Artificial Downtime

UK Evidence on Investor GenAI Reliance and Trading Behaviour from ChatGPT and Claude Outages

Isabella Lejdström

Ellen Wickell

Master Thesis
Stockholm School of Economics
2026

Artificial downtime: UK evidence on investor GenAI reliance and trading behaviour from ChatGPT and Claude outages

Abstract

In this paper, we provide updated evidence on the extent to which investors rely on generative artificial intelligence (GenAI) tools, and what disruptions to these tools reveal about their role in financial markets. We exploit unanticipated outages to ChatGPT and Claude, between March 2025 and February 2026, as plausibly exogenous shocks to investors' access to GenAI tools, and measure changes in trading volume on the London Stock Exchange. We document a significant decline in trading activity during GenAI outage periods. The effect differs systematically across platforms and is primarily driven by Claude outages, with ChatGPT outages producing no statistically significant effect across any specification. We further find suggestive, though not conclusive, evidence that the outage effect is materially stronger on accounting disclosure days than on other regulatory and non-regulatory news days. Finally, a dynamic analysis illustrates that trading activity declines immediately following outages, with the effect intensifying before partially reversing. Taken together, our findings suggest that investors rely on GenAI tools in economically meaningful ways, such that disruptions to a single platform produce observable effects on market activity, even in a more mature multi-tool environment.

Keywords:

Generative AI, ChatGPT, Claude, Trading volume, Information processing, Disclosure processing costs

Authors:

Isabella Lejdström (25520)
Ellen Wickell (25286)

Tutor:

Ting Dong, Assistant Professor, Department of Accounting

Master Thesis
Master Program in Accounting, Valuation and Financial Management
Stockholm School of Economics

Contents

1. Introduction

The widespread adoption of generative artificial intelligence (GenAI) has dramatically expanded the information processing capacity available to anyone with a browser. Tools like ChatGPT and Claude can support a broad range of analytical and professional tasks within a single conversational interface and are increasingly embedded in professional workflows, though evidence on their real-world effects in financial markets remains limited (e.g., Bertomeu et al., 2025; Blankespoor et al., 2026; Bradshaw et al., 2026; Cheng et al., 2025). Cheng et al. (2025) provide early evidence that investors rely on GenAI tools for professional tasks, documenting reduced trading activity during ChatGPT outages on US equity markets. Yet this evidence is limited to a single platform in its early adoption phase and to the US market, leaving open whether investor reliance extends to other platforms, later periods, or other institutional settings. In this paper, we use ChatGPT and Claude outages to provide updated evidence on the following question: *to what extent do investors rely on GenAI tools, and what do disruptions to these tools reveal about their role in financial markets?*

Since OpenAI released ChatGPT in 2022, GenAI adoption has expanded rapidly, with ChatGPT alone reaching approximately 700 million weekly active users by mid 2025 (Chatterji et al., 2025). Competing Large Language Models (LLMs), such as Anthropic's Claude and Google's Gemini, have emerged alongside it, transforming the GenAI landscape from one dominated by a single widely adopted platform into a competitive multi-tool environment (De Vynck and Schaul, 2025). Finance-related occupations are among those most exposed to this new technology (e.g., Cheng et al., 2025; Eisfeldt et al., 2023; Massenkoff and McCrory, 2026), and a growing body of work demonstrates that LLMs can assist with investment-related tasks such as summarising corporate disclosures, forecasting earnings, analysing analyst reports, and processing financial news (e.g., Breitung and Müller, 2025; Li, Mai, et al., 2025; Li, Shen, et al., 2025; Lopez-Lira and Tang, 2023). Anthropic has explicitly targeted this sector through the launch of Claude for Financial Services in July 2025, a dedicated product aimed at investment research, financial modelling, and compliance workflows (Anthropic, 2025).

Investors face limited information processing capacity (Blankespoor et al., 2020; Sims, 2003), and prior research suggests that investors with such constraints often delay or forego trading decisions, resulting in reduced market activity (e.g., deHaan et al., 2017; DellaVigna and Pollet, 2009; Hirshleifer et al., 2009). What makes GenAI particularly relevant to addressing these constraints is arguably not any single task-specific capability but the breadth of tasks a single tool can support, spanning information acquisition, information integration, and routine tasks, and the ease with which investors can adopt it (Cheng et al., 2025). Even when GenAI supports tasks not directly related to investment decisions, time saved on routine tasks frees up capacity for other investment-related activities (Acemoglu, 2025; Cheng et al., 2025). Yet investors' use of GenAI tools is not directly observable, making empirical analysis necessary to address the question motivating this study.

We use unanticipated outages to ChatGPT and Claude, between March 2025 and February 2026, as plausibly exogenous shocks to investors' access to GenAI tools, and examine how trading activity on the London Stock Exchange (LSE) is affected, following the broader outage-based identification literature (e.g., Cheng et al., 2025; Eaton et al., 2022). Outage events are identified using official incident reports published by OpenAI and Anthropic. We identify four ChatGPT outage events and six Claude outage events, a total of ten events during our period of study. We collect intraday trading data from London Stock Exchange Group (LSEG) to compare 5-minute interval trading activity during outage periods with matched non-outage intervals for the same security over the preceding five trading days. The final sample consists of 720,909 security-date-5-minute interval observations across 622 unique stock securities.

To better understand the dynamics and mechanisms underlying the main result, we examine three supporting implications. Our multi-platform design allows us to test whether outage effects generalise across ChatGPT and Claude or differ between platforms. We further explore the cross-sectional distribution of the effect by testing whether outage effects are stronger on accounting disclosure days than on other regulatory and non-regulatory news days. Lastly, we examine the timing of the effect, asking whether trading volume declines immediately at outage onset or emerges with a delay.

We document a significant decline in trading volume during GenAI outages on the LSE, indicating that investors utilise these tools in ways that influence their trading decisions. The drop in trading volume corresponds to approximately 1.41% of the sample standard deviation. The effect differs systematically across platforms and is primarily driven by Claude outages, with ChatGPT outages producing no statistically significant effect across any specification. These findings remain broadly robust across a wide set of sensitivity tests. We further find suggestive, though not conclusive, evidence that the outage effect is materially stronger on accounting disclosure days than on other regulatory and non-regulatory news days, suggesting that investor reliance on GenAI is more closely tied to analytically demanding disclosure settings, and pointing toward information processing as the more economically relevant channel through which GenAI is used. Finally, the dynamic analysis illustrates that trading activity declines immediately following outages, with the effect intensifying before partially reversing, a pattern consistent with GenAI being embedded across multiple stages of the investment workflow. Taken together, our findings suggest that investors rely on GenAI tools in economically meaningful ways and that these tools have become sufficiently embedded in investor workflows such that disruptions to a single platform produce observable effects on market activity, even in a more mature multi-tool environment.

This paper contributes to the emerging literature on the real-world economic effects of GenAI (e.g., Bertomeu et al., 2025; Blankespoor et al., 2026; Bradshaw et al., 2026; Cheng et al., 2025) by extending the existing evidence from an early-adoption setting dominated by a single platform to a more mature multi-tool environment and a different institutional setting. More broadly, we contribute to the literature on technology and information processing in capital markets (e.g., Blankespoor et al., 2019, 2020; Brown et al., 2015; Gao and Huang, 2020) by providing evidence that GenAI tools have become sufficiently embedded in investment activities for disruptions to their availability to produce observable effects on market outcomes. More practically, by showing that disruptions to a single platform produce observable effects on aggregate market activity, our findings are relevant for how financial institutions navigate GenAI adoption decisions and how regulators think about operational resilience in increasingly AI-supported financial markets.

2. Literature review and hypothesis development

This chapter develops the theoretical and empirical foundation for examining to what extent investors rely on GenAI tools and what disruptions to their access reveal about their role in financial markets, as measured through changes in trading volume. We begin by describing the institutional setting in which ChatGPT and Claude have come to occupy a meaningful position in investor workflows. We then turn to the theoretical literature on information frictions and disclosure processing costs, which identifies the constraints that shape how investors translate firm disclosures into trading decisions. Building on this foundation, we situate GenAI tools within the broader trajectory of technological change in disclosure processing and investor workflows and review the still-thin empirical evidence on their real-world market effects. The chapter concludes with the formal development of our hypothesis.

2.1. Institutional background

OpenAI released ChatGPT on November 30, 2022, and the tool reached 100 million users within two months, a pace of adoption that made it the fastest-growing consumer application on record at the time (Milmo, 2023). By mid-2025, weekly active users had grown to approximately 700 million (Chatterji et al., 2025). Built on an LLM, ChatGPT synthesises information, adapts to context, and performs analytical tasks within a single conversational interface. This natural language interface lowers usability barriers regardless of technical expertise, making a wide range of cognitive tasks broadly accessible in ways that earlier computer-based tools did not, supporting more informed decision-making (Blankespoor et al., 2026; Cheng et al., 2025).

These features are particularly relevant for finance-related occupations as recent research ranks those among the most exposed to GenAI (e.g., Cheng et al., 2025; Eisfeldt et al., 2023; Massenkoff and McCrory, 2026). A growing body of work shows that LLMs can perform investment-related tasks that were previously the domain of human investors, including creating business summaries from corporate disclosures, forecasting earnings, analysing analyst reports, and processing business news (e.g., Breitung and Müller, 2025; Li, Mai, et al., 2025; Li, Shen, et al., 2025; Lopez-Lira and Tang, 2023). It is this generality that makes GenAI directly relevant to investor workflows (Cheng et al., 2025). Drawing on a review of anecdotal use cases, Cheng et al. (2025) group investor applications of GenAI, more specifically ChatGPT, into three categories. The first is information acquisition, which involves obtaining and summarising information on specific stocks, industries, or market conditions. The second is information integration, which covers sentiment analysis, forecasting, and portfolio evaluation. The third is routine task execution, such as email drafting, document reviewing, and coding. Insights from our discussions with investment professionals further illustrate the depth of this integration.[1] For example, an equity sales professional at a Nordic financial institution describes feeding four to five research reports daily into GenAI tools to extract key points, compiling morning emails, summarising up to 40 analyst notes, and using GenAI tools to form independent views on companies outside his primary coverage.

The rapid uptake and demonstrated usefulness of ChatGPT attracted substantial investment and intensified competition in the GenAI market, with prominent alternatives emerging including Gemini by Google and Claude by Anthropic, transforming what was initially an environment with a single widely adopted platform into a competitive multi-tool market (De Vynck and Schaul, 2025; Maslej et al., 2025).

Anthropic released Claude in March 2023 as a direct competitor to ChatGPT, positioning the product around reliability, interpretability, and reasoning quality (Anthropic, n.d.). By early 2026, Anthropic had reached a reported valuation of approximately $380bn (Anthropic, 2026), growing its annualised revenue from $1bn at the start of 2025 to more than $9bn by the end of the year (Hammond, 2026). Accordingly, industry analyses indicate that Anthropic overtook OpenAI in the enterprise LLM market share during 2025 (Tully et al., 2025). In July 2025, Anthropic launched Claude for Financial Services, a dedicated product aimed at investment research, financial modelling, and compliance workflows (Anthropic, 2025). Taken together, these developments suggest that Claude is establishing a meaningful footprint in the investment community. Practitioner observations are consistent with this. For example, a derivatives sales professional at a Nordic financial institution describes Claude as widely known and broadly used, with particularly rapid growth at the start of 2025, and that it is now Claude rather than ChatGPT that dominates among the finance professionals he is familiar with.

Adoption of GenAI is neither universal nor uniform, and the picture has evolved rapidly. In early 2024, a survey found that 30% of financial institutions had banned the use of GenAI tools entirely, citing concerns about confidential data leakage and regulatory compliance (Crosman, 2024). By 2025, however, active GenAI use among senior finance leaders had more than doubled relative to the prior year, rising from 30% to 75% across a survey of over 1,000 finance leaders in 20 countries (McAllen et al., 2026), suggesting that many institutions that initially restricted access have since moved toward more structured adoption. Further, research finds that GenAI tools are subject to hallucination, where the model invents incorrect information and presents it as factual (Athaluri et al., 2023; Emsley, 2023). Professionals rate this as their greatest concern in using GenAI (Murgia et al., 2026; Fletcher, 2026), potentially constraining professional uptake particularly for analytical tasks where accuracy is critical. This within-industry unevenness is consistent with practitioner experience. One of the approached investment professionals describes his institution as lagging in GenAI adoption, relying primarily on

[1] Interviews were conducted on a confidential basis in April 2026. Names and institutional affiliations are withheld at participants' request. See Appendix A for further details.

internally developed tools with limited capabilities. In particular, external GenAI systems are prohibited for sensitive trading data and client information, confining use largely to routine tasks such as email drafting rather than analytical investment work.

By 2024-2026, the GenAI market had evolved into a multi-tool environment. ChatGPT remained the most widely used system overall, but Claude had established a credible presence among professionals, and the two platforms are likely to serve overlapping user bases (Kotrotsos, 2026). As such, both tools appear sufficiently adopted and embedded in investor workflows to plausibly affect the constraints that materially shape how investors process and act on information. A long theoretical literature has identified the constraints most likely to be at work.

2.2. Theoretical background

The question of whether GenAI outages affect trading behaviour is, at its core, a question about information frictions. If processing firm disclosures were costless and instantaneous, the temporary loss of a tool that automates portions of that processing would leave trading activity unchanged. If, however, investors operate under binding capacity constraints in processing and acting on information, any factor that materially alters those constraints should leave a measurable mark on trading.

2.2.1. Information frictions: Costly information and limited attention

A central question in financial economics is how and to what extent public information is incorporated into asset prices, and how information frictions shape this process. The efficient market hypothesis (EMH) serves as the theoretical starting point for understanding this. Fama (1970) defines an efficient market as one in which prices reflect available information quickly and without systematic bias, and distinguishes between three forms of efficiency, weak, semi-strong, and strong, depending on the information set reflected in prices. The semi-strong form is most relevant here, as it implies that prices should fully and immediately reflect all publicly available information. Fama identifies three sufficient conditions under which a market would achieve this ideal: zero transaction costs, costless information available to all participants, and homogeneous beliefs regarding the implications of that information.

Fama’s framework has not gone unchallenged. A vast empirical literature has tested its implications across a range of methodological approaches, and while many studies find that prices adjust rapidly to public information, others document patterns that suggest systematic deviations from fully frictionless efficiency (e.g., Bartov et al., 1998; Bernard and Thomas, 1989; Foster et al., 1984).

A related theoretical literature has long emphasised that market efficiency is inherently limited when information is costly. Early noisy rational expectations models show that if acquiring and processing information involves costs, prices cannot be perfectly revealing (Grossman and Stiglitz, 1980; Hellwig, 1980; Verrecchia, 1982). Grossman and Stiglitz (1980) challenge Fama's framework directly, arguing that costless information is not only a sufficient, but also a necessary condition for prices to fully reflect all available information. They demonstrate that if information is costly, perfectly informative prices cannot constitute an equilibrium, as investors would have no incentive to bear the cost of becoming informed if prices already revealed everything they would learn by doing so. As Jensen (1978) argues, market efficiency must therefore be evaluated relative to the costs of obtaining and acting on information. In equilibrium, some degree of information asymmetry must persist, and prices reflect information only to the extent that investors have sufficient incentives to acquire and trade on it.

While noisy rational expectations models emphasise equilibrium outcomes when information is costly, rational inattention theory highlights a distinct and complementary constraint: limits on agents' capacity to process information once available. Sims (2003) introduces the framework in which agents face explicit constraints on their information processing capacity and must optimally allocate limited attention across competing signals. Because cognitive resources are finite and processing information requires effort, agents cannot perfectly absorb and interpret all available data, even when access to that data is, in principle, unrestricted. These frictions bind to every signal investors might engage with, but they become more concrete and empirically accessible in the setting of firm disclosures.

2.2.2. Disclosure processing costs

The frictions described above can be operationalised in disclosure settings as disclosure processing costs. Blankespoor et al. (2019) conceptualise disclosure processing as a sequential economic decision involving three distinct stages, illustrated in Figure 1 Panel A, where a disclosure refers to a specific signal from a firm report or communication. For a disclosure to affect trading behaviour, investors must first learn that the disclosure exists (awareness costs), then obtain the report and extract the relevant disclosure (acquisition costs), and finally analyse the implications of the disclosure for firm value (integration costs). Each stage involves both explicit and opportunity costs, as processing a disclosure consumes resources that could otherwise be allocated to other activities or to processing other disclosures (Blankespoor et al., 2020). Disclosure processing is therefore an optimisation problem in which investors allocate scarce processing resources across competing signals, consistent with the constraints emphasised by the rational inattention literature (Blankespoor et al., 2020; Sims, 2003).

The presence of any material processing costs means that disclosures may not function as fully public information in the traditional sense. Instead, they take the form of costly private information that only a subset of investors will optimally process (Blankespoor et al., 2020), mirroring the equilibrium logic of noisy rational expectation models

(Grossman and Stiglitz, 1980). Because processing resources are scarce and disclosure costs are non-negligible even for publicly available information, investors effectively operate under a binding information processing capacity constraint that determines how much of any given disclosure can be meaningfully engaged with and acted upon. If this constraint is economically meaningful, observed trading activity should respond systematically to any factor that alters disclosure processing costs, a claim that an extensive empirical literature has both tested and supported.

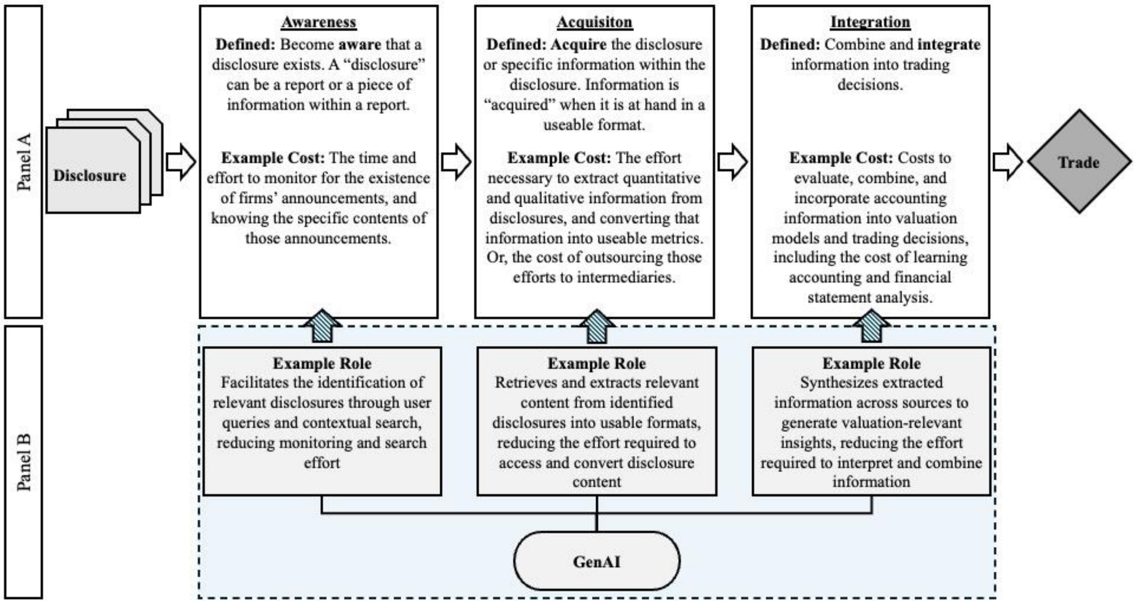

Figure 1. The disclosure processing framework and the role of GenAI

Panel A reproduces the sequential framework of information usage from Blankespoor et al. (2019), illustrating the steps required to incorporate disclosures into trading decisions and examples of the effort required at each stage. Panel B provides illustrative examples of how GenAI tools can affect each stage of the disclosure processing chain by reducing the effort required to identify, extract, and integrate information.

2.2.3. Information processing capacity and trading behaviour

Empirical work has long established that disclosures generate measurable changes in trading activity (e.g., Ball and Brown, 1968; Beaver, 1968; Cready, 1988), and subsequent research documents that disclosure processing costs directly compress investors' available capacity, with observable consequences for trading. Investors disproportionately trade in attention-grabbing stocks rather than distributing cognitive resources evenly across securities (Barber and Odean, 2008), and volume reactions are significantly weaker when disclosure processing costs rise, whether through simultaneous competing announcements (Hirshleifer et al., 2009) or contextual distractions (deHaan et al., 2017; DellaVigna and Pollet, 2009). Disclosure complexity raises these costs directly, with less readable reports reducing investor participation and weakening market reactions (Loughran and McDonald, 2011). Of the three stages in the disclosure processing framework, integration costs appear to impose the greatest constraint. Even when awareness and acquisition costs are substantially reduced, individual investors fail to trade on accounting information, responding instead to less analytically demanding signals (Blankespoor et al., 2019). Quasi-experimental evidence reinforces this picture, showing that exogenous restrictions on an information processing tool, proxied by mobile communication through which investors monitor and process firm disclosures, reduce local trading volume, with effects strongest where investors were most dependent on that tool (Brown et al., 2015).

Taken together, this evidence establishes that trading volume is sensitive to any factor that materially compresses investors' information processing capacity, as investors delay or forgo decisions they lack the capacity to adequately process. Whether GenAI tools function as one such factor, and how strongly, depends on what distinguishes them from the reporting and intermediation technologies that have shaped disclosure processing and investor workflows to date.

2.3. GenAI as an information processing technology

Investor workflows and disclosure processing have been reshaped by successive waves of technological change, each of which has lowered some component of the costs identified above. Understanding what is genuinely new about GenAI requires placing it against this longer trajectory, distinguishing it from the reporting technologies and intermediation services that came before, and reviewing the still-thin empirical evidence on its real-world market effects.

2.3.1. The technological evolution of information processing

Financial reporting technologies have progressively lowered the cost of accessing and processing firm disclosures, but each generation has addressed only a subset of the cost stages identified above. For example, the SEC introduced EDGAR in 1994 as an electronic repository for regulatory filings, principally reducing investors' awareness and acquisition costs (Blankespoor et al., 2020). XBRL, first required in 2009, extended this further by making financial statement data machine-readable, partially addressing acquisition and integration costs (Bhattacharya et al., 2018; Blankespoor et al., 2020). Computational text analysis additionally demonstrated that automated processing of qualitative financial disclosures can extract market-relevant signals (e.g., Li, 2010; Loughran and McDonald, 2011). While each of these innovations reduced the effective cost of processing financial information, their use remained largely concentrated among well-resourced or technically proficient market participants and operated on a subset of the disclosure processing chain (Blankespoor et al. 2020).

Data intermediaries, including journalists, analysts, and data providers, have long reduced disclosure processing costs across all three stages on investors' behalf, reducing awareness and acquisition costs by monitoring and republishing disclosures, and integration costs by synthesising their implications for firm value (Blankespoor et al., 2020). What has changed in recent decades is technology's ability to automate this intermediary function. Blankespoor et al. (2018) show that automated synthesis and distribution of earnings news significantly increased trading volume and liquidity, establishing that automation of the intermediary function measurably affects market participation. Even so, this automated function has remained external to investors' own workflows (Blankespoor et al., 2020), operating through one-way delivery models in which the scope and timing of analysis are determined by the intermediary rather than the investor.

GenAI represents a structural shift in this technological evolution by combining features that prior tools did not. First, it provides end-to-end coverage of the disclosure processing chain. A single conversational interaction with a GenAI tool can surface relevant filings, retrieve and structure their content, and produce valuation-relevant synthesis in natural language (Blankespoor et al., 2026; Cheng et al., 2025), effectively compressing awareness, acquisition, and integration into one process. Figure 1, Panel B illustrates this role of GenAI across the processing chain. Second, GenAI tools provide on-demand analytical support through a conversational interface that requires little specialist training or technical infrastructure (Blankespoor et al., 2026; Cheng et al., 2025), substantially lowering the fixed costs of advanced information processing. Third, these tools are integrated directly into investors' own workflows (Blankespoor et al., 2026; Cheng et al., 2025), rather than operating through externally provided, one-way intermediation services. Investors control the scope, depth, and timing of the analysis themselves.

GenAI's relevance to investors is not, however, confined to information processing. Its versatility extends to complementary professional tasks such as drafting, coding, and routine workflow automation that have no direct analytical function. Disruptions to these supplementary activities can still consume time and cognitive resources that would otherwise be available for processing disclosures, consistent with the task complementarity effect (Acemoglu, 2025; Cheng et al., 2025). Both channels therefore operate through the same binding information processing capacity constraint established above.

As such, unlike earlier technologies that reduced disclosure processing costs only at isolated stages, for limited groups of users, or through externally provided services, GenAI combines end-to-end coverage of the processing chain with on-demand delivery and direct integration into investors' workflows. These features position GenAI tools to materially reduce the costs of both disclosure processing and complementary professional tasks, ultimately expanding investors' effective information processing capacity. This creates an empirical setting in which disruptions to access may compress that capacity and thereby translate into observable changes in trading behaviour.

2.3.2. Real-world market evidence on GenAI

Direct empirical evidence on whether GenAI tools translate this conceptual positioning into observable market effects is still thin, a consequence of the brief period since their release. Much of the developing literature on GenAI in accounting and finance examines ChatGPT as a research tool applied to corporate disclosures, earnings releases, analyst reports, and financial news (e.g., Breitung and Müller, 2025; Li, Mai, et al., 2025; Li, Shen, et al., 2025; Lopez-Lira and Tang, 2023). These studies establish that GenAI tools possess the capabilities relevant to investor decision-making, but they do not speak to whether investors actually rely on those capabilities in ways that influence market outcomes.

A small number of recent studies begin to address this question by exploiting exogenous disruptions to access. Bertomeu et al. (2025) find that the temporary regulatory ban of ChatGPT in Italy led to significant changes in financial information production and analysis among local analysts, indicating that GenAI access shapes how financial information is processed and disseminated more broadly. Cheng et al. (2025) provide the most direct evidence to date. Analysing trading activity during eight major ChatGPT outages in the first half of 2023, they document a statistically and economically significant decline in US equity trading volume during outage periods, indicating that a meaningful number of investors use ChatGPT in ways that influence their trading decisions. They further document that the decline does not appear immediately at outage onset but becomes significant approximately 30 minutes after the disruption begins, a dynamic pattern consistent with ChatGPT supporting the information processing activities that precede trading rather than trade execution itself. Cross-sectional analysis sharpens this interpretation. The volume decline is larger for firms with corporate news released earlier on the same day or during the outage period, indicating that reliance on ChatGPT is most pronounced when investors face heightened information processing demands. The dynamic and cross-sectional patterns therefore point toward information processing being the economically dominant channel through which ChatGPT affects trading behaviour.

Taken together, these papers provide early evidence that GenAI tools have become an economically meaningful technology in investor workflows, especially through information processing, such that disruptions to their access produce observable effects in capital markets. The existing evidence, however, is confined to a single platform during the earliest months of adoption and to specific market contexts. Whether the effect persists in a more mature adoption phase, whether it survives in a multi-tool environment in which investors can in principle substitute across platforms, and whether it generalises to different institutional settings, remain open questions, which the hypothesis developed below is framed to address.

2.4. Hypothesis development

The preceding discussion converges on a single testable prediction. GenAI tools may relax investors' effective information processing constraints by supporting both disclosure processing and related professional workflows. To the extent that these tools have become embedded in investor workflows, disruptions to their availability should constrain that capacity, with observable effects on trading activity. Prior research shows that investors reduce trading activity when information processing capacity becomes constrained (e.g., deHaan et al., 2017; DellaVigna and Pollet, 2009; Hirshleifer et al., 2009), and that disruptions to disclosure processing tools can measurably affect market participation (e.g., Brown et al., 2015). More directly, Cheng et al. (2025) document a decline in US equity trading volume during ChatGPT outages, providing early evidence that investors rely on GenAI tools in economically meaningful ways.

We therefore expect GenAI outages to register as a meaningful shock to investors' access to these tools, disrupting their effective information processing capacity through direct disclosure processing, related workflows, or both, with corresponding effects on trading activity. We state this as a single hypothesis:

> **H1**. Trading volume declines during GenAI outage periods relative to matched non-outage intervals of the same security.

To better understand the underlying dynamics and mechanisms, we examine three supporting implications of H1. The first concerns generalisability across platforms. Our multi-platform design allows us to assess whether outage effects generalise across ChatGPT and Claude or differ between platforms. The second concerns the cross-sectional distribution of the effect. The information processing channel implies that disruptions should fall most heavily on investors facing the greatest processing demands, most directly on days when firms release corporate information requiring analytical interpretation (Blankespoor et al., 2019, 2020; Cheng et al., 2025). We examine this by comparing the outage effect across days without regulatory news, days with non-accounting regulatory news, and days with accounting-related regulatory news, the last of which places the heaviest demand on analytical processing. The third concerns the timing of the effect. If GenAI is more likely to be used for information processing, an outage would result in a delayed volume decline. By contrast, an immediate response would be more consistent with GenAI supporting direct trade execution (Cheng et al., 2025).

3. Method

3.1. Data and sampling

This study employs a quasi-experimental research design to examine to what extent investors rely on GenAI tools and what disruptions to their access reveal about their role in financial markets. Consistent with prior work using service disruptions as natural experiments in financial markets (Cheng et al., 2025; Eaton et al., 2022), the empirical strategy treats unanticipated outages of ChatGPT and Claude as plausibly exogenous shocks to investors' access to these tools and examines how trading activity changes during these periods using intraday panel data.

The outage identification strategy rests on two assumptions. First, the outages must be unanticipated by market participants. If outages arise from idiosyncratic technical failures, such as database instabilities and configuration errors, they are unlikely to be predicted in advance, satisfying this requirement. Second, the outages must not coincide with other events that independently affect trading activity. We address this concern through a validation procedure described in Section 3.1.2.

Even under these assumptions, attributing changes in trading volume to outages requires separating the outage effect from the many other sources of intraday variation that shape trading activity. Following Cheng et al. (2025), the empirical design compares each outage interval to the same intraday interval of the same security on the five preceding trading days. This within-firm, within-time-of-day comparison controls for systematic variation in trading patterns driven by time of day, firm-specific characteristics, and broader market conditions.

3.1.1. Scope

The dataset includes securities listed on the LSE, extending prior US-focused evidence (Cheng et al., 2025) to a European setting. As Europe's largest capital market, the LSE offers a deep pool of liquidity and a broad base of international investors (London Stock Exchange, n.d.), making it a suitable setting for assessing whether GenAI-related trading effects generalise beyond the US. At the same time, focusing on a single exchange ensures a consistent institutional and regulatory environment within the sample. The study adopts a one-year time horizon, spanning March 2025 to February 2026. Using a one-year observation window ensures a sufficient number of outage events for analysis. Also, given the rapid evolution of GenAI tools, a one-year window centred on a more mature adoption phase aims at isolating current investor behaviour more cleanly than a longer horizon that would span distinct phases of the adoption curve.

3.1.2. Outage identification

Building on the approach of Cheng et al. (2025), we identify ChatGPT and Claude outage events occurring between March 2025 and February 2026 based on official reports published by OpenAI and Anthropic, respectively.[2] Both providers have maintained publicly available records of the date and severity of service disruptions since early 2023, along with detailed incident reports outlining the timeline and nature of each event.

Cheng et al. (2025) identify *Major Outages* during which ChatGPT was inaccessible to users. For Claude, we apply the same criterion, using Anthropic's *Major Outage* classification as the primary identification criterion for the claude.ai service, excluding outages solely affecting the underlying API or other Anthropic products. For ChatGPT, the sample period contains far fewer *Major Outage* events than the period examined by Cheng et al. (2025), possibly reflecting improvements in platform stability since its initial launch. We therefore also consider incidents classified as *Partial Outage* and *Degraded Performance* and retain those describing user-facing disruptions to the ChatGPT conversational interface, excluding those affecting only the API, Codex, or other developer-facing products.

To ensure relevance for financial markets, we apply two filters to all candidate events. First, we retain only events that overlap at least partially with LSE trading hours, excluding those occurring entirely outside trading hours or on weekends and public holidays. Second, we drop events with a duration of less than 10 minutes. For each retained event, the start time is defined as the onset of the reported service disruption and the resolution time as the point at which full-service restoration was confirmed, both as recorded in the incident reports. Unlike Cheng et al. (2025) and Eaton et al. (2022), we do not have access to DownDetector, a platform that aggregates real-time user reports of service disruptions, to supplement the timing of outage start. However, Cheng et al. (2025) show that using official OpenAI reported start times instead of DownDetector yields materially unchanged results, supporting the reliability of official incident reports. Since official reports may marginally lag the true start of disruption and lead the true resolution, both timestamps are floored to the nearest 5-minute interval boundary. Thus, the first outage reported from 08:14 to 09:38, is treated as covering the intervals 08:10-09:35.

To confirm that the identified events represent genuine and meaningful service disruptions, we subject each candidate event to a validation procedure drawing on two supplementary sources. First, we use Google Trends, which reports search interest as a relative index normalised within the selected time window, to measure the volume of user searches related to each platform's service status during the reported outages.[3] We filter

---

[2] To see official incident reports go to: https://status.openai.com/history and https://status.claude.com/uptime

[3] To access Google Trends go to: https://trends.google.com/trends/

to the United Kingdom, and use the following search terms, where XXX represents either ChatGPT or Claude: "XXX outage", "XXX service issue", "XXX down", and "is XXX down". We examine annual trends at weekly resolution to identify candidate spike weeks, then narrow to the relevant week to obtain daily resolution. As spikes reflect relative rather than absolute search volume, they are assessed visually rather than against a fixed numerical threshold. We observe clear spikes in service-related search activity coinciding with the reported outages for both platforms. Second, we conduct a thorough search of news sources to assess real-world impact and to screen for overlapping disruptions unrelated to the GenAI services themselves that could contaminate the interpretation of any trading effects. Through this procedure we identify and exclude one event for each platform in November 2025 as well as one event for Claude in December 2025, attributable to broader disruptions involving the network infrastructure provider Cloudflare.[4] As Claude events are restricted to the stricter *Major Outage* classification, official incident records and Google Trends are considered sufficient for validation. For ChatGPT, where candidate events span a broader range of severity classifications, we additionally require that a news report of the disruption exists, retaining only those events satisfying this criterion (see Appendix B).

Following these identification and validation steps, we retain four ChatGPT outage events and six Claude outage events, a total of ten events. The events occur on distinct calendar dates and at varying times of day, reducing the likelihood that any observed trading activity is driven by date- or time-specific market conditions. Durations range from 21 to 510 minutes, with an average of 148 minutes. Appendix B presents an overview of all identified outage events and their key characteristics.

Although the exact causes of all outages are not publicly disclosed, we are able to identify the underlying reasons for four events based on incident reports. These causes vary and include factors such as configuration errors and deployment capacity issues, consistent with outages arising from idiosyncratic technical issues rather than systematic market-related factors.

3.1.3. Financial and news data

Firm-level intraday trading data is obtained from LSEG. Intraday trading volume data, including shares- and trades count, are collected at aggregated 5-minute intervals for all securities traded on the LSE. In addition, we collect daily stock prices, and firm-specific information from LSEG.

We further collect regulatory news announcement data from Investegate, a distributor of announcements released via the London Stock Exchange's Regulatory News Service

[4] For more information on the Cloudflare disruptions go to: https://blog.cloudflare.com/18-november-2025-outage/ and https://blog.cloudflare.com/5-december-2025-outage/

(RNS).[5] The RNS is the primary channel through which UK-listed companies release price-sensitive information to the market, including earnings releases, trading updates, interim and final results, and other regulatory disclosures (LSEG, n.d.). As such, RNS announcements represent the UK equivalent of the SEC filings and earnings press releases examined in the US literature (Cheng et al., 2025). Because Investegate displays announcement listings dynamically, we build a custom browser scraper to collect the relevant data. For each announcement, the scraper records the disclosure time, the issuing company name, the ticker symbol, and the announcement headline. A description of the scraping procedure is provided in Appendix C.1.

3.1.4. Sample construction

We begin with all securities traded on the LSE and their corresponding trading data. For each security and outage event, we retain all 5-minute interval observations within the affected trading period, along with observations from the same time intervals of the preceding five trading days. Our sample selection starts with the maximum number of 3,003,156 possible observations, calculated as 299 5-minute intervals across the ten outage events, multiplied by six days and 1,674 unique securities.

We then apply a series of filters following the sample construction procedure of Cheng et al. (2025). First, we ensure that no observations on control days overlap with other ChatGPT or Claude outages to avoid contaminating the control periods. Second, we exclude securities that are not ordinary common shares. Third, we drop observations for securities with no trading activity on the day. Fourth, we drop observations with closing stock prices below £5 on the day preceding the outage. In addition to these filters, we introduce a liquidity filter that drops observations with fewer than on average 30 active 5-minute intervals with positive trading volume on the control days. This additional filter is motivated by the fact that the US equity markets studied in prior work (Cheng et al., 2025; Eaton et al., 2022) are more liquid than the LSE-listed stocks (e.g., Brogaard et al., 2025; Karolyi et al., 2012), and retaining such securities may introduce noise into the intraday trading volume measures. Lastly, we drop observations with missing control variables.

Our final sample consists of 720,909 security-date-5-minute interval observations from 622 unique stock securities. Table 1 summarises the sample selection process.

[5] To access Investegate go to: https://www.investegate.co.uk/

Table 1. Sample selection
This table presents the sample selection process. The final sample for the trading volume analyses includes 720,909 security-date-5-minute interval observations from 622 unique stock securities during the period from March 2025 to February 2026.

	# Observations	# Securities
Initial sample: 5-minute intraday intervals during the ChatGPT and Claude outages and the same intervals for the same firms over the previous five trading days	3,003,156	1,674
Less:		
Intervals in the control periods that coincide with other outages	(0)	(0)
Non-ordinary common shares	(517,080)	(286)
Observations with no trading activity on the day	(841,355)	(62)
Observations with closing stock prices below £5 on the day preceding the outage	(251,747)	(222)
Observations with fewer than on average 30 active 5-minute intervals on the control days	(647,339)	(480)
Observations with missing control variables	(24,726)	(2)
Final sample	720,909	622

3.2. Research design

3.2.1. Changes in trading volume during GenAI outages

We use the following ordinary least squares (OLS) regression model to examine the impact of GenAI outages on trading volume:

$$\begin{aligned} Trading\ Volume_{i,t,p} = \ & \beta_0 + \beta_1 Outage_{t,p} + \beta_2 Claude_t + \gamma Controls_{i,t} + \\ & + \theta_i + \eta_p + \varepsilon_{i,t,p} \end{aligned} \tag{1}$$

where i, t, and p indicate firm, date, and 5-minute interval, respectively. $Trading\ Volume_{i,t,p}$ is defined as the total number of shares traded in the 5-minute interval p of day t, divided by the total number of shares outstanding for firm i on day t, then multiplied by 10,000 for clearer presentation. Scaling by shares outstanding ensures comparability across securities with different share structures. The primary independent variable, $Outage_{t,p}$, equals one if time interval p of day t falls within an outage period, and zero otherwise. The dummy variable $Claude_t$ equals one for observations associated with Claude outage events (including the associated control period) and zero for ChatGPT outage events.

The coefficient on $Outage_{t,p}$, β_1, captures the average effect of ChatGPT and Claude outages on trading volume by comparing outage intervals to the same intraday intervals on the preceding five trading days, assuming a homogeneous effect across GenAI systems. The coefficient on $Claude_t$, β_2, captures differences in baseline trading activity between ChatGPT and Claude events, ensuring that systematic variation in trading levels across the two sets of events is absorbed. If GenAI outages disrupt investors' trading decisions, we should observe a decrease in trading volume during the outage period and thus a negative coefficient on $Outage_{t,p}$.

Since the baseline specification imposes a homogeneous outage effect across the two GenAI tools, we extend the model by interacting the outage indicator with the model-type indicator to allow for heterogeneity in outage effects across platforms:

$$\begin{aligned} Trading\ Volume_{i,t,p} = \quad & \beta_0 + \beta_1 Outage_{t,p} + \beta_2 (Outage_{t,p} \times Claude_t) + \\ & + \beta_3 Claude_t + \gamma Controls_{i,t} + \theta_i + \eta_p + \varepsilon_{i,t,p} \end{aligned} \tag{2}$$

where β_1 captures the effect of outages for ChatGPT, β_2 captures the differential effect of Claude outages relative to ChatGPT outages. A negative β_2 indicate that Claude outages are associated with lower trading volume relative to ChatGPT outages, while a positive coefficient indicates the opposite. In the absence of differential effects across models, we expect β_2 to be close to zero. The total effect of Claude outages is given by $\beta_1 + \beta_2$. We assess the statistical significance of this total effect using a Wald test of the linear combination of coefficients.

To further assess model-specific effects, we also estimate the baseline specification separately for subsamples of ChatGPT and Claude outage events. These subsample estimates complement the interaction specification in Equation (2) by providing direct model-specific effects based on Equation (1). In these regressions, the model-type indicator is constant and therefore omitted, and the model reduces to:

$$Trading\ Volume_{i,t,p} = \quad \beta_0 + \beta_1 Outage_{t,p} + \gamma Controls_{i,t} + +\theta_i + \eta_p + \varepsilon_{i,t,p} \tag{3}$$

where β_1 captures the effect of outages within each GenAI system in isolation.

Taken together, these specifications allow us to estimate both the average effect of GenAI outages and potential heterogeneity across platforms, while the subsample analyses provide complementary evidence on model-specific effects.

To ensure comparability with Cheng et al. (2025), we adopt their set of control variables that may affect intraday trading volume, to the above specifications: overnight stock returns (*Return Overnight*), return in the previous day (*Return [-1]*), cumulative returns from day *t-5* through day *t-2* (*Return [-5, -2]*), price range in the previous 60 days (*Price Range*), stock prices on day *t-1* (*Price*), and market capitalization on day *t-1* (*Market Cap*). Further, we winsorize all continuous variables, except stock returns and prices, at the 1% tails. Appendix D provides detailed definitions of all variables. We further include firm fixed effects (θ_i) to absorb any time-invariant heterogeneity across securities and 5-

minute interval fixed effects (η_p) to control for systematic intraday variation in market-wide trading activity. Standard errors are clustered at the firm level. Although prior studies (e.g., Cheng et al., 2025; Eaton et al., 2022) use two-way clustering by firm and date, their setting considers a single outage type. In our setting, outage events are split across two GenAI platforms and assigned at the event-date level. As a result, both *Outage* and the *Claude* indicator vary only across dates, and the effective variation is distributed across a limited number of event-date shocks, making date-level clustering unreliable (Petersen, 2009). Clustering at the firm level therefore provides more reliable inference than two-way clustering in this setting.

3.2.2. Sensitivity tests

To ensure robustness, we conduct a series of sensitivity tests building on the procedure of Cheng et al. (2025). First, we adopt two alternative measures of trading volume, the total number of shares traded (*Shares Count*) and the total number of trades (*Trades Count*) in each 5-minute interval and estimate the model using a Poisson Pseudo Maximum Likelihood (PPML) regression to account for the count-based nature of these variables. Second, we employ three sets of more stringent fixed effects: (1) outage × firm and outage × 5-minute interval fixed effects to absorb any firm- or time-of-day effects that vary across outage events; (2) outage × firm × 5-minute interval fixed effects to identify the outage effect purely from day-to-day variation within a given firm, outage event, and 5-minute interval; (3) day-of-week × firm and day-of-week × 5-minute interval fixed effects to additionally control for day-of-week patterns in trading activity. Third, we assess the sensitivity of our results to seasonal variation by implementing alternative sets of seasonal fixed effects, estimating the model with event (outage-date) fixed effects, and with calendar-month fixed effects. Fourth, to verify that our findings are not sensitive to the choice of control period, we re-estimate the baseline model with alternative control windows, using trading days [t-1, t-4], trading days [t-2, t-5], and a tighter two-day window covering trading days [t-1, t-2]. Fifth, to ensure that our results are not driven by outliers or by the inclusion of unusually short or long outage events, we estimate the baseline model using median regression, exclude observations with residuals in the top and bottom one percent and two-and-a-half percent, and sequentially drop the shortest, two shortest, and three shortest, as well as the longest, two longest, and three longest, outage events from the analysis. Sixth, we conduct a falsification test using the day preceding each outage event as a pseudo-outage date, with the four trading days prior to the pseudo-outage date (i.e., days [t-2, t-5]) serving as the corresponding control period. If our baseline results are driven by genuine disruptions to GenAI availability rather than by date-specific market conditions surrounding the outage, we should not observe a similar decline in trading volume during the pseudo-outage periods. Finally, we re-estimate the baseline specification with standard errors two-way clustered by firm and calendar date, mirroring the approach of Cheng et al. (2025) and Eaton et al. (2022).

3.2.3. Mechanism analysis: Accounting news release

We conduct cross-sectional analysis to examine the mechanisms through which GenAI tools support investors' trading activities. The main analysis cannot distinguish whether outages affect trading activity through disruption to direct information processing or through more routine workflow and execution-related activities. Cheng et al. (2025) compile survey and anecdotal evidence suggesting that investors use GenAI primarily for information processing activities. If this channel is the more economically relevant, we would expect outages to have a larger impact on trading volume when investors have a stronger demand for information processing.

Following prior research (Cheng et al., 2025; Drake et al., 2012), we use the presence of corporate news releases as a proxy for such demand. Unlike Cheng et al. (2025), who draw on a broad set of news sources, we focus specifically on RNS announcements and further distinguish between accounting-relevant disclosures and general regulatory announcements. This distinction is theoretically motivated. Accounting disclosures such as earnings releases, profit warnings, and trading updates require investors to process numerical financial data and form updated valuations (Blankespoor et al., 2019, 2020), tasks for which GenAI tools may offer a particular advantage (Blankespoor et al., 2026; Cheng et al., 2025). General regulatory announcements, such as director dealings or administrative notices, are assumed to place lower demands on financial analysis and disclosure processing and are therefore expected to exhibit less dependence on GenAI availability.

We classify announcements using a keyword-based procedure applied to the headline text of each scraped announcement. Headlines are screened for inclusion keywords associated with periodic financial reporting and earnings guidance, including formal reporting terminology such as "final results", "interim results", and "preliminary results", as well as forward-looking earnings communications such as "trading update", "profit warning", and "pre-close trading". Exclusion keywords remove scheduling notices (for example, headlines beginning with "notice of") and sector-specific uses of the word "results" unrelated to financial reporting (for example, clinical trial or drilling results). Appendix C.2 provides the full list of classification keywords. A binary indicator variable, *News Accounting*, is assigned a value of one for all firm-interval observations where the firm released at least one qualifying accounting announcement on that calendar date before or during the outage period, and zero otherwise. In parallel, *News Non-Accounting* takes a value of one when the firm released at least one RNS announcement on that date before or during the outage whose headline does not match the accounting-keyword criteria, and zero otherwise. Where a firm releases both accounting and non-accounting announcements on the same day, the accounting indicator takes priority and *News Non-Accounting* is assigned a value of zero. Both variables are similarly defined for observations from control days. The two indicators are therefore mutually exclusive at the firm-day level and jointly span the RNS announcement universe, allowing us to test whether the outage effect on trading volume increases more broadly in the presence of firm-specific regulatory disclosures or becomes more pronounced specifically in the more analytically demanding setting of accounting disclosures.

To test this mechanism, we estimate the following extension of Equation (1):

$$\begin{aligned} Trading\ Volume_{i,t,p} = \quad & \beta_0 + \beta_1(Outage_{t,p} \times News\ Accounting_{i,t}) + \\ & + \beta_2(Outage_{t,p} \times News\ Non\ Accounting_{i,t}) + \\ & + \beta_3 Outage_{t,p} + \ \beta_4 News\ Accounting_{i,t} + \\ & + \beta_5 News\ Non\ Accounting_{i,t} + \ \beta_6 Claude_t + \\ & + \gamma Controls_{i,t} + \theta_i + \eta_p + \varepsilon_{i,t,p} \end{aligned} \tag{4}$$

The coefficients of primary interest are β_1 and β_2, on the interaction terms *Outage* × *News Accounting* and *Outage* × *News Non-Accounting*, respectively. If the coefficients of both interaction terms are negative and statistically significant, this indicates that GenAI dependence increases more broadly in the presence of firm-specific regulatory news, with a more negative β_1 being consistent with this dependence becoming stronger in the setting of accounting disclosures. If only β_1 is negative and significant, this would be consistent with GenAI dependence being more closely tied to the analytical demands of accounting disclosures. If β_2 is also statistically indistinguishable from zero, this would further suggest that the dependence is specific to accounting-related disclosure processing. We formally test the differential effect with a Wald test of $\beta_1 = \beta_2$.

In addition, β_3 captures the baseline outage effect on days without RNS news, while β_4 *a*nd β_5 capture the general effect of accounting and non-accounting RNS news on trading volume in the absence of an outage. The total outage effect on accounting and non-accounting days is given by $\beta_1 + \beta_3$ and $\beta_2 + \beta_3$, respectively. We assess the statistical significance of these total effects using Wald tests of the corresponding linear combinations.

We include firm and 5-minute interval fixed effects as well as the same set of controls as in Equation (1). Standard errors are clustered at the firm level.

We estimate the specification on the full pooled sample. As a robustness check, we follow Cheng et al. (2025) and deHaan et al. (2023) and estimate a saturated specification in which each control variable and fixed effect is interacted with both *News Accounting* and *News Non-Accounting*, allowing the relationship between firm characteristics and trading volume to differ across accounting-news, non-accounting-news, and non-news days. We further apply the falsification test described in Section 3.2.2 to the news mechanism specification, replacing the true outage indicator with the pseudo-outage date and re-estimating Equation (4). If the results reflect a genuine outage mechanism, the pseudo-outage coefficient and its interactions with both *News Accounting* and *News Non-Accounting* should all be statistically indistinguishable from zero.

3.2.4. Dynamic analysis

To better understand when GenAI outages begin to influence investors' trading activity, we conduct a dynamic analysis following the approach of Cheng et al. (2025) and adapt it to the UK setting. This analysis allows us to examine whether the impact on trading volume arises immediately at the start of an outage or with some delay, and whether any compensatory trading occurs once service is restored. The timing of the effect provides suggestive evidence on the stages of the investment workflow in which GenAI is embedded. A delayed response would be consistent with GenAI tools being used primarily for information processing rather than for direct trade execution, as investors require time to complete information-gathering tasks before translating them into trading decisions, but can continue to trade on information already processed prior to the outage (Cheng et al. 2025). By contrast, an immediate response would be more consistent with GenAI supporting direct trade execution.

We restrict the dynamic sample to outages that begin after 8:30 a.m. and end before 4:00 p.m., ensuring sufficiently long pre- and post-periods within the same trading day given LSE trading hours of 8:00 a.m. to 4:30 p.m.. This narrower inclusion criterion yields three eligible events, all Claude outages. Given the small subsample, this analysis is intended to be illustrative of the timing pattern rather than to provide precise causal estimates.

We estimate the dynamic specification by extending Equation (1), replacing the single outage indicator with a series of interaction terms between the outage dummy and indicators for each 10-minute event-time bin surrounding the outage, to which each underlying security-date-5-minute interval observation is assigned:

$$\begin{aligned} Trading\ Volume_{i,t,p} = \quad & \beta_0 + \sum_{\tau \neq -10} \beta_\tau \left(Interval_\tau \times Outage_{t,p} \right) + \beta_2 Claude_t + \\ & + \ \gamma Controls_{i,t} + \theta_i + \eta_p + \varepsilon_{i,t,p} \end{aligned} \tag{5}$$

where τ indexes 10-minute event-time bins relative to the outage start, with pre-outage bins $\tau \in \{-60, -50, \ldots, -10\}$ and during-outage bins $\tau \in \{0, 10, 20, \ldots, 60\}$ measured from outage start, and post-outage bins $\tau \in \{+10, +20, \ldots, +60\}$ measured from outage end. The first 10-minute interval before the outage (τ =-10) serves as the omitted reference category. As such, each β_τ measures the differential change in trading volume between outage and control days at event-time bin τ, relative to that same differential at the reference bin. While Equation (1) captures the outage effect as a single average coefficient β_1, Equation (5) extends this into a sequence of β_τ estimates, illustrating how the treatment effect changes over the full event window. We include firm and 5-minute interval fixed effects as well as the same set of controls as in Equation (1). Standard errors are clustered at the firm level. Note that because the sample is restricted to only Claude events, the *Claude* indicator is omitted from the specification due to perfect collinearity.

4. Analysis

This section presents the empirical results of the study. We begin with the descriptive statistics for our sample, followed by the main results on changes in trading volume during GenAI outages, both on average and across platforms. Following this, we assess the robustness of these findings through a series of sensitivity tests. We then examine the role of accounting news releases as a mechanism before turning to the dynamic analysis of how the outage effect evolves over the course of a disruption.

4.1. Descriptive statistics

Table 2. Descriptive statistics

This table reports the descriptive statistics of the variables used in the trading volume analyses. The sample includes 720,909 security-date-5-minute interval observations from 622 stock securities during the period from March 2025 to February 2026. All continuous variables except stock returns and prices, are winsorized at the 1% tails. See Appendix D for variable definitions.

Variable	N	Mean	SD	p25	Median	p75
Variables used in the main trading volume analyses						
Trading Volume	720,909	0.591	2.560	0.000	0.023	0.148
Outage	720,909	0.164	0.370	0.000	0.000	0.000
Claude	720,909	0.591	0.492	0.000	1.000	1.000
Return Overnight	720,909	0.002	0.023	-0.004	0.000	0.006
Return [-1]	720,909	0.002	0.030	-0.007	0.001	0.010
Return [-5, -2]	720,909	0.011	0.118	-0.010	0.005	0.023
Price Range	720,909	0.236	0.170	0.129	0.195	0.292
Price	720,909	5.896	1.385	5.061	5.897	6.788
Market Cap	720,909	25.284	1.876	24.134	25.144	26.330
Dependent variables used in additional analyses						
Shares Count	720,909	18,096	301,695	0	681	5,360
Trades Count	598,463	11	33	0	2	10

Note: Trades Count covers fewer observations than other variables, reflecting differences in LSEG's reporting of trade-level data relative to aggregated volume measures.

Table 2 reports the summary statistics of the variables. The mean of *Trading Volume* in a 5-minute interval is 0.591, representing 0.00591% of total shares outstanding. The distribution is right-skewed, with a median of only 0.023, reflecting that most intervals see very little trading activity even after imposing the liquidity filter, consistent with the comparatively lower liquidity of the LSE relative to the US equity markets previously studied (Cheng et al., 2025). The mean of *Outage* is 0.164, indicating that 16.4% of observations fall within outage periods, and the *Claude* indicator has a mean of 0.591, reflecting that 59.1% of observations stem from Claude outage events.

Table 3. Correlation matrix
This table reports pairwise Pearson correlations for the variables used in the trading volume analyses. The sample consists of 720,909 observations from March 2025 to February 2026. All continuous variables except stock returns and prices, are winsorized at the 1% tails. Variable definitions are provided in Appendix D. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
(1) *Trading Volume*	1.000										
(2) *Outage*	-0.007***	1.000									
(3) *Claude*	0.120***	-0.001	1.000								
(4) *Return Overnight*	0.019***	0.015***	0.074***	1.000							
(5) *Return [-1]*	-0.003**	-0.064***	0.010***	0.016***	1.000						
(6) *Return [-5,-2]*	0.042***	-0.013***	0.026***	0.121***	0.102***	1.000					
(7) *Price Range*	0.064***	0.003**	-0.140***	0.163***	0.072***	0.174***	1.000				
(8) *Price*	-0.054***	0.002	-0.022***	-0.195***	-0.031***	-0.075***	-0.343***	1.000			
(9) *Market Cap*	-0.250***	0.007***	-0.136***	-0.146***	-0.036***	-0.076***	-0.244***	0.596***	1.000		
(10) *Shares Count*	0.171***	-0.003***	-0.001	0.010***	-0.002*	0.004***	0.020***	-0.036***	0.034***	1.000	
(11) *Trades Count*	0.218***	-0.007***	-0.019***	-0.032***	-0.007***	-0.012***	-0.028***	0.180***	0.310***	0.142***	1.000

The control variables are broadly similar to those reported by prior research (Cheng et al., 2025). Table 3 reports pairwise Pearson correlations for the variables used in the analyses. *Outage* is negatively and significantly correlated with *Trading Volume*, consistent with the directional prediction of H1. The correlation matrix does not indicate any meaningful multicollinearity concerns, with all pairwise correlations between regressors remaining below 0.60 in absolute value.

4.2. Results and sensitivity tests

4.2.1. Changes in trading volume during GenAI outages

Table 4 reports the main results from Equation (1), which estimates the average effect of GenAI outages on trading volume, and Equation (2), which explores potential heterogeneity in this effect across platforms. Column (1) estimates Equation (1) with firm and 5-minute interval fixed effects only. The coefficient on *Outage* is significantly negative (t = -4.63). Column (2) further includes control variables. The coefficient on *Outage* remains significantly negative (t = -3.21). These results support H1 and indicate a statistically significant decline in trading volume during GenAI outage periods relative to the control periods. The effect is economically significant: the coefficient on *Outage* (-0.036) in Column (2) suggests a relative decrease of 1.41% (= -0.036 / 2.560) of the sample standard deviation of *Trading Volume*.

Columns (3) and (4) report the results from Equation (2), examining whether the outage effect on trading volume differs across ChatGPT and Claude events. In Column (3), which only includes firm and 5-minute interval fixed effects, the coefficient on *Outage*, which captures the effect of ChatGPT outages, is negative and statistically significant (t = -2.97), and the interaction term *Outage* × *Claude* is negative but not statistically significant (t = -0.52). The total effect of Claude outages, given by the sum of these coefficients, is negative and statistically significant (t = -3.71). In Column (4), which additionally includes control variables, the coefficient on *Outage* remains negative but statistically insignificant (t = -0.46), while the interaction term *Outage* × *Claude* remains negative and becomes significant (t = -2.27). The total effect of Claude outages remains negative and statistically significant (t = -3.26). This indicates that the outage effect on *Trading Volume* differs significantly across models, with the effect being significantly larger for Claude outages than for ChatGPT outages relative to their respective matched control intervals, while no statistically significant outage effect can be found for ChatGPT outages. Economically, the interaction coefficient (-0.049) in column (4) implies that the estimated trading volume effect of Claude outages is 1.91% (= -0.049 / 2.560) of the sample standard deviation more negative than that of ChatGPT outages. Correspondingly, the total effect of Claude outages (-0.055) corresponds to a reduction of 2.15% (= -0.055 / 2.560) of the sample standard deviation of *Trading Volume*. This suggests that the negative average effect observed in Equation (1) is primarily attributable to Claude outage events.

Table 4. GenAI outages and change in trading volume, pooled results

This table reports the OLS regression results of the analysis of the average impact of ChatGPT and Claude outages on trading volume at 5-minute intervals, and the differential outage effect between the two models, using Equation (1) and (2), respectively. The sample includes 720,909 security-date-5-minute interval observations from 622 stock securities during the period from March 2025 to February 2026. See Appendix D for variable definitions. Standard errors are clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Dependent Variable =	*Trading Volume*			
	(1)	(2)	(3)	(4)
Outage	-0.057***	-0.036***	-0.050***	-0.006
	(-4.63)	(-3.21)	(-2.97)	(-0.46)
Outage × *Claude*			-0.012	-0.049**
			(-0.52)	(-2.27)
Return Overnight		-1.206		-1.206
		(-1.07)		(-1.07)
Return [-1]		0.596		0.596
		(0.87)		(0.87)
Return [-5, -2]		0.541***		0.540***
		(4.42)		(4.42)
Price Range		1.031***		1.032***
		(3.25)		(3.26)
Price		0.412***		0.412***
		(2.61)		(2.61)
Market Cap		-1.061***		-1.062***
		(-9.49)		(-9.49)
Claude	0.717***	0.349***	0.719***	0.357***
	(7.73)	(6.27)	(7.79)	(6.39)
Firm FE	YES	YES	YES	YES
Time Interval FE	YES	YES	YES	YES
N	720,909	720,909	720,909	720,909
Adjusted R^2	0.204	0.368	0.204	0.368
Total outage effect ($\beta_{Outage} + \beta_x$)				
Claude			-0.062***	-0.055***
			(-3.71)	(-3.26)

Note: Statistical significance for the reported total outage effect is assessed using a Wald test.

Table 5 reports the results from Equation (3), estimating the outage effect on trading volume separately for ChatGPT and Claude events. Columns (1) and (2) present the ChatGPT subsample without and with controls, where the coefficient on *Outage* is negative but statistically insignificant (t = -0.84 and t = -1.55, respectively). In contrast, Columns (3) and (4), which report the Claude subsample without and with controls, show a consistently negative and statistically significant coefficient on *Outage* (t = -3.68 and t = -3.13, respectively). These results confirm the absence of a statistically significant effect for ChatGPT outages and a significant decline in trading volume during Claude outage periods relative to the matched control periods.

Table 5. GenAI outages and change in trading volume, model specific results

This table reports the OLS regression results of the analysis of the impact of ChatGPT and Claude outages on trading volume at 5-min intervals, estimated on model specific subsamples, using Equation (3). See Appendix D for variable definitions. Standard errors are clustered at the firm level. ***, **, and * indicate significance at the 1 %, 5 %, and 10 % levels, respectively, based on two-sided tests.

Dependent Variable =	*Trading Volume*			
	ChatGPT		Claude	
	(1)	(2)	(3)	(4)
Outage	-0.011	-0.013	-0.062***	-0.053***
	(-0.84)	(-1.55)	(-3.68)	(-3.13)
Return Overnight		-1.186		-1.173
		(-0.88)		(-0.69)
Return [-1]		-0.513		0.995
		(-0.49)		(1.39)
Return [-5, -2]		0.841***		0.438***
		(2.66)		(6.93)
Price Range		0.334*		1.110**
		(1.88)		(2.25)
Price		0.617***		-0.231
		(3.05)		(-0.79)
Market Cap		-0.680*		-1.016***
		(-1.94)		(-9.42)
Firm FE	YES	YES	YES	YES
Time Interval FE	YES	YES	YES	YES
N	295,053	295,053	425,856	425,856
Adjusted R^2	0.231	0.236	0.321	0.438

4.2.2. Sensitivity tests

Table 6 reports a series of sensitivity tests of the main results reported in Columns (2) and (4) of Table 4. Column (1) reports the coefficient on *Outage* from Equation (1), with its corresponding test statistics in Column (2). Columns (3) and (4) report the coefficient on *Outage* × *Claude* interaction term from Equation (2) and its corresponding test statistic. This holds for all panels except in Panel F, where Columns (1) and (3) report the coefficients on *Pseudo-Outage* and *Pseudo-Outage* × *Claude*, respectively. The results are largely robust across all specifications, with some variation in Panel A and no significance in Panel G.

Panel A yields mixed results across alternative measures of trading activity. For *Shares Count*, neither the *Outage* coefficient from Equation (1) nor the *Outage* × *Claude* interaction reaches significance, although the direction remains the same as in the main specification. For *Trades Count*, both the *Outage* coefficient from Equation (1) and the *Outage* × *Claude* interaction remains negative and significant at the 1% level.

Panel B shows that the findings hold under more stringent fixed effects. The *Outage* coefficient from Equation (1) remains negative and significant at the 1% level across all three fixed effects structures. The *Outage* × *Claude* interaction remains negative and significant at the 1% and 5% levels under outage × firm and outage × time-interval fixed effects, though it loses significance under day-of-week fixed effects.

Panel C confirms the results are not sensitive to seasonal fixed effects. Both the *Outage* coefficient from Equation (1) and the *Outage* × *Claude* interaction remains significant and quantitatively similar.

Panel D shows the findings are insensitive to the choice of control window. Across all three alternative control periods, both the *Outage* coefficient from Equation (1) and the *Outage* × *Claude* interaction remains negative and significant.

Panel E shows that the findings are not driven by outliers or extreme outage events. Median regression preserves both the *Outage* effect from Equation (1) and the *Outage* × *Claude* interaction, although both weaken. Excluding residual outliers at the 1% and 2.5% tails maintains the significance of the *Outage* coefficient from Equation (1), though the interaction term loses significance, and again, both weaken. Sequentially dropping the shortest outage events leaves both coefficients significant and largely unchanged, while excluding the longest outage events strengthens both the main effect from Equation (1) and the interaction.

Panel F rules out any date-specific market-wide conditions surrounding the outage days as a driver of the main results. Replacing actual outage dates with pseudo-outages yields no significant coefficients on either *Pseudo-Outage* or *Pseudo-Outage* × *Claude*.

Panel G reports two-way clustering by firm and date. Point estimates are unchanged, but test statistics decrease and both the *Outage* effect from Equation (1) and the *Outage* × *Claude* interaction lose significance.

Table 6. GenAI outages and change in trading volume, sensitivity tests

This table presents the sensitivity tests for the average impact of ChatGPT and Claude outages on trading volume and the differential outage effect between the two models, as reported in Columns (2) and (4) of Table 4. The baseline results refer to the specification in Column (2) and (4) of Table 4. For conciseness, we only report the coefficients on *Outage* and *Outage* × *Claude*, except in Panel F, where we report the coefficient on *Pseudo-Outage* and *Pseudo-Outage* × *Claude*. Panel A reports the PPML regression using two alternative measures of trading volume (*Shares Count* and *Trades Count*). Panel B reports the OLS regression results when controlling for alternative (stricter) fixed effects. Panel C reports the results using seasonal fixed effects. Panel D reports the results using alternative control periods. Panel E reports the results using a median regression or using subsamples. Panel F represents the falsification test using day *t-1* as the pseudo-outage date and days *t-2* through *t-5* as the corresponding control days. Panel G reports the results using two-way clustering. See Appendix D for variable definitions. Standard errors are clustered at the firm level, except in Panel G where it is clustered on firm and date. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

	(1)	(2)	(3)	(4)
	Coeff. on Outage	*t/z-statistic*	*Coeff. on Outage × Claude*	*t/z-statistic*
Baseline results (Table 4, Columns (2) and (4))	-0.036***	-3.21	-0.049**	-2.27
Panel A: Alternative dependent variables				
Shares Count (PPML regression)	-0.077	-0.99	-0.259	-1.53
Trades Count (PPML regression)	-0.046***	-2.58	-0.102***	-2.52
Panel B: Alternative (stricter) fixed effects				
Outage × Firm FE and Outage × Time-Interval FE	-0.034***	-3.23	-0.056***	-2.63
Outage × Firm × Time-Interval FE	-0.034***	-2.95	-0.056**	-2.40
Day-of-Week × Firm FE and Day-of-Week × Time-Interval FE	-0.040***	-3.09	-0.023	-1.04
Panel C: Seasonal fixed effects				
Interval FE + Event FE	-0.038***	-3.41	-0.043**	-2.01
Interval FE + Calendar-Month FE	-0.038***	-3.41	-0.043**	-2.00
Panel D: Alternative control periods				
Trading days [t-1, t-4]	-0.031***	-2.80	-0.050**	-2.19
Trading days [t-2, t-5]	-0.039***	-3.20	-0.054**	-2.27
Trading days [t-1, t-2]	-0.026**	-2.13	-0.059**	-2.54
Panel E: Alternative model or using subsamples				
Median regression	-0.003**	-2.10	-0.010***	-3.04
Excluding observations with residuals at the 1% tails	-0.013**	-2.31	-0.004	-0.35
Excluding observations with residuals at the 2.5% tails	-0.007**	-2.47	-0.001	-0.15
Excluding the shortest outage	-0.037***	-3.25	-0.046**	-2.07
Excluding the two shortest outages	-0.036***	-3.08	-0.050**	-2.22

(*Table 6 continued on next page*)

(*Table 6 continued*)

	(1) *Coeff. on Outage*	(2) *t/z-statistic*	(3) *Coeff. on Outage × Claude*	(4) *t/z-statistic*
Excluding the three shortest outages	-0.036***	-3.09	-0.053**	-2.28
Excluding the longest outage	-0.041***	-2.77	-0.076***	-3.76
Excluding the two longest outages	-0.060***	-2.76	-0.105***	-3.70
Excluding the three longest outages	-0.071***	-3.05	-0.106***	-3.03
Panel F: Falsification test				
Using same periods of the previous trading day as pseudo-outage periods	-0.013	-1.19	-0.024	-1.15
Panel G: Two-way clustering				
Firm + date clustering	-0.036	-0.96	-0.049	-0.95

4.2.3. Mechanism analysis: Accounting news release

Table 7. Descriptive statistics of news variables

This table reports the descriptive statistics for the news variables used in the news mechanism analysis. Panel A presents the number of firm-day observations with any RNS news (*News All*), accounting RNS news (*News Accounting*), and non-accounting RNS news (*News Non-Accounting*) released on the same calendar date as the outage or matched control day and before or during the relevant outage period, along with their shares relative to all RNS news and their occurrence during GenAI outage days. Panel B reports summary statistics for the same variables as well as their interactions with the Outage indicator at the 5-minute interval level. The sample includes 720,909 security-date-5-minute interval observations from 622 stock securities over the period March 2025 to February 2026. All variables are defined in Appendix D.

Panel A: Firm-day news composition and overlap with outage days

	# Firm-day N	% of News All	# Firm-day N during outage days	% of category on outage days
News All	6,488	100%	1,148	17.8%
News Accounting	230	3.6%	38	16.5%
News Non-Accounting	6,218	96.4%	1,110	17.9%

Note: Of the 230 *News Accounting* firm-days, 133 (57.8%) are pure accounting releases and 97 (42.2%) coincide with at least one non-accounting RNS release on the same day.

Panel B: Descriptive statistics

Variable	N	Mean	SD	p25	Median	p75
News All	720,909	0.270	0.444	0.000	0.000	1.000
News Accounting	720,909	0.009	0.095	0.000	0.000	0.000
News Non-Accounting	720,909	0.261	0.439	0.000	0.000	1.000

(*Table 7 continued on next page*)

(*Table 7 continued*)

Variable	N	Mean	SD	p25	Median	p75
Outage × News All	720,909	0.045	0.208	0.000	0.000	0.000
Outage × News Accounting	720,909	0.002	0.040	0.000	0.000	0.000
Outage × News Non-Accounting	720,909	0.044	0.204	0.000	0.000	0.000

Table 7 reports descriptive statistics for the news variables. Panel A shows that, among firm-day observations with regulatory news, non-accounting news constitutes the vast majority at 96.4%, while accounting news represents only 3.6%. Approximately 17.8% of news events occur on outage days, with similar proportions across accounting and non-accounting news. Panel B shows that news is present in 27.9% of security-date-5-minute interval observations, while accounting news is relatively rare, occurring in less than 1% of observations.

Table 8 reports the regression results from Equation (4), which estimates the differential impact of GenAI outages on trading volume conditional on the presence of RNS announcements, further distinguished between accounting and non-accounting news. Column (1) presents the baseline specification including controls and fixed effects. The coefficient on *Outage* × *News Accounting* is significantly negative (t = -2.43), while the coefficient on *Outage* × *News Non-Accounting* is small and statistically insignificant (t = 0.00). A Wald test rejects the null that $\beta_1 = \beta_2$, confirming that the outage effect is significantly stronger on accounting-news days than on non-accounting-news days. The total outage effect on accounting news days, given by the sum of the *Outage* and *Outage* × *News Accounting* coefficients, is negative and statistically significant (t = -2.54), whereas the corresponding total effect on non-accounting news days is statistically insignificant (t = -1.05). The coefficient on *Outage* is significantly negative (t = -2.10) and confirms a baseline outage effect that applies to non-news days quantitatively similar to the main results. The coefficients on *News Accounting* and *News Non-Accounting* are both significantly positive (t = 5.30 and t = 2.10), with a larger magnitude for *News Accounting*.

The results suggest that the outage effect differs across news types and is amplified during periods with accounting news. In contrast, we do not find evidence of a significantly stronger outage effect on RNS announcement days more broadly relative to non-announcement days, as indicated by the insignificant *Outage* × *News All* coefficient reported in Table E1, nor do we observe a differential outage effect for non-accounting news days relative to non-RNS periods. This pattern is observed across both the ChatGPT and Claude subsamples (see Table E2).

Economically, the incremental outage effect on accounting-news days, given by the coefficient on *Outage* × *News Accounting* (-0.742) from Column (1), corresponds to an additional relative decline of 29.0% (= -0.742/2.560) of the sample standard deviation of

Trading Volume compared to non-announcement days. The total outage effect conditional on accounting news releases, combining the coefficients (-0.030 and -0.742) from Column (1), implies a relative decrease of 30.2% (= -0.772/2.560) of the sample standard deviation of *Trading Volume*.

Because 42.2% of firm-days with accounting news also feature coincident non-accounting RNS releases, we further distinguish between pure accounting releases and co-release days. The incremental outage effects are nearly identical across the two groups (-0.767 and -0.723 respectively), and we cannot reject equality between the coefficients. This suggests that the result is driven by accounting content rather than coincident non-accounting news (see results in Table E1).

Table 8. GenAI outages and change in trading volume: Conditional on accounting and non-accounting news releases

This table reports the OLS regression results examining whether the effect of GenAI outages on trading volume varies with the presence and type of RNS news, using Equation (4). Column (1) presents the baseline specification, while Column (2) reports the saturated specification. The sample includes 720,909 security-date-5-minute interval observations from 622 stock securities during the period from March 2025 to February 2026. See Appendix D for variable definitions. Standard errors are clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Dependent Variable =	*Trading Volume*	
	(1)	(2)
Outage × *News Accounting*	-0.742**	0.428
	(-2.43)	(0.83)
Outage × *News Non-Accounting*	0.000	-0.017
	(0.00)	(-0.46)
Outage	-0.030**	-0.028**
	(-2.10)	(-2.08)
News Accounting	1.054***	
	(5.30)	
News Non-Accounting	0.080**	
	(2.10)	
Claude	0.349***	0.324***
	(6.28)	(5.58)
Controls	YES	YES
Firm FE	YES	YES
Time Interval FE	YES	YES
Controls × News Category	NO	YES
Firm FE × News Category	NO	YES
Time Interval FE × News Category	NO	YES
N	720,909	720,909
Adjusted R^2	0.369	0.391

(*Table 8 continued on next page*)

(*Table 8 continued*)

Dependent Variable =	*Trading Volume*	
	(1)	(2)
Interaction-effect differences ($\beta_{\text{Outage} \times \text{A}} - \beta_{\text{Outage} \times \text{B}}$)		
News Accounting vs. *News Non-Accounting*	-0.743**	0.445
	(-2.44)	(0.86)
Total outage effect ($\beta_{\text{Outage}} + \beta_{\text{x}}$)		
News Accounting	-0.772**	0.399
	(-2.54)	(0.77)
News Non-Accounting	-0.030	-0.050
	(-1.05)	(-1.43)

Note: Statistical significance for reported interaction-effect differences and total outage effects is assessed using Wald tests.

Column (2) presents the saturated specification as a robustness check. The coefficient on *Outage* remains significant and qualitatively unchanged (t = -2.08). However, the coefficient on *Outage* × *News Accounting* loses statistical significance (t = 0.83) and reverses in sign relative to Column (1), indicating that the outage effect conditioned on accounting news does not hold once controlling for news-interacted fixed effects and controls. The *Outage* × *News Non-Accounting* remains statistically insignificant, and neither total outage effect is statistically distinguishable from zero.

Table 9. Falsification test: Pseudo-outage news specification

This table presents the falsification test of the impact of GenAI outages on trading volume conditional on the presence and type of RNS news, as reported in Column (1) in Table 8. The test replaces the *Outage* indicator with a *Pseudo-Outage* indicator, using day *t-1* as the pseudo-outage date and days *t-2* through *t-5* as the corresponding control days. Standard errors are clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Dependent Variable =	*Trading Volume*
	(1)
Pseudo Outage × *News Accounting*	0.021
	(0.06)
Pseudo Outage × *News Non-Accounting*	-0.014
	(-0.38)
Pseudo Outage	-0.010
	(-0.73)
Controls	YES
Firm FE	YES
Time Interval FE	YES
N	602,742
Adjusted R^2	0.370

Table 9 reports the results from a falsification test of the news-conditioned specification reported in Column (1) of Table 8. The test rules out any date-specific market-wide conditions surrounding outage days as a driver of the news mechanism results. Replacing actual outage dates with pseudo-outages yields no statistically significant coefficients on either the *Pseudo-Outage* indicator or its interactions with *News Accounting* and *News Non-Accounting*.

4.2.4. Dynamic analysis

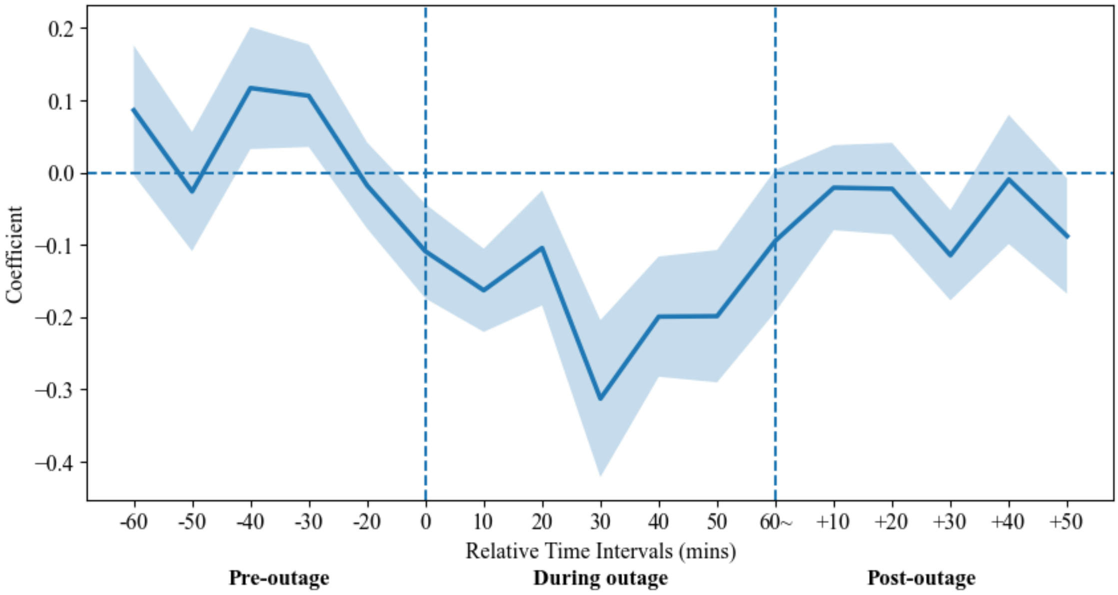

Figure 2. Changes in trading volume surrounding GenAI outages

This figure illustrates estimated outage effects on trading volume in the 10-minute windows surrounding GenAI outages. We focus on the three outages that begin after 8:30 a.m. and end before 4:00 p.m. to ensure sufficient pre- and post-periods within trading hours of the same day. This limits us to only Claude outage events. The solid line represents the estimated changes in trading volume at each 10-minute window surrounding the outage, derived from the interaction term in Equation (5), relative to the 10-minute interval immediately before the outage started (τ = -10), which serves as the baseline. The shaded area indicates 90% confidence intervals. The left and right vertical dash lines mark the start and end times of the outages, respectively. See Table E3 for full regression results.

Figure 2 illustrates the dynamic outage effects estimated from Equation (5), showing how the effect of GenAI outages on trading volume evolves over the course of a disruption. The negative effect on trading volume becomes significant immediately at outage onset, with during-outage coefficients being negative and significant at the 1% or 5% level up to 50 minutes post outage start time. The effect becomes stronger over the course of the outage and reaches its largest magnitude after approximately 30 minutes. Pre-outage coefficients show no systematic directional trend, though *Pre 40* and *Pre 30* are individually positive and significant. Post-outage coefficients are consistently negative but mostly small in magnitude and insignificant, with the exception of *Post 30* and *Post 50* which reach significance. See Table E3 for full regression results.

5. Discussion

In this section, we discuss our main empirical findings: (1) GenAI outages are associated with a significant decline in trading volume, with the effect differing systematically across platforms and primarily driven by Claude outages; (2) the outage effect is amplified on accounting disclosure days, suggesting that information processing is the more economically relevant channel through which GenAI is used; and (3) trading volume declines immediately at outage onset, with the effect intensifying before partially reversing, a pattern consistent with GenAI being embedded across multiple stages of the investment workflow. We also discuss our data and methodology, including choices made and limitations.

5.1. Main findings

5.1.1. GenAI outages reduce trading volume, driven by Claude outages

The main results support H1. GenAI outages are associated with a statistically significant decline in trading volume relative to matched non-outage intervals of the same security, indicating that investors utilise these tools in ways that influence their trading decisions. The finding is robust to the inclusion of controls and alternative specifications. In economic terms, the drop in trading volume corresponds to a relative decline of approximately 1.4% of the sample standard deviation, indicating a modest reduction in trading activity during GenAI outages.

This finding is broadly consistent with Cheng et al. (2025) and suggests that the impact of GenAI on investor trading is not confined to an early-adoption, single-platform setting, but extends to a more mature multi-tool environment in which investors continue to rely on GenAI tools in economically meaningful ways. The result is consistent with exogenous disruptions to GenAI access tightening investors' effective information processing capacity (Blankespoor et al., 2020; Sims, 2003), with corresponding effects on trading activity. This disruption may operate through GenAI's direct support of the disclosure processing chain, compressing awareness, acquisition, and integration costs (Blankespoor et al., 2019, 2020; Cheng et al., 2025), or through its support of complementary workflow activities consistent with the task complementarity effect (Acemoglu, 2025; Cheng et al., 2025). The main analysis cannot distinguish between these channels, and the observed decline in trading volume is consistent with either mechanism or a combination of both.

Our discussions with finance professionals shed light on this variation in how GenAI is used and relied upon. One professional described using GenAI daily to process reports and compile analyst summaries, noting that an outage would require reverting to manual summarisation, possible but considerably more effortful. Another described GenAI use at his institution as confined to routine administrative tasks due to internal restrictions, with core trading decisions unaffected by tool availability. This variation in usage and reliance across professionals and institutions mirrors the moderate average GenAI outage effect documented in the main analysis.

The average GenAI outage effect, however, is not evenly distributed across platforms. The evidence indicates that the effect differs systematically across GenAI systems, with a significant decline in trading volume during Claude outages, while ChatGPT outages do not show a detectable effect in this setting. This pattern is consistent across both the interaction specification and the model-specific subsample regressions. In economic terms, the outage effect for Claude corresponds to an additional reduction in trading activity of approximately 1.9% of the sample standard deviation relative to ChatGPT outages. The total effect of Claude outages corresponds to a decline in trading activity of just over 2% of the sample standard deviation, indicating a somewhat larger reduction in trading activity than implied by the average GenAI outage effect. Taken together, this suggests that the average effect of GenAI outages is primarily driven by Claude outages.

We consider two possible explanations for this pattern. The first relates to Claude's positioning within professional financial workflows. For a GenAI outage to produce a detectable effect on trading volume, the disrupted platform must be both sufficiently adopted among active investors and sufficiently embedded in their workflows that it cannot easily be substituted during the disruption. Following the launch of Claude for Financial Services in mid-2025 and its increasing enterprise adoption, Claude may be particularly well-suited for financial analytical tasks. This, in turn, may lead to greater adoption among active investors. Further, a platform embedded in activities such as investment research, financial modelling, and compliance review is likely harder to substitute away from during an outage than one used more interchangeably alongside other tools. Our results cannot establish which of these mechanisms plays the larger role, but they do suggest that disruptions to Claude are associated with a more pronounced effect in trading activity than disruptions to ChatGPT in this setting. The second explanation relates to the design of the study rather than to any genuine difference in how investors use the two platforms. In particular, the outage events in our sample differ in severity across platforms, with Claude outages generally being more disruptive. This may contribute to the observed difference and should be considered when interpreting platform-specific effects, as discussed in Section 5.2.

Relative to Cheng et al. (2025), the outage effects documented in this study are economically smaller. One possible explanation is that investors in 2025-2026 can, in principle, substitute across GenAI platforms during outages in ways that were less feasible in the earlier single-platform-dominated environment studied by Cheng et al. (2025). To the extent that such substitution occurs, the effects estimated in this study may understate overall investor reliance on GenAI more broadly, and a simultaneous disruption across multiple platforms would likely produce a larger effect. The fact that a significant trading volume effect nonetheless remains detectable shows that substitution across platforms is not frictionless, and that reliance on at least one GenAI platform has become sufficiently embedded in investor workflows to affect trading activity even in a mature multi-tool environment.

The average outage effect and the differential effect across platforms are broadly robust across the set of sensitivity specifications presented in Table 6, holding under more stringent fixed effects (Panel B), seasonal fixed effects (Panel C), and alternative control windows (Panel D), and are not explained by date-specific market conditions as confirmed by the falsification test (Panel F). Further, the results are not driven by outliers or extreme outage events (Panel E), as median regression and outlier exclusions preserve the direction of the effects, while dropping both the shortest and longest outage events leaves the results largely unchanged, with some strengthening when excluding the longest events.

At the same time, some variation across specifications should be noted. Alternative trading volume measures yield more mixed results (Panel A, Table 6). For *Trades Count*, both the *Outage* coefficient and the *Outage* × *Claude* interaction term remain statistically significant at the 1% level, consistent with the baseline findings, although the regressions are estimated on a smaller sample due to data availability constraints. For *Shares Count*, however, neither coefficient remains significant, although the directional pattern is preserved. This may reflect the fact that raw share counts retain variation that the baseline specification normalises for through scaling. Under day-of-week fixed effects (Panel B, Table 6), the interaction term loses significance, indicating some sensitivity of the platform-differential effect to this more restrictive specification. The main sensitivity arises in Panel G of Table 6, where two-way clustering by firm and date causes both the main effect and the differential effect to lose statistical significance, despite similar point estimates. This is a meaningful robustness consideration, and the primary inference therefore rests on firm-clustered standard errors, which we argue are more appropriate given the structure of the identifying variation. We return to this in Section 5.2.

5.1.2. The outage effect is amplified on accounting disclosure days

The mechanism analysis suggests that the effect of GenAI outages on trading volume differs systematically across news types. Specifically, the outage effect is amplified on days with accounting news relative to the outage effect on days without RNS announcements, and this amplification is significantly larger than for non-accounting regulatory news. At the same time, we do not find clear evidence of a differential outage effect for non-accounting regulatory news days relative to non-RNS announcement days and therefore do not draw conclusions for this category. Economically, the total outage effect on accounting-news days corresponds to a decline in trading activity of approximately 30% of the sample standard deviation, almost entirely attributable to the incremental outage effect associated with accounting disclosures, given that the corresponding outage effect on non-RNS announcement days is approximately 1.2% of the sample standard deviation by comparison. This indicates that the concurrence of outages and accounting disclosures generates materially stronger disruptions to trading activity.

These results refine the evidence of Cheng et al. (2025), who document a stronger outage effect in the presence of firm-specific corporate news and interpret this as evidence that GenAI supports investors' processing of firm-specific information more broadly. By distinguishing between accounting and non-accounting regulatory news, our findings add precision to this interpretation. While both accounting and non-accounting regulatory news are firm-specific, only accounting disclosures generate a detectable amplification of the outage effect. This suggests that not all firm-specific content requires GenAI in equal measure, rather, investors' reliance on GenAI for information processing is more closely tied to settings where analytical demands are high enough to make it a meaningful input. Non-accounting regulatory news, though equally firm-specific, does not appear to generate the same processing demands, and accordingly does not produce the same pattern of reliance.

Practitioner observations point in the same direction. One of the finance professionals we spoke with described the time pressure around earnings announcements as particularly acute, noting that prices adjust the moment corporate news is released, making speed in processing decisive and GenAI's absence most costly in exactly these settings.

One potential concern is that the absence of a differential outage effect for non-accounting regulatory news days relative to non-RNS announcement days reflects a lack of trading relevance rather than differences in information processing demands. However, both accounting and non-accounting news are associated with significantly higher trading volume in the absence of outages, indicating that both types of disclosures are relevant for trading activity. The absence of an additional outage effect for non-accounting news therefore does not reflect a lack of market relevance but rather suggests that these disclosures do not rely on GenAI-supported processing to the same extent as accounting news.

More broadly, the news-type heterogeneity provides suggestive evidence on the mechanism underlying the main outage effect. The main analysis alone cannot distinguish whether outages affect trading activity primarily through disruption to direct information processing or through more routine workflow and execution-related channels. However, accounting disclosures are precisely the settings where information processing constraints are most likely to bind (Blankespoor et al., 2019, 2020; Sims, 2003). If the outage effect operated predominantly through execution or workflow channels, which do not vary systematically with disclosure complexity, no differential amplification would be predicted across news types. The concentration of the amplification on accounting news days, where processing demands are highest, is therefore genuinely informative, it points toward information processing as the more economically relevant channel through which GenAI affects investors' trading-related activity, even if the precise mechanism within that channel remains unidentified. The small but significant baseline outage effect on non-RNS announcement days does not contradict this interpretation. It is consistent with task complementarity or with the processing of firm-specific, industry-, or market-wide news not captured by the RNS classification.

The mechanism evidence should, however, be interpreted with caution. In the saturated specification the interaction between outages and accounting news becomes statistically insignificant and reverses signs. We attribute this to data constraints as simultaneous outage and accounting news observations account for approximately 0.2% of the sample, leaving the saturated framework with insufficient variation to identify the interaction reliably. Additional sensitivity tests confirm that the result is not driven by coincident non-accounting news and does not appear under pseudo-outages dates. The keyword-based classification of accounting disclosures may nonetheless introduce measurement error in borderline cases.

Taken together, the results are consistent with GenAI supporting investors' processing of analytically demanding disclosures and give more support to an information processing mechanism than to a task complementarity explanation, though they do not allow for definitive conclusions about the underlying channel within that mechanism. At the same time, the evidence should be interpreted as suggestive rather than conclusive given the limited number of simultaneous outage and accounting-news observations, which constrains the precision with which the magnitude of the amplification effect can be estimated.

5.1.3. Trading declines immediately, intensifying before partially reversing

The dynamic analysis provides suggestive evidence that GenAI outages produce an immediate and sustained reduction in trading activity. Trading volume declines immediately at the start of the outage and remains negative, with the effect becoming stronger over the first 30 minutes, where it reaches its largest magnitude. Though, given that the analysis is based on only three outage events, the patterns should be interpreted as illustrative rather than conclusive.

This timing pattern differs from Cheng et al. (2025), who find that the effect of ChatGPT outages becomes significant only after approximately 30 minutes and interpret this delay as evidence that GenAI is used primarily for information processing. In contrast, we observe both an immediate decline and a subsequent deepening. This pattern does not align with a purely information processing constraint explanation, under which trading would be expected to decline only with a delay, as investors can still trade on information processed prior to the outage but are unable to process new information during the disruption. Instead, the pattern is consistent with GenAI being embedded across multiple stages of the investment workflow, from pre-trade information processing to execution-stage activities and mid-workflow tasks. One potential explanation for this difference relative to Cheng et al. (2025) is that our dynamic results are driven by Claude outages in a later time-period, whereas Cheng et al. study ChatGPT outages in early 2023. The difference may therefore reflect changes in how investors use GenAI over time, differences across platforms, or both, and our setting does not allow us to isolate these explanations.

Further, the pre-outage coefficients show no systematic directional trend, providing reassurance that trading activity on outage days was not already declining before the disruption began. Although the coefficients at 40 and 30 minutes prior to the outage are individually positive and significant, they do not form part of a consistent pattern. This weakens a reverse-causality interpretation in which elevated trading activity causes outages through increased system load, as such a mechanism would predict a sustained build-up in trading volume prior to the disruption rather than two isolated positive coefficients surrounded by insignificant ones. For the subset of outages where the cause of disruption is disclosed, reverse causality can also be ruled out more directly, as the disruptions are attributed to internal technical failures such as configuration errors and capacity issues, unrelated to user-side demand.

Finally, we find no clear evidence of compensatory trading after the outage ends. Post-outage coefficients remain negative and are generally small and statistically insignificant, suggesting that foregone trades are not fully made up once service is restored. Perhaps reflecting that some trading opportunities are time-sensitive and cannot be recovered after a delay.

Taken together, the dynamic evidence is consistent with investors relying on GenAI tools in both information processing and related professional workflows.

5.2. Limitations in data and method

The greatest vulnerability in our study lies in the measurement of outage events as shocks to investors' access to GenAI tools. Three issues arise. First, the timing, severity, and scope of any individual event are difficult to pin down precisely. Official reports may lag the true onset and lead the true resolution, severity classifications vary across providers, and the precise scope of individual events is not always clearly defined. We address these challenges through the procedures set out in Section 3.1.2, though some residual imprecision remains. The retained sample also remains heterogeneous in duration, severity, time of day, day of week, and season, which Section 4.2.2 examines through sensitivity tests. Second, while we attempt to isolate disruptions affecting the user-facing conversational interfaces most relevant to investors, some retained events may also have affected related products or services within the same platform ecosystem, making it difficult to fully isolate the effects of conversational-interface outages alone. Third, we cannot observe which investors rely on ChatGPT or Claude, or how they use these tools. The design assumes that aggregate trading volume on outage days reflects the behaviour of an unobserved subset of users with unobserved use patterns. This is not a limitation confined to our study, but an assumption that is inherent to tool-disruption designs (e.g., Brown et al., 2015; Cheng et al., 2025; Eaton et al., 2022). The mitigations described above and the consistency of our results across the sensitivity tests support reasonable confidence in the headline finding. The third concern, regarding unobservable variation in user dependence, is more likely to bias estimates toward zero than away from it, since investors with no GenAI exposure contribute noise rather than signal. The average GenAI outage effect should therefore be read as a lower bound on the response among investors who actually rely on these tools.

The multi-tool nature of our setting raises concerns that do not arise in single-tool designs. Cheng et al. (2025) study ChatGPT in isolation, during a period of early adoption. As set out in Section 2.1, the GenAI market has since evolved into an environment in which ChatGPT and Claude operate as competing and complementary tools with overlapping user bases, and adoption continues to grow. This complicates inference about platform-specific reliance in two ways. First, investors may substitute between tools in ways we cannot directly observe, which we discuss in Section 5.1.1. Second, the available outage events also differ systematically in severity across platforms, which is the more material concern in our data. All six Claude events involve a *Major Outage*, while all four ChatGPT events are classified as *Partial Outage* or *Degraded Performance* only. The differential outage effect documented may therefore reflect the more disruptive nature of Claude events in our sample period, rather than a difference in user dependence. We cannot fully separate these explanations within the data available. As such, the platform-specific interpretation requires more caution, even though the direction of the average outage effect remains well-supported. Claims about whether Claude users respond more strongly than ChatGPT users should be treated with care, given the asymmetric severity of available events and the unobservable substitution between platforms.

Our setting also limits generalisability across markets and time. We study the LSE rather than the US market examined by Cheng et al. (2025). As established in Section 3.1.4, the LSE is structurally less liquid than US markets, and the liquidity filter applied removes 480 securities from the initial sample. Thus, the residual consists of more liquid securities and may differ from the full LSE universe in ways that affect results of GenAI dependence. The composition of the LSE investor base may also shape our results, as outage effects scale with the level of GenAI dependence in the underlying investor population. Our findings are therefore most informative for liquid LSE securities, and may not generalise cleanly to thinly traded stocks or to other markets where investor composition may differ.

Further, the twelve-month window from March 2025 to February 2026 covers a specific phase of GenAI adoption. As set out in Section 2.1, GenAI adoption among investors continued to evolve through this period and is likely to keep changing. This evolution also creates heterogeneity within our sample, since later outages likely affected a larger and potentially different investor base than earlier ones, particularly for Claude, whose adoption expanded materially over the period (Hammond, 2026). Our estimates therefore capture an average across a window of evolving investor dependence. The magnitude of the effect may shift further as adoption deepens and as competition between platforms develops, and our findings are most reliably read as evidence for the existence and direction of a GenAI outage effect on trading volume in this setting, rather than as a fixed estimate of its magnitude.

Lastly, statistical inference is constrained by two related limitations. The first is the small number of outage events. Our sample contains ten events, slightly more than Cheng et al. (2025), but split unevenly with six Claude and four ChatGPT events. Combined with the severity asymmetry described above, this limits the precision of platform-specific estimates. The dynamic analysis is the most constrained, as it relies on only three Claude outages. The second is sensitivity to the choice of standard error clustering. We cluster at the firm level (Section 3.2.1), while previous research (Cheng et al., 2025; Eaton et al., 2022) clusters two-way by firm and date, supported by larger samples. We report two-way clustering as a robustness check in Section 4.2.2. Point estimates are unchanged, but both the average outage effect and the platform differential effect lose statistical significance under this alternative. We attribute this to the limited within-date variation in our regressors and the small number of distinct outage dates in our multi-tool sample. The point estimate of the average GenAI outage effect is therefore well-identified and stable across most specifications, while its statistical precision is sensitive to the small event count. Findings that rest on smaller subsets of events, including the dynamic pattern and any standalone ChatGPT effect, should be treated as suggestive rather than definitive.

As a final note, all regressions report VIF values below 5, indicating no material multicollinearity concerns.

6. Conclusion

The growing use of GenAI in financial markets raises an empirical question that is important yet difficult to answer directly: *to what extent do investors rely on GenAI tools, and what do disruptions to these tools reveal about their role in financial markets?* While Cheng et al. (2025) provide the first direct evidence that investors rely on GenAI tools in ways such that disruptions to their access reduce trading volume in US equity markets, we examine whether this reliance persists in a more mature multi-tool environment and generalise across platforms and institutional settings.

Using unanticipated service disruptions to ChatGPT and Claude between March 2025 and February 2026, we document a statistically significant decline in trading volume during GenAI outages on the LSE, indicating that investors utilise these tools in ways that influence their trading decisions. The drop in trading volume corresponds to approximately 1.41% of the sample standard deviation. The effect differs systematically across platforms and is primarily driven by Claude outages, with ChatGPT outages producing no statistically significant effect across any specification. We further find suggestive, though not conclusive, evidence that the outage effect is materially stronger on accounting disclosure days than on days with other regulatory announcements or no regulatory news at all, suggesting that investor reliance on GenAI is more closely tied to analytically demanding disclosure settings, and pointing toward information processing as the more economically relevant channel through which GenAI is used. Finally, the dynamic analysis illustrates that trading activity declines immediately following an outage, with the effect intensifying before partially reversing, a pattern consistent with GenAI being embedded across multiple stages of the investment workflow.

Taken together, the findings suggest that investors rely on GenAI tools in economically meaningful ways and that these tools have become sufficiently embedded in investor workflows such that disruptions to a single platform produce observable effects on market activity, even in a competitive multi-tool environment where substitutes exist. At the same time, the results should be interpreted with appropriate caution. Platform-specific effects are partly limited by differences in outage severity and unobservable substitution across tools, and statistical precision remains constrained by the limited number of outage events. The findings are therefore more reliably interpreted as evidence of the existence and direction of a GenAI outage effect on trading activity in this setting, rather than as a precise estimate of its magnitude or a definitive account of the underlying mechanism. More broadly, the results suggest that GenAI has begun to function as part of the infrastructure through which investors process information and make trading decisions.

7. Future research

As the adoption of GenAI tools continues, understanding the scope, systemic implications, and underlying mechanisms of investor reliance on these tools becomes increasingly important. Building on the evidence presented in our study, we propose several directions for future research.

First, our analysis is confined to a single market and two platforms, leaving open whether the findings generalise beyond this setting. Extending the outage-based design to other widely used platforms such as Gemini and Microsoft Copilot would clarify whether the documented reliance reflects features of specific tools or a more general pattern of investor's integration of GenAI. Examining whether comparable effects arise in markets with varying levels of institutional participation and financial intermediation would similarly help determine how broadly the findings apply. How these effects evolve over time is a further dimension worth tracking, as the nature and intensity of reliance may shift as adoption deepens and the competitive environment matures. Repeated analysis across different periods would allow researchers to observe whether market sensitivity to outages grows or diminishes.

Second, our finding that the outage effect appears primarily driven by Claude rather than ChatGPT raises questions the current design cannot address. A key open question is whether investors treat platforms as close substitutes or whether different tools serve distinct functions within the investment process. If platforms are used for different parts of the workflow, for example more continuous analytical tasks versus context-dependent ones, disrupting one platform removes a specific capability rather than simply redirecting users to an alternative. The news mechanism analysis could naturally be extended to the platform level, but the limited number of concurrent outage and accounting news observations precluded a reliable test, making this a promising direction for future work with richer data. Examining this directly, through platform-level usage data or natural experiments involving platform-specific disruptions, would help establish whether the concentration of effects reflects switching costs, specialisation, or differences in user bases, and whether growing reliance on specific platforms constitutes a meaningful source of market fragility.

Third, the central limitation of the outage-based design is that it establishes economically meaningful investor reliance without identifying how investors use GenAI. Our discussion acknowledges that the observed trading effects are consistent with both direct information processing support and more routine workflow assistance. While the news-based heterogeneity results provide suggestive evidence more consistent with the former, the design cannot conclusively distinguish between these channels. Survey or field data linking individual investors to their actual usage patterns would allow researchers to identify whether reliance is concentrated in direct information processing or routine task execution, and within the former, at which stage of the disclosure processing chain GenAI generates the most value. This would help resolve the ambiguity left by the current study and shed light on where investor reliance on GenAI is most concentrated. Examining how effects vary across firm complexity and information environments, or differ systematically across investor types, are further open questions that the current design cannot address but that future work could explore.

References

Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, 40(121), 13–58. https://doi.org/10.1093/epolic/eiae042

Anthropic. (2026, February 12). *Anthropic raises $30 billion in Series G funding at $380 billion post-money valuation*. https://www.anthropic.com/news/anthropic-raises-30-billion-series-g-funding-380-billion-post-money-valuation

Anthropic. (2025, July 15). *Claude for financial services*. https://www.anthropic.com/news/claude-for-financial-services

Anthropic. (n.d.). *Claude's constitution*. https://www.anthropic.com/constitution

Athaluri, S. A., Manthena, S. V., Kesapragada, V. S. R. K. M., Yarlagadda, V., Dave, T., & Duddumpudi, R. T. S. (2023). Exploring the boundaries of reality: Investigating the phenomenon of artificial intelligence hallucination in scientific writing through ChatGPT references. *Cureus*, *15*(4), e37432. https://doi.org/10.7759/cureus.37432

Ball, R., & Brown, P. (1968). An empirical evaluation of accounting income numbers. *Journal of Accounting Research*, 6(2), 159–178. https://doi.org/10.2307/2490232

Barber, B. M., & Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies*, 21(2), 785–818. https://doi.org/10.1093/rfs/hhm079

Bartov, E., Lindahl, F. W., & Ricks, W. E. (1998). Stock price behavior around announcements of write-offs. *Review of Accounting Studies*, *3*, 327–346. https://doi.org/10.1023/A:1009644800963

Beaver, W. H. (1968). The information content of annual earnings announcements. *Journal of Accounting Research*, *6*, 67–92. https://doi.org/10.2307/2490070

Bhattacharya, N., Cho, Y. J., & Kim, J. B. (2018). Leveling the playing field between large and small institutions: Evidence from the SEC's XBRL mandate. *The Accounting Review*, *93*(5), 51–71. https://doi.org/10.2308/accr-52000

Bernard, V. L., & Thomas, J. K. (1989). Post-earnings-announcement drift: Delayed price response or risk premium? *Journal of Accounting Research*, 27(Supplement), 1–36. https://doi.org/10.2307/2491062

Bertomeu, J., Lin, Y., Liu, Y., & Ni, Z. (2025). The impact of generative AI on information processing: Evidence from the ban of ChatGPT in Italy. *Journal of Accounting and Economics*, 80(1), 101782. https://doi.org/10.1016/j.jacceco.2025.101782

Blankespoor, E., Croom, J., & Grant, S. M. (2026). *Generative AI and investor processing of financial information* (SSRN Working Paper No. 5053905). https://doi.org/10.2139/ssrn.5053905

Blankespoor, E., deHaan, E., & Marinovic, I. (2020). Disclosure processing costs, investors' information choice, and equity market outcomes: A review. *Journal of*

Accounting and Economics, 70(2–3), 101344. https://doi.org/10.1016/j.jacceco.2020.101344

Blankespoor, E., deHaan, E., & Zhu, C. (2018). Capital market effects of media synthesis and dissemination: Evidence from robo-journalism. *Review of Accounting Studies*, *23*(1), 1–36. https://doi.org/10.1007/s11142-017-9422-2

Blankespoor, E., deHaan, E., Wertz, J., & Zhu, C. (2019). Why do individual investors disregard accounting information? The roles of information awareness and acquisition costs. *Journal of Accounting Research*, *57*(1), 53–84. https://doi.org/10.1111/1475-679X.12248

Bradshaw, M. T., Ma, C., Yost, B. P., & Zou, Y. (2026). Generative AI use by capital market information intermediaries: Evidence from Seeking Alpha. *Journal of Accounting Research*. Advance online publication. https://doi.org/10.1111/1475-679x.70053

Breitung, C., & Müller, S. (2025). *Global business networks*. *Journal of Financial Economics, 166*, 104007. https://doi.org/10.1016/j.jfineco.2025.104007

Brogaard, J., Hagströmer, B., & Zareei, A. (2025). *The European liquidity gap* (SSRN Working Paper No. 5782262). https://doi.org/10.2139/ssrn.5782262

Brown, N. C., Stice, H., & White, R. M. (2015). Mobile communication and local information flow: Evidence from distracted driving laws. *Journal of Accounting Research*, *53*(2), 275–329. https://doi.org/10.1111/1475-679X.12077

Chatterji, A., Cunningham, T., Deming, D. J., Hitzig, Z., Ong, C., Shan, C. Y., & Wadman, K. (2025). *How people use ChatGPT* (NBER Working Paper No. 34255). National Bureau of Economic Research. https://doi.org/10.3386/w34255

Cheng, Q., Lin, P., & Zhao, Y. (2025). Does generative AI facilitate investor trading? Early evidence from ChatGPT outages. *Journal of Accounting and Economics*. Advance online publication. https://doi.org/10.1016/j.jacceco.2025.101821

Cready, W. M. (1988). Information value and investor wealth: The case of earnings announcements. *Journal of Accounting Research*, *26*(1), 1–27. https://doi.org/10.2307/2491111

Crosman, P. (2024, March 18). A third of banks ban employees from using gen AI: Here's why. *American Banker*. https://www.americanbanker.com/news/a-third-of-banks-ban-employees-from-using-gen-ai-heres-why

De Vynck, G., & Schaul, K. (2025, December 5). ChatGPT started the AI race. Now its lead is looking shaky. *The Washington Post*. https://www.washingtonpost.com/technology/2025/12/05/chatgpt-ai-gemini-competition/

deHaan, E., Madsen, J., & Piotroski, J. D. (2017). Do weather-induced moods affect the processing of earnings news? *Journal of Accounting Research*, *55*(3), 509–550. https://doi.org/10.1111/1475-679X.12160

deHaan, E., Moon, J. R., Shipman, J. E., Swanquist, Q. T., & Whited, R. L. (2023). Control variables in interactive models. *Journal of Financial Reporting*, *8*(2), 77–85. https://doi.org/10.2308/JFR-2021-023

DellaVigna, S., & Pollet, J. M. (2009). Investor inattention and Friday earnings announcements. *Journal of Finance*, 64(2), 709–749. https://doi.org/10.1111/j.1540-6261.2009.01447.x

Drake, M. S., Roulstone, D. T., & Thornock, J. R. (2012). Investor information demand: Evidence from Google searches around earnings announcements. *Journal of Accounting Research*, *50*(4), 1001–1040. https://doi.org/10.1111/j.1475-679X.2012.00443.x

Eaton, G. W., Green, T. C., Roseman, B. S., & Wu, Y. (2022). Retail trader sophistication and stock market quality: Evidence from brokerage outages. *Journal of Financial Economics*, *146*(2), 502–528. https://doi.org/10.1016/j.jfineco.2022.08.002

Eisfeldt, A. L., Schubert, G., & Zhang, M. B. (2023). *Generative AI and firm values* (NBER Working Paper No. 31222). National Bureau of Economic Research. https://doi.org/10.3386/w31222

Emsley, R. (2023). ChatGPT: these are not hallucinations – they're fabrications and falsifications. *Schizophrenia*, 9(1), 52. https://doi.org/10.1038/s41537-023-00379-4

Fama, E. F. (1970). *Efficient capital markets: A review of theory and empirical work*. *The Journal of Finance, 25*(2), 383–417. https://www.jstor.org/stable/2325486

Fletcher, L. (2026, May 4). Hedge funds seek an edge by using AI's speed. *Financial Times*. https://www.ft.com/content/0feb5743-ecf3-48f3-8425-faabea4b6f86

Foster, G., Olsen, C., & Shevlin, T. (1984). Earnings releases, anomalies, and the behavior of security returns. *The Accounting Review*, *59*(4), 574–603. https://www.jstor.org/stable/247321

Gao, M., & Huang, J. (2020). Informing the market: The effect of modern information technologies on information production. *The Review of Financial Studies*, *33*(4), 1367–1411. https://doi.org/10.1093/rfs/hhz100

Grossman, S. J., & Stiglitz, J. E. (1980). *On the impossibility of informationally efficient markets*. *The American Economic Review, 70*(3), 393–408. https://www.jstor.org/stable/1805228

Hammond, G. (2026, February 6). Anthropic's breakout moment: How Claude won business and shook markets. *Financial Times*. https://www.ft.com/content/a75555a6-24c3-4468-aba9-7fe12b5def31

Hellwig, M. F. (1980). *On the aggregation of information in competitive markets*. *Journal of Economic Theory, 22*(3), 477–498. https://doi.org/10.1016/0022-0531(80)90056-3

Hirshleifer, D., Lim, S. S., & Teoh, S. H. (2009). Driven to distraction: Extraneous events and underreaction to earnings news. *Journal of Finance*, 64(5), 2289–2325. https://doi.org/10.1111/j.1540-6261.2009.01501.x

Jensen, M. C. (1978). *Some anomalous evidence regarding market efficiency*. *Journal of Financial Economics, 6*(2-3), 95–101. https://doi.org/10.1016/0304-405X(78)90025-9

Karolyi, G. A., Lee, K.-H., & van Dijk, M. A. (2012). Understanding commonality in liquidity around the world. *Journal of Financial Economics*, *105*(1), 82–112. https://doi.org/10.1016/j.jfineco.2011.12.008

Kotrotsos, M. (2026, March 16). The AI market share numbers are wild right now. *Medium*. https://kotrotsos.medium.com/the-ai-market-share-numbers-are-wild-right-now-e8755cec734e

Li, E., Shen, M., Tu, Z., & Zhou, D. (2025). *The promise and peril of generative AI: Evidence from GPT as sell-side analysts* (SSRN Working Paper No. 4480947). https://doi.org/10.2139/ssrn.4480947

Li, F. (2010). The information content of forward-looking statements in corporate filings-A naïve Bayesian machine learning approach. *Journal of Accounting Research*, *48*(5), 1049–1102. https://doi.org/10.1111/j.1475-679X.2010.00382.x

Li, K., Mai, F., Shen, R., Yang, C., & Zhang, T. (2025). *Dissecting corporate culture using generative AI* (SSRN Working Paper No. 4558295). https://doi.org/10.2139/ssrn.4558295

London Stock Exchange Group. (n.d.). *Regulatory news services (RNS)*. https://www.lseg.com/en/capital-markets/regulatory-news-service

London Stock Exchange. (n.d.). *The world's international exchange*. https://www.londonstockexchange.com/discover

Lopez-Lira, A., & Tang, Y. (2023). Can ChatGPT forecast stock price movements? Return predictability and large language models. *arXiv preprint*. https://doi.org/10.48550/arXiv.2304.07619

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance*, 66(1), 35–65. https://doi.org/10.1111/j.1540-6261.2010.01625.x

Massenkoff, M., & McCrory, P. (2026). *Labor market impacts of AI: A new measure and early evidence*. Anthropic. https://www.anthropic.com/research/labor-market-impacts

Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., Betts Lotufo, J., Rome, A., Shi, A., & Oak, S. (2025). *The AI index 2025 annual report*. AI Index Steering Committee, Institute for Human-Centered AI, Stanford University. https://hai.stanford.edu/ai-index/2025-ai-index-report

McAllen, N., Stöckle, S., and Stender, C. (2026, May). *KPMG global AI in finance 2026: The decision advantage*. KPMG International. https://assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2026/05/global-ai-in-finance-report.pdf

Milmo, D. (2023, February 2). ChatGPT reaches 100 million users two months after launch. *The Guardian*. https://www.theguardian.com/technology/2023/feb/02/chatgpt-100-million-users-open-ai-fastest-growing-app

Murgia, M., Murray, C., & Clark, D. (2026, March 22). AI hallucinations haunt users more than job losses. *Financial Times*. https://www.ft.com/content/e074d3a9-7fd8-447d-ac0a-e0de756ac5c5

Petersen, M. A. (2009). Estimating standard errors in finance panel data sets: Comparing approaches. *Review of Financial Studies*, *22*(1), 435–480. https://doi.org/10.1093/rfs/hhn053

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3), 665–690. https://doi.org/10.1016/S0304-3932(03)00029-1

Tully, T., Redfern, J., Das, D., & Xiao, D. (2025, December 9). *2025: The state of generative AI in the enterprise*. Menlo Ventures. https://menlovc.com/perspective/2025-the-state-of-generative-ai-in-the-enterprise/

Verrecchia, R. E. (1982). Information acquisition in a noisy rational expectations economy. *Econometrica*, *50*(6), 1415–1430. https://doi.org/10.2307/1913389

Appendices

Appendix A. Practitioner interviews

To supplement the quantitative analysis, we conducted informal interviews with two finance professionals during April 2026. Both participants consented to their responses being referenced in this thesis but requested that their names and institutional affiliations remain confidential. Quotes and observations are therefore attributed by role only. The interviews were unstructured and exploratory in nature, and the responses are used as illustrative context rather than primary evidence.

Appendix B. Outage events

This appendix lists the ChatGPT and Claude outage events included in our sample. Outage severity and timing are collected from the official incident reports published by OpenAI and Anthropic, respectively, with start and end times defined as the point at which the service disruption began and the point at which full-service restoration was confirmed. For outage events with timespans extending beyond LSE trading hours (8:00 to 16:30), we replace the start time and end time with 8:00 and 16:30, respectively. Since official reports may marginally lag the true start of disruption and lead the true resolution, both timestamps are floored to the nearest 5-minute interval boundary. Thus, the first outage reported from 08:14 to 09:38, is treated as covering the intervals 08:10-09:35. ChatGPT events additionally required corroboration from a published news report, indicated by footnotes below.

Date	GenAI - model	Outage severity	Start time	End time	Duration (mins)	Number of 5-minute intervals	Day of week
18/03/2025	Claude	Major	08:14	09:38	84	17	Tuesday
24/03/2025	Claude	Major & partial	14:00	16:30	150	30	Monday
22/04/2025[1]	ChatGPT	Degraded performance	16:09	16:30	21	5	Tuesday
10/06/2025[2]	ChatGPT	Partial	08:00	16:30	510	102	Tuesday
21/07/2025	Claude	Major	08:12	09:57	105	21	Monday
03/09/2025[3]	ChatGPT	Partial	08:00	10:25	145	29	Wednesday

(*continued on next page*)

(*continued*)

Date	GenAI - model	Outage severity	Start time	End time	Duration (mins)	Number of 5-minute intervals	Day of week
23/10/2025[4]	ChatGPT	Degraded performance	16:09	16:30	21	5	Thursday
31/10/2025	Claude	Major	09:17	10:36	79	16	Friday
14/01/2026	Claude	Major & partial	08:34	14:17	343	69	Wednesday
03/02/2026	Claude	Major	15:32	15:56	24	5	Tuesday

[1] Tom's Guide, "ChatGPT was down in the UK — the outage as it happened"; Yahoo News, "ChatGPT down as thousands of users fume 'this cannot be happening' amid global issues"

[2] Forbes, "ChatGPT Outage: OpenAI's Fix Applied As 'Degraded Performance' Affects Chatbot"; Tom's Guide, "ChatGPT is back following global outage — here's what happened"

[3] Tom's Guide, "ChatGPT was down — OpenAI confirms status of outage"

[4] Tom's Guide, "ChatGPT was down — live updates on chatbot outage affecting the UK right now"

Appendix C. Supplementary data and method details

C.1 News announcement data collection

We use a browser-based scraper written in Python using Selenium and a Chrome WebDriver to collect announcement listings from Investegate's advanced-search page and extract structured data from the returned HTML using BeautifulSoup. For each target date, the scraper searches the page with the date restricted to that day, the source limited to RNS announcements, and non-security announcements (for example, fund NAVs) excluded. For each announcement, the scraper records the disclosure time, company name, ticker symbol, and headline, and constructs the corresponding Refinitiv Identification Code (RIC) by appending ".L" to the ticker symbol. We collect announcements for each outage date in the sample and for the five trading days preceding each outage. The tool accesses only publicly available information and is designed to avoid placing excessive load on the website.

C.2 Accounting-news classification

A headline is classified as accounting-relevant if it contains at least one of the inclusion keywords listed in this appendix and none of the exclusion keywords. Matching is case-insensitive and performed on the headline text only.

Category	Keywords
Inclusion keywords	
Periodic financial results and reports	final results, interim results, half year results, half-year results, half year report, half-year report, half year financial report, half-year financial report, full year results, full-year results, full-year report, full-year financial report, annual results, annual financial results, annual financial report, preliminary results, quarterly results, quarterly report, interim report, interim financial, unaudited interim, unaudited financial, financial results, financial report
Period-specific trading	full year trading, full-year trading, half year trading, half-year trading, year end trading, year-end trading, pre-close trading, quarterly activities, h1 trading, h2 trading, q1 trading, q2 trading, q3 trading, q4 trading
Trading updates and earnings guidance	trading update, trading statement, management statement, AGM trading, profit warning, profit alert, revenue warning
Exclusion keywords	
	notice of, posting of annual report, publication of annual report, annual report and notice of AGM, annual financial report and notice, delay to, revised date of, request of extension, order re extension, short delay to, phase 1, phase 2, phase 3, clinical, drilling, flow testing

Note: Headlines beginning with "notice of" are excluded to avoid classifying announcements that merely schedule upcoming results as substantive disclosures. The sector-specific exclusions "phase 1", "phase 2", "phase 3", "clinical", "drilling", and "flow testing" remove non-financial uses of the word "results", such as clinical trial or resource extraction announcements.

Appendix D. Variable definitions

Variable definitions follow Cheng et al. (2025), adapted where necessary to reflect the context of our study.

Variable name	Definition
Trading volume analyses	
Dependent variables	
Trading Volume	The total number of shares traded in a 5-minute interval, divided by the total number of shares outstanding, then multiplied by 10,000
Shares Count	The total number of shares traded in a 5-minute interval
Trades Count	The total number of trades in a 5-minute interval
Independent variables	

(*continued on next page*)

(*continued*)

Variable name	Definiton
Outage	An indicator variable that equals one if the trading interval coincides with a ChatGPT or Claude outage period, and zero otherwise
Pseudo-Outage	An indicator variable that equals one if the trading interval coincides with a ChatGPT or Claude outage period on the day before the outage date, and zero otherwise
Control variables	
Return Overnight	Stock return measured from the closing price in day *t-1* to the opening price on day *t*
Return [-1]	Stock return on day *t-1*
Return [-5, -2]	Cumulative stock return from day *t-5* through *t-2*
Price Range	The logarithm of the highest closing price minus the logarithm of the lowest closing price during the previous 60-day period
Price	The logarithm of one plus the stock price on day *t-1*
Market Cap	The logarithm of one plus the product of the stock price and the total number of shares outstanding on day *t-1*
Cross-sectional variables	
Claude	An indicator variable that equals one if the trading interval coincides with an outage event caused by a disruption to Claude, and zero if caused by ChatGPT. It is similarly defined for observations of the corresponding control days
News Accounting	An indicator variable that equals one if there is at least one qualifying accounting RNS announcement about the firm released on that calendar date before or during the outage period, and zero otherwise. It is similarly defined for observations from control days
News Non-Accounting	An indicator variable that equals one if there is at least one RNS announcement about the firm released on that calendar date before or during the outage period, provided that none of the announcements are classified as accounting news, and zero otherwise. It is similarly defined for observations from control days
News All	An indicator variable that equals one if there is at least one RNS announcement about the firm released on that calendar date before of during the outage period, and zero otherwise. It is similarly defined for observations from control days

(*continued on next page*)

(*continued*)

Variable name	Definiton
News Pure Accounting	An indicator variable that equals one if there is at least one RNS announcement about the firm released on that calendar date before or during the outage period, and all such announcements are qualified accounting news, and zero otherwise. It is similarly defined for observations from control days
News Mixed	An indicator variable that equals one if there is at least one qualifying accounting RNS announcement and at least one non-accounting RNS announcement about the firm released on that calendar date before or during the outage period, and zero otherwise. It is similarly defined for observations from control days.

Appendix E. Supplementary results

E.1 News mechanism analysis

Table E1. GenAI outages and change in trading volume: Conditional on RNS news releases and type

This table reports the OLS regression results examining whether the effect of GenAI outages on trading volume varies with the presence and type of RNS news, using Equation (4). Column (1) presents the specification using all RNS news observations, while Column (2) examine pure accounting, mixed, and non-accounting only news classification. The sample includes 720,909 security-date-5-minute interval observations from 622 stock securities during the period from March 2025 to February 2026. See Appendix D for variable definitions. Standard errors are clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Dependent variable =	*Trading volume*	
	(1)	(2)
Outage × *News All*	-0.024	
	(-0.65)	
Outage × *News Pure Accounting*		-0.766*
		(-1.75)
Outage × *News Mixed*		-0.723**
		(-1.98)
Outage × *News Non-Accounting*		0.000
		(0.00)
Outage	-0.030**	-0.030**
	(-2.11)	(-2.10)
News All	0.128***	
	(3.24)	

(*Table E1 continued on next page*)

(*Table E1 continued*)

Dependent variable =	*Trading Volume*	
	(1)	(2)
News Pure Accounting		1.099***
		(4.19)
News Mixed		1.000***
		(3.34)
News Non-Accounting		0.080**
		(2.06)
Claude	0.352***	0.349***
	(6.32)	(6.27)
Controls	YES	YES
Firm FE	YES	YES
Time Interval FE	YES	YES
N	720,909	720,909
Adjusted R^2	0.368	0.369
Interaction-effect differences ($\beta_{\text{Outage} \times A} - \beta_{\text{Outage} \times B}$)		
News Pure Accounting vs. *News Mixed*		-0.043
		(-0.08)
News Pure Accounting vs. *News Non-Accounting*		-0.767*
		(-1.75)
News Mixed vs. *News Non-Accounting*		-0.723**
		(-1.97)
Total outage effect ($\beta_{\text{Outage}} + \beta_x$)		
News All	-0.054*	
	(-1.82)	
News Pure Accounting		-0.796*
		(1.82)
News Mixed		-0.753**
		(-2.06)
News Non-Accounting		-0.030
		(-1.05)

Note: Statistical significance for reported interaction-effect differences and total outage effects is assessed using Wald tests.

Table E2. GenAI outages and change in trading volume: Conditional on accounting and non-accounting news releases, model-specific results

This table reports the OLS regression results examining whether the effect of GenAI outages on trading volume varies with the presence and type of RNS news, using Equation (4) and estimated on model specific subsamples. See Appendix D for variable definitions. Standard errors are clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Dependent variable =	*Trading volume*	
	ChatGPT	Claude
	(1)	(3)
Outage × *News Accounting*	-0.367**	-0.719*
	(-2.14)	(-1.65)
Outage × *News Non-Accounting*	0.049	0.033
	(1.25)	(0.67)
Outage	-0.026**	-0.054**
	(-2.26)	(-2.63)
News Accounting	0.430**	1.266***
	(2.45)	(4.78)
News Non-Accounting	0.024	0.123
	(1.32)	(1.54)
Controls	YES	YES
Firm FE	YES	YES
Time Interval FE	YES	YES
N	295,053	425,856
Adjusted R^2	0.237	0.440
Interaction-effect differences ($\beta_{\text{Outage} \times \text{A}} - \beta_{\text{Outage} \times \text{B}}$)		
News Accounting vs. *News Non-Accounting*	-0.416**	-0.752*
	(-2.31)	(-1.72)
Total effect ($\beta_{\text{Outage}} + \beta_{x}$)		
News Accounting	-0.393**	-0.773*
	(-2.28)	(-1.78)
News Non-Accounting	0.024	-0.021
	(0.73)	(-0.53)

Note: Statistical significance for reported interaction-effect differences and total outage effects is assessed using Wald tests.

E.2 Dynamic analysis

Table E3. Regression estimates for the dynamic analysis of GenAI outages

This table reports the OLS regression estimates underlying Figure 2. Coefficients are estimated from the dynamic specification in Equation (5), which examines changes in trading volume across 10-minute intervals surrounding GenAI outage periods. The sample includes 248,844 security-date-5-minute interval observations from 552 stock securities during the period from March 2025 to February 2026. Standard errors are clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Dependent Variable =	*Trading Volume*	
	(1)	(2)
Event time interval (τ)	Coef.	SE
Pre 60 min	0.086	(0.055)
Pre 50 min	-0.026	(0.050)
Pre 40 min	0.117**	(0.051)
Pre 30 min	0.106**	(0.043)
Pre 20 min	-0.018	(0.036)
Pre 10 min		
During 0 min	-0.109***	(0.040)
During 10 min	-0.163***	(0.035)
During 20 min	-0.104**	(0.048)
During 30 min	-0.313***	(0.066)
During 40 min	-0.199***	(0.051)
During 50 min	-0.199***	(0.056)
During 60 min	-0.095	(0.060)
Post 10 min	-0.021	(0.036)
Post 20 min	-0.023	(0.039)
Post 30 min	-0.115***	(0.038)
Post 40 min	-0.010	(0.054)
Post 50 min	-0.088*	(0.049)
Post 60 min		
Controls	YES	
Firm FE	YES	
5-min interval FE	YES	
N	248,844	
Adjusted R^2	0.3662	

Appendix F. AI use disclosure

Throughout the process of writing the thesis AI has been used as support, primarily Claude and ChatGPT, as well as ResearchRabbit for literature discovery. AI was used exclusively to support the writing and coding process and all theoretical reasoning, methodological choices, empirical tests, and conclusions are our own.

AI assistance took four forms. First, AI was used for brainstorming, mainly during the early stages of the thesis helping to e.g. identify potential research angles and think through how different parts of the thesis connected to each other. Second, AI was used to clarify theoretical concepts mostly encountered during the literature review phase, supporting our understanding before forming independent views. Third, AI was used to give feedback on drafts, including suggestions on clarity, structure, and phrasing. In all cases, suggestions were of course critically evaluated based on our own judgement. Fourth, AI was used as support throughout the development of the Python code underlying the empirical analysis, including data cleaning, web scraping, data processing, regression implementation, and debugging. All code was reviewed, tested, and verified by us, making sure we understood and agreed with the suggestions made. In addition, ResearchRabbit was used to identify relevant academic literature and ensure that important contributions within the research field were not overlooked. All sources identified through ResearchRabbit were read and evaluated independently before being included in the thesis.

The primary risk associated with AI use is factual inaccuracy, given the tendency of language models to produce plausible but incorrect information. To mitigate this, any substantive claims, citations, or empirical facts were verified independently against primary sources, and nothing was accepted solely based on AI output.

The experience reinforced that AI tools can meaningfully support efficiency in tasks such as brainstorming, literature discovery, coding assistance, and writing feedback, while analytical reasoning, methodological decisions, and interpretation of results have remained grounded in our own judgement. AI contributed to how ideas were communicated and organised, not to the ideas themselves.