
MktCap ($B)	226	97.263	29.590	251.761	0.631	2,232.280
Leverage (%)	226	44.823	36.985	62.144	0.000	616.770
Beta	226	1.279	1.161	0.713	−0.111	7.098
Volatility	226	0.416	0.366	0.198	0.175	1.353
ETF Share	226	0.184	0.100	0.132	0.100	0.800
ETF Weight	226	0.0044	0.0009	0.0152	0.0000	0.1325
Panel C: Non-AI ETF Firms (N = 2,136)						
MktCap ($B)	2,132	9.624	2.218	26.251	0.062	387.335
Leverage (%)	2,126	39.487	36.595	46.000	−347.560	1,207.600
Beta	2,122	1.322	1.252	0.862	−9.450	6.597
Volatility	2,126	0.471	0.409	0.241	0.012	4.995
ETF Share	—	—	—	—	—	—
ETF Weight	—	—	—	—	—	—

We retrieve abnormal returns through the WRDS Event Study tool (2026), which builds on the CRSP database. For each firm-event observation, the tool returns raw returns over symmetric event windows of up to ± 10 trading days around each event date. Abnormal returns are estimated using the Fama–French three-factor model, with parameters estimated over a 255-day pre-event window ending 46 days before the event date. The event study output is then combined with our AI exposure measures in Python to produce the final cross-sectional results.

V. Empirical Analysis

A. Constructing AI Exposure Measures

We construct three measures of AI exposure, each capturing a different dimension of a firm's presence in AI-themed ETFs.

The first measure, which we call *AI ETF*, is a binary indicator equal to one if a firm appears as a constituent in at least one of the 10 identified AI ETFs, and zero otherwise:

$$\text{AI ETF}_i = \mathbf{1}\left\{\sum_{j=1}^{J} \mathbf{1}\{i \in \mathcal{H}_j\} > 0\right\}, \tag{1}$$

where $J = 10$ is the number of AI-themed ETFs, $\mathcal{H}_j$ is the set of firms held by ETF j, and $\mathbf{1}\{i \in \mathcal{H}_j\}$ equals one if firm i is held by ETF j, and zero otherwise. Of the approximately 2,400 Russell 3000 constituents in our sample, 226 are classified as AI-exposed. The intuition is straightforward: inclusion in an AI-themed ETF reflects a classification, by either ETF managers or the index providers they track, that a firm is part of the AI theme. The measure thus uses the aggregated classifications of ETF issuers and index providers to draw a coarse boundary between firms the market treats as AI-relevant and those it does not.

The second measure, which we call *ETF Share*, builds on the binary classification by measuring the intensity of AI exposure. Rather than treating all ETF-included firms symmetrically, it assigns each firm a score equal to the fraction of the 10 ETFs in which it appears:

$$\text{ETF Share}_i = \frac{1}{J}\sum_{j=1}^{J} \mathbf{1}\{i \in \mathcal{H}_j\}, \tag{2}$$

where $J = 10$ is the number of AI-themed ETFs, $\mathcal{H}_j$ is the set of firms held by ETF j, and $\mathbf{1}\{i \in \mathcal{H}_j\}$ equals one if firm i is held by ETF j, and zero otherwise. A firm held by all 10 ETFs receives a score of one, while a firm held by one ETF receives a score of 1/10. The intuition is that firms included by more AI ETFs are more consistently classified by the market as AI-relevant, whereas inclusion in only one ETF may reflect a more idiosyncratic classification. Nvidia, for example, appears in 8 of the 10 ETFs, yielding an ETF Share score of 8/10. Appendix Table IX reports the 15 firms with the highest ETF Share scores.

The third measure, which we call *ETF Weight*, incorporates the size of each firm's position within AI ETFs. We construct a synthetic AI ETF by pooling the holdings of all 10 funds and using each reported dollar holding value as of December 31, 2020. We then define firm i's exposure as its aggregate holding value across all AI ETFs, scaled by the total value of all holdings in the synthetic portfolio:

$$\text{ETF Weight}_i = \frac{\sum_{j=1}^{J} V_{ij}}{\sum_{j=1}^{J}\sum_{k\in\mathcal{H}_j} V_{kj}}, \tag{3}$$

where V_{ij} is the dollar value of ETF j's position in firm i, with $V_{ij} = 0$ if ETF j does not hold the firm. This measure captures not only whether a firm is held by AI ETFs, but also how prominently it appears in aggregate AI-themed capital allocation. Nvidia again receives the highest score, accounting for approximately 13% of the synthetic portfolio's total assets. The advantage of ETF Weight over ETF Share is that it reflects the economic size of each position. A firm that receives large allocations across several ETFs is treated as more AI-exposed than a firm that is included only at negligible weights. The tradeoff is that ETF Weight also reflects fund-level portfolio construction choices, which may introduce variation unrelated to the firm's underlying AI relevance.

Together, the three measures form a hierarchy of granularity: *AI ETF* provides a discrete classification of AI exposure, *ETF Share* captures the consistency with which market participants identify a firm as AI-relevant, and *ETF Weight* further incorporates the economic magnitude of that exposure. Using all three allows us to assess whether our findings are robust to the choice of measurement strategy and to draw more nuanced conclusions about the cross-sectional structure of AI exposure.

B. Event Study Results

To validate our AI exposure measures, we conduct an event study around five key legislative milestones of the EU AI Act over the period 2020-2024, as well as the launch of ChatGPT in November 2022. For each event, we first present the difference in cumulative abnormal returns (CARs) between AI-exposed and non-AI-exposed firms using the *AI ETF* dummy, before turning to the regression-based estimates using all three exposure measures.

We compute CARs for each firm over an event window of $[-1, +10]$ relative to each event date. The choice of window reflects a deliberate tradeoff. A sufficiently wide window is necessary to allow investors to fully digest the regulatory implications of each legislative milestone, particularly given that the EU AI Act's risk-based compliance framework involves complex, firm-specific consequences that may not be immediately apparent to market participants. At the same time, an overly wide window risks contamination from unrelated AI-related news or macroeconomic developments occurring in close proximity to the event dates. The $[-1, +10]$ window balances these concerns, providing sufficient time for the market to absorb the informational content of each event while remaining narrow enough to limit confounding noise. Results for the $[-1, +5]$ window are reported in Appendix Table XI as a robustness check.

A natural concern with our identification strategy is that AI-themed ETFs disproportionately hold firms in technology-related sectors (semiconductors, software publishing, data processing) and that these sectors react to macroeconomic news in ways unrelated to AI specifically. To address this, we estimate all regressions both with and without 3-digit NAICS sector fixed effects. Standard errors are clustered at the NAICS-3 level, since firms within the same broad industry may have correlated abnormal returns around the events. Heteroskedasticity-robust (HC1) standard errors are reported in Appendix Table XII as a robustness check. Of the 81 unique 3-digit NAICS codes in our sample, 26 contain both AI-classified and non-AI-classified firms and contribute to fixed-effects identification; the AI coefficient under FE is identified off within-sector variation between AI and non-AI firms within these codes. This isolates the AI-specific component of the return reaction from sector-wide co-movement.

We note that the NAICS-3 fixed-effect specification does not show that our measure captures AI exposure independently of broader sector characteristics. It shows only that our results are not entirely driven by the concentration of AI firms in technology-adjacent sectors. The ETF-based measure likely reflects both AI exposure and broader exposure to technology, growth firms, and semiconductors, since AI-themed ETFs invest heavily in these related areas. We therefore interpret the measure as capturing market-perceived AI relevance, which may be correlated with technology and growth exposure, rather than as a pure AI signal orthogonal to all sector characteristics.

B.1 *The EU AI Act*

Table V reports market reactions to the five EU AI Act events, as well as the release of ChatGPT.

Table V

Cumulative Abnormal Returns: AI vs. Non-AI Firms

This table reports the difference in mean cumulative abnormal returns (CARs) between AI-classified and non-AI-classified Russell 3000 firms over the $[-1, +10]$ event window. AI classification is based on inclusion in any of the 10 AI-themed ETFs identified in our sample. Abnormal returns are estimated using the Fama–French three-factor model with a 255-day estimation window ending 46 trading days before the event. The Unconditional columns report differences in mean CAR with heteroskedasticity-robust (t-test) standard errors. The NAICS-3 FE columns report the coefficient on the AI ETF dummy from a regression that includes 3-digit NAICS sector fixed effects, with standard errors clustered at the NAICS-3 level. The FE specification is identified off within-sector variation between AI and non-AI firms within the 26 NAICS-3 codes that contain both groups. t-statistics in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

Event	**Unconditional** AI − Non-AI	**NAICS-3 FE** AI − Non-AI
Panel A: EU AI Act Legislative Milestones		
Event 1: Commission Proposal (21 Apr 2021)	−1.991%***	−0.807%
	(−3.74)	(−1.62)
Event 2: Council General Approach (6 Dec 2022)	2.552%**	0.095%
	(2.24)	(0.14)
Event 3: Parliament Position (14 Jun 2023)	1.032%	2.149%
	(0.87)	(1.06)
Event 4: Provisional Agreement (11 Dec 2023)	0.263%	−1.640%
	(0.43)	(−1.25)
Event 5: Formal Approval (13 Mar 2024)	−1.874%**	−0.948%**
	(−2.04)	(−2.09)
Panel B: ChatGPT Launch		
ChatGPT Launch (30 Nov 2022)	1.836%**	0.883%
	(2.20)	(1.30)

Notes: Event 4 (Provisional Agreement) occurred on Saturday, 9 December 2023; because financial markets were closed, the event window begins on the following trading day, Monday, 11 December 2023.

Table VI

Cross-Sectional Regressions of CARs on AI Exposure Measures

This table reports separate cross-sectional regressions of cumulative abnormal returns over the [−1, +10] event window on three measures of AI exposure: *AI ETF*, a binary indicator equal to one if the firm is held by at least one AI-themed ETF; *ETF Share*, the share of the 10 AI ETFs in which the firm is included; and *ETF Weight*, the firm's value-weighted share of the synthetic AI portfolio formed by pooling all 10 ETFs. Each cell reports the coefficient on the indicated AI exposure measure from a separate regression. The *Unconditional* columns use heteroskedasticity-robust (HC1) standard errors. The *NAICS-3 FE* columns include 3-digit NAICS sector fixed effects with standard errors clustered at the NAICS-3 level; the AI exposure coefficient is identified off within-sector variation. t-statistics in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	Unconditional			**NAICS-3 FE**		
Event	*AI ETF*	*ETF Share*	*ETF Weight*	*AI ETF*	*ETF Share*	*ETF Weight*
Panel A: EU AI Act Legislative Milestones						
Event 1	−0.020***	−0.077***	−0.322	−0.008	−0.021	0.263**
	(−3.74)	(−3.20)	(−1.01)	(−1.62)	(−1.15)	(2.19)
Event 2	0.026**	0.116***	1.786*	0.001	0.019	0.687***
	(2.24)	(2.91)	(1.82)	(0.14)	(0.67)	(6.29)
Event 3	0.010	0.004	−0.133	0.021	0.023	−0.121
	(0.88)	(0.16)	(−0.40)	(1.06)	(0.50)	(−0.53)
Event 4	0.003	−0.002	0.042	−0.016	−0.064	−0.278
	(0.43)	(−0.08)	(0.16)	(−1.25)	(−1.39)	(−1.07)
Event 5	−0.019**	−0.081***	−0.410	−0.009**	−0.040*	−0.087
	(−2.04)	(−3.08)	(−1.10)	(−2.09)	(−1.77)	(−0.22)
Panel B: ChatGPT Launch						
ChatGPT	0.018**	0.071**	0.981***	0.009	0.035	0.894***
	(2.21)	(2.15)	(3.18)	(1.30)	(1.57)	(3.12)
NAICS-3 FE	No	No	No	Yes	Yes	Yes
Clustered SE	No	No	No	Yes	Yes	Yes

Event 5 produces the strongest and most robust result. Under the unconditional specification, AI-exposed firms underperform their non-AI peers by 1.9% over the event window. The continuous regressions confirm this pattern: *ETF Share* yields a coefficient of -0.08 with high significance, implying that a firm held by all ten AI ETFs would underperform a firm held by none by approximately 0.8%. Under NAICS-3 FE, the binary differential halves to −0.9% but remains statistically significant at the 5% level, and *ETF Share* remains significant at the 10% level. The fact that the within-sector AI effect survives the most demanding specification establishes Event 5 as the cleanest identification opportunity in our sample, consistent with its role as the legislative milestone that most decisively eliminated remaining regulatory uncertainty.

Event 1 shows directionally consistent but less robust results. Unconditionally, AI-exposed firms underperform by 2.0%, with *ETF Share* again significant at conventional levels. Under NAICS-3 FE, however, the binary differential narrows to −0.8% and loses statistical significance, and the *ETF Share* coefficient is similarly weakened. The directional sign aligns with our predictions, but the FE results suggest that the unconditional underperformance was largely driven by sector-wide reactions among technology firms rather than by AI-specific firm-level pricing within sectors.

Event 2 yields a positive and significant differential, with AI-exposed firms outperforming by 2.6%. Events 3 and 4 yield insignificant unconditional results under the binary measure. Their interpretation requires accounting for two considerations: whether each event was genuinely unanticipated, and whether confounding developments occurred within the event window. The latter is particularly important in our setting. With only five legislative events, a single confounding episode can materially influence the inference drawn from any individual event, unlike a standard event study with many observations where confounders tend to average out.

For Event 2, the public release of ChatGPT eight days earlier likely generated positive momentum for AI-related equities, which would mechanically inflate AI firm CARs around the event window and plausibly accounts for the observed outperformance (OpenAI, 2022). The NAICS-3 FE specification provides direct support for this interpretation: the unconditional differential collapses from +2.6% to essentially zero (+0.1%) once sector composition is absorbed. This is consistent with the ChatGPT spillover operating as a tech-sector-wide phenomenon rather than an AI-specific firm-level effect, and the disappearance of the result under FE strengthens our confounding-factor interpretation. For Event 3, the U.S. Federal Reserve announced a shift toward further interest rate increases on June 15 2023, one trading day after the event date (Federal Reserve, 2023). Since AI firms tend to be classified as growth stocks and growth valuations are particularly sensitive to discount rate changes, this announcement would be expected to exert downward pressure on AI firms independent of any EU regulatory news, complicating identification. For Event 4, Google's announcement of Gemini approximately one week earlier introduced AI-related noise into the event window, plausibly explaining the insignificant results (Pichai and Hassabis, 2023).

Taken together, the pattern of results is consistent with our exposure measures working as intended: they produce strong, directionally consistent effects when events are clean and salient, and null effects when events are either expected or contaminated by confounding news. The Google Trends evidence in Section III reinforces this interpretation. Events 1 and 5 generated discrete attention spikes and were free of confounders, and these are the events where our measures produce statistically significant negative CARs of the predicted magnitude. Events 3 and 4 generated investor attention but coincided with confounding macroeconomic and AI-related news. The null results on these events should therefore not be read as evidence against our measures, but as a reminder that event studies require both surprise and clean identification, neither of which holds uniformly across the EU AI Act timeline. Having documented the EU AI Act's cross-sectional effects, we now turn to a contrasting AI shock: the launch of ChatGPT.

B.2 The ChatGPT Launch

The ChatGPT launch on 30 November 2022 provides arguably the cleanest shock to AI-exposed firms in our sample, as it was a discrete, globally unanticipated event with no obvious confounders within the event window. AI-exposed firms outperform their non-AI peers by 1.8% over the $[-1, +10]$ window under the unconditional specification.

The continuous regressions reinforce this finding across all three exposure measures. *ETF Share* yields a coefficient of 0.07, implying that a firm held by all ten ETFs earns a CAR approximately 0.7% higher than a firm held by none. *ETF Weight* produces a coefficient of 0.98 and is highly significant, indicating that not merely a firm's inclusion but also the economic magnitude of its presence in AI-themed funds was priced around this event. Notably, this is the one event in our sample where *ETF Weight* achieves strong significance, suggesting that the market reaction to ChatGPT was particularly concentrated among the largest AI holdings.

Under NAICS-3 fixed effects, *AI ETF* and *ETF Share* lose statistical significance, with *AI ETF* narrowing to 0.9% and *ETF Share* to 0.035. The *ETF Weight* coefficient, however, remains highly significant, with a coefficient of 0.89, confirming that the value-weighted reaction operates within sectors rather than across them. This pattern is informative: the market's positive reaction to ChatGPT was concentrated in a small number of large AI-exposed firms whose performance diverged from their sector peers, rather than uniformly affecting all firms classified as AI-exposed by ETF inclusion.

Combined with the negative within-sector reaction at Event 5, these results provide bidirectional validation that our measures identify firms whose valuations are sensitive to AI-related news in either direction.

C. Validation Against Babina et al.'s (2024) AI Investment Measure

To further validate our exposure measures, we compare our classifications against those of Babina et al. (2024), who construct a firm-level AI exposure measure using data on AI-related job postings.

Babina et al.'s (2024) measure covers approximately 13,000 firms as of December 31, 2020, of which roughly 1,350 are featured in the Russell 3000 at that point in time. We restrict the comparison to these 1,350 firms to ensure that differences in results reflect variation in classification methodology rather than sample composition. Babina et al.'s (2024) measure is continuous by construction and captures the intensity of a firm's investment in AI-skilled human capital. To facilitate a direct comparison with our binary *AI ETF* indicator, we transform the AI Investment score into a dummy variable using a top-decile cutoff. Firms in the top 10% of the AI Investment score distribution are classified as AI-exposed and assigned a value of one; all remaining firms are classified as non-AI-exposed and assigned a value of zero. This yields 130 AI-exposed firms under Babina et al.'s (2024) definition. Of our 226 AI-classified firms in the full sample, 150 fall within Babina et al.'s (2024) coverage universe and form the AI-exposed group in this restricted comparison.

Table VII
Validation: Babina et al. (2024) vs. Our ETF-Based Measures

This table compares cumulative abnormal returns over the $[-1, +10]$ event window across two independently constructed measures of firm-level AI exposure: the labor-based AI Investment score of Babina et al. (Panel A) and our ETF-based measures (Panel B). The sample is restricted to the 1,350 Russell 3000 firms covered by both datasets. Babina et al.'s continuous score is transformed into a binary indicator using a top-decile cutoff, yielding approximately the same number of AI-classified firms as our *AI ETF* measure. The *binary measure* column reports the mean CAR difference between AI and non-AI firms; the *continuous measure* column reports the regression coefficient on the corresponding continuous variable. All specifications use heteroskedasticity-robust (HC1) standard errors. To match the original methodology of Babina et al., specifications do not include sector fixed effects; sector-controlled results for our measures are reported in Tables V and VI. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

Event	Binary measure	Continuous measure
Panel A: Babina et al. Measure (top-decile dummy and AI Investment score)		
Event 1	−2.372%***	−1.824***
Event 2	1.974%**	1.853***
Event 3	−1.221%	−1.036*
Event 4	0.592%	0.445
Event 5	−1.992%**	−1.644***
ChatGPT	1.488%	1.798***
Panel B: Our Measures (AI ETF dummy and ETF Share)		
Event 1	−1.824%***	−0.080***
Event 2	1.438%*	0.076**
Event 3	−0.088%	−0.019
Event 4	0.687%	0.017
Event 5	−2.736%***	−0.081***
ChatGPT	0.917%	0.054

Table VII shows that the binary measures, our *AI ETF* indicator and Babina et al.'s (2024) measure, yield directionally consistent results across all five EU AI Act events.

For the two cleanest events, the magnitudes are close but not identical: for Event 1, the differential between AI-exposed and non-AI-exposed firms is −2.4% under our measure and −1.8% under Babina et al.'s (2024); for Event 5, the figures are −2.0% and −2.7%, respectively. The modest differences in magnitude are unsurprising given that the two measures are constructed from fundamentally different data sources and may classify marginal firms differently. What matters for validation purposes is that both measures identify the same directional effect on the same events, which they do.

For the ChatGPT event, both measures point in the same positive direction, 1.5% and 0.9% respectively, though neither achieves statistical significance in this restricted sample. We attribute this to the reduction in cross-sectional variation that results from limiting the sample to Babina et al.'s coverage universe rather than to any fundamental disagreement between the measures.

Turning to the continuous measures, *ETF Share* produces results consistent with the full-sample findings reported in Table VI. Babina et al.'s (2024) continuous AI Investment score achieves significance on four of the five EU AI Act events; *ETF Share* reaches significance on three EU AI Act events. Babina et al.'s (2024) measure achieves significance on more events than ours, likely reflecting the richer firm-level information embedded in their continuous score. Even so, the broad consistency across events supports the interpretation that our measures capture a related underlying construct.

This interpretation is further supported by the correlations between measures shown in Table VIII. *ETF Share* and Babina et al.'s (2024) continuous variable exhibit a correlation of 0.55, while *ETF Weight* correlates at 0.37. We view these as evidence of moderate and meaningful agreement. A higher correlation would not necessarily be expected given that the two approaches are built on entirely different data: ours reflects the portfolio decisions of ETF managers, while Babina et al.'s (2024) reflects firm-level investment in AI capabilities. That they agree as much as they do lends credibility to both.

Table VIII

Correlation Between AI Exposure Measures

This table reports pairwise Pearson correlations between three firm-level AI exposure measures. *ETF Share* is the fraction of shares outstanding held by AI-focused thematic ETFs. *ETF Weight* is the average portfolio weight assigned to the firm across all AI ETFs holding it. *AI Investment* is the resume-based firm-level AI investment measure from Babina et al. (2024), constructed from employee resumes as the share of AI-related occupations in a firm's workforce. The sample consists of Russell 3000 firms for which all three measures are available.

	ETF Share	**ETF Weight**	**AI Investment**
ETF Share	1.00		
ETF Weight	0.48	1.00	
AI Investment	0.55	0.37	1.00

VI. Discussion

This study's primary contribution lies in the development and validation of a novel, market-perceived measure of firm-level AI exposure. By classifying firms based on their inclusion and weighting in AI-themed ETFs, our approach offers a fresh perspective that complements existing measures grounded in internal firm characteristics. Importantly, our methodology is not confined to AI; it can in principle be applied to any thematic exposure for which dedicated ETFs exist, given its reliance on publicly available holdings data and a transparent classification logic. The validation of our measures, particularly against the legislative milestones of the EU AI Act and the launch of ChatGPT, supports their effectiveness and underscores their relevance for empirical research on thematic investing.

A particularly notable finding is that AI-themed ETF holdings successfully identify firms whose valuations are sensitive to AI-related news in both directions: AI-classified firms exhibit significantly negative abnormal returns around the cleanest legislative events of the EU AI Act, and significantly positive abnormal returns around the ChatGPT launch. The convergence between our market-perceived measures and the labor-based classification of Babina et al. (2024), together with moderate correlations of 0.55 *ETF Share* and 0.37 *ETF Weight* with their AI Investment score, indicates that our approach captures a meaningful classification of AI exposure rather than purely thematic noise. These findings suggest that thematic ETF inclusion can serve as a credible proxy for market-perceived AI relevance.

A. Implications for Theory

The findings of this study have several theoretical implications for the literature on thematic investing, AI classification, and market reactions to regulatory and technological shocks.

First, our research contributes a novel market-perceived measure of AI exposure to a literature that has so far relied predominantly on internal firm characteristics. Existing approaches measure AI exposure through occupational task composition (Eisfeldt et al., 2023), AI-skilled human capital (Babina et al., 2024), or textual disclosures (Ante and Saggu, 2024). Each captures a meaningful but partial dimension of how firms relate to AI. Our ETF-based approach asks a fundamentally different question: not how AI-intensive a firm is from the inside, but which firms the market itself classifies as AI-exposed. The moderate correlations with Babina et al.'s (2024) measure indicate that our approach captures a related but independently informative signal, suggesting that market-perceived and fundamentals-based measures are complementary rather than substitutable.

Second, this paper contributes to the literature on thematic investing (Blitz, 2021; Ben-David, Franzoni, Kim, and Moussawi, 2023), which emphasizes the speculative and narrative-driven nature of thematic funds. Because AI ETFs are a prominent vehicle through which investors obtain thematic exposure to artificial intelligence, their holdings provide a natural setting for testing whether thematic fund classifications contain information about the firms linked to the stated investment theme. We show that, at least in the context of AI, ETF-classified firms respond consistently to both regulatory and technological AI shocks, suggesting that thematic ETF inclusion contains economically meaningful information about market-perceived AI exposure.

Third, most validation exercises in this literature rely on positive technological shocks, particularly the ChatGPT release. By demonstrating that the same firms respond negatively to regulatory shocks and positively to technological ones, our work shows that exposure measures should ideally be validated bidirectionally. This methodological point generalises to any setting where researchers seek to distinguish genuine thematic sensitivity from one-sided sectoral co-movement.

Finally, our methodology can readily be extended to other thematic domains. Any investment theme with dedicated ETFs (such as cybersecurity, clean energy, semiconductors, or quantum computing) could in principle be analysed using the same framework. This generalisability connects our work to the broader discourse on how data-driven, transparent classification methods can enhance the study of thematic investing across emerging technologies.

B. Implications for Practice

Our findings also have practical implications for investors, fund managers, researchers, and risk managers.

First, our approach offers a transparent and low-cost first pass for identifying AI-relevant firms. ETF holdings data is publicly available, easy to retrieve, and replicable across researchers. Compared to constructing labor-based measures from proprietary job-postings data or scraping and processing 10-K filings at scale, our approach offers a substantial reduction in barriers to entry.

Second, our findings have implications for ETF issuers and index providers. The evidence that thematic AI ETFs collectively identify firms whose valuations are sensitive to AI-related news supports the credibility of their thematic classifications, even if individual fund performance varies. At the same time, the variation in holdings across the ten ETFs we examine, combined with the sensitivity of value-weighted measures like ETF Weight to dominant individual holdings such as NVIDIA, suggests that issuers could benefit from greater methodological transparency about their inclusion criteria.

Third, our measure provides a diagnostic tool for portfolio and risk managers. By aggregating *ETF Share* or *ETF Weight* scores across portfolio holdings, managers can measure portfolio exposure to the AI theme. This can support portfolio construction, performance attribution, and risk assessment around AI-related regulatory, technological, and sentiment shocks.

Fourth, our approach is well suited to questions where the relevant construct is investor perception rather than firm fundamentals. For investors and analysts seeking to identify which firms the market perceives as AI-exposed, and which are therefore likely to react to AI-related news, ETF-based classification offers a direct and economically meaningful signal.

Finally, our measure can serve as a complement to fundamentals-based measures for institutional investment and risk management. Because it captures market-perceived rather than fundamentals-based exposure, it can be used alongside internal AI investment metrics to triangulate a firm's thematic profile. A firm that scores high on Babina et al.'s (2024) measure but low on ours has substantial internal AI capabilities that the market does not yet recognise, a potentially underpriced AI play. Conversely, a firm with high ETF inclusion but low internal investment may be overvalued on AI hype. This two-dimensional framework supports more sophisticated investment and risk management decisions than either measure alone.

C. Limitations and Future Research

This study is subject to several limitations that suggest avenues for future research.

First, our AI classification is based on ETF holdings observed during 2016-2020, which predates both the generative AI wave triggered by ChatGPT and the legislative events we study. While this temporal separation mitigates concerns that fund managers adjusted their holdings in response to the regulation, it also means that our classification may not fully capture the AI landscape as it existed at the time of each event. Firms that became prominent AI players after 2020 but were not yet included in thematic ETFs would be classified as non-AI-exposed in our framework and thus enter our benchmark sample. Future research could address this by using time-varying ETF holdings with a sufficiently long lag (e.g., 12 to 18 months before each event), allowing the classification to evolve with the AI landscape.

Second, our measure is best understood as identifying the most visibly AI-exposed firms rather than as a complete ranking of AI exposure across the Russell 3000. Some of the firms in our benchmark sample likely had some degree of AI activity that simply did not meet the inclusion threshold of any thematic ETF manager. Future work could refine this by comparing a firm's representation in AI-themed ETFs against its representation in broad-market ETFs, which would generate a continuous over- or under-representation score.

Third, the inclusion criteria of thematic ETFs are not fully transparent. Fund managers and index providers rarely disclose the precise weights placed on factors such as revenue exposure, patent activity, or advisory committee assessments, and our measures therefore inherit whatever biases exist in these selection processes. While the consistency of our results across the three exposure measures and their alignment with Babina et al. (2024) suggests that such biases do not dominate, they cannot be ruled out entirely. Future research could examine AI-themed ETFs more directly, for instance by studying whether included firms subsequently exhibit stronger AI investment, innovation output, or operating performance than comparable non-included firms.

Finally, our ETF-based measure may partly reflect broader exposure to other sectors such as technology and semiconductor, rather than AI exposure in a narrow sense. AI-themed ETFs draw heavily from these adjacent categories, and although our 3-digit NAICS sector controls help isolate within-sector AI variation, they cannot fully separate AI exposure from other tech-adjacent firm characteristics. Our measure should therefore be interpreted as capturing market-perceived AI relevance rather than exposure in an operational sense, a construct that inevitably overlaps with broader technology and semiconductor exposure.

VII. Conclusion

This thesis develops and validates a market-based measure of firm-level AI exposure using the holdings of AI-themed ETFs. Existing measures capture AI exposure through firms' internal characteristics, including AI-skilled human capital (Babina et al., 2024), occupational task composition (Eisfeldt et al., 2023), and textual disclosures (Ante and Saggu, 2024). Our approach instead measures which firms financial markets classify as AI-relevant. Using the holdings of ten AI-themed ETFs over 2016-2020, we construct three measures of AI exposure: a binary indicator, *AI ETF*; a frequency-based measure, *ETF Share*; and a value-weighted measure, *ETF Weight*.

We validate these measures using stock-price reactions to two types of AI-related shocks. The first is the EU AI Act, which increased expected compliance costs and regulatory risk for AI-exposed firms. The second is the launch of ChatGPT in November 2022, which increased the perceived commercial value of generative AI. Around the two cleanest legislative events (the Commission Proposal in April 2021 and the European Parliament's formal approval in March 2024) AI-exposed firms underperform non-AI firms by 2.0% and 1.9% respectively, with the March 2024 reaction surviving 3-digit NAICS sector fixed effects. Around the ChatGPT launch, AI-exposed firms outperform non-AI firms by about 1.8%, with the ETF Weight measure remaining significant under sector controls. The opposite signs across regulatory and commercial shocks indicate that our measures identify firms whose valuations are sensitive to AI-related news in either direction.

Comparison with the labor-based measure of Babina et al. (2024) provides additional validation. Our binary classifications produce directionally similar return patterns across all five EU AI Act events. Moreover, ETF Share is moderately correlated with their AI Investment score, with a correlation of 0.55. This suggests that the two measures capture related, but distinct, dimensions of firm-level AI exposure. This is expected, since the measures are constructed from different data sources.

The measures are useful for both research and practice. For empirical researchers, ETF holdings provide a transparent and low-cost way to classify thematically exposed firms using publicly available data. For investors and risk managers, the measures provide a simple diagnostic for measuring portfolio exposure to the AI theme and for assessing sensitivity to AI-related news. The methodology can also be applied to other themes, such as cybersecurity, clean energy, semiconductors, or quantum computing, whenever dedicated thematic ETFs exist. More broadly, the ETF-based approach complements fundamentals-based measures. It captures how financial markets perceive firms' AI relevance, while previous measures capture how firms are operationally exposed to AI.

References

Amplify ETFs. (2020). *Amplify Bloomberg AI Value Chain ETF (AIVC)*. Retrieved from https://amplifyetfs.com/aivc/

Amundi ETF. (2020). *Amundi MSCI Robotics & AI UCITS ETF Acc (GOAI)*. Retrieved from https://www.amundietf.se/en/professional/products/equity/amundi-msci-robotics-ai-ucits-etf-acc/lu1861132840

Ante, L., & Saggu, A. (2024). Quantifying a firm's AI engagement: Constructing objective, data-driven AI stock indices using 10-K filings. *Technological Forecasting and Social Change*, 212, 123965.

Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151.

Ben-David, I., Franzoni, F., Kim, B., & Moussawi, R. (2023). Competition for attention in the ETF space. *Review of Financial Studies*, 36(3), 987–1042.

Blitz, D. (2021). Betting against quant: Examining the factor exposures of thematic indexes. *Journal of Index Investing*, 12, 5–14.

Claus, I., & Krippner, L. (2018). Contemporary topics in finance: a collection of literature surveys. *Journal of Economic Surveys*, 32(5), 1221–1228. https://doi.org/10.1111/joes.12298

Da, Z., Engelberg, J., & Gao, P. (2011). In search of attention. *Journal of Finance*, 66(5), 1461–1499.

Eisfeldt, A. L., Schubert, G., Zhang, M. B., & Taska, B. (2023). Generative AI and firm values. *Journal of Finance*.

EU Artificial Intelligence Act. (2026a). *EU AI Act Explorer*. Retrieved from https://artificialintelligenceact.eu/ai-act-explorer/

EU Artificial Intelligence Act. (2026b). *EU AI Act: Developments*. Retrieved from https://artificialintelligenceact.eu/developments/

European Commission. (2021). *Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)* (COM/2021/206 final). Retrieved from https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206

European Commission. (2023). *Council of the EU: General Approach on the AI Act*. Retrieved from https://futurium.ec.europa.eu/en/european-ai-alliance/community-content/council-eu-general-approach-ai-act

European Council. (2023). *Artificial intelligence act: Council and Parliament strike a deal on the first worldwide rules for AI* [Press release]. Retrieved from https://www.consilium.europa.eu/en/press/press-releases/2023/12/09/artific ial-intelligence-act-council-and-parliament-strike-a-deal-on-the-first-worldwide-rules-for-ai/

European Parliament. (2023). *Artificial Intelligence Act: At a Glance* (PE 747.926). Retrieved from https://www.europarl.europa.eu/RegData/etudes/ATAG/2023/747926/EPRS_ATA(2023)747926_EN.pdf

European Parliament. (2024a). *Artificial Intelligence Act: European Parliament legislative resolution of 13 March 2024* (P9_TA(2024)0138). Retrieved from https://www.europarl.europa.eu/doceo/document/TA-9-2024-0138_EN.html

European Parliament. (2024b). *Artificial Intelligence Act: MEPs adopt landmark law* [Press release]. Retrieved from https://www.europarl.europa.eu/news/en/press-room/20240308IPR19015/artificial-intelligence-act-meps-adopt-landmark-law

FactSet. (2018). *ETF Classification System: Rules and Methodology*. Retrieved from https://www.etf.com/docs/FactSet_ETF_Classification_System_Rules_and_Methodology_June2018.pdf

Federal Reserve. (2023). *Federal Reserve issues FOMC statement* [Press release]. Board of Governors of the Federal Reserve System. Retrieved from https://www.federalreserve.gov/newsevents/pressreleases/monetary20230614a1.htm

First Trust. (2020). *First Trust Nasdaq Artificial Intelligence and Robotics ETF (ROBT)*. Retrieved from https://www.ftportfolios.com/Retail/Etf/EtfHoldings.aspx?Ticker=ROBT

FTSE Russell. (2020). *Russell 3000 Index factsheet*. Retrieved from https://www.lseg.com/ftse-russell

Global X ETFs. (2020a). *Global X Artificial Intelligence & Technology ETF (AIQ)*. Retrieved from https://www.globalxetfs.com/funds/aiq

Global X ETFs. (2020b). *Global X Robotics & Artificial Intelligence ETF (BOTZ)*. Retrieved from https://www.globalxetfs.com/funds/botz

Global X ETFs Canada. (2020). *Global X Robotics & AI Index ETF (RBOT)*. Retrieved from https://www.globalx.ca/product/rbot

Google. (2026). *Google Trends: Search interest for "EU AI Act", "AI Act", and "EU AI regulation/law", 1 January 2021–31 December 2024*. Retrieved from https://trends.google.com/trends/explore?date=2021-01-01%202024-12-31&q=eu%20ai%20act%20%2B%20the%20eu%20ai%20act,%22ai%20act%22%20%2B%20%22the%20ai%20act%22,eu%20ai%20regulation%20%2B%20eu%20ai%20law&hl=sv&legacy

iShares. (2020). *iShares Future AI & Tech ETF (ARTY)*. Retrieved from https://www.ishares.com/us/products/297905/ishares-future-ai-tech-etf

Legal & General. (2020). *L&G Artificial Intelligence UCITS ETF USD Acc (AIAG).* Retrieved from https://fundcentres.landg.com/en/se/institutional/fund-centre/ETF/Artificial-Intelligence/

LSEG Workspace. (2026). *Vanguard Russell 3000 ETF (VTHR) holdings as of December 31, 2020* [Dataset]. London Stock Exchange Group. Retrieved from https://www.lseg.com/en/data-analytics/products/workspace

OpenAI. (2022). *Introducing ChatGPT.* Retrieved from https://openai.com/index/chatgpt/

Pichai, S., & Hassabis, D. (2023). *Introducing Gemini: our largest and most capable AI model.* Google. Retrieved from https://blog.google/innovation-and-ai/technology/ai/google-gemini-ai/

Somefun, K., Perchet, R., Yin, C., & Leote de Carvalho, R. (2023). Allocating to thematic investments. *Financial Analysts Journal*, 79(1), 18–36. https://doi.org/10.1080/0015198X.2022.2112895

TCW. (2020). *TCW Artificial Intelligence ETF (AIFD).* Retrieved from https://www.tcw.com/products/etfs/tcw-artificial-intelligence-etf/aifd

Vanguard. (2026). *What are index ETFs and how do they track their benchmarks?* Retrieved from https://www.vanguard.co.uk/professional/vanguard-365/investment-knowledge/etf-knowledge/what-are-Index-ETFs-and-how-do-they-track-their-benchmarks

Wharton Research Data Services. (2026). *WRDS.* Retrieved May 10, 2026, from https://wrds.wharton.upenn.edu

Wu, C.-C., & Chen, W.-P. (2022). What's an AI name worth? The impact of AI ETFs on their underlying stocks. *Finance Research Letters*, 46, 102474.

Xtrackers (DWS). (2020). *Xtrackers Artificial Intelligence & Big Data UCITS ETF 1C (XAID).* Retrieved from https://etf.dws.com/en-gb/IE00BGV5VN51-artificial-intelligence-big-data-ucits-etf-1c/

Appendix

Table IX
Firms with Highest ETF Share and ETF Weight

This table lists the 15 firms with the highest ETF Share in Panel A and the 15 firms with the highest ETF Weight in Panel B. *ETF Share* is the fraction of a firm's shares outstanding held by ETFs. *ETF Weight* is the average weight of the firm across all ETFs holding it. *MktCap* is the firm's market capitalization in $ billions. *Industry* is classified at the NAICS 3-digit level.

Panel A: Top 15 firms with the highest ETF Share

Firm Name	**ETF Share**	**MktCap**	**Industry**
Nvidia Corp	0.80	323	Computer & Electronic Prod. Mfg.
Alphabet Inc Class A	0.70	527	Web Search Portals & Info. Svcs.
Microsoft Corp	0.70	1,682	Publishing Industries
Amazon.com Inc	0.60	1,634	Sporting Goods & Misc. Retailers
Autodesk Inc	0.60	67	Publishing Industries
Salesforce Inc	0.60	204	Professional & Technical Svcs
Intl. Business Machines Corp	0.50	112	Professional & Technical Svcs
Adobe Inc	0.50	240	Publishing Industries
Meta Platforms Inc	0.50	657	Web Search Portals & Info. Svcs.
Intuitive Surgical Inc	0.50	96	Computer & Electronic Prod. Mfg.
FARO Technologies Inc	0.50	1	Computer & Electronic Prod. Mfg.
ServiceNow Inc	0.40	107	Professional & Technical Svcs
Advanced Micro Devices Inc	0.40	110	Computer & Electronic Prod. Mfg.
Netflix Inc	0.40	239	Web Search Portals & Info. Svcs.
QUALCOMM Inc	0.40	172	Computer & Electronic Prod. Mfg.

Panel B: Top 15 firms with the highest ETF Weight

Firm Name	**ETF Weight (%)**	**MktCap**	**Industry**
Nvidia Corp	13.25	323	Computer & Electronic Prod. Mfg.
Intuitive Surgical Inc	12.77	96	Computer & Electronic Prod. Mfg.
Azenta Inc	7.32	5	Professional & Technical Svcs
Cerence Inc	5.33	4	Publishing Industries
JBT Marel Corp	5.29	4	Machinery Mfg.
AeroVironment Inc	3.37	2	Transp. Equipment Mfg.
FARO Technologies Inc	3.23	1	Computer & Electronic Prod. Mfg.
Alphabet Inc Class A	2.06	527	Web Search Portals & Info. Svcs.
Helix Energy Solutions Group	0.94	1	Support Activities for Mining
Microsoft Corp	0.91	1,682	Publishing Industries
Amazon.com Inc	0.89	1,634	Sporting Goods & Misc. Retailers
ServiceNow Inc	0.86	107	Professional & Technical Svcs
Salesforce Inc	0.82	204	Professional & Technical Svcs
Intl. Business Machines Corp	0.78	112	Professional & Technical Svcs
Advanced Micro Devices Inc	0.76	110	Computer & Electronic Prod. Mfg.

Table X

AI Exposure Scores by Industry

This table reports average AI exposure scores for the 10 NAICS 3-digit industries with the largest number of firms in the sample, sorted by average ETF Share in descending order. *Avg. ETF Share* is the average fraction of ETFs holding at least one firm in the industry. *Avg. ETF Weight* is the average portfolio weight assigned to industry firms across all ETFs holding them. *# Firms* is the number of firms in the industry.

Industry	**Avg. ETF Share**	**Avg. ETF Weight**	**# Firms**
Other Information Svcs	0.2436	0.0013	13
Web Search Portals, Libraries & Archives	0.2156	0.0014	46
Publishing Industries	0.1812	0.0023	69
Computer & Electronic Prod. Mfg.	0.1667	0.0012	47
Professional & Technical Svcs	0.1458	0.0002	8
Data Processing, Hosting & Related Svcs	0.1439	0.0015	11
Transp. Equipment Mfg.	0.1417	0.0003	10
Credit Intermediation & Related Activities	0.1250	0.0015	18
Machinery Mfg.	0.1071	0.0001	14
Miscellaneous Mfg.	0.1054	0.0001	34

Table XI

Robustness: Cumulative Abnormal Returns over the [−1, +5] Window

This table replicates Table V using a shorter [−1, +5] event window. All other specification details (sample, *AI ETF* dummy classification, abnormal return estimation via the Fama–French three-factor model, NAICS-3 sector fixed effects, and standard errors clustered at the NAICS-3 level in the FE column) are identical to those in Table V. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	Unconditional	**NAICS-3 FE**
Event	AI − Non-AI	AI − Non-AI
Panel A: EU AI Act Legislative Milestones		
Event 1: Commission Proposal (21 Apr 2021)	−0.578%	−0.556%*
	(−1.50)	(−1.75)
Event 2: Council General Approach (6 Dec 2022)	1.423%**	0.592%
	(2.46)	(0.83)
Event 3: Parliament Position (14 Jun 2023)	−1.190%***	−0.313%
	(−2.66)	(−0.74)
Event 4: Provisional Agreement (11 Dec 2023)	0.225%	−1.030%
	(0.47)	(−1.63)
Event 5: Formal Approval (13 Mar 2024)	−2.047%***	−0.784%*
	(−3.22)	(−1.86)
Panel B: ChatGPT Launch		
ChatGPT Launch (30 Nov 2022)	0.303%	0.282%
	(0.56)	(0.61)

Table XII

Robustness: Heteroskedasticity-Robust (HC1) Standard Errors

This table replicates Table V using heteroskedasticity-robust (HC1) standard errors in the NAICS-3 fixed-effects specification, in place of the NAICS-3-clustered standard errors used in the main analysis. Point estimates are unchanged from Table V; only the standard errors (and therefore t-statistics and significance stars) differ. The Unconditional column is identical to Table V, since both use HC1 standard errors. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

Event	**Unconditional** AI − Non-AI	**NAICS-3 FE** AI − Non-AI
Panel A: EU AI Act Legislative Milestones		
Event 1: Commission Proposal (21 Apr 2021)	−1.991%***	−0.807%
	(−3.74)	(−1.22)
Event 2: Council General Approach (6 Dec 2022)	2.552%**	0.095%
	(2.24)	(0.10)
Event 3: Parliament Position (14 Jun 2023)	1.032%	2.149%
	(0.87)	(1.51)
Event 4: Provisional Agreement (11 Dec 2023)	0.263%	−1.640%**
	(0.43)	(−2.05)
Event 5: Formal Approval (13 Mar 2024)	−1.874%**	−0.948%
	(−2.04)	(−0.83)
Panel B: ChatGPT Launch		
ChatGPT Launch (30 Nov 2022)	1.836%**	0.883%
	(2.20)	(0.94)

We report two robustness checks in Tables XI and XII. Table XI replicates the binary specification using a shorter $[-1, +5]$ event window. The headline results are largely preserved: Event 5 remains negative and significant at the 10% level under NAICS-3 FE (-0.78%, $t = -1.86$), while Events 1 and 2 retain their directional signs and Event 2's positive effect again weakens substantially under sector controls.

Table XII replicates the main specification using heteroskedasticity-robust (HC1) standard errors in place of the NAICS-3-clustered standard errors used in our main analysis. Point estimates are unchanged by construction; only the standard errors differ. The unconditional results (already estimated with HC1 in Table 4) are identical, but the FE significance levels shift: Event 5's within-sector effect is no longer statistically significant, while Event 4 reaches significance at the 5% level ($p = 0.04$). This sensitivity reflects the well-known fact that clustered and HC1 standard errors make different assumptions about the residual correlation structure. Clustering at NAICS-3 is appropriate when firms within the same broad industry have correlated abnormal returns around the events, which is plausible in our setting; HC1 treats observations as independent. We adopt clustered standard errors in our main analysis for this reason, but acknowledge that the significance of the Event 5 within-sector effect is sensitive to this choice. The directional pattern of results (negative reactions at the cleanest legislative events, positive reaction at ChatGPT) is preserved across both specifications.

AI Recognition

We have made use of the AI tools ChatGPT, Grammarly and Claude for assistance with grammar, Python coding, and LaTeX formatting. However, all analyses, interpretations, and conclusions are entirely our own.