# AI-Powered Corporate Governance

Jiung Lee, Seungjoon Oh, Joanna Wang*

June 20, 2026

**Abstract**

We study how generative AI adoption by institutional investors reshapes corporate governance. Exploiting the November 2022 release of ChatGPT as a quasi-experimental shock to AI adoption costs, we show that firms with greater pre-shock exposure to machine-based investors experience valuation gains and strengthened governance, including increased CEO pay-performance sensitivity, greater shareholder activism, and fewer financial restatements. These firms also undertake fewer but higher-quality acquisitions and more valuable patents, indicating improved investment efficiency. We trace these effects to two mechanisms: generative AI strengthens exit-based discipline through coordinated trading, and facilitates coalition formation with activist investors, amplifying voice-based governance.

Keywords: Generative AI, institutional investors, corporate governance, machine-based investors, information acquisition costs, shareholder activism, exit and voice, investment efficiency, ChatGPT

JEL Classification: G14, G23, G31, G34

* Lee is from the HSBC Business School, Peking University (jiung.lee@phbs.pku.edu.cn), Oh is from the HSBC Business School, Peking University (sjoonoh@phbs.pku.edu.cn), and Wang is from the HSBC Business School, Peking University (joannawang@phbs.pku.edu.cn). We are grateful to Jan Bena, William Cong, Bingfeng Cen, Jian Feng, Jarrad Harford, Seil Kim, Jaemin Lee, Sung Kwan Lee, Chen Lyu, Fangyuan Ma, Han Ma, Karol Mazur, Xianhua Peng, Wang Qi, Alberto Rossi, Yuan Shi, Wei Wang, Hang Zhou, Jiaqi Zheng, Xiaodi Zhang and seminar participants at the Chinese University of Hong Kong Shenzhen, Hong Kong Polytechnic University, Peking University HSBC Business School, and Shanghai University of Finance and Economics, Renmin University, and conference participants at the KIF-KCMI Joint Symposium, 2026 PHBS Finance Symposium, 6th Greater China Area Finance Conference, and 10th PKU-NUS Annual International Conference for valuable comments and suggestions. All errors are our own.

## 1. Introduction

Artificial intelligence is rapidly transforming how institutional investors acquire, process, and act on corporate information. Advances in machine learning, natural language processing, and generative AI have enabled asset managers and financial analysts to process vast quantities of corporate disclosures with unprecedented speed and scale, augmenting human judgment with algorithmic systems that extract and interpret information from financial filings, earnings calls, and regulatory documents.[1] A growing literature examines how AI adoption shapes firm productivity, innovation, and valuation, as well as price efficiency (Dugast and Foucault, 2018; D'Acunto et al., 2019; Bartram, Branke, and Motahari, 2020; Acemoglu et al., 2022; Cao et al., 2023; Babina et al., 2024; Eisfeldt et al., 2025; Brynjolfsson, Li, and Raymond, 2025). Recent evidence further suggests that generative AI reshapes investor trading behavior and performance, with broader implications for market efficiency (Dou, Goldstein, and Ji, 2025; Sheng et al., 2026). Despite this progress, how AI adoption by investors affects the corporate governance of the firms they own and subsequently shapes real corporate decisions remains largely unexplored.

AI adoption alters investor behavior by reducing the cost of processing information and expanding the range of information that can be effectively analyzed. By relaxing traditional attention constraints (Kacperczyk, Van Nieuwerburgh, and Veldkamp, 2016), AI enables investors to evaluate a broader set of firms with greater speed and scalability.[2] It allows investors to process

[1] According to a survey by the Alternative Investment Management Association, 58% of fund managers expect increased use of generative AI in investment processes over the next year (AIMA, Front Office Gen AI Adoption Shifts from 'If' to 'When' for Leading Fund Managers, 2025, available at https://www.aima.org/article/press-release-front-office-gen-ai-adoption-shifts-from-if-to-when-for-leading-fund-managers-aima-research-finds.html, accessed Apr. 2026).

Major asset managers including BlackRock, Morgan Stanley, Goldman Sachs, Fidelity, and Vanguard are already actively deploying AI across research, risk analysis, and portfolio workflows, collectively managing trillions of dollars under AI-augmented strategies, with generative AI assistants increasingly used to digest earnings calls and regulatory filings at scale (7 Top Investment Firms Using AI for Asset Management, U.S. News & World Report, 2025, available at https://money.usnews.com/investing/articles/7-top-investment-firms-using-ai-for-asset-management, accessed Apr. 2026).

[2] Prior work documents that institutional investors face binding attention constraints that limit portfolio breadth and engagement depth (Sims, 2003; Peng and Xiong, 2006), and that reductions in information acquisition costs lead to meaningful portfolio reallocation (Grossman and Stiglitz, 1980; Cohen, Frazzini, and Malloy, 2008). AI-driven investment systems can improve the speed, scale, and consistency of decision-making by automating information extraction, summarization, and reporting tasks. Prior studies document that AI enhances pattern recognition and prediction in financial applications (e.g., Gu, Kelly, and Xiu, 2020), facilitates sentiment extraction from textual data, and improves operational efficiency in tasks such as underwriting and fraud detection (Bao et al., 2019; Jansen, Nguyen, and Shams, 2025). These developments collectively reduce the cognitive burden on investors and enable more scalable analysis of firm-specific information.

unstructured, firm-specific information from diverse sources (e.g., regulatory filings, corporate news, earnings calls) at scale and at substantially lower cost than conventional NLP tools (e.g., machine learning algorithm). In Internet Appendix Section A, we present a broad range of evidence on institutional investors' use of generative AI in processing information about portfolio firms, and categorize these applications along multiple dimensions. Consistent with the evidence from both industry practice and empirical studies (e.g., Sheng et al., 2026), AI adoption enhances investors' ability to extract and act on idiosyncratic information. These changes have direct implications for portfolio behavior: as information processing costs decrease, investors can monitor a broader set of firms, shorten decision cycles, and respond more rapidly to governance-relevant signals, potentially reshaping the agency relationship between managers and shareholders.[3] How investor-side AI adoption ultimately affects corporate governance and real firm decisions thus becomes a central question.

However, whether these changes translate into improved corporate governance is ex ante ambiguous. On one hand, AI-enhanced information processing may strengthen governance through more effective monitoring. By continuously processing firm disclosures and extracting structured signals, AI-based systems can improve the detection of accounting irregularities, facilitate benchmarking of executive compensation, and enable more systematic evaluation of firms' capital allocation decisions (Bao et al., 2020). Such monitoring can be implemented at scale and in near real time, increasing the timeliness and consistency of governance-relevant assessments. These improvements in monitoring may also reinforce traditional governance mechanisms through exit. Investors who can more efficiently assess firm quality are more credible sellers when performance deteriorates (Edmans, 2009; Edmans and Manso, 2011). In addition, the resulting price pressure can create opportunities for activist investors to accumulate positions and intervene, linking exit to voice through coordination.

[3] The economic relevance of this channel is increasingly evident in practice. A recent McKinsey Global Survey reports that 88% of organizations now regularly use AI in at least one business function, with financial services among the earliest and most intensive adopters (McKinsey & Company, The State of AI: McKinsey Global Survey, 2024, available at https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai, accessed Apr. 2026). In parallel, J.P. Morgan Asset & Wealth Management has reportedly replaced third-party proxy advisory firms, including ISS and Glass Lewis, with an internal AI-based platform for proxy voting, becoming one of the first major asset managers to internalize this governance function using AI (JPMorgan Chase Cuts Ties with Proxy Advisory Firms, Reuters, Jan. 7, 2026, available at https://www.reuters.com/business/finance/jpmorgan-chase-cuts-ties-with-proxy-advisory-firms-wsj-reports-2026-01-07/, accessed Apr. 2026).

On the other hand, AI-powered investing may weaken governance along several dimensions. Most directly, AI-driven investing may crowd out the qualitative, relationship-based engagement that supports long-term oversight (Becht et al., 2009; McCahery, Sautner, and Starks, 2016; Liberti and Petersen, 2019). As algorithmic strategies emphasize rapid, signal-driven trading, they may reduce incentives to provide patient capital and sustain the firm-specific monitoring that underlies effective long-horizon governance. In addition, investors' financial decision-making cannot be fully automated by AI due to legal and fiduciary accountability, limiting the extent to which judgment can be delegated to algorithmic systems (Gennaioli, Shleifer, and Vishny, 2015; Goldstein, Jiang, and Karolyi, 2019). AI also introduces model risks, including sensitivity to historical data patterns, vulnerability to structural breaks, and the potential for errors to propagate more rapidly and at greater scale than in traditional investment processes, raising concerns about systemic fragility (Daniel, Hirshleifer, and Sun, 2020; Giglio, Maggiori, Stroebel, and Utkus, 2021; Dou et al., 2025). The net effect on governance therefore depends on the relative strength of coordination-driven discipline and the potential erosion of long-horizon monitoring, and is ultimately an empirical question.

To examine these questions, we exploit the November 2022 release of ChatGPT as a quasi-experimental shock to the cost of generative AI adoption among institutional investors. Our empirical design employs a difference-in-differences (DiD) framework in which treatment intensity is defined at the firm level based on pre-shock ownership by machine-based investors. We identify machine-based investors as those whose SEC EDGAR filing access patterns are consistent with automated, algorithmic information acquisition in the pre-shock period, following Ryans (2021) and Drake, Roulstone, and Thornock (2015). Specifically, using SEC EDGAR log data from 2015 to 2017, we classify an investor as machine-based if its IP addresses download filings from more than 50 firms within a single day, and we map IP addresses to institutional identities following Chen et al. (2020). Machine-based ownership is then defined as the average equity ownership held by these investors over the three-year pre-shock period (2019Q3–2022Q3). This classification is highly persistent, with year-to-year transition rates exceeding 97%, indicating that machine-based information acquisition reflects stable organizational infrastructure rather than transitory behavior,[4] and machine-based investors in our sample exhibit average holding durations of approximately 1.5 years, distinguishing them from high-frequency traders.[5]

Our identification strategy rests on the premise that investors with pre-existing machine-based infrastructure are better positioned to benefit from the reduction of information processing costs, realizing stronger complementarities between generative AI tools and their superior information acquisition capability. This mechanism is consistent with evidence that pre-existing AI talent predicts post-ChatGPT alpha generation among hedge funds (Sheng et al., 2026), and that human-machine complementarities make the marginal value of AI tools increasing in investors' pre-existing technological capabilities (Cao et al., 2024). As further validation, we document that machine-based investors earn significantly higher portfolio returns and exhibit higher reliance on generative AI tools following the ChatGPT release.

Our empirical analyses reveal several novel findings. First, firms with greater pre-shock exposure to machine-based investors experience significant increases in firm value following the ChatGPT release: a one standard deviation increase in machine-based ownership implies a 3.7% increase in Tobin's Q and an 11.7% increase in Total Q. These results are robust to a battery of robustness tests, including pre-trend validation, alternative AI-based investor classifications based on SEC Form ADV filings, and a propensity-score matched DiD model. Cross-sectional tests reveal that valuation effects are concentrated where monitoring frictions are greater: among operationally complex firms with high inventory volatility or intangible intensity, firms with fewer analyst coverage, firms in less competitive industries, and firms reporting negative earnings surprises. The evidence is consistent with the view that AI is more valuable where firm-specific information is more costly to process.

Second, we provide direct evidence of strengthened governance across multiple channels. Treated firms exhibit stronger CEO pay-performance and turnover-performance sensitivity. We also document improvements in accounting quality: treated firms experience fewer financial

---

[4] Appendix Table IA.1 shows that machine-based investor classification is highly persistent (with 97.4% of machine investors and 99.8% of non-machine investors retaining their status year-to-year) and remains virtually unchanged after controlling for investor characteristics, supporting the validity of our pre-shock measure and alleviating concerns about measurement error.

[5] On average, machine-based investors in our sample exhibit a holding duration of approximately 1.5 years (6.116 quarters), which is inconsistent with the short-term trading behavior of high-frequency traders. Consistent with this distinction, the top 20 machine-based investors in our sample include institutions such as JPMorgan Chase, T. Rowe Price, and GMO LLC (i.e., entities not typically characterized as algorithmic or high-frequency trading firms). See Appendix Table IA.2 for the full list, along with average holding periods and assets under management.

restatements and smaller restatement corrections, consistent with machine-based investors' comparative advantage in algorithmically detecting earnings management and reporting irregularities (Bao et al., 2020; Healy and Wahlen, 1999). Treated firms also experience significantly higher rates of activist campaigns, proxy fights, and shareholder proposals in the post-ChatGPT period. These effects are concentrated in firms with greater activist ownership, suggesting that machine-based investors strengthen voice-based governance indirectly by attracting and enabling external activists.

Third, AI-powered monitoring reshapes real corporate investment decisions in ways that reflect a shift toward investment efficiency. Treated firms undertake fewer acquisitions overall, but acquisition quality improves: the fraction of value-destroying deals declines by approximately 18% relative to the sample mean, and announcement returns improve by about 0.67 percentage point. Treated firms also reduce long-term investment, R&D spending, and patent filings. A one standard deviation increase in machine-based ownership reduces long-term investment, R&D, and patent counts by 6.0%, 5.8%, and 13.8%, respectively. However, the patents filed by treated firms are of significantly higher value, with a one standard deviation increase in exposure associated with a 6.2% increase in patent value. These patterns suggest that AI-powered governance disciplines marginal investments while concentrating resources on higher-impact projects.

Lastly, we document the investor- and firm-level behavioral changes that underlie these governance outcomes. At the investor level, the ChatGPT shock leads machine-based investors to hold broader portfolios across more industries, exhibit higher portfolio turnover, shorten investment horizons, and reduce passive ETF exposure, consistent with greater reliance on firm-specific, idiosyncratic information. At the firm level, treated firms attract more machine-based investors following the shock, while average ownership per investor declines. Selling and buying dispersion both fall significantly, indicating increased synchronization in trading decisions across machine-based investors. Our findings suggest that shared AI-based analytical frameworks lead to more convergent firm assessments, strengthening exit-based discipline. Beyond exit, we show that treated firms attract significantly greater activist ownership following the shock. This increase is not driven by direct engagement from machine-based investors, whose involvement remains limited. Instead, treated firms attract more experienced activists who are more likely to initiate campaigns and achieve successful outcomes, including settlements and board representation. These patterns suggest that AI-powered investors facilitate activist intervention indirectly by improving the information environment and reshaping the shareholder base, enabling more effective voice-based governance through coalition formation.

Our study contributes to several strands of literature. First, we contribute to the growing literature on artificial intelligence and corporate finance. Recent work examines how firms adopt AI and its consequences for productivity and innovation (Babina et al., 2024; Brynjolfsson et al., 2025), employment and valuation (Alekseeva et al., 2021; Acemoglu et al., 2022; Eisfeldt et al., 2025; Chen and Wang, 2025a), financial misconduct (Khalid et al., 2025), audit quality (Fedyk et al., 2022), venture capital (Gofman and Jin, 2024), and entrepreneurship (Gupta et al., 2025; Bena, Bian, and Giannetti 2026). A closely related strand examines how AI shapes corporate disclosure: Cao et al. (2023) show that growing machine readership leads firms to strategically reshape their disclosure style and sentiment, documenting a feedback effect running from investor-side AI adoption to firm-level disclosure behavior; Chen and Wang (2025b) further examine how generative AI affects disclosure quality and informativeness. Less is known, however, about how AI adoption by institutional investors — rather than by firms — affects corporate outcomes. We contribute by showing that investor-side AI adoption has real consequences for firm governance and value, complementing Cao et al. (2023) by documenting the other side of this relationship: while they study how firms adapt their disclosure to machine readers, we show how AI enables investors to more effectively monitor and discipline the firms they own.

Second, we contribute to the literature on information technology and financial markets. Prior work examines how algorithmic trading affects market quality and price discovery (Hendershott, Jones, and Menkveld, 2011; Brogaard, Hendershott, and Riordan, 2014; Boehmer, Fong, and Wu, 2021), how machine learning improves asset pricing (Gu, Kelly, and Xiu, 2020), and how textual analysis of corporate disclosures generates trading signals (Tetlock, 2007; Loughran and McDonald, 2011; Hoberg and Phillips, 2016). A growing body of work studies AI adoption across financial markets, including equity investment (Lopez-Lira and Tang, 2023), analyst behavior (Cao et al., 2024; Bertomeu et al., 2025), banking and credit markets (Ben-David et al., 2026), financial advising (D'Acunto et al., 2019), and asset management (Bartram et al., 2020; Sheng et al., 2026). Specifically, Dou et al. (2025) show theoretically and empirically that AI-powered trading can generate correlated signals across investors, raising concerns about algorithmic collusion and price efficiency, and Sheng et al. (2026) document that AI-powered investors generate higher abnormal returns following the ChatGPT release. We extend this literature by documenting that generative AI increases synchronization in information processing and trading behavior among institutional investors as it lowers the cost of acquiring idiosyncratic information, thereby leading to shorter investment horizons and more coordinated governance discipline that ultimately generates valuation gains for firms.

Third, we contribute to the literature on institutional investors and corporate governance. A large body of work establishes that institutional investors affect governance through both voice and exit mechanisms (Shleifer and Vishny, 1986; Admati, Pfleiderer, and Zechner, 1994; Edmans, 2009; Edmans and Manso, 2011; Appel, Gormley, and Keim, 2016), while a related literature emphasizes the role of investor horizons in shaping monitoring incentives (Bushee, 1998; Gaspar, Massa, and Matos, 2005). We contribute by showing that AI adoption shifts investors toward shorter investment horizons while simultaneously strengthening both exit- and voice-based governance. The shift toward shorter horizons, driven by rule-based exit threats and trading synchronization, does not crowd out firm value creation; rather, AI-powered investors discipline marginal investments and concentrate resources on higher-impact projects, improving overall investment efficiency. Coalition formation between machine-based and activist investors further amplifies voice-based governance without requiring direct engagement by machine-based investors themselves. Our findings contribute to ongoing debates about short-termism by suggesting that AI-induced horizon shortening need not conflict with long-term value creation when accompanied by improved information processing and coordinated investor action.

## 2. Identification Strategy

Our identification strategy exploits the November 2022 release of ChatGPT as an exogenous shock to the cost of adopting AI-based information processing technologies. The timing and nature of the release satisfy key requirements for a valid natural experiment. First, the release was largely unexpected by market participants. While OpenAI had been developing large language models for years, the timing of ChatGPT's public launch was not announced in advance, and its capabilities surprised even AI researchers and industry experts. The rapid adoption—reaching 100 million users within two months—further suggests that the release represented a discrete technological breakthrough rather than a predictable incremental improvement. Second, the shock was plausibly orthogonal to firm-specific characteristics or performance trajectories. Unlike regulatory changes or industry-specific technological developments, ChatGPT's availability was broad and simultaneous across sectors and firm types. Third, the shock was persistent, with subsequent model improvements and competing products reinforcing the shift toward accessible AI tools.

However, using a single economy-wide shock presents empirical challenges. The release coincides with other potentially confounding events, including the SEC's share repurchase disclosure modernization, increased regulatory scrutiny following several accounting scandals in 2022, and macroeconomic uncertainty related to post-pandemic recovery and rising interest rates. In addition, a uniform shock provides strong time-series variation but limited cross-sectional variation. To address these concerns, we combine the ChatGPT shock with firm-level heterogeneity, following recent work (Bertomeu et al., 2025; Eisfeldt et al., 2025; Chen and Wang, 2025b).

While ChatGPT dramatically lowered technical barriers to generative AI for the broader market, its impact is not uniform across investors. Institutional investors with pre-existing machine-based infrastructure, automated data pipelines, algorithmic filing access routines, and organizational capacity to integrate new AI tools are better positioned to exploit the ChatGPT shock. For these investors, generative AI functions as a powerful complement to existing systems, enabling them to process firm-specific disclosures at greater scale, extract richer signals from unstructured text, and act on governance-relevant information more rapidly than before. Investors without such foundations, by contrast, face a steeper path from ChatGPT access to meaningful analytical deployment: the marginal value of a generative AI interface is substantially lower when the surrounding data infrastructure, processing routines, and institutional workflows are absent. This mechanism is consistent with evidence from Sheng et al. (2026), who document that hedge funds with greater pre-ChatGPT AI talent earn significantly higher abnormal returns following the release of ChatGPT. This is also consistent with evidence from Cao et al. (2024) that human-machine complementarities make the marginal value of AI tools increase in investors' existing technological capabilities. The shock therefore amplifies, rather than equalizes, the informational advantage of machine-based investors. Accordingly, we define treatment intensity at the firm level by pre-shock machine-based investor ownership, with details of the variable construction procedure provided in Section 3.1. Our main difference-in-differences specification is:

$$Y_{i,t} = \beta_1\, Post_t \times Machine\text{-}based\ Ownership_i + \beta_2\, X_{i,t-1} + \alpha_i + \gamma_t + \varepsilon_{i,t}$$

where $Y_{i,t}$ represents various outcome variables for firm $i$ in quarter $t$, and $Post_t$ is an indicator equal to one for quarters after 2022Q4. *Machine-based Ownership*$_i$ is the average equity ownership held by machine-based investors over the three-year pre-shock window (2019Q3–2022Q3), ensuring our treatment measure captures persistent rather than transitory investor interest. $X_{i,t-1}$ is a vector of lagged firm-level controls including size (i.e., log-transformed assets), leverage, market-to-book ratio, R&D/assets, tangible assets/total assets, and firm-level AI adoption proxied by average AI-skilled employment in the past five years. We include firm fixed effects to absorb time-invariant firm characteristics and year-quarter fixed effects to control for common time trends. The coefficient of interest is $\beta_1$, which captures the differential change in outcomes for firms with higher pre-shock machine-based investor ownership following the ChatGPT release. Firm fixed effects ensure our estimates are not driven by time-invariant differences between treated and control firms, and the inclusion of firm-level AI adoption controls allows us to separate the governance effects of external monitoring from the operational effects of firms' own AI adoption.

We provide additional validation for our identification strategy using investor-level evidence. First, we show that machine-based investors earn significantly higher risk-adjusted returns following the ChatGPT shock, as measured by Fama–French five-factor alphas. This result is consistent with the interpretation that these investors are better able to exploit AI-driven improvements in information processing. Second, we directly examine changes in generative AI usage by constructing an investor-level measure of generative AI reliance using the methodology of Sheng et al. (2026). We find that machine-based investors exhibit a significantly greater increase in reliance on generative AI following the ChatGPT release. Together, these findings support our identification assumption that pre-shock exposure to machine-based investors captures heterogeneity in the ability to benefit from generative AI, and that firms with higher exposure are more affected by the ChatGPT shock. Detailed empirical settings and results are provided in Section 7.1.

## 3. Data and Summary Statistics

### *3.1 Data and Variable Construction*

Our analysis combines data from multiple sources to construct a comprehensive dataset spanning 2019Q3-2025Q3 and covering all U.S.-listed firms. We obtain institutional ownership data from Thomson Reuters 13F filings, which provide quarterly snapshots of equity holdings by institutional investment managers with at least $100 million in assets under management. These filings cover a broad set of investor types, including hedge funds, mutual funds, pension funds, and other institutional investors. Following standard practice in the literature (e.g., Ben-David et al., 2021), we aggregate holdings across multiple filing entities belonging to the same institutional investor family to avoid double-counting.

To identify machine-based institutional investors, we utilize EDGAR server log data from the SEC between 2015 and 2017, which records detailed information about document access patterns, including IP addresses, timestamps, and file identifiers. Following the literature (Drake et al., 2015; Cao et al., 2023), we classify an investor as machine-based if it exhibits EDGAR access consistent with automated information acquisition. Specifically, an IP address is classified as a machine visitor if it downloads filings from more than 50 unique firms within a single day. We then classify an institutional investor as machine-based if any associated IP address meets this criterion during the sample period. Since the EDGAR log file anonymizes IP users by encrypting the last digit of the IP address, we reconstruct full IP addresses flagged as machines in the EDGAR logs following the procedure in Chen et al. (2020). We then map reconstructed IP addresses to registered organizations using the American Registry for Internet Numbers (ARIN).[6] Next, we link identified organizations to institutional investors with matching or closely corresponding names in institutional ownership databases. We employ the RapidFuzz library for matching, given its competitive performance in terms of computational efficiency and matching accuracy (Elmobark, 2025), and retain candidate matches with similarity scores above 65. To mitigate false positives arising from IP reassignment, shared hosting infrastructure, or service providers (e.g., cloud vendors or data intermediaries), we manually verify each candidate match by cross-checking organizational descriptions, business activities, and investor identity information. Institutional investors that satisfy both (i) algorithm-consistent EDGAR access patterns and (ii) validated organizational matches are classified as machine-based investors. We then construct our primary treatment variable, *Machine-based Ownership*, as the average equity ownership held by machine-based investors over the three-year period preceding the ChatGPT shock (2019Q3–2022Q3).

[6] https://www.arin.net/reference/research/bulkwhois/

A potential concern with our identification strategy is measurement error in our classification of machine-based investors in the pre-shock period. Specifically, some institutional investors may have already adopted AI technologies but not yet exhibited the observable behavioral patterns we use for classification. This measurement error would attenuate our treatment intensity measure and bias our estimates toward zero, making it harder to detect true effects. A related concern is that investors classified as machine-based in the pre-period may subsequently revert to less automated information acquisition behavior, which could weaken our identification. Appendix Table IA.1 addresses this concern by reporting both observed and fitted one-year transition matrices for machine-based investor status during 2015–2017. The results reveal strong persistence in both investor types: among investors classified as machine-based in year t, 97.44% retain that classification in year t+1, with a reversion rate of only 2.56%.[7] The complementary pattern holds for non-machine investors, who transition to machine-based status at a rate of just 0.24%. This persistent classification supports the validity of our pre-shock treatment measure and suggests that measurement error in our classification is unlikely to drive the results. Section 7.1 provides additional investor-level analysis serving as both validation of the shock's effect on investor behavior and mechanism tests linking portfolio changes to the firm-level governance outcomes we document.

To construct the firm-level AI adoption control, we use a text-based classification measure of AI-skilled workers drawn from employees' job titles in the Revelio Labs database. We apply a curated dictionary of AI-related keywords capturing both role-based designations (e.g., "AI engineer," "machine learning scientist," "AI researcher") and technical capabilities (e.g., "machine learning," "deep learning," "neural network," "natural language processing," "large language model," and "generative AI"). The dictionary further incorporates terms related to the deployment and production of AI systems (e.g., "MLOps," "model deployment," and "model serving") to capture workers involved in implementation and scaling, in addition to those engaged in model development. To ensure measurement precision, we implement regular expression matching with word boundaries to avoid spurious matches (e.g., excluding partial matches such as "air" for "AI"). An employee is classified as AI-skilled if their job title contains at least one predefined keyword,

[7] The fitted transition matrix, which controls for observable investor characteristics including portfolio size, turnover, breadth, average holdings, and large-block concentration, yields almost identical estimates (97.42% and 0.23%, respectively).

and at the firm-year level, we aggregate these classifications to calculate the yearly count of AI-skilled employees.[8]

Firm-level financial data come from Compustat, including variables for firm size, leverage, market-to-book ratio, R&D intensity, and tangible asset ratios. Stock return data are obtained from CRSP. We obtain information on CEO compensation and turnover events from Execucomp. Accounting quality, measured by the incidence of financial restatements and the magnitude of restatement corrections, is obtained from Audit Analytics. To measure shareholder activism, we obtain data on activist campaigns and proxy fights from *FactSet SharkWatch*, which provides comprehensive coverage of activist interventions, including campaign initiations, proxy solicitations, and their outcomes. We construct three variables at the firm-quarter level: an indicator for whether the firm is targeted by at least one activist campaign, an indicator for whether any campaign escalates to a definitive proxy fight, and the total number of shareholder proposals filed against the firm. As measures of real corporate investment decisions, patent data are obtained from Kogan et al. (2017),[9] and M&A activity data are collected from the SDC database. Our final sample consists of 68,325 firm-quarter observations for 5,063 unique firms over the period 2019Q3-2025Q3.[10]

### *3.2 Summary Statistics*

Table 1 reports summary statistics. The key treatment variable, *Machine-based Ownership*, has a mean of 0.143 with an interquartile range of 0.045 to 0.226, confirming that treatment intensity is well-dispersed across the sample rather than concentrated in a small set of firms. Notably, machine-based investors account for 14.3% of total equity ownership on average, representing a substantial fraction of the 57.5% average institutional ownership, underscoring the economic relevance of this investor group as a governance force. The mean *Tobin's Q* is 1.753, and the mean *Total Q* is 1.961, both exhibiting considerable dispersion. Turning to governance outcomes, financial restatements are relatively common at 3.6% of firm-quarters, providing

[8] The full job list is available upon request.

[9] The patent data are obtained from the extended dataset provided by Kogan et al. (2017), available at: https://github.com/KPSS2017/Technological-Innovation-Resource-Allocation-and-Growth-Extended-Data

[10] Our results remain unchanged when excluding firms in financial (SIC codes 6000-6999) and utilities (SIC codes 4900-4999) industries.

meaningful variation for our accounting quality tests. Shareholder activism events are infrequent but nontrivial: 1.9% of firm-quarters experience an activist campaign, and 1.8% involve a proxy fight. M&A activity occurs in 14.8% of firm-quarters, of which approximately 44% are value-destroying as measured by negative announcement returns. Innovation activity is heavily right-skewed, with the median firm filing zero patents per quarter. Panel B reports investor-level statistics for 46,742 investor-quarter observations. Machine-based investors constitute 5.1% of the investor sample but hold disproportionately large and actively managed positions, with an average Fama-French five-factor alpha of 4.3% and a mean holding duration of 6.8 quarters.

## 4. Investor AI Adoption and Firm Valuation

### *4.1 AI-Powered Governance and Firm Valuation*

A central question in corporate governance research is whether institutional monitoring creates firm value. Theoretical models predict that effective monitoring disciplines managerial behavior, reduces agency costs, and thereby raises firm valuations (Shleifer and Vishny, 1986; Edmans, 2009). If AI-powered investors constitute a new and more effective class of monitor, treated firms should exhibit higher market valuations following the ChatGPT shock. We measure firm valuation using both Tobin's Q and Total Q. *Tobin's Q* is calculated as the market value of assets divided by the book value of assets, representing the market's overall assessment of the firm relative to its replacement cost. *Total Q*, following Peters and Taylor (2017), incorporates intangible capital by adding the capitalized value of R&D and other intangibles to the denominator, providing a more comprehensive valuation measure that accounts for knowledge assets particularly relevant for innovative firms.

Table 2 reports the valuation consequences of AI adoption among machine-based institutional investors. Firms with higher pre-period machine-based investor exposure experience statistically significant increases in both Tobin's Q and Total Q following the ChatGPT shock. Economically, a one-standard-deviation increase in pre-shock machine-based ownership is associated with an increase of 0.065 in Tobin's Q, representing approximately a 3.7% premium relative to the sample mean of 1.753. Results are qualitatively similar for Total Q, with treated firms exhibiting a larger increase in this broader valuation measure that accounts for intangible capital: a one-standard-deviation increase in machine-based ownership is associated with an increase of approximately 0.23, corresponding to an 11.7% premium relative to the sample mean of 1.961. Importantly, the valuation effect remains statistically significant after controlling for firm-level AI adoption, suggesting that the results are more plausibly attributed to the external monitoring channel, rather than driven by firms' own operational use of AI technologies. Instead, the evidence points to improved external monitoring and governance as the primary channel through which valuations increase.

To visualize the dynamic treatment effects and assess the validity of the parallel trends assumption, Figure 1 plots the estimated coefficients and 95% confidence intervals from a yearly event-time specification, where the horizontal axis indicates years relative to the ChatGPT shock in 2022. Panel A reports results for Tobin's Q and Panel B for Total Q. In both panels, the pre-treatment coefficients are statistically indistinguishable from zero, providing no evidence of differential valuation trends between high- and low-exposure firms prior to the shock. This pattern supports the identifying assumption that, absent the ChatGPT release, firms with varying levels of machine-based investor exposure would have followed parallel valuation trajectories. The coefficients become positive and statistically significant in the post-shock period, with the magnitude increasing over time, consistent with a gradual strengthening of AI-powered governance as investors progressively integrate generative AI tools into their monitoring infrastructure. The absence of pre-trends combined with the sharp post-treatment divergence reinforces a causal interpretation of our baseline estimates.

Our main analysis treats machine-based ownership as a continuous treatment intensity variable, which exploits all available variation in the data. To assess robustness to potential outliers in the ownership distribution, we re-estimate all specifications using a binary treatment indicator equal to one if a firm had any machine-based investor ownership in the pre-shock period and zero otherwise. Table IA.3 presents these results. The findings remain qualitatively consistent with our main results. In the binary specification, firms with any pre-shock machine-based investor exposure exhibit a Tobin's Q that is 0.261 higher than otherwise comparable firms following the ChatGPT shock, corresponding to a 14.9% premium relative to the sample mean.

We then conduct a battery of robustness checks on firm valuation effects. First, it is possible that firms with greater machine-based investor exposure may independently adopt AI technologies at higher rates, confounding the governance channel with firms' own operational AI adoption. Table 3 Panel A addresses this concern through a pre-trend analysis that regresses firm-level AI adoption (i.e., the logarithm of AI-related employment) on machine-based ownership in the pre-shock period. The coefficient is economically small and statistically insignificant at both the year-quarter level and the annual level, indicating that pre-shock machine-based investor exposure does not predict firms' own AI investment trajectories. This result alleviates the concern that our treatment variable captures firms' internal technological capabilities rather than external monitoring intensity, supporting the interpretation that the valuation effects we document operate through the investor monitoring channel.

We present additional robustness tests for our baseline results in Table 3. Panel B employs an alternative classification of AI-powered investors based on *GenAI Reliance*, following Sheng, et al., (2026). Under this approach, investors are classified as AI-powered if their *GenAI Reliance* measure is statistically significant (p-value < 0.01), capturing the extent to which institutional investors rely on ChatGPT-generated signals, beyond standard firm fundamentals, to explain their portfolio holding changes. We find consistent and statistically significant positive valuation effects of *AI-powered ownership* under this alternative classification, with a coefficient of 0.552 for Tobin's Q and 2.633 for Total Q, both significant at the 1% level. This result confirms that our main findings are not an artifact of the specific AI-powered investor classification used in the baseline analysis.

Panel C employs an alternative classification of machine-based investors based on textual analysis of SEC regulatory filings. Specifically, we classify investors as machine-based if their disclosures contain keywords indicative of machine-based strategies. To implement this approach, we extract text from Item 8 ("Methods of Analysis, Investment Strategies, and Risk of Loss") of ADV Part 2A brochures and from the full text of Form N-CSR shareholder reports, and apply a predefined keyword taxonomy designed to capture machine-based (AI-related) investment processes, including terms such as "artificial intelligence," "machine learning," and "large language models". The keyword taxonomy is constructed to distinguish AI-related methods from traditional quantitative approaches and is applied consistently across both data sources. An investor is classified as machine-based if at least one such keyword is present in the disclosure. Under the alternative machine-based investor classification, we find consistent positive valuation effects. Panel C augments the baseline specification with industry-by-year-quarter fixed effects to absorb industry-specific time trends that might confound identification; the results remain qualitatively similar. Another concern is that treated firms, which by construction have high pre-period machine-based investor exposure, could be more liquid, more technology-oriented, and more information-intensive. In Panel D we implement a propensity score matched difference-in-differences design and find comparable estimates, alleviating concerns about covariate imbalance driving our results.

Lastly, we re-estimate our baseline regressions after excluding the "Magnificent Seven" firms, Apple, Microsoft, Alphabet, Amazon, NVIDIA, Meta, and Tesla. This exclusion addresses several concerns. First, these firms are extreme outliers in size, liquidity, institutional ownership, and investor attention, which may give them disproportionate leverage in firm-level estimates. Second, their exceptional valuation gains during the sample period may partly reflect elevated AI-related expectations, raising the possibility that an AI-hype premium drives the observed effects. Third, machine-based investors may hold these stocks more passively, for example through index-replication strategies, generating ownership dynamics that differ from the information-driven monitoring channel our framework emphasizes. The results in Panel E show that our main estimates remain quantitatively similar and statistically significant after excluding these firms, suggesting that our findings are not driven by the non-random exposure of a small set of mega-cap stocks.

### *4.2 Heterogeneity on Governance Frictions*

We next examine cross-sectional variation in the valuation effects documented in the previous section. If AI-augmented monitoring operates through a governance channel, its impact should be stronger precisely where traditional monitoring mechanisms are less effective and where the wedge between managerial discretion and shareholder interests is larger. We focus on two broad dimensions along which monitoring frictions vary: operational complexity and information environment quality.

Firms with high operational complexity are inherently more difficult to monitor using conventional financial analysis. In these settings, managerial actions are less transparent, performance attribution is more ambiguous, and traditional monitoring technologies are relatively blunt. AI-powered monitoring reduces the cost of processing unstructured, high-frequency, and forward-looking information. As a result, the effective monitoring frontier shifts outward more for firms whose operational structure generates greater informational opacity. Under this mechanism, we expect the valuation effect of machine-based ownership after the ChatGPT shock to be concentrated among operationally complex firms.

Table 4 provides evidence consistent with this prediction. We use inventory volatility and intangible asset intensity as proxies of operational complexity. Inventory volatility introduces rapid state changes that are difficult to track in real time, while intangible-intensive firms derive value from assets that are only partially reflected in accounting statements and are subject to greater measurement error. Panel A splits firms by pre-period inventory volatility, measured as the standard deviation of quarterly inventory over two years (eight quarters). The coefficient *on Post × Machine-based Ownership* is statistically significant for Tobin's Q only in the high-volatility subsample, with no significant effect among low-volatility firms. A parallel pattern emerges for Total Q. Panel B repeats the analysis using intangible asset intensity and yields similar results: only firms above the median in intangible intensity exhibit statistically significant valuation effects. These findings indicate that AI-driven reductions in information acquisition costs generate the largest marginal gains in settings where conventional monitoring is most constrained.

We next examine whether the valuation response varies with the quality of the firm's external information environment. Firms with thin analyst coverage or limited external benchmarking operate under weaker informational discipline, creating greater scope for managerial slack. Two complementary mechanisms may operate. First, in poor information environments, AI-powered investors may substitute for missing external monitoring by generating independent signals. Second, in concentrated industries where product market competition provides fewer external performance benchmarks and weaker discipline, AI monitoring may substitute for the disciplinary role that competition would otherwise play.

Panel C splits firms by analyst coverage. The valuation effect is more than twice as large among low-coverage firms, with statistical significance concentrated in this subsample. This pattern is consistent with a substitution mechanism: AI-enhanced monitoring is particularly valuable where third-party information production is weak. Panel D examines industry concentration, measured as the combined market share of the top eight firms. Consistent with a substitution interpretation, treated firms in concentrated industries experience significantly larger increases in Tobin's Q, while the coefficient in competitive industries is statistically insignificant.

We further examine heterogeneity based on earnings surprises. Governance discipline is most relevant following disappointing performance, when investors must determine whether poor results reflect transitory shocks or persistent managerial underperformance. Negative earnings surprises represent precisely such moments of heightened scrutiny. We split the sample based on the sign of standardized unexpected earnings (SUE). Panel E shows that valuation effects are concentrated among firms reporting negative earnings surprises. The coefficients on *Post* × *Machine-based Ownership* are statistically significant only in the negative-SUE subsample. This asymmetry is consistent with a downside-discipline mechanism: when performance disappoints, the credibility of the exit threat increases, and the value of timely monitoring rises. AI-powered investors are better positioned to distinguish noise from persistent underperformance, strengthening managerial discipline precisely when governance intervention is most consequential.

Taken together, the cross-sectional evidence strongly supports a monitoring-based interpretation. The valuation effects of AI-powered governance are systematically concentrated in environments where monitoring is intrinsically more difficult (operational complexity), where external information production is limited (thin analyst coverage and concentrated industries), and when disciplining managers is most salient (negative earnings surprises). These patterns are difficult to reconcile with purely mechanical trading or price discovery explanations[11] and instead point to an economically meaningful enhancement of governance efficiency.

## 5. AI-Powered Investors and Corporate Governance Outcomes

### *5.1 Executive Discipline: Turnover and Compensation*

The executive labor market literature predicts that institutional monitoring strengthens the pay-for-performance contract between shareholders and managers (Jensen and Murphy, 1990; Core, Holthausen, and Larcker, 1999). AI-powered investors, with their superior ability to process standardized performance metrics and detect underperformance in real time through algorithmic screening of earnings releases, regulatory filings, and audit reports, should strengthen both performance-turnover and pay-performance sensitivities following the shock. We predict that firms with a larger proportion of machine-based investors will exhibit both more responsive CEO turnover (higher turnover-performance sensitivity) and more tightly aligned CEO compensation (higher pay-performance sensitivity) following the ChatGPT shock.

[11] One potential alternative interpretation is that ChatGPT and related tools simply enhance investors' information processing capacity, with the post-treatment effect reflecting faster signal extraction and improved price discovery rather than any change in monitoring intensity or governance-relevant behavior. In untabulated results, we find our results qualitatively robust after controlling for price efficiency proxies such as Amihud illiquidity and bid-ask spreads. To further address this concern, in the next section, we provide direct evidence of monitoring enhancement through executive discipline, financial reporting quality, and corporate policy outcomes.

Table 5 investigates whether AI-augmented investor monitoring strengthens executive discipline, examining both CEO turnover-performance sensitivity and pay-performance sensitivity. Column (1) estimates a linear probability model in which CEO turnover is regressed on operating performance over two-year rolling windows and other firm characteristics. The key coefficient of interest is the triple interaction *Post* × *Machine-based Ownership* × *ROA*. Following the ChatGPT shock, firms with greater machine-based investor exposure exhibit stronger sensitivity of CEO turnover to poor performance: the coefficient on *Post* × *Machine-based Ownership* × *ROA* is negative and statistically significant, implying that worse ROA leads to higher turnover probability more strongly in treated firms. Column (2) estimates a linear probability model in which CEO total compensation is regressed on operating performance over two-year rolling windows and other firm characteristics. The results shows that CEO compensation becomes more closely tied to firm performance post-shock, indicating an increase in pay-performance sensitivity. Together, these findings provide direct evidence that machine-based investors enhance governance by more effectively disciplining managers rather than merely influencing market valuations, consistent with reduced monitoring costs enabling machine-based investors to maintain effective performance-contingent discipline through algorithmic performance benchmarking and exit-threat mechanisms.

### *5.2 Managerial Misconduct in Financial Reporting*

Managerial misconduct in financial reporting, such as financial restatement, is a direct target of institutional monitoring. Agency theory predicts that managers with private information about firm performance have incentives to manipulate reported earnings upward, smooth volatility, or delay the recognition of bad news (Jones, 1991; Dechow, Sloan, and Sweeney, 1995; Healy and Wahlen, 1999). Institutional investors with strong monitoring capabilities can deter these behaviors through two mechanisms. First, more accurate detection of irregularities in reported numbers, such as comparing actual accruals to expected accruals, cross-checking disclosure consistency across filings, and benchmarking performance against industry peers, raises the probability that manipulation is caught and penalized. Second, the credible threat of exit on news of accounting misconduct increases the reputational and market-value costs of misreporting. AI-powered investors enjoy a strong comparative advantage in the detection dimension: algorithmic processing of structured accounting data can identify anomalies in accrual patterns and flag unusual discretionary items across the full portfolio in real time. We therefore predict that higher machine-based investor exposure will be associated with better accounting quality following the ChatGPT shock.

Table 6 examines the effects of AI-powered monitoring on financial reporting quality outcomes. We measure accounting quality through two complementary metrics: the incidence of financial restatements (binary indicator for any restatement in the quarter), and the magnitude of restatement corrections. These measures capture different aspects of financial reporting quality, from outright errors requiring correction to more subtle earnings management. Column (1) shows that firms with higher pre-period machine-based investor exposure experience a significant decline in financial restatements following the ChatGPT shock. The coefficient on *Post* × *Machine-based Ownership* is negative and statistically significant in the logit specification, implying a meaningful reduction in the probability of financial restatements. Column (2) shows that not only does the incidence of restatements decline, but the magnitude of restatement corrections also falls significantly, suggesting that the deterrence effect operates on the severity of misreporting, not merely its frequency. These results are robust to the inclusion of firm-level AI adoption controls, confirming that the improvements in reporting quality arise from external monitoring pressure rather than from internal operational changes. Taken together, the accounting quality evidence suggests that algorithmic monitoring improves both the detection and deterrence of financial misreporting, consistent with AI-powered investors' comparative advantage in processing complex accounting information.

### *5.3 Shareholder Activism*

AI-powered monitoring may strengthen governance not only through direct exit threats but also by facilitating more active shareholder engagement. We expect that the likelihood and effectiveness of activism increase when investors face lower costs of identifying governance deficiencies, evaluating strategic alternatives, and coordinating collective action (Brav et al., 2008; Edmans and Holderness, 2017). We examine whether machine-based institutional ownership shapes the broader governance ecosystem by attracting external activist pressure, using three complementary measures of activism intensity: the incidence of activist campaigns, proxy fights, and shareholder proposal submissions.

Table 7 presents the results. In Panel A, firms with greater pre-shock exposure to machine-based investors experience significantly more activism following the ChatGPT shock. The interaction coefficient on *Post* × *Machine-based Ownership* is positive and significant across all three measures: the probability of facing any activist campaign increases by 3.8 percentage points (Column 1), the probability of a definitive proxy fight increases by 3.7 percentage points (Column 2), and the number of shareholder proposals filed increases by 0.115 per quarter (Column 3). In Panel B, we split the sample by activist ownership above and below the sample median. The effects are concentrated among firms with greater activist ownership, suggesting that machine-based investors strengthen voice-based governance indirectly by attracting and enabling external activists rather than engaging themselves directly.

The mechanism behind these results operates through the exit channel and coalition facilitation rather than machine-based investors' direct engagement. When machine-based investors algorithmically identify governance deficiencies (e.g., inefficient investment, earnings manipulation, or value-destroying managerial decisions), their typical response is to reduce or exit their positions. This exit behavior, particularly when coordinated across multiple algorithmic investors responding to the same disclosures, generates downward price pressure that depresses market valuation relative to fundamental value. The resulting undervaluation creates precisely the conditions that attract external activist investors: a cheap entry point, an identifiable governance problem, and a shareholder base increasingly tilted toward short-term investors, lowering the coalition-building threshold needed to win a proxy contest. In untabulated results, we find that machine-based investors are typically not themselves filing campaigns or proxy fights, confirming that their role is catalytic rather than direct. Their algorithmic trading behavior creates the conditions that make external activism both attractive and feasible.

## 6. AI-Powered Governance and Real Corporate Decisions

### *6.1 M&A Activity*

Corporate takeovers represent one of the most consequential discretionary investments available to managers. Agency theory predicts that entrenched executives may pursue empire-building M&A strategies to increase firm size, managerial compensation, and personal prestige at the expense of shareholder value (Jensen, 1986). Institutional investors with strong monitoring capabilities can prevent such behavior by increasing the probability of post-deal accountability through exit or voice mechanisms. Machine-based investors, equipped with generative AI tools to evaluate deal valuations and synergy projections in real time, should be particularly effective at identifying and discouraging value-destroying deals. The more informed and timely monitoring approach enforces capital discipline, constraining opportunistic managerial behavior and improving the average quality of deals that do proceed. In addition, their shift in attention toward exit-based discipline and their shorter investment horizons may also lead them to vote against long-term investment, including strategically sound but uncertain M&A deals. We therefore predict that increased AI-powered monitoring will reduce (value-destroying) M&A activities but increase the value creation among announced deals.

Table 8 analyzes merger and acquisition activity. First, we find that the likelihood of engaging in M&A activity is significantly reduced for treated firms in the post-shock period, with the coefficient on *Post* × *Machine-based Ownership* indicating a decline of 10.1% relative to the sample base rate, suggesting that AI-powered monitoring extends less support for long-term investment. Next, we find that firms with greater exposure to machine-based investors are less likely to engage in value-destroying M&A transactions following the ChatGPT shock. Among completed deals, the fraction of value-destroying transactions falls by 18.2% relative to the sample mean.[12] Furthermore, conditional on deal announcement, the cumulative abnormal returns (*CARs*) over the [−1, +1] announcement window improve by approximately 0.67 percentage points for treated firms relative to controls, suggesting that deals that do proceed are of higher quality on average. These patterns indicate that AI-powered investors not only discourage long-term risky projects but also constrain managerial empire-building through acquisitions that destroy shareholder value, consistent with institutional investors acting as a disciplinary force on managers' most visible capital allocation decisions.

### *6.2 Corporate Long-term Investment and Innovation*

The relationship between institutional monitoring and corporate investment is theoretically contested. On one hand, active monitoring can mitigate managerial short-termism and

[12] As an additional robustness check, Table IA.4 re-estimates our baseline specifications for binary outcome variables using probit models rather than linear probability models. The probit coefficients on *Post* × *Machine-based Ownership* are -0.815 for financial restatements, -0.647 for any M&A deal, and -0.608 for any value-destroying deal, all statistically significant at the 1% level. These results confirm that the reductions in restatement incidence and M&A activity documented in our main analysis are robust to the nonlinear structure of binary dependent variables and are not driven by the linear probability model's well-known limitations near the boundaries of the unit interval.

overinvestment driven by empire-building motives (Shleifer and Vishny, 1986; Edmans, 2009), thereby improving the efficiency of capital allocation. On the other hand, certain forms of institutional oversight may exacerbate myopic behavior if investors primarily respond to quantifiable, short-horizon performance signals (Bushee, 1998). AI-powered investors present a distinctive case: their reliance on algorithmically extractable metrics and their potential shift toward shorter holding periods and exit-based governance may systematically disadvantage investment projects characterized by long payback periods, high uncertainty, and limited near-term measurability. R&D expenditures and long-term capital investment are precisely such projects. If AI-augmented investors become more impatient and short-term oriented (e.g., disliking or not advising long-term risky investments), they may broadly reduce investment intensity as managers respond to the heightened threat of exit.

Table 9 Panel A examines real investment and innovation quantity across columns (1)-(3). Following the ChatGPT shock, firms with higher machine-based investor exposure significantly reduce long-term investment [13] and R&D intensity. Economically, a one-standard-deviation increase in pre-shock machine-based ownership is associated with a decline of approximately 0.003 in long-term investment scaled by assets (6.0% of the sample mean) and a reduction of approximately 0.001 in R&D intensity (5.8% of the sample mean). We also measure innovation quantity using the number of patents filed by the firm in each quarter, adjusted for filing lags using application dates rather than grant dates to ensure contemporaneous measurement. The number of patents filed declines significantly, with treated firms filing roughly 13.8% fewer patents per quarter compared to the control group. These effects are not mechanically driven by changes in firm size or profitability, as the specifications control for *ROA, leverage,* and *market-to-book ratio*. Moreover, the results persist after controlling for firms' own AI adoption, confirming that investment reductions stem from external monitoring pressure rather than from internal technological transformation. This evidence is consistent with our interpretation that machine-based investors become more impatient and short-term oriented, suggesting that they discourage firms' long-term risky investments. As AI-powered investors shorten their horizons and rely on

[13] Long-term investment is measured as Compustat quarterly item IVLTQ (Investments and Advances, Long-Term) scaled by total assets (ATQ). IVLTQ captures the book value of long-term investment commitments recorded on the balance sheet, including equity method investments, long-term advances, and other non-current investment holdings. This measure captures a broader set of long-horizon capital commitments than capital expenditures (CAPX) alone, which records only periodic cash outflows for property, plant, and equipment, and is therefore better suited to capturing the full scope of managerial investment discretion over long-horizon projects.

exit-based rather than voice-based governance, managers face heightened pressure to deliver measurable near-term outcomes, leading them to scale back uncertain, long-horizon projects even when those projects may have positive expected value.

Earlier results show that AI-powered monitoring reduces the quantity of innovation; it is possible that tighter monitoring simultaneously improves the quality of investment selection. Agency theory frameworks predict that managerial slack may manifest not only as overinvestment but also as undisciplined project selection. For instance, managers may scatter resources across too many marginal projects rather than concentrating on the most promising ones (Holmstrom and Tirole, 1997). If AI-powered investors, through strict monitoring, enforce greater project selectivity, firms may respond by pruning low-quality innovation while doubling down on high-potential endeavors. This "innovation efficiency" hypothesis predicts opposite directional effects on patent quantity and patent quality following enhanced monitoring. At the same time, AI-powered investors' shorter horizons and impatience toward long-horizon uncertainty may induce a tilt away from exploratory, frontier innovation (which generates highly cited breakthrough patents) toward exploitative, incremental innovation with more predictable near-term commercialization. The net effect on patent quality is therefore an empirical question. We expect a positive effect on patent value if the innovation efficiency channel dominates, accompanied by a shift in the composition of patenting activity toward more commercially focused projects.

Table 9 Panel B examines innovation quality rather than quantity, using cumulative future average patent value per patent over 4-, 8-, and 12-quarter windows following Kogan et al. (2017). The significantly positive coefficients on *Post* × *Machine-based Ownership* across all horizons suggest that treated firms experience meaningfully higher value creation from patent grants following the ChatGPT shock. For nominal patent value, the coefficients are 1.838, 1.618, and 1.532 at the 4-, 8-, and 12-quarter horizons, respectively, all significant at the 1% level. These estimates indicate that a one-standard-deviation increase in pre-shock machine-based ownership is associated with economically meaningful increases in per-patent value over both near- and medium-term horizons. Results are qualitatively robust to inflation adjustment, with real patent value coefficients of 1.416, 1.272, and 1.215 at the 4-, 8-, and 12-quarter horizons, respectively, all significant at the 1% level. The consistency across nominal and real measures suggests that the quality improvement is not a mechanical consequence of reduced patent counts or inflation effects. This divergence between patent quantity and patent value provides direct support for the innovation efficiency hypothesis: AI-powered investors strictly monitor investment behavior, which helps increase innovation efficiency. Rather than spreading resources thin across many marginal projects, firms under AI-powered governance concentrate their innovative efforts on fewer but higher-impact inventions. This is consistent with the previous results on M&A outcomes: this evidence collectively indicates that AI-powered governance reshapes corporate resource allocation by constraining low-quality, high-uncertainty commitments while sharpening focus on economically significant projects.

## 7. Investor AI Adoption and Institutional Investors' Portfolio Changes

### *7.1 Investor AI Adoption and Portfolio Changes*

At the investor level, we examine how the ChatGPT shock affects institutional investors' portfolio structure and trading patterns. Information acquisition costs determine the scope and intensity of shareholder oversight (Kacperczyk et al., 2016). When it is costly to analyze each additional firm, investors face a binding attention constraint that forces them to concentrate their portfolios and monitoring effort on a limited set of positions. Generative AI, as a technology shock, can dramatically reduce the marginal cost of processing corporate disclosures and relax this constraint along two margins simultaneously. First, investors can expand the breadth of their coverage to a larger set of firms while maintaining monitoring quality. Second, with more firms under active surveillance, exit becomes a more credible and actionable governance mechanism, shortening optimal holding periods and increasing portfolio turnover. This scale effect of cheaper information is expected to produce a distinctive portfolio structure: broader portfolios across industries and shorter investment horizons among investors who adopt AI-augmented analysis. We test this prediction directly using investor-level data.

Table 10 examines how institutional investors adjust their portfolio strategies following the ChatGPT shock. Panel A reports results for portfolio breadth, where the dependent variables are the number of newly entered industries and the industry-level HHI. Panel B examines investment horizon and trading activity, with dependent variables for average holding duration and portfolio turnover. We find that machine-based investors hold more diversified portfolios across industries, consistent with a decline in information acquisition costs, and exhibit significantly higher portfolio turnover, indicating shorter holding periods and greater reliance on exit-based governance.

Economically, machine-based investors exhibit approximately a 3.9% reduction in industry-level HHI and a 19.0% increase in portfolio turnover following the ChatGPT shock.

Panel C of Table 10 further reveals that machine-based investors earn significantly higher risk-adjusted returns, as measured by Fama-French five-factor alpha, while simultaneously reducing their passive ETF exposure. The positive alpha effect indicates that AI-powered information processing translates into genuine investment performance gains, not merely portfolio reshuffling. The decline in ETF share is consistent with a shift from passive, index-tracking strategies toward active, firm-specific analysis enabled by lower information-processing costs. These results suggest that generative AI adoption enables institutional investors to scale monitoring across a broader portfolio while shifting toward more liquid, active, and information-intensive trading strategies.

Panel D provides additional validation for our DiD design by examining whether machine-based investors meaningfully increase their reliance on generative AI following the ChatGPT shock. Following Sheng, et al. (2026), we construct a measure of *GenAI Reliance* that captures the degree to which an investor's portfolio changes are driven by ChatGPT-generated signals beyond standard firm fundamentals. Column (1) shows that machine-based investors exhibit significantly higher *GenAI Reliance* following the shock, consistent with the ChatGPT release representing a genuine shift in the information processing behavior of treated investors. Column (2) confirms that this pattern persists in the post-shock period, where machine-based investors continue to display meaningfully higher *GenAI Reliance* relative to other institutional investors. These results confirm that our treatment variable captures real changes in AI-powered information processing rather than a spurious correlation with the shock timing.

*7.2 Investor AI Adoption and Firm-level Institutional Ownership Changes*

The investor-level evidence establishes that machine-based investors expand their portfolio breadth and shorten their investment horizons following the ChatGPT shock. A natural implication at the firm level is that firms with pre-existing exposure to machine-based investors should experience an increase in the number of such investors following the shock. Cheaper information processing lowers the fixed cost of monitoring any given firm, inducing entry by investors who were previously unable to justify the cost. However, because the cost reduction enables each investor to simultaneously cover more firms, individual ownership stakes need not increase and may in fact decline as each investor spreads attention more broadly. The net effect on aggregate machine-based ownership, therefore, reflects a tension between extensive-margin entry and intensive-margin dilution. We test these predictions using firm-level decompositions of institutional ownership.

Table 11 analyzes how institutional ownership responds at the firm level. Consistent with the scale effect, firms with higher pre-shock exposure to machine-based investors experience a statistically significant post-shock increase in the total number of machine-based investors. At the same time, average ownership per machine-based investor declines significantly, indicating that the effects of the new investor entry outweigh the diluted ownership stakes by existing investors. Together, these patterns suggest that the ChatGPT shock expanded machine-based investor coverage breadth due to the reduction of information processing cost resulted from AI adoption.

Interestingly, beyond changes in ownership levels, we find that machine-based investors become significantly more synchronized in their trading behavior following the shock. Columns (3) and (4) examine the dispersion of selling and buying decisions across machine-based investors holding the same firm. Following the shock, both selling dispersion and buying dispersion decline significantly among treated firms, indicating that machine-based investors converge in their portfolio decisions after the ChatGPT shock. This synchronization is consistent with investors sharing similar AI-based analytical frameworks and information processing routines, which naturally leads to correlated assessments of firm fundamentals and coordinated trading responses. Importantly, this coordination arises without explicit communication: investors relying on similar algorithms and processing the same publicly available information independently converge on similar decisions. The result suggests a more collectively enforced, rule-based governance mechanism that can amplify the disciplinary effect of any individual investor's exit threat. Therefore, we expect that when machine-based investors move in concert, the credibility and market impact of exit-based discipline is substantially strengthened, giving rise to a more powerful and scalable form of governance.

### *7.3 Investor AI Adoption and Active Voice*

Our findings in Section 5.3, showing enhanced shareholder activism following the ChatGPT shock in treated firms, raise a natural question: how does the shortening of investor horizons induced by AI-powered information processing enhance shareholder activism? Shorter horizons may weaken incentives for direct engagement by machine-based investors, as activism is costly and typically requires time to realize its benefits. Instead, we hypothesize that machine-based investors facilitate coalition formation with activist shareholders through two channels. First, enhanced information processing improves the precision of target identification, reducing activists' screening costs and sharpening the detection of persistent governance deficiencies. Second, the correlated beliefs and synchronized trading behavior of machine-based investors reduce coordination frictions among dispersed shareholders, making it easier for activists to build coalitions and secure voting support.[14] In this section, we examine changes in activist presence and campaign effectiveness using our DiD setting.

Table 12 examines whether the ChatGPT shock affects the composition of activist investors and the effectiveness of their campaigns once initiated. The results in columns (1)-(3) indicate that treated firms attract a more experienced and capable pool of activist investors following the shock. Column (1) measures the presence of experienced activists, defined as investors who have engaged in at least one activism campaign in U.S.-listed firms since 2015 up to a given firm-quarter. We find a significant increase in the number of experienced activists in treated firms after the shock, consistent with improved targeting and selection. Column (2) shows that activist ownership increases by 1.2 percentage points in treated firms following the shock. Column (3) further indicates that newly entering activists in treated firms are more likely to initiate campaigns within a given firm-quarter. In untabulated results, we find little evidence that machine-based investors themselves engage directly in activism.

Importantly, these activists are not only more numerous but also more effective once campaigns are initiated. Columns (4) and (5) show that the likelihood of reaching a settlement increases by 0.9 percentage points, and the probability of securing board representation also rises significantly. These findings suggest that activism in treated firms is more likely to generate substantive governance outcomes. Our evidence is consistent with an indirect channel: by reshaping the information environment, AI-powered investors make governance problems more

[14] Although machine-based investor classification is not directly observable, activists can infer the likely presence of such investors from publicly available information. Institutional holdings disclosed through 13F filings reveal the identity and composition of a firm's shareholder base, allowing activists to identify quantitative and systematically oriented investors as plausible proxies for machine-based investment behavior. In addition, trading patterns, such as high turnover, synchronized rebalancing, and strong price responsiveness to disclosures, provide further signals of model-driven investment behavior. Together, these observable characteristics enable activists to screen for firms with shareholder bases that are more responsive and predictable, lowering coordination uncertainty and increasing the expected payoff to intervention.

visible and firms more attractive to specialized activists. At the same time, synchronized trading and portfolio reallocation by machine-based investors generate signals that may attract activist attention and increase the willingness of other institutional shareholders to support activist proposals. This, in turn, lowers the coordination costs that typically constrain activist success.

Taken together, these findings point to a governance structure in which AI-powered investors and activist shareholders play complementary roles. Machine-based investors strengthen exit-based discipline through trading, while activist investors provide voice-based discipline through direct engagement. By facilitating coordination between these two channels, AI adoption enhances the effectiveness of shareholder oversight without requiring machine-based investors to engage directly in activism.

## 8. Conclusion

This paper examines how the rise of AI-powered institutional investors reshapes corporate governance. Exploiting the November 2022 release of ChatGPT as a quasi-experimental shock to generative AI adoption costs, we document that exposure to machine-based investors generates significant valuation gains. These effects are concentrated among firms facing greater ex-ante monitoring frictions, suggesting that AI-enhanced monitoring improves governance efficiency precisely where information-processing constraints previously limited oversight.

To understand the sources of these valuation gains, we examine governance outcomes across multiple dimensions. Firms with greater machine-based investor exposure exhibit stronger pay-performance sensitivity, greater CEO turnover-performance sensitivity, improved accounting quality, and fewer financial restatements. Shareholder activism also intensifies following the shock, with these firms experiencing higher activist campaign incidence, more proxy fights, and greater shareholder proposal submissions. These patterns confirm that the valuation improvements reflect genuine improvements in governance quality. Importantly, these results remain robust after controlling for firms' own AI adoption, confirming that the governance effects operate through external monitoring rather than internal technological transformation.

AI-powered governance also reshapes corporate investment. Managers facing heightened exit threats scale back uncertain, long-horizon investment projects. Firms with greater machine-based investor exposure reduce overall M&A activity, long-term investment, R&D intensity, and patent counts. Yet the deals and patents that survive are of higher quality: we observe lower incidence of value-destroying acquisitions, higher announcement returns, and greater patent value. These findings suggest that AI-powered investors improve capital allocation efficiency by screening out marginal projects and concentrating resources on higher-impact investments.

We then examine the mechanisms underlying these outcomes at both the investor and firm levels. At the investor level, reduced information acquisition costs lead AI-oriented institutional investors to expand portfolio breadth while simultaneously shortening investment horizons and increasing their reliance on generative AI tools following the shock. These shifts are consistent with a transition toward scalable, signal-driven monitoring and greater reliance on exit-based discipline rather than sustained engagement. At the firm level, exposure to machine-based investors reshapes ownership structure and monitoring intensity. Firms with greater machine-based investor exposure experience an increase in the number of machine-oriented investors, a decline in ownership concentration per investor, and greater synchronization in trading decisions, strengthening collective discipline and amplifying market-based governance mechanisms. The governance improvements operate through two complementary channels: machine-based investors discipline managers directly through synchronized trading and credible exit threats, and indirectly by reshaping the information environment and ownership composition in ways that attract and empower external activist investors.

As AI adoption costs continue to decline, machine-based monitoring is likely to become increasingly pervasive, reshaping the institutional investor landscape. Future research should further examine how AI adoption by investors interacts with firms' strategic investment, disclosure design, and innovation trajectories. Understanding the long-run equilibrium consequences of algorithmic monitoring, particularly its impact on growth, risk-taking, and stakeholder welfare, remains an important frontier in corporate finance.

## References


Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P., 2022. Artificial intelligence and jobs: Evidence from online vacancies. Journal of Labor Economics, 40(S1), pp.S293-S340.

Admati, A. R., and P. Pfleiderer, 2009, The "wall street walk" and shareholder activism: Exit as a form of voice, Review of Financial Studies 22, 2645-2685.

Admati, A. R., P. Pfleiderer, and J. Zechner, 1994, Large shareholder activism, risk sharing, and financial market equilibrium, Journal of Political Economy 102, 1097-1130.

Alekseeva, L., Azar, J., Giné, M., Samila, S. and Taska, B., 2021. The demand for AI skills in the labor market. Labour Economics, 71, p.102002.

Appel, I. R., T. A. Gormley, and D. B. Keim, 2016, Passive investors, not passive owners, Journal of Financial Economics 121, 111-141.

Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager, 2021, Synthetic difference-in-differences, American Economic Review 111, 4088-4118.

Babina, T., A. Fedyk, A. He, and J. Hodson, 2024, Artificial intelligence, firm growth, and product innovation, Journal of Financial Economics 151, 103745.

Bao, Y., B. Ke, B. Li, J. Yu, and J. Zhang, 2020, Detecting accounting fraud in publicly traded U.S. firms using a machine learning approach, Journal of Accounting Research 58, 199-235.

Bartram, S.M., Branke, J. and Motahari, M., 2020. Artificial intelligence in asset management. CFA Institute Research Foundation.

Becht, M., J. Franks, C. Mayer, and S. Rossi, 2009, Returns to shareholder activism: Evidence from a clinical study of the Hermes UK Focus Fund, Review of Financial Studies 22, 3093-3129.

Ben-David, D., Mintz, I., and Sade, O., 2026. Using AI and behavioral finance to cope with limited attention and reduce overdraft fees. Management Science, 72(1), pp.204-222.

Ben-David, I., Franzoni, F., Moussawi, R., and Sedunov, J., 2021, The granular nature of large institutional investors, Management Science, 67(11), 6629-6659.

Bena, J., Bian, B., and Giannetti, M. (2026). Prompted to Start: How Generative AI is Transforming Entrepreneurship. HKU Jockey Club Enterprise Sustainability Global Research Institute Paper, (2025/182).

Bertomeu, J., Y. Lin, Y. Liu, and Z. Ni, 2025, The impact of generative AI on information processing: Evidence from the ban of ChatGPT in Italy, Journal of Accounting and Economics, forthcoming.

Boehmer, E., K. Y. L. Fong, and J. Wu, 2021, Algorithmic trading and market quality: International evidence, Journal of Financial and Quantitative Analysis 56, 2659-2686.

Brav, A., W., Jiang, F., Partnoy, and R., Thomas, 2008, Hedge fund activism, corporate governance, and firm performance. The Journal of Finance, 63(4), 1729-1775.

Brogaard, J., T. Hendershott, and R. Riordan, 2014, High-frequency trading and price discovery, Review of Financial Studies 27, 2267-2306.

Brynjolfsson, E., D. Li, and L. Raymond, 2025, Generative AI at work, Quarterly Journal of Economics 140, 889–942.

Bushee, B. J., 1998, The influence of institutional investors on myopic R&D investment behavior, Accounting Review 73, 305-333.

Cao, S., Jiang, W., Wang, J. and Yang, B., 2024. From man vs. machine to man+ machine: The art and AI of stock analyses. Journal of Financial Economics, 160, 103910.

Cao, S., W. Jiang, B. Yang, and A. L. Zhang, 2023, How to talk when a machine is listening: Corporate disclosure in the age of AI, Review of Financial Studies 36, 3603-3642.

Chen, H., Cohen, L., Gurun, U., Lou, D., and Malloy, C., 2020, IQ from IP: Simplifying search in portfolio choice. Journal of Financial Economics, 138(1), 118-137.

Chen, M.A. and Wang, J., 2025a. Displacement or augmentation? The effects of AI innovation on workforce dynamics and firm value. Available at SSRN 5246388.

Chen, M.A. and Wang, J.X., 2025b. More Than Human: The Informational Effects of AI-Powered Corporate Disclosure. Available at SSRN 5959496.

Cohen, L., A. Frazzini, and C. Malloy, 2008, The small world of investing: Board connections and mutual fund returns, Journal of Political Economy 116, 951-979.

Core, J. E., R. W. Holthausen, and D. F. Larcker, 1999, Corporate governance, chief executive officer compensation, and firm performance, Journal of Financial Economics 51, 371-406.

D'Acunto, F., Prabhala, N. and Rossi, A.G., 2019. The promises and pitfalls of robo-advising. The Review of Financial Studies, 32(5), 1983-2020.

Daniel, K., D., Hirshleifer, and L. Sun, 2020, Short- and long-horizon behavioral factors. The Review of Financial Studies, 33(4), 1673-1736.

Dechow, P. M., Sloan, R. G., and Sweeney, A. P. 1995, Detecting earnings management. The Accounting Review, 70(2), 193–225.

Dou, W. W., I., Goldstein, and Y., Ji, 2025, AI-powered trading, algorithmic collusion, and price efficiency. National Bureau of Economic Research.

Drake, M. S., D. T. Roulstone, and J. R. Thornock, 2015, The determinants and consequences of information acquisition via EDGAR, Contemporary Accounting Research 32, 1128-1161.

Dugast, J., and T. Foucault, 2018, Data abundance and asset price informativeness, Journal of Financial Economics 130, 367-391.

Edmans, A., 2009, Blockholder trading, market efficiency, and managerial myopia, Journal of Finance 64, 2481-2513.

Edmans, A., and C. G. Holderness, 2017, Blockholders: A survey of theory and evidence. The handbook of the economics of corporate governance, 1, 541-636.

Edmans, A., and G. Manso, 2011, Governance through trading and intervention: A theory of multiple blockholders, Review of Financial Studies 24, 2395-2428.

Eisfeldt, A. L., G. Schubert, M. B. Zhang, and B. Taska, 2025, The labor impact of generative AI on firm values, Working paper.

Elmobark, N., 2025, A comparative analysis of Python text matching libraries: A multilingual evaluation of capabilities, performance and resource utilization, International Journal of Environment, Engineering and Education, 7(1), 48-60.

Fedyk, A., Hodson, J., Khimich, N. and Fedyk, T., 2022. Is artificial intelligence improving the audit process? A. Fedyk et al. Review of Accounting Studies, 27(3), 938-985.

Gaspar, J., M. Massa, and P. Matos, 2005, Shareholder investment horizons and the market for corporate control, Journal of Financial Economics 76, 135-165.

Gennaioli, N., A., Shleifer, and R., Vishny, 2015, Neglected risks: The psychology of financial crises. American Economic Review, 105(5), 310-314.

Giglio, S., M., Maggiori, J., Stroebel, and S., Utkus, 2021, Five facts about beliefs and portfolios. American Economic Review, 111(5), 1481-1522.

Gofman, M. and Jin, Z., 2024. Artificial intelligence, education, and entrepreneurship. The Journal of Finance, 79(1), 631-667.

Goldstein, I., W. Jiang, and G. A. Karolyi, 2019, To FinTech and beyond, Review of Financial Studies 32, 1647-1661.

Grossman, S. J., and J. E. Stiglitz, 1980, On the impossibility of informationally efficient markets, American Economic Review 70, 393-408.

Gu, S., B. Kelly, and D. Xiu, 2020, Empirical asset pricing via machine learning, Review of Financial Studies 33, 2223-2273.

Gupta, A., F., Qian, E., Simintzi, and Y., Sun, 2025, Generative AI and Entrepreneurship, Available at SSRN 5971474.

Healy, P. M., and J. M. Wahlen, 1999, A review of the earnings management literature and its implications for standard setting, Accounting Horizons 13, 365-383.

Hendershott, T., C. M. Jones, and A. J. Menkveld, 2011, Does algorithmic trading improve liquidity?, Journal of Finance 66, 1-33.

Hoberg, G., and G. Phillips, 2016, Text-based network industries and endogenous product differentiation, Journal of Political Economy 124, 1423-1465.

Holmstrom, B., and J. Tirole, 1997, Financial intermediation, loanable funds, and the real sector, Quarterly Journal of Economics 112, 663-691.

Jensen, M. C., 1986, Agency costs of free cash flow, corporate finance, and takeovers, American Economic Review 76, 323-329.

Jensen, M. C., and K. J. Murphy, 1990, Performance pay and top-management incentives, Journal of Political Economy 98, 225-264.

Jones, K. K., 1991, Earnings management during import relief investigations, Journal of Accounting Research 29, 193-228.

Kacperczyk, M., S. Van Nieuwerburgh, and L. Veldkamp, 2016, A rational theory of mutual funds' attention allocation, Econometrica 84, 571-626.

Khalid, F., Srivastava, M. and Toumi, F., 2025. AI technology adoption and corporate financial misconduct. Journal of Financial Reporting and Accounting, 1-23.

Kogan, L., Papanikolaou, D., Seru, A. and Stoffman, N., 2017. Technological innovation, resource allocation, and growth. The Quarterly Journal of Economics, 132(2), 665-712.

Liberti, J. M., and M. A. Petersen, 2019, Information: Hard and soft, Review of Corporate Finance Studies 8, 1-41.

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

Loughran, T., and B. McDonald, 2011, When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks, Journal of Finance 66, 35-65.

McCahery, J. A., Z. Sautner, and L. T. Starks, 2016, Behind the scenes: The corporate governance preferences of institutional investors, Journal of Finance 71, 2905-2932.

Peng, L., and W. Xiong, 2006, Investor attention, overconfidence and category learning, Journal of Financial Economics 80, 563-602.

Peters, R. H., and L. A. Taylor, 2017, Intangible capital and the investment-q relation, Journal of Financial Economics 123, 251-272.

Ryans, J. P., 2021, Textual classification of SEC comment letters: JP Ryans. Review of Accounting Studies, 26(1), 37-80.

Sheng, J., Sun, Z., Yang, B. and Zhang, A.L., 2026. Generative AI and asset management. Review of Financial Studies, forthcoming.

Shleifer, A., and R. W. Vishny, 1986, Large shareholders and corporate control, Journal of Political Economy 94, 461-488.

Sims, C. A., 2003, Implications of rational inattention, Journal of Monetary Economics 50, 665-690.

Tetlock, P. C., 2007, Giving content to investor sentiment: The role of media in the stock market, Journal of Finance 62, 1139-1168.

## Appendix A. Variable Definition

| Variable | Explanation |
|---|---|
| Panel A. Firm-level variables | |
| Activist Ownership | The share of ownership held by institutional investors associated with activist campaign. |
| Any Campaign | An indicator variable equal to one if the firm is targeted by at least one activist campaign during the quarter, and zero otherwise. The data are sourced from *FactSet SharkWatch*. |
| Any M&A deal | An indicator variable equal to one if the firm announces any M&A deals as an acquirer in a given quarter, as reported in the *SDC* database. |
| Any Proxy Fight | An indicator variable equal to one if at least one activist campaign against the firm results in a definitive proxy fight during the quarter, and zero otherwise. |
| Any Restatement | An indicator variable equal to one if the firm announces any financial restatement in a given quarter. |
| Any Value-destroying Deal | An indicator variable equal to one if the firm announces any M&As deals as an acquirer and experiences a negative announcement return (CAR) in a [-1, +1] day window around the deal announcement. |
| Average Ownership of Machine Investor | Total machine-based ownership divided by the number of machine-based investors. |
| Board Seat Won | The number of board seats secured by activists at the firm in a given quarter. |
| Buying Dispersion of Machine Investor | Firm-quarter cross-sectional variance across machine-based investors of the variable that equal to one if a machine-based investor executes at least one buy transaction in a firm's stock during a given quarter and zero otherwise. |
| CAR [-1, +1] | The cumulative abnormal return is estimated from the market model in a [-1, +1] day window around the deal announcement. |
| Cash | Total cash holdings divided by total assets. |
| CEO Ownership | Percentage of firm stock shares owned by the CEO in a given firm-year. |
| CEO Tenure | The number of years the CEO on position. |
| CEO Total Compensation | The total CEO compensation, including salary, bonus, other annual compensation, and restricted stock grants in a given firm-year, as reported in the *Execucomp* database. |
| CEO Turnover | An indicator variable equal to one if the CEO turnover event occurs in a given firm-quarter. |
| Experienced Activists | The number of investors who have engaged in at least one activism campaign in U.S.-listed firms since 2015 up to a given firm-quarter. |
| Firm AI Adoption | Logarithm of one plus the five-year average number of employees with AI-related skills at the firm-year level, where AI-skilled employees are identified from job titles in the Revelio Labs database using a keyword-based classification that captures AI roles, technical skills, and deployment activities. |
| Institutional Ownership | The fraction of shares outstanding held by institutional investors. |
| Leverage | Total liabilities divided by total assets. |
| Log (Number of Patents) | The log-transformed number of patents obtained from Kogan et al. (2017). |
| Log (R&D expenditure) | The log-formed R&D expenditures as reported in Compustat in a given firm-quarter. |
| Long-term Investment | The total long-term investment (Compustat quarterly item IVLTQ), which captures the book value of long-term investment commitments recorded on the balance sheet, including equity method investments, long-term advances, and other non-current investment holdings, scaled by total assets in a given firm-quarter. |
| Machine-based Ownership | Average of the fraction of shares outstanding held by Machine-based institutional investors for three years before the rollout of ChatGPT. |
| Market-to-Book | Market capitalization relative to its book value of equity. |

| | |
|---|---|
| New Activists | The number of newly entering investors that initiate an activist campaign in a given firm-quarter |
| Percentage of Value-destroying Deals | The percentage of announced deals in a given firm-quarter that experience a negative announcement return (CAR) in a [-1, +1] day window around the deal announcement. |
| Patent value (dollar adjusted) | The log-transformed average value per patent per firm-quarter deflated to 1982 (million) dollars using the CPI obtained from Kogan et al. (2017). |
| Patent value (nominal) | The log-transformed average value per patent per firm-quarter, deflated to 1982 dollars using the CPI, measured in millions of real dollars, obtained from Kogan et al. (2017). |
| Post | An indicator variable equal to one if the year-quarter is after 2022 Q4, after the rollout of ChatGPT on November 30, 2022, and zero otherwise. |
| R&D | The total R&D expenditures scaled by total assets in a given firm-quarter as reported in Compustat. |
| R&D missing | An indicator variable equals 1 if the firm has a missing value in R&D expenditures as reported in Compustat in a given firm-quarter. |
| Restatement Correction | The magnitude of financial restatement corrections (for net income) announced in a given firm-quarter. |
| ROA | The ratio of net income to total assets. |
| Selling Dispersion of Machine Investor | Firm-quarter cross-sectional variance across machine-based investors of the variable that equal to one if a machine-based investor executes at least one sell transaction in a firm's stock during a given quarter and zero otherwise. |
| Settled Campaign | The number of activist campaigns associated with a firm that reach a settlement within a given quarter. |
| Shareholder Proposals | The total number of shareholder proposals filed against the firm during the quarter. |
| Size | Logarithm of book assets. |
| Tangibility | Tangible assets divided by total assets. |
| Tobin's Q | Firm's market value, measured as the average quarterly share price times shares outstanding plus total debt, divided by total assets. |
| Total Number of Machine Investor | The number of machine-based institutional investors. |
| Total Q | The ratio of the market value of the firm to the replacement cost of total capital, measured as the sum of tangible assets and estimated intangible capital (Peters and Taylor 2017), where intangible capital is depreciated uniformly in each quarter. |

| Panel B. Investor-level variables | |
|---|---|
| Alpha (FF5) | Risk-adjusted abnormal return estimated using the Fama-French Five-Factor Model. |
| Average Holding Duration | Average length of time (in quarters) that the investor holds a stock position. |
| Earnings Growth Share | Proportion of holdings in firms with positive earnings growth. |
| ETF Share | Proportion of ETF among the total market value of holdings. |
| GenAI Reliance | A measure of institutional investor reliance on generative AI information constructed following Sheng, et al., (2026). For each institutional investor in each quarter, GenAI Reliance is defined as the incremental *R*-squared from adding 14 ChatGPT-generated signals derived from earnings conference call transcripts to a regression of quarterly portfolio holding changes on standard firm fundamentals, capturing the marginal contribution of AI-generated information to the investor's portfolio decisions beyond what is explained by firm fundamentals alone. |
| Industry-level HHI | Sum of square of the proportion of shares in each industry (classified by 2-digit SIC code) held by the investor. |
| Log(Portfolio Concentration) | Logarithm of the total market value of holdings divided by the number of stocks held. |
| Log(Total Portfolio Value) | Logarithm of the total market value of holdings by the investor. |
| Log(Total Shares) | Logarithm of the total number of shares held by the investor |

| | |
|---|---|
| Long-Term Holdings Share | Proportion of shares held continuously for more than eight quarters |
| Machine-based Investor (indicator) | An indicator variable equals 1 if the investor exhibits “robot-like” filings download activities in the SEC EDGAR web traffic data. |
| New-Industries | The number of industries newly entered by the investor (based on 2-digit SIC classification). |
| Past Portfolio Return | Value-weighted return of the investor’s portfolio in the previous quarter. |
| Portfolio Turnover (Count) | Sum of absolute changes in the number of shares across all stocks between two quarters, scaled by the sum of total holding amount in the current and previous quarters. |
| Portfolio Turnover (Value) | Sum of absolute changes in portfolio weights across all stocks between two quarters, scaled by the sum of portfolio weights in the current and previous quarters. |
| Return Volatility | Portfolio-level return volatility computed as the weighted average of individual stock return volatilities, using portfolio weights. |

## Figure 1. Dynamic Effect of AI-Powered Governance on Firm Valuation

The figure plots the estimated coefficients and their associated two-tailed 95% confidence intervals from a yearly event-time specification, where the horizontal axis indicates years relative to the ChatGPT shock in 2022. Panel A reports results for Tobin's Q, defined as the ratio of the market value of assets to the book value of assets. Panel B presents results for Total Q, which extends Tobin's Q by incorporating intangible capital following Peters and Taylor (2017). Both measures are constructed as yearly averages.

Panel A. Dynamic Analysis for Tobin's Q

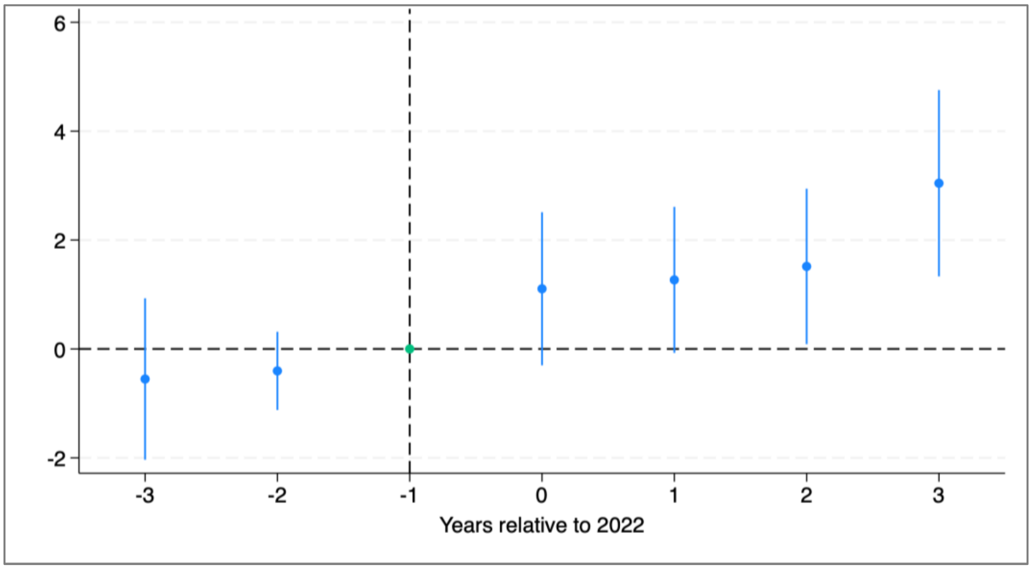


Panel B. Dynamic Analysis for Total Q

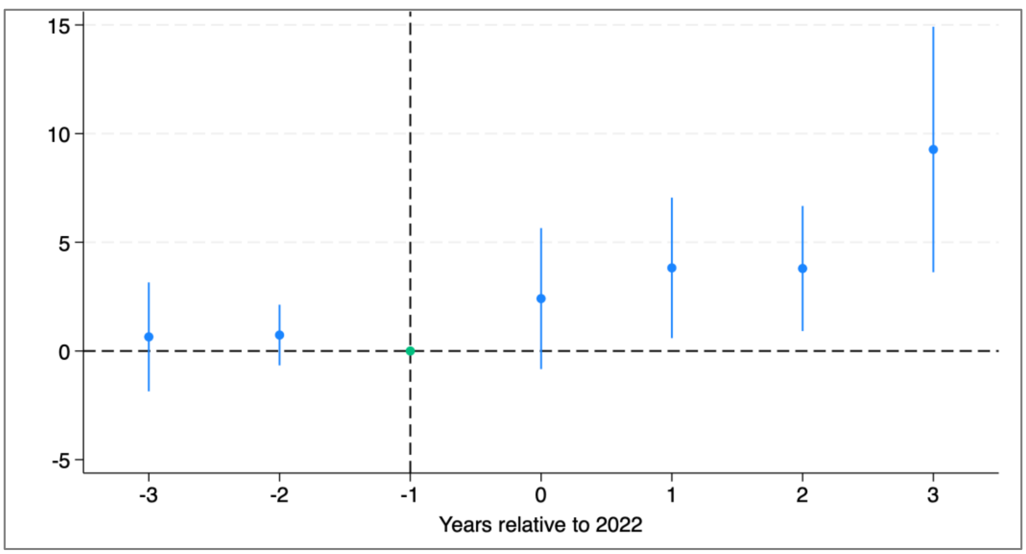

## Table 1. Summary Statistics

The table presents summary statistics for our sample. Panel A reports the firm-level statistics. The firm-level sample consists of 68,325 firm-year-quarter observations for 5,063 unique firms, covering all US-listed firms over the sample period from 2019 Q3 to 2025 Q3. Panel B reports the investor-level statistics. The investor-level sample consists of 46,742 investor-year-quarter observations for 2,529 investors. Appendix A provides detailed variable descriptions.

Panel A. Firm-level statistics

| | Mean | Std. Dev | 25th | 50th | 75th |
|---|---|---|---|---|---|
| Activist Share | 0.022 | 0.032 | 0.000 | 0.010 | 0.029 |
| Any Campaign | 0.019 | 0.137 | 0.000 | 0.000 | 0.000 |
| Any M&A deal | 0.148 | 0.355 | 0.000 | 0.000 | 0.000 |
| Any Proxy Fight | 0.018 | 0.133 | 0.000 | 0.000 | 0.000 |
| Any Restatement | 0.036 | 0.187 | 0.000 | 0.000 | 0.000 |
| Any Value-destroying Deal | 0.065 | 0.246 | 0.000 | 0.000 | 0.000 |
| Average Ownership of Machine Investor | 0.235 | 0.438 | 0.106 | 0.172 | 0.267 |
| Board Seat Won | 0.006 | 0.159 | 0.000 | 0.000 | 0.000 |
| Buying Dispersion of Machine Investor | 0.234 | 0.051 | 0.231 | 0.249 | 0.255 |
| CAR [-1, +1] | 0.014 | 0.083 | -0.020 | 0.002 | 0.039 |
| Cash | 0.245 | 0.276 | 0.043 | 0.120 | 0.357 |
| CEO Ownership | 2.253 | 5.684 | 0.097 | 0.415 | 1.588 |
| CEO Tenure | 7.225 | 8.317 | 2.000 | 5.000 | 10.000 |
| CEO Total Compensation (USD bn) | 0.010 | 0.015 | 0.003 | 0.006 | 0.012 |
| CEO Turnover (Indicator) | 0.099 | 0.299 | 0.000 | 0.000 | 0.000 |
| Experienced Activist | 3.651 | 3.964 | 1.000 | 3.000 | 5.000 |
| Firm AI Adoption | 1.235 | 1.086 | 0.470 | 0.936 | 1.686 |
| Institutional Ownership | 0.575 | 0.304 | 0.314 | 0.643 | 0.843 |
| Leverage | 0.259 | 0.241 | 0.049 | 0.194 | 0.416 |
| Log (Number of Patents) | 0.161 | 0.646 | 0.000 | 0.000 | 0.000 |
| Log (R&D expenditure) | 15.558 | 61.578 | 0.000 | 0.000 | 4.711 |
| Long-term Investment | 0.045 | 0.146 | 0.000 | 0.000 | 0.013 |
| Machine-based Ownership | 0.143 | 0.106 | 0.045 | 0.134 | 0.226 |
| Market-to-Book | 2.248 | 2.186 | 1.020 | 1.422 | 2.501 |
| New Activist | 0.001 | 0.034 | 0.000 | 0.000 | 0.000 |
| Patent value (dollar adjusted) | 2.176 | 1.469 | 0.899 | 2.041 | 3.205 |
| Patent value (nominal) | 3.091 | 1.692 | 1.712 | 3.092 | 4.318 |
| Percentage of Value-destroying Deals | 0.031 | 0.169 | 0.000 | 0.000 | 0.000 |
| Post | 0.388 | 0.487 | 0.000 | 0.000 | 1.000 |
| R&D | 0.024 | 0.053 | 0.000 | 0.000 | 0.024 |
| R&D missing (Indicator) | 0.499 | 0.500 | 0.000 | 0.000 | 1.000 |
| Restatement Correction | -0.024 | 7.823 | 0.000 | 0.000 | 0.000 |
| ROA | -0.005 | 0.072 | -0.010 | 0.013 | 0.031 |
| Selling Dispersion of Machine Investor | 0.243 | 0.041 | 0.242 | 0.252 | 0.256 |
| Settled Campaign | 0.001 | 0.051 | 0.000 | 0.000 | 0.000 |
| Shareholder Proposals | 0.014 | 0.264 | 0.000 | 0.000 | 0.000 |
| Size | 7.157 | 2.244 | 5.612 | 7.171 | 8.702 |
| Tangibility | 0.217 | 0.248 | 0.027 | 0.112 | 0.324 |
| Tobin's Q | 1.753 | 2.221 | 0.533 | 1.031 | 2.075 |
| Total Number of Machine Investor | 86.867 | 80.911 | 26.000 | 63.000 | 119.000 |
| Total Q | 1.961 | 3.581 | 0.435 | 0.891 | 1.893 |

*Table 1 (Continued)*

Panel B. Investor-level statistics

| | Mean | Std. Dev | 25th | 50th | 75th |
|---|---|---|---|---|---|
| Alpha (FF5) | 0.043 | 0.051 | 0.015 | 0.042 | 0.075 |
| Average Holding Duration | 6.797 | 4.093 | 3.425 | 6.123 | 9.591 |
| Earnings Growth Share | 0.049 | 0.054 | 0.006 | 0.030 | 0.072 |
| ETF Share | 0.176 | 0.265 | 0.000 | 0.030 | 0.253 |
| GenAI Reliance | 0.255 | 0.204 | 0.092 | 0.202 | 0.373 |
| Industry-level HHI | 0.230 | 0.213 | 0.092 | 0.139 | 0.272 |
| Log(Portfolio Concentration) | 11.272 | 3.593 | 8.346 | 9.755 | 14.955 |
| Log(Total Portfolio Value) | 4.923 | 1.288 | 4.025 | 4.812 | 5.673 |
| Log(Total Shares) | 16.187 | 3.756 | 13.057 | 14.758 | 19.780 |
| Long-Term Holdings Share | 0.499 | 0.280 | 0.286 | 0.577 | 0.727 |
| Machine-based Investor (indicator) | 0.051 | 0.220 | 0.000 | 0.000 | 0.000 |
| New-Industries | 1.615 | 2.480 | 0.000 | 1.000 | 2.000 |
| Past Portfolio Return | 0.031 | 0.092 | -0.015 | 0.048 | 0.086 |
| Portfolio Turnover (Count) | 0.151 | 0.165 | 0.047 | 0.083 | 0.177 |
| Portfolio Turnover (Value) | 0.158 | 0.200 | 0.062 | 0.093 | 0.153 |
| Post | 0.339 | 0.473 | 0.000 | 0.000 | 1.000 |
| Return Volatility | 0.019 | 0.010 | 0.013 | 0.016 | 0.022 |

## Table 2. Investor AI Adoption and Firm Valuation

This table reports results from firm-level regressions examining the effects of AI-powered governance on firm valuation following the ChatGPT shock. The sample consists of 58,661 firm-year-quarter observations for 4,464 unique firms for the sample period from 2019 Q3 to 2025 Q3, constructed from 13F institutional ownership data. In columns (1) and (2), the dependent variables are Tobin's Q and Total Q, respectively, which measure firm market valuation. Tobin's Q is defined as the ratio of the market value of assets to the book value of assets, while Total Q incorporates intangible capital following Peters and Taylor (2017). The key independent variable is Post × Machine-based Ownership, which interacts an indicator for the post-ChatGPT period with firms' pre-period exposure to Machine-based institutional investors. Control variables include firm size, return on assets (or lagged return on assets where applicable), leverage, cash holdings, asset tangibility, and institutional ownership. To account for firms' own adoption of artificial intelligence, all specifications additionally control for Firm AI Adoption, measured as the average log-transformed AI-skilled workers in the past five years. All regressions include firm fixed effects and year-quarter fixed effects. Standard errors reported in parentheses are robust and clustered by industry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

| | Tobin's Q<br>(1) | Total Q<br>(2) |
|---|---|---|
| Post × Machine-based Ownership | 0.613* | 2.168** |
| | (0.318) | (0.850) |
| Firm AI Adoption | -0.097 | -0.212 |
| | (0.099) | (0.215) |
| Size | -0.662*** | -0.509** |
| | (0.129) | (0.235) |
| ROA | 0.392 | 5.903*** |
| | (0.931) | (1.057) |
| Leverage | -2.070*** | -2.462*** |
| | (0.251) | (0.296) |
| Cash | 0.940*** | 4.170*** |
| | (0.220) | (0.986) |
| Tangibility | -0.084 | -1.796** |
| | (0.448) | (0.778) |
| Institutional Ownership | 0.245** | -0.458** |
| | (0.093) | (0.178) |
| | | |
| Firm fixed effects | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes |
| Observations | 58,661 | 40,074 |
| Adjusted R-squared | 0.768 | 0.678 |

## Table 3. Investor AI Adoption and Firm Valuation: Robustness Tests

This table reports results from robustness tests. Panel A presents a pre-trend analysis of firm-level AI adoption. The dependent variable is firm AI adoption, measured by the logarithm of the number of AI employees at the year-quarter and annual levels, and the independent variable is *Machine-based Ownership*. For the remaining panels, the dependent variables are Tobin's Q and Total Q, which measure firm market valuation. Tobin's Q is defined as the ratio of the market value of assets to the book value of assets, while Total Q incorporates intangible capital following Peters and Taylor (2017). The key independent variable is *Post* × *Machine-based Ownership*, which interacts an indicator for the post-ChatGPT period with firms' pre-period exposure to machine-based institutional investors. Panels B–F report robustness tests related to variable identification and specification. Panel B uses an alternative identification of AI-powered investors based on *GenAI Reliance* in the pre-period. *GenAI Reliance* is constructed following Sheng, Sun, Yang, and Zhang (2026) and measures the incremental explanatory power of ChatGPT-generated signals, derived from earnings conference call transcripts, for explaining institutional investors' portfolio holding changes beyond what is accounted for by standard firm fundamentals. Specifically, *GenAI Reliance* is estimated as the difference in R-squared between a regression of portfolio holding changes on both firm fundamentals and ChatGPT-generated signals and a regression on firm fundamentals alone, following a two-step procedure analogous to Kacperczyk and Seru (2007). We classify institutional investors into AI-powered investors if their GenAI Reliance is statistically significant (p-value <0.01). The key independent variable in Panel B is *Post* × *AI-powered Ownership (GenAI Reliance)*, which interacts the post-ChatGPT indicator with firms' pre-period ownership by GenAI Reliance-based AI-powered investors. Panel C uses an alternative identification of machine-based investors based on ADV reports, classifying investors as machine-based if their disclosures contain keywords associated with machine-based strategies (e.g., "artificial intelligence," "machine learning," and "large language model"). Panel D augments the baseline specification by including industry × year-quarter fixed effects to absorb time-varying industry-level shocks. Panel E reports results from propensity score matched difference-in-differences estimations. Panel F presents results excluding the "Magnificent Seven" firms (Apple, Microsoft, Alphabet, Amazon, NVIDIA, Meta, and Tesla), which may represent outliers due to their size, valuation dynamics, and ownership structure. Standard errors reported in parentheses are robust and clustered at the industry level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Pre-trend analysis of firm-level AI adoption

| | Firm AI Adoption | |
|---|---|---|
| | Year-Quarter | Annual |
| | (1) | (2) |
| Machine-based Ownership | -0.043 | -0.035 |
| | (0.434) | (0.403) |
| | | |
| Firm Controls in Table 2 except Firm AI Adoption | Yes | Yes |
| Time fixed effects | Yes | Yes |
| Observations | 42,953 | 14,104 |
| Adjusted R-squared | 0.246 | 0.241 |

Panel B. Alternative identification of AI-powered investors based on Gen AI Reliance

| | Tobin's Q | Total Q |
|---|---|---|
| | (1) | (2) |
| Post × AI-powered Ownership (Gen AI Reliance) | 0.552*** | 2.633*** |
| | (0.193) | (0.472) |
| | | |
| Firm Controls in Table 2 | Yes | Yes |
| Firm fixed effects | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes |
| Observations | 56,461 | 38,521 |
| Adjusted R-squared | 0.679 | 0.609 |

*Continued*

*Table 3 (Continued)*

Panel C. Alternative identification of AI-powered investors based on ADV

| | Tobin's Q | Total Q |
|---|---|---|
| | (1) | (2) |
| Post ×Machine-based Ownership (ADV) | 0.250* | 0.859*** |
| | (0.135) | (0.252) |
| | | |
| Firm Controls in Table 2 | Yes | Yes |
| Firm fixed effects | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes |
| Observations | 60,895 | 41,610 |
| Adjusted R-squared | 0.771 | 0.678 |

Panel D. Inclusion of Industry × Year-quarter fixed effects

| | Tobin's Q | Total Q |
|---|---|---|
| | (1) | (2) |
| Post × Machine-based Ownership | 0.633* | 1.951*** |
| | (0.356) | (0.613) |
| | | |
| Firm Controls in Table 2 | Yes | Yes |
| Firm fixed effects | Yes | Yes |
| Industry × Year-Quarter fixed effects | Yes | Yes |
| Observations | 58,598 | 39,963 |
| Adjusted R-squared | 0.774 | 0.690 |

Panel E. Propensity score matched DiD

| | Tobin's Q | Total Q |
|---|---|---|
| | (1) | (2) |
| Post × Machine-based Ownership | 0.987* | 2.962** |
| | (0.586) | (1.166) |
| | | |
| Firm Controls in Table 2 | Yes | Yes |
| Firm fixed effects | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes |
| Observations | 10,882 | 7,416 |
| Adjusted R-squared | 0.712 | 0.555 |

Panel F. Excluding the "Magnificent Seven" firms

| | Tobin's Q | Total Q |
|---|---|---|
| | (1) | (2) |
| Post × Machine-based Ownership | 0.593* | 2.141** |
| | (0.311) | (0.860) |
| | | |
| Firm Controls in Table 2 | Yes | Yes |
| Firm fixed effects | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes |
| Observations | 58513 | 39956 |
| Adjusted R-squared | 0.766 | 0.676 |

## Table 4. Investor AI Adoption and Firm Valuation: Cross-Sectional Analysis

This table reports results from cross-sectional subsample regressions examining how the valuation effects of AI-powered governance vary with firm complexity, information environment quality, and earnings news. The sample consists of 58,661 firm-year-quarter observations for 4,464 unique firms over the period 2019 Q3 to 2025 Q3, constructed from 13F institutional ownership data. The dependent variables are Tobin's Q and Total Q, which measure firm market valuation. Tobin's Q is defined as the ratio of market value of assets to book value of assets, while Total Q incorporates intangible capital following Peters and Taylor (2017). The key independent variable is Post × Machine-based Ownership, which interacts an indicator for the post-ChatGPT period with firms' pre-period exposure to machine-based institutional investors. Each panel splits the sample into two subgroups based on a firm- or industry-level characteristic measured in the pre-period. Panel A splits by inventory volatility, measured as the standard deviation of quarterly inventory over two years, where high (low) inventory volatility is defined as above (below) the sample median in the pre-period. Panel B splits by intangible asset intensity, measured as the total intangible assets divided by total assets, where high (low) intangibles is defined as above (below) the sample median in the pre-period. Panel C splits by analyst coverage, measured as the natural logarithm of the number of analysts covering the firm in the quarter, where low (high) coverage is defined as below (above) the sample median in the pre-period. Panel D splits by industry concentration, measured as the sum of market share of the top 8 firms in the industry; concentrated (competitive) industries are those above (below) the sample median. Panel E splits by earnings surprise, measured as the standardized unexpected earnings (SUE), and the sample is divided based on the sign of SUE. Control variables include firm size, return on assets, leverage, market-to-book ratio, cash holdings, asset tangibility, institutional ownership, and firm-level AI adoption. All regressions include firm fixed effects and year-quarter fixed effects, absorbing time-invariant firm characteristics and common macroeconomic shocks. Standard errors reported in parentheses are robust and clustered by industry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Inventory Volatility

| | Low Inventory Volatility | | High Inventory Volatility | |
|---|---|---|---|---|
| | Tobin's Q | Total Q | Tobin's Q | Total Q |
| | (1) | (2) | (3) | (4) |
| Post × Machine-based Ownership | 0.444 | 0.260 | 0.767** | 1.257** |
| | (0.960) | (0.997) | (0.379) | (0.578) |
| Firm controls | Yes | Yes | Yes | Yes |
| Firm fixed effects | Yes | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes | Yes |
| Observations | 25,113 | 16,289 | 32,850 | 23,201 |
| Adjusted R-squared | 0.737 | 0.667 | 0.776 | 0.647 |

Panel B. Intangible Asset Intensity

| | Low Intangible Assets | | High Intangible Assets | |
|---|---|---|---|---|
| | Tobin's Q | Total Q | Tobin's Q | Total Q |
| | (1) | (2) | (3) | (4) |
| Post × Machine-based Ownership | 0.088 | 0.983 | 0.519** | 1.240*** |
| | (0.623) | (1.343) | (0.207) | (0.380) |
| Firm controls | Yes | Yes | Yes | Yes |
| Firm fixed effects | Yes | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes | Yes |
| Observations | 25,215 | 16,800 | 33,111 | 22,904 |
| Adjusted R-squared | 0.764 | 0.676 | 0.824 | 0.767 |

*Continued*

*Table 4 (Continued)*

Panel C. Analyst Coverage

| | Low Analyst Coverage | | High Analyst Coverage | |
|---|---|---|---|---|
| | Tobin's Q | Total Q | Tobin's Q | Total Q |
| | (1) | (2) | (3) | (4) |
| Post × Machine-based Ownership | 1.253*** | 1.519** | 0.327 | 0.574 |
| | (0.276) | (0.696) | (0.325) | (0.683) |
| Firm controls | Yes | Yes | Yes | Yes |
| Firm fixed effects | Yes | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes | Yes |
| Observations | 17,934 | 12,129 | 22,802 | 15,777 |
| Adjusted R-squared | 0.741 | 0.707 | 0.809 | 0.728 |

Panel D. Industry Competition

| | High Competition | | Low Competition | |
|---|---|---|---|---|
| | Tobin's Q | Total Q | Tobin's Q | Total Q |
| | (1) | (2) | (3) | (4) |
| Post × Machine-based Ownership | 0.356 | 2.240** | 1.075*** | 1.929*** |
| | (0.508) | (0.926) | (0.358) | (0.663) |
| Firm controls | Yes | Yes | Yes | Yes |
| Firm fixed effects | Yes | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes | Yes |
| Observations | 29,879 | 20,352 | 28,298 | 19,172 |
| Adjusted R-squared | 0.749 | 0.644 | 0.796 | 0.725 |

Panel E. Earnings Surprise

| | Negative SUE | | Positive SUE | |
|---|---|---|---|---|
| | Tobin's Q | Total Q | Tobin's Q | Total Q |
| | (1) | (2) | (3) | (4) |
| Post × Machine-based Ownership | 0.879** | 1.565** | 0.213 | -0.052 |
| | (0.430) | (0.782) | (0.250) | (0.545) |
| Firm controls | Yes | Yes | Yes | Yes |
| Firm fixed effects | Yes | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes | Yes |
| Observations | 17,489 | 11,658 | 32,537 | 22,081 |
| Adjusted R-squared | 0.656 | 0.567 | 0.745 | 0.690 |

## Table 5. AI-Powered Governance and Contracting: CEO Turnover and Compensation

This table reports results from firm-level regressions examining how AI-powered investors affect executive discipline following the ChatGPT shock. The sample consists of 12,361 firm-quarter observations for 799 unique firms with CEO information from ExecuComp from 2019 Q3 to 2025 Q3. The unit of observation is a firm-quarter. Column (1) examines CEO turnover, where the dependent variable is an indicator equal to one if the CEO is replaced in a given quarter and zero otherwise. Column (2) examines CEO total compensation, defined as the total CEO compensation, including salary, bonus, other annual compensation, and restricted stock grants, measured as of the year-end quarter. The key independent variable is *Post* × *Machine-based Ownership* × *ROA*, which captures changes in the sensitivity of executive discipline to firm performance following the ChatGPT shock for firms with greater pre-period exposure to machine-based institutional investors. We additionally include all lower-order interaction terms and main effects of *Post*, *Machine-based Ownership*, and *ROA* to allow for flexible heterogeneity in performance sensitivity. $ROA_{(t-1,t-2)}$ is calculated as the average return on assets in the past two years. Control variables include firm size, return on assets (*ROA*), leverage, market-to-book ratio, cash holdings, asset tangibility, institutional ownership, firm-level AI adoption, CEO tenure, and CEO ownership. All regressions include firm fixed effects and year-quarter fixed effects, absorbing time-invariant firm characteristics and common macroeconomic shocks. Standard errors reported in parentheses are robust and clustered by industry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

| | CEO Turnover | CEO Total Compensation |
|---|---|---|
| | (1) | (2) |
| Post × Machine-based Ownership × $ROA_{(t-1,t-2)}$ | -2.231** | 0.059* |
| | (0.981) | (0.035) |
| $ROA_{(t-1,t-2)}$ | 0.278* | -0.001 |
| | (0.147) | (0.001) |
| Post ×$ROA_{(t-1,t-2)}$ | 0.277 | -0.007 |
| | (0.198) | (0.007) |
| Post × Machine-based Ownership | 0.277* | 0.000 |
| | (0.153) | (0.000) |
| Machine-based Ownership ×$ROA_{(t-1,t-2)}$ | -1.016 | -0.033** |
| | (0.748) | (0.016) |
| Firm AI Adoption | 0.025 | 0.001 |
| | (0.022) | (0.001) |
| Size | 0.012 | 0.005*** |
| | (0.018) | (0.001) |
| ROA | -0.330** | 0.001 |
| | (0.137) | (0.005) |
| Leverage | 0.051 | -0.005*** |
| | (0.067) | (0.002) |
| Market-to-Book | -0.003 | 0.002*** |
| | (0.003) | (0.000) |
| Cash | -0.036 | -0.003 |
| | (0.062) | (0.006) |
| Tangibility | -0.181 | 0.001 |
| | (0.131) | (0.003) |
| Institutional Ownership | -0.009 | 0.000 |
| | (0.031) | (0.000) |
| CEO Tenure | -0.023*** | 0.001*** |
| | (0.006) | (0.000) |
| CEO Ownership | -0.007 | 0.000 |
| | (0.005) | (0.000) |
| Firm fixed effects | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes |
| Observations | 12,361 | 12,332 |
| Adjusted R-squared | 0.230 | 0.633 |

## Table 6. AI-Powered Governance, Accounting Quality, and Financial Restatements

This table reports results from firm-level regressions examining the effects of AI-powered investors on accounting quality following the ChatGPT shock. The sample consists of 60,752 firm-quarter observations for 4,395 unique firms with financial restatement information obtained from Audit Analytics from 2019 Q3 to 2025 Q3. The unit of observation is a firm-quarter. Columns (1) and (2) examine the incidence of financial restatements. In column (1), the dependent variable is *Any Restatement*, an indicator equal to one if the firm issues any financial restatement in a given quarter and zero otherwise, estimated using a logit model. Column (2) examines *Restatement Correction*, which captures the magnitude or severity of restatements, estimated using an OLS specification. The key independent variable is *Post × Machine-based Ownership*, which measures whether firms with greater pre-period exposure to machine-based institutional investors experience changes in accounting quality following the ChatGPT shock. We additionally include the main effects of the post-period indicator and machine-based ownership to flexibly capture level effects. Control variables include firm size, return on assets (ROA), leverage, market-to-book ratio, cash holdings, asset tangibility, institutional ownership, and firm-level AI adoption. All regressions include firm fixed effects and year-quarter fixed effects, absorbing time-invariant firm characteristics and common macroeconomic shocks. Standard errors reported in parentheses are robust and clustered by industry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

| | Any Restatement<br>(1) | Restatement Correction<br>(2) |
|---|---|---|
| Post ×Machine-based Ownership | -1.414** | -1.463* |
| | (0.688) | (0.751) |
| Firm AI Adoption | -0.125 | -0.027 |
| | (0.092) | (0.045) |
| Size | 0.267*** | -0.243 |
| | (0.064) | (0.207) |
| ROA | -1.370*** | -0.230 |
| | (0.456) | (0.411) |
| Leverage | 0.237 | 0.337 |
| | (0.212) | (0.253) |
| Market-to-Book | -0.001 | -0.053* |
| | (0.027) | (0.030) |
| Cash | -0.285 | -0.359 |
| | (0.234) | (0.280) |
| Tangibility | -1.022*** | -1.915 |
| | (0.388) | (1.705) |
| Institutional Ownership | 0.178 | -0.098 |
| | (0.162) | (0.097) |
| | | |
| Sample | Full sample | Restatement sample |
| Model | Logit | OLS |
| Firm fixed effects | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes |
| Observations | 64,986 | 64,797 |
| Adjusted R-squared | | -0.013 |

## Table 7. AI-Powered Governance and Activism Intensity

This table reports results from firm-level regressions examining the effects of AI-powered investors on activism intensity following the ChatGPT shock. The sample consists of 68,105 firm-quarter observations for 4,226 unique firms with activism data obtained from FactSet SharkWatch from 2019 Q3 to 2025 Q3. In Panel A, we use the full sample, and in Panel B, we split the sample by activist ownership above (below) the sample median. The unit of observation is a firm-quarter. *Any Campaign* is an indicator variable equal to one if the firm is targeted by at least one activist campaign during the quarter, and zero otherwise. The data are sourced from FactSet SharkWatch. *Any Proxy Fight* is an indicator variable equal to one if at least one activist campaign against the firm results in a definitive proxy fight during the quarter, and zero otherwise. *Shareholder Proposals* is the total number of shareholder proposals filed against the firm during the quarter. The key independent variable is *Post × Machine-based Ownership*, which measures whether firms with greater pre-period exposure to machine-based institutional investors experience changes in activism intensity following the ChatGPT shock. We additionally include the main effects of the post-period indicator and machine-based ownership to flexibly capture level effects. Control variables include firm size, return on assets (ROA), leverage, market-to-book ratio, cash holdings, asset tangibility, institutional ownership, and firm-level AI adoption. All regressions include firm fixed effects and year-quarter fixed effects, absorbing time-invariant firm characteristics and common macroeconomic shocks. Standard errors reported in parentheses are robust and clustered by industry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Activism Intensity using Full Sample

| | Any Campaign | Any Proxy Fight | Shareholder Proposals |
|---|---|---|---|
| | (1) | (2) | (3) |
| Post ×Machine-based Ownership | 0.038*** | 0.037*** | 0.115*** |
| | (0.011) | (0.011) | (0.032) |
| Firm AI Adoption | 0.005* | 0.004 | 0.008 |
| | (0.003) | (0.003) | (0.006) |
| Size | 0.004** | 0.003* | -0.001 |
| | (0.002) | (0.002) | (0.002) |
| ROA | -0.005 | -0.002 | 0.011 |
| | (0.013) | (0.015) | (0.011) |
| Leverage | 0.005 | 0.003 | 0.004 |
| | (0.004) | (0.004) | (0.004) |
| Market-to-Book | 0.002*** | 0.002*** | 0.000 |
| | (0.001) | (0.001) | (0.000) |
| Cash | 0.011** | 0.010** | 0.005 |
| | (0.005) | (0.005) | (0.006) |
| Tangibility | -0.012* | -0.017** | -0.005 |
| | (0.007) | (0.006) | (0.008) |
| Institutional Ownership | -0.007 | -0.007* | -0.002 |
| | (0.005) | (0.004) | (0.009) |
| | | | |
| Firm fixed effects | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes |
| Observations | 68,105 | 68,105 | 68,105 |
| Adjusted R-squared | 0.090 | 0.089 | 0.143 |

*Continued*

*Table 7 (Continued)*

Panel B. Subsample by Activist Ownership

| | High Activist Share Sample | | | Low Activist Share Sample | | |
|---|---|---|---|---|---|---|
| | Any Campaign (1) | Any Proxy Fight (2) | Shareholder Proposal (3) | Any Campaign (4) | Any Proxy Fight (5) | Shareholder Proposal (6) |
| Post × Machine-based Ownership | 0.060** | 0.010** | 0.139*** | 0.019 | 0.000 | 0.012 |
| | (0.025) | (0.004) | (0.044) | (0.016) | (0.004) | (0.014) |
| Size | 0.001 | 0.001* | 0.003 | 0.001 | 0.001** | -0.006* |
| | (0.006) | (0.001) | (0.008) | (0.004) | (0.000) | (0.003) |
| ROA | 0.011 | 0.002 | 0.051 | -0.018 | -0.004 | 0.027 |
| | (0.043) | (0.006) | (0.054) | (0.022) | (0.004) | (0.021) |
| Leverage | 0.030*** | 0.000 | -0.005 | 0.004 | 0.005** | 0.000 |
| | (0.011) | (0.002) | (0.012) | (0.009) | (0.002) | (0.003) |
| Market-to-Book | -0.000 | -0.000* | 0.001 | 0.003*** | 0.000 | -0.000 |
| | (0.001) | (0.000) | (0.001) | (0.001) | (0.000) | (0.000) |
| Cash | 0.006 | -0.002 | 0.012 | 0.010 | 0.001 | 0.016 |
| | (0.013) | (0.003) | (0.014) | (0.009) | (0.002) | (0.013) |
| Tangibility | -0.025 | 0.002 | 0.026 | -0.020 | -0.002 | 0.001 |
| | (0.030) | (0.006) | (0.038) | (0.015) | (0.002) | (0.014) |
| Firm AI Adoption | 0.006 | 0.001 | 0.015* | -0.002 | -0.000 | 0.000 |
| | (0.005) | (0.001) | (0.009) | (0.003) | (0.000) | (0.002) |
| Institutional Ownership | 0.008 | 0.002 | -0.014 | -0.006 | -0.000 | 0.014 |
| | (0.011) | (0.002) | (0.017) | (0.008) | (0.002) | (0.019) |
| Firm fixed effects | Yes | Yes | Yes | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 34,404 | 34,404 | 34,404 | 33,308 | 33,308 | 33,308 |
| Adjusted R-squared | 0.083 | 0.032 | 0.110 | 0.083 | 0.032 | 0.110 |

## Table 8. AI-Powered Governance and M&A Activity

This table reports results from firm-level regressions examining how AI-powered governance affects firms' merger and acquisition (M&A) activity following the ChatGPT shock. The sample consists of 68,270 firm-quarter observations for 4,226 unique firms with M&A information obtained from SDC Platinum from 2019 Q3 to 2025 Q3. The unit of observation is a firm-quarter. Column (1) examines the likelihood of engaging in any M&A activity, estimated using a logit specification. The dependent variable in column (1) is *Any M&A Deal*, an indicator equal to one if the firm completes at least one M&A transaction in a given quarter and zero otherwise. Column (2) focuses on value-destroying acquisitions, estimated using a logit specification. The dependent variable is *Any Value-Destroying Deal*, an indicator equal to one if the firm completes an acquisition associated with negative announcement returns and zero otherwise. Column (3) examines the *Percentage of Value-Destroying Deals*, defined as the fraction of a firm's M&A transactions in a quarter that are value-destroying, estimated using an OLS specification. Column (4) examines market reactions, where the dependent variable is the cumulative abnormal return (*CAR*) over the $[-1, +1]$ event window around M&A announcements, estimated using an OLS specification and restricted to the sample of M&A events. The key independent variable is *Post × Machine-based Ownership*, which captures whether firms with greater pre-period exposure to machine-based institutional investors adjust their M&A behavior following the ChatGPT shock. Control variables include firm size, return on assets (ROA), leverage, market-to-book ratio, cash holdings, asset tangibility, and firm-level AI adoption. All specifications include firm fixed effects and year-quarter fixed effects. Standard errors reported in parentheses are robust and clustered by industry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

| | Any M&A deal | Any Value-destroying Deal | Percentage of Value-destroying Deals | CAR [-1, +1] |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Post × Machine-based Ownership | -1.149*** | -1.801*** | -0.063*** | 0.087*** |
| | (0.331) | (0.473) | (0.020) | (0.026) |
| Firm AI Adoption | -0.029 | 0.030 | 0.004 | 0.004 |
| | (0.061) | (0.088) | (0.003) | (0.004) |
| Size | -0.061 | 0.109 | -0.000 | -0.026*** |
| | (0.056) | (0.084) | (0.000) | (0.009) |
| ROA | 0.798* | 0.470 | 0.003 | 0.093 |
| | (0.441) | (0.691) | (0.009) | (0.080) |
| Leverage | -2.077*** | -2.669*** | -0.053*** | 0.061** |
| | (0.176) | (0.271) | (0.007) | (0.027) |
| Market-to-Book | 0.006 | 0.039** | 0.001 | -0.006*** |
| | (0.012) | (0.019) | (0.001) | (0.002) |
| Cash | 1.726*** | 1.049*** | 0.018** | 0.063*** |
| | (0.195) | (0.297) | (0.007) | (0.021) |
| Tangibility | -0.290 | -1.551*** | -0.010 | 0.144*** |
| | (0.353) | (0.540) | (0.012) | (0.038) |
| Institutional Ownership | 0.259** | 0.432*** | 0.011** | 0.004 |
| | (0.112) | (0.163) | (0.005) | (0.009) |
| Sample | Full | Full | Full | M&A events |
| Model | Logit | Logit | OLS | OLS |
| Firm fixed effects | Yes | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes | Yes |
| Observations | 36,879 | 23,809 | 68,105 | 4,688 |
| Adjusted R-squared | | | 0.065 | 0.106 |

## Table 9. AI-Powered Governance and Innovation Outcomes

This table reports results from firm-level regressions examining how AI-powered governance affects corporate investment and innovation outcomes following the ChatGPT shock. In Panel A, we examine the effect on corporate long-term investment and innovation quantity, and the sample consists of 68,073 firm-quarter observations for 4,224 unique firms over the period 2019 Q3 to 2025 Q3. Columns (1) and (2) examine firms' real investment decisions. *Long-term Investment* is the firm-level long-term investment scaled by total assets. *R&D Intensity* is measured as research and development expenditures scaled by total assets. *Log (Number of Patents)* is defined as the logarithm of one plus the number of patents filed by the firm in a given quarter. In Panel B, we examine the effect on innovation quality, and the sample consists of 5,794 firm-quarter observations for 760 unique firms with patent-level information aggregated at the firm level. The dependent variable is cumulative future average patent value per patent over 4-, 8-, and 12-quarter windows. Nominal patent value is measured using nominal Kogan et al. (2017) patent value, while real patent value is inflation-adjusted. The key independent variable is *Post × Machine-based Ownership*, which captures changes in investment and innovation behavior for firms with greater pre-period exposure to machine-based institutional investors following the ChatGPT shock. Control variables include firm size, return on assets (ROA), leverage, market-to-book ratio, cash holdings, asset tangibility, institutional ownership, and firm-level AI adoption. In Panel B, we additionally control for *Log (Number of Patents)*, *Log (R&D expenditure)*, and indicators for missing R&D data to separate changes in innovation quality from changes in innovation quantity and reporting. All regressions include firm fixed effects and year-quarter fixed effects, absorbing time-invariant firm characteristics and common macroeconomic shocks. Standard errors reported in parentheses are robust and clustered by industry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Investment and Innovation Quantity

| | Long-term Investment | R&D | Log (Number of Patents) |
|---|---|---|---|
| | (1) | (2) | (3) |
| Post × Machine-based Ownership | -0.025*** | -0.013* | -1.405*** |
| | (0.009) | (0.007) | (0.284) |
| Firm AI Adoption | 0.001 | 0.001* | -0.017 |
| | (0.002) | (0.001) | (0.027) |
| Size | -0.004** | -0.008*** | 0.025** |
| | (0.002) | (0.002) | (0.010) |
| ROA | 0.021** | -0.143*** | -0.009 |
| | (0.010) | (0.043) | (0.046) |
| Leverage | -0.000 | -0.006* | -0.002 |
| | (0.000) | (0.003) | (0.043) |
| Market-to-Book | -0.001*** | -0.001 | 0.003 |
| | (0.000) | (0.001) | (0.002) |
| Cash | -0.083*** | -0.001 | 0.091*** |
| | (0.014) | (0.003) | (0.026) |
| Tangibility | -0.147*** | 0.012*** | 0.104*** |
| | (0.034) | (0.004) | (0.036) |
| Institutional Ownership | 0.004 | 0.001 | -0.079*** |
| | (0.003) | (0.002) | (0.021) |
| Observations | 68,073 | 68,073 | 65,301 |
| Adjusted R-squared | 0.922 | 0.783 | 0.716 |
| Firm fixed effects | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes |

*Continued*

*Table 9 (Continued)*

Panel B. Average Patent Value per Patent

| | Nominal Patent Value | | | Real Patent Value | | |
|---|---|---|---|---|---|---|
| | 4-quarter | 8-quarter | 12-quarter | 4-quarter | 8-quarter | 12-quarter |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Post × Machine-based Ownership | 1.838*** | 1.618*** | 1.532*** | 1.416*** | 1.272*** | 1.215*** |
| | (0.392) | (0.319) | (0.285) | (0.407) | (0.340) | (0.312) |
| Firm AI Adoption | 0.082 | 0.121** | 0.133*** | 0.074 | 0.108** | 0.116*** |
| | (0.073) | (0.052) | (0.049) | (0.055) | (0.041) | (0.039) |
| Size | -0.139*** | -0.080*** | -0.070*** | -0.113*** | -0.061*** | -0.053*** |
| | (0.035) | (0.026) | (0.023) | (0.030) | (0.023) | (0.019) |
| ROA | 0.191 | 0.100 | 0.108 | 0.203 | 0.111 | 0.105* |
| | (0.143) | (0.087) | (0.067) | (0.130) | (0.076) | (0.061) |
| Leverage | 0.258* | 0.157 | 0.128 | 0.204* | 0.132 | 0.105 |
| | (0.132) | (0.123) | (0.096) | (0.108) | (0.092) | (0.077) |
| Market-to-Book | 0.001 | 0.002 | 0.003 | -0.001 | 0.001 | 0.001 |
| | (0.005) | (0.005) | (0.004) | (0.004) | (0.004) | (0.003) |
| Cash | -0.256*** | -0.167** | -0.117* | -0.209** | -0.147** | -0.103* |
| | (0.095) | (0.073) | (0.061) | (0.080) | (0.067) | (0.055) |
| Tangibility | -0.242 | -0.300* | -0.224 | -0.185 | -0.211 | -0.162 |
| | (0.241) | (0.177) | (0.142) | (0.202) | (0.149) | (0.125) |
| Log(Number of Patents) | -0.208*** | -0.174*** | -0.158*** | -0.160*** | -0.134*** | -0.121*** |
| | (0.025) | (0.019) | (0.017) | (0.023) | (0.017) | (0.015) |
| Log(R&D expenditure) | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| | (0.000) | (0.000) | (0.000) | (0.000) | (0.000) | (0.000) |
| R&D missing | 0.011 | 0.045 | 0.045 | 0.015 | 0.043 | 0.041 |
| | (0.087) | (0.066) | (0.062) | (0.086) | (0.063) | (0.060) |
| Observations | 8,072 | 8,805 | 9,003 | 8,072 | 8,805 | 9,003 |
| Adjusted R-squared | 0.921 | 0.947 | 0.959 | 0.913 | 0.942 | 0.954 |
| Firm fixed effects | Yes | Yes | Yes | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes | Yes | Yes | Yes |

## Table 10. Investor AI Adoption and Machine-based Investor Portfolio Change

This table reports results from investor-level regressions examining how AI-powered investors adjust portfolio structure, investment horizon, and performance following the ChatGPT shock. The sample consists of 46,742 investor-year-quarter observations for 2,529 unique investors, constructed from 13F filings over the period 2019Q3 to 2025Q3. The unit of observation is an institutional investor-quarter. Panel A examines portfolio breadth, where the dependent variables in column (1), *New Industries*, is defined as the number of industries newly entered by the investor, and the dependent variables in columns (2), *Industry-level HHI*, is defined as the sum of squared portfolio weights across industries (2-digit SIC). In Panel B, columns (1)-(3) examine investment horizon and trading activity, where the dependent variables are *Average Holding Duration*, defined as the average number of quarters a stock is held, *Portfolio Turnover (Value)*, defined as the sum of absolute changes in portfolio weights across holdings between two quarters scaled by total portfolio weight, and *Portfolio Turnover (Count)*, defined as the sum of absolute changes in share counts across holdings scaled by total shares held. In Panel C, columns (1) and (2) examine investor performance and allocation, where the dependent variables are *Alpha (FF5)*, defined as risk-adjusted abnormal return estimated using the Fama-French five-factor model, and *ETF Share*, defined as the proportion of portfolio market value invested in ETFs. Panel D examines the extent to which machine-based investors rely on generative AI tools following the ChatGPT shock. The dependent variable, *GenAI Reliance*, is constructed following Sheng, Sun, Yang, and Zhang (2026) and measures the incremental explanatory power of ChatGPT-generated signals, derived from earnings conference call transcripts, for explaining institutional investors' portfolio holding changes beyond what is accounted for by standard firm fundamentals. Specifically, *GenAI Reliance* is estimated as the difference in R-squared between a regression of portfolio holding changes on both firm fundamentals and ChatGPT-generated signals and a regression on firm fundamentals alone, following a two-step procedure analogous to Kacperczyk and Seru (2007). The key independent variable is *Post × Machine-based Investor*, an interaction between an indicator for the post-ChatGPT period and an indicator for investors classified as machine-based prior to the shock. Control variables include the log of total portfolio value, the log of the number of shares held, the log of portfolio concentration, return volatility, earnings growth share, long-term holdings share, and past portfolio return. All regressions include investor fixed effects and year-quarter fixed effects, absorbing time-invariant investor characteristics and common macroeconomic shocks. Standard errors reported in parentheses are robust. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Portfolio Breadth

| | New-Industries | Industry-HHI |
|---|---|---|
| | (1) | (2) |
| Post × Machine-investor | 0.207** | -0.009* |
| | (0.090) | (0.005) |
| Log(Total Portfolio Value) | 0.015 | -0.002 |
| | (0.037) | (0.004) |
| Log(Total Shares) | 0.835*** | -0.061*** |
| | (0.075) | (0.007) |
| Log(Portfolio Concentration) | -0.021 | 0.004 |
| | (0.036) | (0.004) |
| Return Volatility | -4.218 | -3.020*** |
| | (3.221) | (0.375) |
| Earnings Growth Share | -0.423** | -0.133*** |
| | (0.195) | (0.014) |
| Long-Term Holdings Share | -1.320*** | -0.010*** |
| | (0.058) | (0.004) |
| Past Portfolio Return | 0.426** | -0.072*** |
| | (0.189) | (0.011) |
| | | |
| Investor fixed effects | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes |
| Observations | 46,742 | 46,742 |
| Adjusted R-squared | 0.356 | 0.926 |

*Continued*

*Table 10 (Continued)*

Panel B. Investment Horizon and Portfolio Turnover

| | Average Holding Duration | Portfolio Turnover (Value) | Portfolio Turnover (Count) |
|---|---|---|---|
| | (1) | (2) | (3) |
| Post × Machine-investor | -0.761*** | 0.030** | 0.019*** |
| | (0.278) | (0.012) | (0.007) |
| Log(Total Portfolio Value) | 0.565*** | 0.011** | -0.005 |
| | (0.148) | (0.005) | (0.003) |
| Log(Total Shares) | -1.861*** | -0.008 | 0.050*** |
| | (0.197) | (0.006) | (0.005) |
| Log(Portfolio Concentration) | -0.561*** | 0.015*** | 0.003 |
| | (0.148) | (0.005) | (0.003) |
| Return Volatility | -4.700 | 1.477*** | -0.029 |
| | (5.759) | (0.273) | (0.267) |
| Earnings Growth Share | 1.310*** | -0.012 | 0.018 |
| | (0.405) | (0.017) | (0.016) |
| Long-Term Holdings Share | -4.255*** | -0.033*** | -0.105*** |
| | (0.124) | (0.004) | (0.004) |
| Past Portfolio Return | 0.888*** | 0.024 | -0.060*** |
| | (0.243) | (0.016) | (0.014) |
| | | | |
| Investor fixed effects | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes |
| Observations | 46,742 | 46,742 | 46,742 |
| Adjusted R-squared | 0.705 | 0.566 | 0.566 |

Panel C. Investor Performance and Portfolio Allocation

| | Alpha (FF5) | ETF Share |
|---|---|---|
| | (1) | (2) |
| Post × Machine-investor | 0.003** | -0.011** |
| | (0.001) | (0.005) |
| Log(Total Portfolio Value) | 0.000 | -0.004 |
| | (0.001) | (0.003) |
| Log(Total Shares) | 0.000 | -0.014** |
| | (0.001) | (0.006) |
| Log(Portfolio Concentration) | -0.000 | 0.006** |
| | (0.001) | (0.003) |
| Return Volatility | -0.414*** | -4.518*** |
| | (0.131) | (0.366) |
| Earnings Growth Share | 0.026*** | -0.148*** |
| | (0.005) | (0.013) |
| Long-Term Holdings Share | -0.001 | -0.014*** |
| | (0.001) | (0.004) |
| Past Portfolio Return | 0.126*** | -0.014* |
| | (0.007) | (0.008) |
| | | |
| Investor fixed effects | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes |
| Observations | 46,713 | 46,742 |
| Adjusted R-squared | 0.524 | 0.955 |

*Continued*

*Table 10 (Continued)*

Panel D. Generative AI Reliance

| | GenAI Reliance | |
|---|---|---|
| | Full Sample | Post GPT Shock |
| | (1) | (2) |
| Post × Machine-investor | 0.011*** | |
| | (0.004) | |
| Machine-investor | | 0.041*** |
| | | (0.009) |
| Log(Total Portfolio Value) | 0.000 | -0.035*** |
| | (0.003) | (0.007) |
| Log(Total Shares) | -0.096*** | -0.103*** |
| | (0.006) | (0.006) |
| Log(Portfolio Concentration) | 0.000 | 0.038*** |
| | (0.003) | (0.006) |
| Return Volatility | -0.449 | 2.870*** |
| | (0.535) | (0.647) |
| Earnings Growth Share | -0.026 | -0.415*** |
| | (0.025) | (0.066) |
| Long-Term Holdings Share | -0.012*** | -0.094*** |
| | (0.004) | (0.010) |
| Past Portfolio Return | 0.016 | 0.052 |
| | (0.029) | (0.075) |
| | | |
| Investor fixed effects | Yes | No |
| Year-Quarter fixed effects | Yes | Yes |
| Observations | 33,089 | 11,727 |
| Adjusted R-squared | 0.637 | 0.507 |

**Table 11. Investor AI Adoption and Firm-level Machine-based Investor Ownership**

This table reports results from firm-level regressions examining how machine-based investors' institutional ownership and portfolio patterns respond to the ChatGPT shock. The sample consists of 68,325 firm-year-quarter observations for 5,063 unique firms for the sample period from 2019 Q3 to 2025 Q3, constructed from 13F institutional ownership data. The unit of observation is a firm-quarter. In column (1), the dependent variable is *Total Number of Machine-based Investors* holding the firm. In column (2), the dependent variable is *Average Ownership per Machine-based Investor*, defined as total machine-based ownership divided by the number of machine-based investors. In columns (3) and (4), the dependent variable is *Selling (Buying) Dispersion*, defined as the cross-sectional variance across machine-based investors of an indicator equal to one if the investor executes at least one sell (buy) transaction in a portfolio firm's stock during the quarter and zero otherwise. The key independent variable is *Post × Machine-based Ownership*, which interacts an indicator for the post-ChatGPT period with firms' pre-period exposure to machine-based investors. Control variables include firm size, return on assets (ROA), leverage, market-to-book ratio, cash holdings, asset tangibility, institutional ownership, and firm-level AI adoption. All regressions include firm fixed effects and year-quarter fixed effects. Standard errors reported in parentheses are robust and clustered by industry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

| | Total Number | Average Ownership | Selling Dispersion | Buying Dispersion |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Post × Machine-based Ownership | 37.184*** | -0.245* | -0.011** | -0.020*** |
| | (6.499) | (0.138) | (0.005) | (0.005) |
| Firm AI Adoption | 2.026* | 0.011** | -0.001 | 0.001 |
| | (1.209) | (0.004) | (0.001) | (0.001) |
| Size | 13.555*** | -0.048** | -0.000 | 0.004** |
| | (1.544) | (0.021) | (0.001) | (0.001) |
| ROA | 5.392 | 0.044 | 0.007 | 0.016 |
| | (8.075) | (0.048) | (0.006) | (0.010) |
| Leverage | -32.058*** | 0.154** | -0.003 | -0.023*** |
| | (3.016) | (0.072) | (0.003) | (0.005) |
| Market-to-Book | 2.342*** | 0.001 | -0.000** | 0.001*** |
| | (0.142) | (0.003) | (0.000) | (0.000) |
| Cash | -6.150*** | 0.006 | 0.004 | 0.002 |
| | (1.401) | (0.026) | (0.003) | (0.003) |
| Tangibility | -3.311 | 0.001 | 0.012** | -0.004 |
| | (2.782) | (0.049) | (0.005) | (0.007) |
| Institutional Ownership | 58.767*** | 0.306** | 0.024*** | 0.031*** |
| | (4.991) | (0.122) | (0.005) | (0.004) |
| | | | | |
| Firm fixed effects | Yes | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes | Yes |
| Observations | 68,325 | 67,719 | 67,501 | 67,501 |
| Adjusted R-squared | 0.972 | 0.588 | 0.138 | 0.190 |

## Table 12. Investor AI Adoption and Investor Coalition

This table reports results from firm-level regressions examining how activist ownership and coalitions of activists respond to the ChatGPT shock. The full sample consists of 68,325 firm-year-quarter observations for 5,063 unique firms over the period from 2019 Q3 to 2025 Q3. The unit of observation is a firm-quarter. In columns (1)-(3), we examine the activist participation dynamics. *Experienced Activists* is defined as the number of investors with any prior history of participation in activist campaigns. *Activist Ownership* is defined as the fraction of shares outstanding held by institutional investors classified as an activist. *New Activists* is defined as the number of newly entering investors that initiate an activist campaign in a given firm-quarter. In column (4)-(5), we examine activist campaign outcomes. *Settled Campaign* is defined as the number of activist campaigns associated with a firm that reach a settlement within a given quarter. *Board Seat Won* is defined as the number of board seats secured by activists at the firm in a given quarter. The key independent variable is *Post × Machine-based Ownership*, which measures whether firms with greater pre-period exposure to machine-based institutional investors experience changes in activism intensity following the ChatGPT shock. We additionally include the main effects of the post-period indicator and machine-based ownership to flexibly capture level effects. Control variables include firm size, return on assets (ROA), leverage, market-to-book ratio, cash holdings, asset tangibility, institutional ownership, and firm-level AI adoption. All regressions include firm fixed effects and year-quarter fixed effects, absorbing time-invariant firm characteristics and common macroeconomic shocks. Standard errors reported in parentheses are robust and clustered by industry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

| | Activist Participation | | | Activist Campaign Outcomes | |
|---|---|---|---|---|---|
| | Experienced Activists | Activist Ownership | New Activists | Settled Campaign | Board Seat Won |
| | (1) | (2) | (3) | (4) | (5) |
| Post ×Machine-based Ownership | 3.240*** | 0.012*** | 0.004* | 0.009*** | 0.022* |
| | (0.438) | (0.004) | (0.002) | (0.003) | (0.012) |
| Firm AI Adoption | 0.670*** | 0.001 | -0.000 | 0.000 | 0.002 |
| | (0.107) | (0.001) | (0.001) | (0.001) | (0.002) |
| Size | -0.267 | 0.002*** | 0.004 | 0.002** | 0.009*** |
| | (0.279) | (0.001) | (0.005) | (0.003) | (0.012) |
| ROA | -1.075*** | -0.003 | 0.008*** | -0.004 | -0.026** |
| | (0.156) | (0.002) | (0.002) | (0.003) | (0.012) |
| Leverage | 0.120*** | 0.007** | -0.000 | 0.005 | 0.019** |
| | (0.009) | (0.003) | (0.000) | (0.003) | (0.007) |
| Market-to-Book | -0.513*** | -0.000 | 0.003 | -0.000 | -0.000 |
| | (0.120) | (0.000) | (0.002) | (0.000) | (0.000) |
| Cash | -0.290 | 0.002 | -0.001 | 0.003 | 0.001 |
| | (0.248) | (0.002) | (0.002) | (0.002) | (0.007) |
| Tangibility | 0.380*** | 0.003 | -0.001*** | 0.003 | 0.018 |
| | (0.131) | (0.004) | (0.000) | (0.004) | (0.013) |
| Institutional Ownership | 2.493*** | 0.032*** | 0.003 | -0.000 | 0.001 |
| | (0.216) | (0.003) | (0.002) | (0.001) | (0.005) |
| Firm fixed effects | Yes | Yes | Yes | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes | Yes | Yes | Yes |
| Observations | 68,325 | 68,325 | 68,325 | 68,325 | 68,325 |
| Adjusted R-squared | 0.885 | 0.579 | 0.009 | 0.004 | 0.000 |

*Internet Appendix for*

# "AI-Powered Corporate Governance"

Jiung Lee, Seungjoon Oh, Joanna Wang

Section A. Generative AI Adoption by Institutional Investors

## Section A. Generative AI Adoption by Institutional Investors

To understand how fund managers actually employ GenAI in information processing, we draw on documented practice at major institutions as reported by credible industry sources, such as the Alternative Investment Management Association (AIMA), proprietary research from Resonanz Capital, and public statements from BlackRock. We organize these practices into seven distinct information-processing channels, each reflecting a category of unstructured text for which LLMs offer a structurally novel processing advantage over both human analysts and conventional NLP tools. The seven channels described share a common structural feature: each involves reducing the cost of processing a category of unstructured, text-based information that was previously accessible through costly process. This observation suggests a unified mechanism through which GenAI may reduce the marginal cost of information processing, allowing fund managers to cover wider securities universes, conduct deeper due diligence within their existing coverage, or identify signals in information types (e.g., legal covenants, central bank language, ESG disclosures) that were previously costly to systematically analyze.

### Earnings Call Transcript Parsing

Earnings calls are among the richest and most time-sensitive sources of forward-looking management information, yet their unstructured, conversational format makes systematic comparison across companies and periods labor-intensive. LLMs are particularly well-suited to this task because they can parse the discourse-level properties of speech (e.g., hedging, evasion, tonal shifts, deviations from scripted language) that carry incremental information beyond the numerical announcements themselves.

BlackRock has developed a proprietary LLM trained on “more than 400,000 earnings call transcripts covering over 17,000 public firms, combined with two decades of historical market data.” The model generates forecasts of market reactions following earnings announcements and, in BlackRock’s own comparative testing, outperforms general-purpose LLMs on financial prediction tasks. BlackRock describes its broader systematic process as synthesizing information from “analyst reports, corporate earnings call transcripts, news articles, and social media to help inform investment forecasts and uncover potential alpha opportunities” (BlackRock, 2024).

At Point72, a partnership with an external AI platform enables the firm to “process earnings calls in real-time, extracting sentiment and identifying linguistic patterns that human analysts might miss” (Resonanz Capital, 2024). This real-time pipeline is qualitatively different from legacy sentiment scorers, which apply fixed dictionaries of positive and negative words to call text. LLMs can instead capture contextual meaning: a CFO’s unusual hesitation when discussing inventory, a departure from the standard guidance script, or management’s implicit acknowledgement of competitive pressure without stating it directly.

Balyasny Asset Management has built a system, BAMChatGPT, that ingests ten distinct data streams, including “transcripts, sales and sell-side commentaries, and broker research, with the goal of creating bots that can proactively push relevant information to portfolio managers and business teams." The system employs specialized financial embeddings that achieve "60 percent accuracy in internal testing compared to OpenAI's sub-40 percent performance for financial document retrieval," reflecting the practical value of domain-specific fine-tuning (Resonanz Capital, 2024).

### Regulatory Filing Analysis (10-K, 10-Q, 8-K)

Annual and quarterly SEC filings contain the most legally consequential disclosures that a public company makes. Identifying material changes in risk factor language, management discussion tone, or footnote accounting policy across successive filings is valuable but manually burdensome at scale. LLMs can perform cross-period comparisons automatically, flagging linguistically significant deviations for portfolio manager review.

Balyasny operates a dedicated micro-agent that "flags 10-K wording changes" in SEC filings as they are released. According to Resonanz Capital's account, a process that previously required two days from a senior analyst is reduced to thirty minutes. The same system generates morning notes and pushes automated notifications on "breaking-news moves, discrepancies in filings, ESG controversies, and other unknown-unknowns" (Resonanz Capital, 2024).

An AIMA survey of 157 hedge fund managers found that 67 percent of respondents were using AI for investment idea generation and 58 percent for portfolio construction, suggesting that filing analysis feeds directly into position-level decisions rather than remaining a background screening step (AIMA, 2024).

### Credit and Legal Document Processing

Bond indentures, loan agreements, and offering prospectuses are among the densest and most technically demanding texts that credit analysts encounter. Covenant packages, restricted payment baskets, and change-of-control provisions can run to hundreds of pages, and their economic implications are non-trivially encoded in legal boilerplate. LLMs fine-tuned on legal and financial corpora can extract economically relevant clause-level summaries reliably, covering an entire market segment with a lean team.

Resonanz Capital documents a credit-focused hedge fund that used GenAI "to ingest and summarize thousands of bond prospectuses and indentures, allowing their analysts to cover more securities with greater speed and consistency, even with a lean team." This application illustrates a structural advantage of LLMs over legacy search tools: unlike keyword-based retrieval, a language model can recognize that a clause is economically equivalent to a known covenant form even when drafted in unfamiliar language (Resonanz Capital, 2024).

Balyasny's Deep Research bot, which "combs five million documents and answers PM questions in minutes," is designed to serve portfolio managers directly—not to produce analyst deliverables for later review, but to insert structured answers into the PM's decision workflow in near-real time.

### Cross-Document Thematic Extraction

A structurally distinctive capability of LLMs relative to both human analysts and legacy NLP pipelines is the ability to answer a structured question simultaneously across hundreds of documents, returning results in a consistent, comparable format. This enables research at scale: rather than covering a handful of securities in depth, a fund manager can obtain a standardized view of a specific KPI, theme, or risk factor across an entire sector.

BlackRock's "Thematic Robot" illustrates this channel. The tool "blends the power of LLMs with proprietary data to build long/short or long-only equity baskets" by using language models combined with conference-call text analysis to "uncover investment exposures across multiple companies around emerging market themes like GLP-1 pharmaceuticals." Portfolio manager expertise is, in BlackRock's own description, "engrained in every part of this process—from defining the specific theme or scenario, customizing the analysis with human determined priors." (BlackRock, 2024)

Two Sigma articulates the structural shift this creates for the research process: "The research funnel is inverting: Large language models are widening the top, shifting the bottleneck from 'we need more ideas' to 'we need to evaluate ideas faster.'" The firm further notes that "models are learning to read, joining text and structured data into unified representations. The economics of data are being rewritten through automated document understanding at near-zero marginal cost." (Two Sigma, 2026)

**Internal Research Retrieval via Retrieval-Augmented Generation (RAG)**

A substantial fraction of the analytical capital in an asset management firm exists as tacit, unstructured knowledge: prior research notes, analyst memos, model documentation, and internal commentary accumulated over years. Retrieval-Augmented Generation (RAG) systems allow fund managers to query this proprietary archive in natural language, surfacing relevant prior work without requiring the user to know exactly where or how it was documented.

Resonanz Capital describes this channel as enabling investment teams to "search and query internal research archives more intuitively—surfacing insights that would otherwise remain buried in PDFs or internal wikis." Balyasny's BAMChatGPT system is explicitly designed so that bots "proactively push relevant information to portfolio managers and business teams," inverting the conventional model in which the PM initiates a research request and waits for analyst delivery (Resonanz Capital, 2024).

D.E. Shaw's infrastructure illustrates the governance architecture required to make RAG deployments both effective and compliant. The firm's LLM Gateway "brokers calls to two dozen external models only after stripping personally identifiable information and logging queries," while its DocLab system "tags confidence scores and audit hashes with every retrieval." This combination of data sanitization, query logging, and source attribution is the institutional response to concerns about model hallucination and regulatory accountability (Resonanz Capital, 2024).

Two Sigma frames the expectation at the firm level: "Everyone at Two Sigma is expected to use LLMs to accelerate their work and consider workflow improvements for their teams." The firm elaborates: "It's like each of us getting a team of highly capable, tireless assistants that can process a vast amount of information, synthesize insights, and automate mundane tasks." (Two Sigma, 2026)

### Hypothesis Generation and Research Automation

Beyond processing existing documents, LLMs are increasingly deployed to generate new investment hypotheses—scanning research papers, financial databases, and historical market observations to propose relationships or strategies that human researchers then evaluate. This channel is most advanced at quantitatively oriented managers, where the hypothesis-to-backtest pipeline can itself be automated.

Man Group's AlphaGPT system (developed in partnership with Anthropic and powered by Claude) functions as what the firm describes as a "digital three-person research team that never sleeps." Its three-stage architecture includes an idea generator that "generates testable propositions…explores subtle relationships between seemingly unrelated market factors;" a coder that "writes production-grade Python code leveraging Man Group's research tools and is capable of interacting with the proprietary databases;" and an evaluator that "rigorously evaluates whether the research works…applies the same strict criteria that human-generated research must pass." As of 2025, AlphaGPT has produced several dozen investment signals approved for live trading (Man Group, 2025).

Man Group is explicit that the system "doesn't replace human judgment but amplifies it—Humans provide strategic direction, market context, and final decision-making, while AlphaGPT handles the heavy lifting of data processing, hypothesis generation, and initial analysis." Through its Anthropic partnership, the firm states that "investment teams can analyze a significantly wider universe of securities and their associated financial risk models," with Claude "accessing data through code execution, then summarizing and synthesizing insights for use by portfolio management teams globally" (Man Group–Anthropic Partnership, 2025).

Two Sigma has identified a countervailing disciplinary risk: because LLMs can generate investment hypotheses rapidly, they also raise the risk of overfitting. The firm observes that "with AI agents, researchers can easily generate a large number of hypotheses and backtest them, which can exacerbate the overfitting issue," a challenge the firm views as becoming "more difficult with powerful AI" (Two Sigma, 2026). This observation points to a genuine tension at the frontier of quantitative research: the same tools that lower the cost of idea generation may also lower the cost of generating spurious ideas.

### Macroeconomic and Policy Document Interpretation

Central bank communications, government budget documents, and international regulatory releases are information-dense texts in which subtle wording changes can carry large market implications. Parsing these documents manually in real time is challenging, particularly for global macro managers monitoring multiple jurisdictions simultaneously. LLMs offer a scalable approach to first-pass interpretation, with human analysts providing the contextual calibration.

Resonanz Capital documents a macro fund that deployed GenAI to “summarize central bank statements and interpret policy shifts.” Notably, the same case study records an instance where the tool missed subtle language changes that a human analyst subsequently caught—an important illustration of the current ceiling on this channel and the continued necessity of human oversight for high-stakes macro calls (Resonanz Capital, 2024).

Man Group describes its AI copilot as capable of helping “draft trade rationales and surface anomalies in alternative data before they hit production signals,” compressing the cycle from macro observation to formalized investment thesis (Resonanz Capital, 2024).

**Processing Alternative Disclosure on Special Topics**

Some alternative disclosure on special topics, such as ESG disclosures, present a distinctive information-processing challenge: they are non-standardized, jurisdiction-specific, and scattered across sustainability reports, proxy filings, and voluntary frameworks (TCFD, GRI, SASB) that vary substantially in structure and coverage. Manually normalizing ESG data across large investment universes is extremely resource-intensive, making it a natural candidate for LLM-assisted processing.

Resonanz Capital describes a fund that deployed GenAI to “analyze proxy voting data and sustainability disclosures across dozens of jurisdictions, enabling a faster test of whether ESG screens could be alpha-generating in their specific investment universe.” The ability to comparably process dozens of non-standardized disclosure formats simultaneously is a capability that is structurally difficult to replicate with either human analysts or legacy rule-based extraction tools (Resonanz Capital, 2024).

## Table IA.1. Transition Matrix: Machine-based Investor vs Non-machine-based Investor

This table reports one-year transition matrices for machine-based investor classification, constructed from consecutive investor-year observations. The transition probability $P_{i,j} = P(X_{t+1} = j|X_t = i)$, where $X_t \in \{0,1\}$ indicates non-machine and machine-based investing behavior, respectively. Panel A reports the observed transition matrix, computed as the empirical frequency of transitions from state i to state j across adjacent years. Panel B reports the fitted transition matrix, constructed from a logit model in which next-year machine status is regressed on current machine status and investor characteristics, including industry-level HHI, portfolio turnover, log total portfolio value, log number of stocks owned, portfolio breadth, average percentage holdings, and the percentage of holdings in large blocks, and averaging the predicted conditional probabilities within each originating state.

Panel A. Observed Transition Matrix

| | | Year t+1 | |
|---|---|---|---|
| | | Machine-based Investor | Non-machine Investor |
| Year t | Machine-based Investor | 97.44 | 2.56 |
| | Non-machine Investor | 0.24 | 99.76 |

Panel B. Fitted Transition Matrix

| | | Year t+1 | |
|---|---|---|---|
| | | Machine-based Investor | Non-machine Investor |
| Year t | Machine-based Investor | 97.42 | 2.58 |
| | Non-machine Investor | 0.23 | 99.77 |

## Table IA.2. Top 20 Machine-Investors and Average Holding Duration

This table reports the top 20 institutional investors ranked by the total number of machine-based EDGAR filing downloads during the 2015–2017 SEC server log sample period. For each investor, we report the institution name, total number of machine-classified downloads, and the average holding duration measured in quarters. Holding duration is calculated as the average number of consecutive quarters each investor holds a given stock position, following Gaspar, Massa, and Matos (2005). The average holding duration across the top 20 machine-based investors is approximately 6.1 quarters (1.5 years), indicating that machine-based investors in our sample are predominantly fundamental or quantitative institutional investors with multi-quarter holding horizons rather than high-frequency traders. The list includes major asset managers such as T. Rowe Price, JPMorgan Chase, and Grantham Mayo Van Otterloo (GMO), alongside quantitative firms such as Renaissance Technologies and Two Sigma Investments.

| Rank | Name of Institution | # Machine Downloads | Average Holding Duration (Quarter) |
|---|---|---|---|
| 1 | T. Rowe Price Associates, INC. | 548,954 | 5.784 |
| 2 | Grantham Mayo Van Otterloo | 377,337 | 5.356 |
| 3 | Wells Fargo Bank, N.A. | 215,514 | 17.227 |
| 4 | Panagora | 147,465 | 8.703 |
| 5 | Renaissance Technologies Corp. | 107,513 | 5.597 |
| 6 | AGF Management Limited | 91,040 | 10.825 |
| 7 | Paloma Partners | 83,720 | 2.335 |
| 8 | Invesco Group Services, Inc. | 76,768 | 14.249 |
| 9 | Two Sigma Investments, LP | 75,599 | 5.925 |
| 10 | Zacks Investment Research Inc. | 49,775 | 3.989 |
| 11 | Soros Fund Management, LLC | 43,551 | 2.372 |
| 12 | Morgan Stanley | 37,492 | 2.667 |
| 13 | Tower Research Capital LLC | 32,374 | 4.587 |
| 14 | SquarePoint Ops LLC | 29,657 | 2.939 |
| 15 | Point72 Asset Management, LP | 25,733 | 2.507 |
| 16 | JPMorgan Chase | 23,835 | 10.118 |
| 17 | Deutsche Bank AG | 17,715 | 9.444 |
| 18 | Abacus Wealth Partners, LLC | 12,123 | 2.500 |
| 19 | Guggenheim Partners | 10,466 | 11.193 |
| 20 | SIG Growth Equity | 8,458 | 4.000 |

## Table IA.3. AI-Powered Governance and Firm Valuation: Binary Treatment

This table reports results from firm-level regressions examining the effects of AI-powered governance on firm valuation and operating performance following the ChatGPT shock. The sample consists of 58,661 firm-year-quarter observations for 4,464 unique firms for the sample period from 2019 Q3 to 2025 Q3, constructed from 13F institutional ownership data. In columns (1) and (2), the dependent variables are *Tobin's Q* and *Total Q*, respectively, which measure firm market valuation. *Tobin's Q* is defined as the ratio of the market value of assets to the book value of assets, while *Total Q* incorporates intangible capital following Peters and Taylor (2017). The key independent variable is *Post × Machine-based Investor Dummy*, which interacts an indicator for the post-ChatGPT period with firms' pre-period exposure to Machine-based institutional investors. Control variables include firm size, return on assets (or lagged return on assets where applicable), leverage, market-to-book ratio, cash holdings, institutional ownership, and asset tangibility. To account for firms' own adoption of artificial intelligence, we additionally control for *Firm AI Adoption*, measured using the AI employment intensity variable from Babina et al. (2024). All regressions include firm fixed effects and year-quarter fixed effects, absorbing time-invariant firm characteristics and common macroeconomic shocks. Standard errors reported in parentheses are robust and clustered by industry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

| | Tobin's Q<br>(1) | Total Q<br>(2) |
|---|---|---|
| Post × Machine-based Investor Dummy | 0.261*** | 1.327*** |
| | (0.095) | (0.435) |
| Firm AI Adoption | -0.091 | -0.195 |
| | (0.097) | (0.214) |
| Size | -0.664*** | -0.530** |
| | (0.128) | (0.234) |
| ROA | 0.421 | 6.055*** |
| | (0.926) | (1.077) |
| Leverage | -2.065*** | -2.410*** |
| | (0.249) | (0.305) |
| Cash | 0.948*** | 4.209*** |
| | (0.228) | (0.959) |
| Tangibility | -0.106 | -1.873** |
| | (0.441) | (0.797) |
| Institutional Ownership | 0.244** | -0.419* |
| | (0.103) | (0.222) |
| | | |
| Firm fixed effects | Yes | Yes |
| Year-Quarter fixed effects | Yes | Yes |
| Observations | 58,661 | 40,074 |
| Adjusted R-squared | 0.767 | 0.678 |

## Table IA.4. Probit Regression Results

This table reports results from probit regressions examining the robustness of the baseline linear probability model estimates for binary outcome variables. Column (1) examines the incidence of financial restatements, where the dependent variable is an indicator equal to one if the firm issues any financial restatement in a given quarter and zero otherwise. Column (2) examines the likelihood of M&A activity, where the dependent variable is an indicator equal to one if the firm completes any M&A deal in a given quarter and zero otherwise. Column (3) examines the likelihood of value-destroying M&A activity, where the dependent variable is an indicator equal to one if the firm completes a deal with negative announcement-period abnormal returns and zero otherwise. The key independent variable is *Post × Machine-based Ownership*, which interacts an indicator for the post-ChatGPT period with firms' pre-period exposure to machine-based institutional investors. Control variables include firm size, return on assets (ROA), leverage, market-to-book ratio, cash holdings, asset tangibility, institutional ownership, and firm-level AI adoption. All specifications include year-quarter fixed effects. Standard errors reported in parentheses are robust and clustered by industry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

| | Any Restatement | Any M&A deal | Any Value-destroying Deal |
|---|---|---|---|
| | (1) | (2) | (3) |
| Post × Machine-based Ownership | -0.815*** | -0.647*** | -0.608*** |
| | (0.293) | (0.147) | (0.191) |
| Machine-based Ownership | -0.625** | 0.733*** | 0.577*** |
| | (0.287) | (0.122) | (0.159) |
| Firm AI Adoption | 0.038* | 0.076*** | 0.065*** |
| | (0.021) | (0.007) | (0.009) |
| Size | -0.044** | 0.083*** | 0.098*** |
| | (0.020) | (0.005) | (0.006) |
| ROA | -0.088 | 0.610*** | 0.345** |
| | (0.202) | (0.113) | (0.158) |
| Leverage | 0.297*** | -0.465*** | -0.453*** |
| | (0.098) | (0.038) | (0.051) |
| Market-to-Book | -0.005 | 0.021*** | 0.027*** |
| | (0.005) | (0.003) | (0.004) |
| Cash | -0.254 | -0.481*** | -0.450*** |
| | (0.181) | (0.038) | (0.053) |
| Tangibility | -0.163 | -0.293*** | -0.269*** |
| | (0.201) | (0.032) | (0.042) |
| Institutional Ownership | 0.250*** | 0.122*** | 0.213*** |
| | (0.080) | (0.033) | (0.044) |
| | | | |
| Sample | Full | Full | Full |
| Model | Probit | Probit | Probit |
| Firm fixed effects | No | No | No |
| Year-Quarter fixed effects | Yes | Yes | Yes |
| Observations | 64,986 | 68,278 | 68,270 |