

When ChatGPT Stops Talking:

GenAI, Retail Trading Coordination, and Market Liquidity

Xiaoxue He

*Department of Accounting and Finance, Lancaster University Management School,
Lancaster, LA1 4YX, UK*

x.he21@lancaster.ac.uk

THIS VERSION: SEPTEMBER 2026

ABSTRACT

Using 14 unexpected ChatGPT outages, we examine whether GenAI coordinates retail trading and how it affects trading informativeness. When ChatGPT is available, retail order flow is more correlated, and a larger share of information-related order flow comoves across stocks. This common component shows little evidence of return predictability, while it appears to substitute for the informative individual component. The informativeness effect of ChatGPT is concentrated in that individual component and runs in opposite directions: it makes retail selling less informative while making retail buying more informative. This asymmetry suggests that the informational effect of such a tool may depend on how it interacts with investors' existing information sets. Such ChatGPT-related coordination also leads to higher liquidity commonality, implying greater exposure to common liquidity shocks.

Keywords: Generative AI; Retail Trading; Correlated Trading; Informativeness; Liquidity Commonality

1 Introduction

Generative artificial intelligence (GenAI) has rapidly democratized complex financial information processing, significantly lowering the barriers to informed market participation. Publicly available GenAI platforms such as ChatGPT have great potential to reshape how retail investors interact with financial markets, particularly for those traditionally constrained by high information costs and limited access to sophisticated analytical tools. Survey evidence already points to meaningful GenAI adoption among retail investors: over 40% of retail investors report using GenAI tools to inform their investment decisions (eToro, 2023; Oliver Wyman, 2024; Blankespoor et al., 2026), and projections suggest these applications will become the dominant source of retail investment advice by 2027, reaching 78% of users by 2028 (Deloitte Center, 2024). Therefore, the key question is no longer whether GenAI affects retail investors, but how it changes their trading behavior and what this implies for market quality.

Emerging research has begun to answer parts of this question. Several studies find that GenAI fosters retail trading activity and supports more informed decisions (e.g., Cheng et al., 2025; Hansen and Lee, 2025; Chang et al., 2026; Ecker et al., 2026). However, existing research largely examines retail investors as individual decision makers or as an aggregate group. Much less is known about how GenAI affects the relation among them. This distinction matters because GenAI models operate through a centralized architecture, and risk compressing diverse information sets into homogenized outputs, thereby delivering similar signals to many users and increasing correlation in their trading decisions. Regulators have already flagged this as a burgeoning systemic risk, referring to it as "GenAI-induced herding" (ESMA, 2023; BIS, 2024; FSB, 2024; IMF, 2024; Bank of England, 2025), while direct empirical evidence remains scarce.

The theoretical foundation for this concern is well established, but direct empirical evidence on whether it actually occurs remains scarce, and its effects are still unclear. Investors with limited attention may rely heavily on salient AI outputs (Hirshleifer and Teoh, 2003; Cao et al., 2026), and such low-cost, widely shared signals can coordinate

beliefs and generate comovement in trading and asset prices (Veldkamp, 2006). Whether this improves or distorts investment decisions remains an open empirical question with important welfare implications (Cao et al., 2026). If shared signals are informative, the resulting correlated trading may reflect convergence toward value-relevant information and contribute to price discovery. Alternatively, shared signals may crowd out independent information acquisition and processing, reducing price efficiency when they are less informative than the personal information investors would otherwise acquire or rely on (Grossman and Stiglitz, 1980; Morris and Shin, 2002; Dugast and Foucault, 2018). Thereby, coordination reflects correlated uninformed trading, and it can generate nondiversifiable noise-trader risk (De Long et al., 1990; Lee et al., 1991).

To explore this question, we exploit 14 unexpected ChatGPT full outages between 2023 and 2024 that overlap U.S. equity trading hours as plausibly exogenous shocks to GenAI availability for retail investors. ChatGPT is the natural focus of our study. Its publicly accessible and user-friendly interface requires little technical expertise, making it particularly relevant to the broad retail investor population. ChatGPT surpassed 100 million users within two months of its launch in November 2022 (Reuters, 2023) and has remained the dominant GenAI platform, accounting for 77% of traffic across the top 60 GenAI tools as of April 2025 (Liu et al., 2025). Given its broad adoption, ChatGPT provides a setting that more closely reflects the behavior of typical retail GenAI users than any alternative tool could offer (Chatterji et al., 2025).

We combine these outage events with high-frequency trade and quote data from the NYSE TAQ database and identify retail trades following Barber et al. (2024). For each outage, we compare outcomes during the outage window with the same clock-time five-minute intervals for the same stock on matched control days. This design absorbs time-invariant firm characteristics and intraday seasonality, allowing us to identify the effect of ChatGPT availability.

Our analysis centers on three questions: Does ChatGPT increase correlation in retail trading? Does the resulting correlated order flow reflect price discovery or correlated noise? And what are the broader market consequences of such coordination?

Consistent with Cheng et al. (2025), we first confirm that ChatGPT plays a meaningful role in retail trading: during outages, retail trade count falls by 6.9% and retail trade volume by 12.2%. This result provides the empirical foundation for the analyses that follow.

Our first main finding is that ChatGPT availability is associated with greater correlation in retail trading, both within and across stocks. During outages, within stocks, retail trades are less directionally concentrated, and order imbalance is more volatile and less persistent over time; across stocks, retail order flow comoves less with that of industry peers; the market-level evidence is weaker but points in the same direction.

Our second finding is that greater coordination does not necessarily represent improved price discovery, and ChatGPT has asymmetric effects on trading informativeness across retail selling and buying. We examine buying and selling separately because they face different information constraints. Selling decisions depend more on investors' information processing about the limited set of securities they already hold, whereas buying requires searching across a much larger opportunity set, where investors face greater attention constraints and thinner prior information (Barber and Odean, 2008; Barber et al., 2022). An information tool may have different marginal effects on the two sides. Our results show that retail selling becomes more informative during outages, whereas the informativeness of retail buying decreases.

To shed light on the mechanism underlying this asymmetry, we follow Boehmer et al. (2021) and decompose retail order imbalance. The outage effect on return predictability is concentrated in the information-related component. We further separate this component into common and individual parts. During outages, the common component declines while the individual component increases, with little change in their combined magnitude. This pattern is consistent with common signals substituting for more heterogeneous individual inputs when ChatGPT is available. However, the common

component itself provides little evidence of return predictability. Instead, the change in informativeness is concentrated in the individual component and moves in opposite directions across selling and buying. One plausible explanation is that, on the sell side, a low-cost common signal from ChatGPT substitutes for investors' own more informative heterogeneous information. On the buy side, where investors possess less prior information about potential purchases, ChatGPT may instead facilitate information acquisition and processing, improving the informativeness of individual trading. The evidence is consistent with a crowding-out effect on retail selling while a crowding-in effect on retail buying. Thus, ChatGPT's informational effect appears to depend on how a shared, low-cost information tool interacts with investors' existing information sets.

Our third finding concerns market liquidity. When ChatGPT is available, price impact is 12.5% higher, while realized spreads remain unchanged, indicating that the increase in trading costs is concentrated in post-trade price movements rather than liquidity-provider compensation. More importantly, liquidity commonality is also stronger when ChatGPT is available, both within industries and, to a somewhat weaker extent, across the broader market. These findings point to a potential microstructure channel through which widespread reliance on a common GenAI technology may increase exposure to common liquidity shocks.

Our paper makes several contributions. First, we add to the growing literature on GenAI and financial markets. We extend this literature from examining how GenAI affects retail investors individually or as an aggregate group (e.g., Cheng et al., 2025; Even-Tov et al., 2025; Hansen and Lee, 2025; Chang et al., 2026; Ecker et al., 2026) to examining how it affects the relation among their trading decisions. Our evidence on ChatGPT-induced correlation in retail trading suggests that GenAI can shape not only individual trading decisions but also the collective behavior of retail investors.

Second, we contribute to the literature on information technology and price discovery. Prior literature provides competing predictions on whether modern information

technologies crowd in or crowd out market participants' own information production (e.g., Grossman and Stiglitz, 1980; Morris and Shin, 2002; Gao and Huang, 2020). Existing empirical studies generally examine which force dominates on average. Our evidence points to the possibility that both forces may operate simultaneously within the same investor group and suggests that the effect of information technology on price discovery may depend on how a shared information tool interacts with investors' existing information sets and the information constraints associated with different trading decisions.

Third, we provide new evidence on demand-side sources of liquidity commonality. Prior research has explored demand-side sources driven by correlated institutional trading activity (Kamara et al., 2008; Koch et al., 2016). We show that GenAI reliance can operate through a related channel: ChatGPT coordinates retail order flow and further makes liquidity comove more strongly. These findings also provide early empirical evidence relevant to regulatory concerns that widespread GenAI adoption may foster “GenAI-induced herding” and increase exposure to common liquidity shocks (ESMA, 2023; BIS, 2024; FSB, 2024; IMF, 2024; Bank of England, 2025).

The remainder of the paper proceeds as follows. Section 2 discusses the relevant literature. Section 3 describes the data and sample construction. Section 4 presents empirical results. Section 5 concludes.

2 Literature Review

2.1 Generative AI and Financial Markets

The literature on GenAI in capital markets is growing rapidly. An expanding strand of this literature examines how GenAI affects information production, dissemination, and processing by investors (e.g., Bertomeu et al., 2025; Blankespoor et al., 2026; Bradshaw et al., 2026; Cao et al., 2026; Choi and Xie, 2026).

Recent studies provide evidence that GenAI reduces information-processing frictions and affects trading. Cheng et al. (2025), whose empirical setting is closest to ours, exploit unexpected ChatGPT outages and document lower trading activity, short-run price impact, and return variance during outages, together with evidence that GenAI improves longer-term price informativeness. Chang et al. (2026) show that retail trading becomes more aligned with AI-sentiment signals following ChatGPT's deployment. Ecker et al. (2026), using investor queries submitted to a GenAI platform, show that investors use GenAI primarily to integrate and interpret financial information and document associations between GenAI use, informed trading, and market liquidity.

Several studies also point to potential costs of GenAI use. Croom (2026) confirms that GenAI interactivity blurs the attribution of information processing skill and strengthens overconfidence-like behavior in a laboratory setting. Hansen and Lee (2025) show experimentally that AI agents can still be induced to herd when explicitly guided toward profit-maximizing decisions. Moreover, AI hallucinations can introduce errors into AI-generated outputs that are difficult for users to detect (Choi et al., 2021; Cao et al., 2026).

Despite this growing evidence, much less is known about how GenAI can affect the relation among investors' trading decisions. Moreover, the net welfare implications of reliance on AI-generated information without adequate verification remain an open empirical question (Cao et al., 2026). Our study attempts to shed light on these gaps by examining how unexpected disruptions to ChatGPT affect correlation in retail trading, the informativeness of the resulting order flow, and its broader market consequences.

2.2 Information Technology and Price Discovery

Modern information technologies reduce the costs of accessing, disseminating, and processing financial information. However, their effect on market participants' own information production remains ambiguous.

Modern information technologies can either crowd in or crowd out private information production, with different implications for price discovery. The crowding-in effect arises when easier access to information and lower processing costs increase the net benefits of producing information, thereby encouraging market participants' information production and improving pricing efficiency (e.g., Verrecchia, 1982; Kim and Verrecchia, 1994; Gao and Huang, 2020). The crowding-out effect arises when easily accessible common information substitutes for costly independent information production, potentially reducing price efficiency when the displaced information is more precise (Grossman and Stiglitz, 1980; Morris and Shin, 2002; Dugast and Foucault, 2018).

Theoretically, crowding-in and crowding-out need not be mutually exclusive. Broader theoretical work on public information shows that both forces can arise depending on the information environment and signal precision (Cai and Dong, 2023; Aoyagi, 2025). Although these studies do not examine information technology directly, they provide theoretical support for the possibility that the two effects can coexist. However, existing empirical studies of information technology largely examine which force dominates on average. Much less is known about whether the same information technology can simultaneously crowd in and crowd out information production across different decisions within the same investor group. Our research attempts to fill the gap.

2.3 Correlated Trading and Liquidity Commonality

Retail investors' trades are systematically correlated (Kumar and Lee, 2006; Barber et al., 2008). Correlated trading can arise from common sentiment, attention, location, or information (e.g., Feng and Seasholes, 2004; Kumar and Lee, 2006; Colla and Mele, 2010; Chelley-Steeley et al., 2019). For example, Veldkamp (2006) shows that the high fixed costs and low replication costs of information production make widely shared information relatively inexpensive. Investors therefore have incentives to acquire similar information, generating commonality in beliefs and ultimately comovement in asset prices.

One important market consequence of correlated trading is liquidity commonality (Chordia et al., 2000). The literature identifies both supply-side and demand-side sources of liquidity commonality, with correlated trading widely recognized as an important demand-side source. Kamara et al. (2008) link differences in liquidity commonality to institutional ownership, while Karolyi et al. (2012) show that liquidity commonality is stronger when trading activity is more correlated. Moreover, Koch et al. (2016) also show that stocks with greater mutual fund ownership exhibit stronger liquidity comovement and trace this relation to common trading by mutual funds.

Another important market consequence links to the informational content of correlated trading. When correlated trading reflects correlated noise, the resulting “noise trader risk” is non-diversifiable and can be priced (De Long et al., 1990; Shleifer and Summers, 1990; Lee et al., 1991). This mechanism is particularly relevant to our setting because reliance on a common GenAI platform may generate information-induced correlated retail demand. We therefore examine whether greater correlation in retail order flow associated with ChatGPT is also accompanied by greater liquidity commonality and whether the correlated trading is informative.

3 Data and Sample

We identify ChatGPT outage events in 2023–2024 from OpenAI’s publicly disclosed incident reports, which provide the official records of incident timelines and severity for service disruptions.¹ We focus on incidents classified as full outages and confirm that ChatGPT was broadly inaccessible to users based on the root-cause reports. We retain events that overlap at least partially with U.S. equity trading hours (9:30–16:00 ET). For the exact outage windows, we adopt the adjusted outage timestamps from Cheng et al. (2025), who refine OpenAI’s reported incident timelines using user reports compiled by Downdetector, a platform that tracks online service disruptions. Appendix A provides the complete list of outage events used in our analysis. The final event set

¹ See <https://status.openai.com/history> for OpenAI’s incident history.

contains 14 full outages, with durations ranging from 18 to 240 minutes (1,482 minutes in total; mean duration of about 106 minutes).

Our intraday trade and quote data come from the NYSE's TAQ database. We obtain daily stock prices and returns from CRSP, firm characteristics from Compustat, firm-specific information events from Capital IQ Key Developments, and analyst coverage from I/B/E/S.

We begin with U.S. ordinary common stocks listed on NYSE, AMEX, or Nasdaq and exclude non-ordinary shares and stocks priced below \$1. Because our setting is retail-focused, we identify retail trades (purchases and sales) following Boehmer et al. (2021) and Barber et al. (2024), which exploits institutional features of retail execution.² Retail orders are often executed off-exchange and receive small price improvement relative to the national best bid and offer (NBBO). Specifically, we classify retail trades as those reported with TAQ exchange code “D” (off-exchange) and with prices 0.1–0.4 cents below a round penny for buys and 0.1–0.4 cents above a round penny for sells. We do not classify trades priced exactly at a round penny or near the half-penny (0.4–0.6 cents, inclusive) to reduce misclassification.

Following prior work (e.g., Eaton et al., 2022; Cheng et al., 2025), we aggregate retail transactions to 5-minute intervals. For each outage, we include all 5-minute observations that fall within the outage window. We construct the matched control sample using the same clock-time intervals from the preceding five trading days for the same stock. This matching procedure holds both the stock and intraday trading period fixed, thereby reducing concerns arising from persistent firm characteristics and systematic intraday trading patterns. We exclude any control-day interval that overlaps with another ChatGPT outage.³ We further require each outage interval to have at least

² BJZZ’s classification approach has been widely adopted in recent studies of retail trading (e.g., Eaton et al., 2022; Cheng et al., 2025; Havakhor et al., 2025).

³ For example, the ChatGPT outage on March 20, 2023, ran from 12:20 to 22:54 ET. Since the outage extends beyond market close, we truncate the window at 16:00 ET, yielding 44 5-minute intervals within

three eligible control intervals. The final sample contains 3,052,058 stock–date–5-minute observations from 3,628 unique stocks. The number of observations could vary across subsequent tests due to specific variable construction.

Table 1 reports summary statistics for the main sample. To minimize the influence of outliers, we winsorize all continuous variables at the 1st and 99th percentiles. Variable definitions are provided in Appendix B. The table shows that retail trading is economically active at the 5-minute level: a typical interval contains around 21 retail trades, corresponding to roughly 3,100 shares per stock-interval on average.

[Insert Table 1 here]

4 Research Design and Empirical Results

Our empirical analysis proceeds in four steps. We first establish that ChatGPT meaningfully affects retail trading intensity (Section 4.1). We then examine whether ChatGPT availability increases coordination in retail trading within and across stocks (Section 4.2), whether such coordination improves or reduces trading informativeness (Section 4.3), and whether it has broader consequences for market liquidity and liquidity commonality (Section 4.4).

4.1 The Effect of ChatGPT Outages on Retail Trading Intensity

We begin by examining whether retail trading activity declines when investors temporarily lose access to ChatGPT. If ChatGPT meaningfully supports retail trading decisions, we expect a decrease in retail trading activity during outage periods. We estimate:

outage window on March 20. The matched control sample covers the same 44 intervals on each of the five preceding trading days (March 13, 14, 15, 16, and 17, 2023). We exclude March 15 from the control window because a separate ChatGPT outage occurred on that day.

$$Retail\ Trade_{i,t,p} = \alpha + \beta_1 Outage_{t,p} + \gamma Controls_{i,t} + \gamma_i + \delta_{t,p} + \varepsilon_{i,t,p} \quad (1)$$

where i , t , and p indicate stock, date, and 5-min interval, respectively. We measure retail trading intensity using Ln Retail Trade Count and Ln Retail Trade Volume. The independent variable, $Outage_{t,p}$, equals one if interval p on day t falls within a ChatGPT outage window. and zero otherwise. The coefficient β_1 captures the difference in retail trading intensity between outage intervals and the matched control intervals for a stock. Following prior work (e.g., Eaton et al., 2022; Cheng et al., 2025), we control for past returns (overnight, prior-day, and days $t-5$ to $t-2$), stock price, market capitalization, book-to-market ratio, intraday volatility, and analyst coverage. We include firm fixed effects and event-by-interval fixed effects to absorb time-invariant firm heterogeneity and time-varying market-wide conditions within each event and time-of-day bin. Standard errors are two-way clustered at the firm and event-by-date levels.

Table 2 shows that retail trading activity declines significantly when ChatGPT is unavailable, consistent with Cheng et al. (2025). Specifically, the coefficient on $Outage$ is -0.071 for Ln Retail Trade Count ($t = -2.33$), corresponding to a 6.9% decrease in retail trade count, and the coefficient on $Outage$ is -0.130 for Ln Retail Trade Volume ($t = -2.98$), implying a 12.2% decline. These estimates suggest that access to ChatGPT meaningfully raises retail participation even at short intraday horizons.

[Insert Table 2 here]

A primary identification concern is whether the documented outage effects are truly exogenous. We address this in two steps. First, we identify the root causes of outages from OpenAI's incident reports and find no systematic evidence that outages are related

to market conditions.⁴ Second, we conduct a battery of robustness tests to ensure our results are not spurious, summarized in Table 3.

[Insert Table 3 here]

Panel A excludes observations coinciding with firm-specific news events. The results remain similar for both Ln Retail Trade Count (-0.072 , $t = -2.35$) and Ln Retail Trade Volume (-0.134 , $t = -3.01$). Panel B drops events during which scheduled macroeconomic announcements, specifically U.S. Consumer Price Index (CPI), Producer Price Index (PPI), and the Federal Open Market Committee (FOMC) minutes releases, are made before or during the outage window, as these announcements are known to trigger significant market reactions.⁵ The results still hold (-0.071 , $t = -2.00$; -0.141 , $t = -2.79$).⁶ Panel C excludes outages shorter than 30 minutes to rule out the effect of minor disruptions, with stable coefficients (-0.074 , $t = -2.37$; -0.132 , $t = -2.90$).⁷ To mitigate potential confounding from after-hours news or elevated opening-period volatility, Panel D drops outages beginning before 10:00 AM, and the results remain statistically and economically significant (-0.093 , $t = -2.08$; -0.171 , $t = -2.70$).⁸ Finally, Panel E presents a placebo test by assigning the outage indicator to day $t-1$ rather than the actual outage day. The placebo coefficients are statistically insignificant ($t = 1.66$; $t = 1.15$), confirming that the observed outage effect is unique to actual outage periods and not a mechanical artifact of the research design. We report similar

⁴ The root causes include database instability, node provisioning failures, underlying subsystem failures, software bugs, authentication issues, cloud provider data-center power failures, and deployment changes. None of these causes are plausibly related to stock market conditions.

⁵ Prior research documents that CPI, PPI, and FOMC announcements are major macroeconomic news that can lead to strong stock market reaction (e.g., Savor and Wilson, 2013; Lucca and Moench, 2015; Fisher et al., 2022). CPI and PPI announcement dates are obtained from the Bureau of Labor Statistics (https://www.bls.gov/bls/archived_sched.htm), and FOMC minutes dates are from the Federal Reserve Board (<https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm>).

⁶ The four excluded events are: February 14, 2023 (CPI announcement), March 15, 2023 (PPI announcement), May 24, 2023 (FOMC minutes release), and September 13, 2023 (CPI announcement).

⁷ Two outage events are excluded: March 15, 2023, and August 21, 2024.

⁸ Five outage events are excluded: February 27, August 29, August 31, September 13, and November 8, 2023.

robustness checks for subsequent analyses in Appendix C. Taken together, these tests substantially mitigate identification concerns that the documented outage effects are driven by observable information events, particular outage characteristics, or the matching procedure.

One additional concern is that retail investors may partially switch to competing AI tools, such as Google Gemini or Anthropic's Claude, during ChatGPT outages.⁹ To the extent that such substitution occurs, it would tend to attenuate the estimated effect of ChatGPT outages on retail trading activity. More broadly, our findings speak to regulatory concerns about concentration risk in AI services (e.g., ESMA, 2023; BIS, 2024; FSB, 2024; IMF, 2024; Bank of England, 2025): if a single dominant platform can reshape retail investors' trading behavior, any technical disruption or systematic bias in its outputs could have outsized effects on financial markets.

4.2 The Effect of ChatGPT Outages on Correlated Retail Trading

Having established that ChatGPT plays a meaningful role in retail trading, we now investigate whether ChatGPT availability is associated with more correlated retail trading. As a widely used centralized information-processing tool, ChatGPT may compress diverse information sets into homogenized outputs and deliver similar signals to users. If retail investors rely on such shared signals, their trading decisions may become more correlated. We examine both within-stock and cross-stock correlation.

4.2.1 Within-Stock Correlated Trading

We start by constructing the herding measure developed by Lakonishok et al. (1992), which has been applied by Dorn et al. (2008) and Barber et al. (2008) to measure systematic correlation in retail trading. The LSV measure is defined as:

⁹ We attempted to conduct a parallel analysis using outage events for these competing platforms but found that neither maintains a complete and reliable public incident history comparable to OpenAI's structured status page.

$$LSV_{i,t,p} = |p_{i,t,p} - E[p_{i,t,p}]| - E|p_{i,t,p} - E[p_{i,t,p}]| \quad (2)$$

where $p_{i,t,p}$ denotes the proportion of retail trades classified as buys for stock i during interval p on date t . $E[p_{i,t,p}]$ represents the expected buy proportion, estimated from matched control intervals. The second term, $E|p_{i,t,p} - E[p_{i,t,p}]|$, captures the expected absolute deviation of the realized buy proportion from the expected buy proportion that would arise mechanically from random trading, given the finite number of trades. To ensure reliable measurement, we require at least five retail trades in each interval, and at least 10 trades in the matched-control benchmark. A higher LSV measure indicates greater directional concentration of retail trades beyond what random trading would produce. If ChatGPT coordinates retail trading, we expect to see a decrease in LSV during outages. We estimate:

$$LSV_{i,t,p} = \alpha + \beta_1 Outage_{t,p} + \gamma Controls_{i,t} + \gamma_i + \delta_{t,p} + \varepsilon_{i,t,p} \quad (3)$$

Table 4 reports a negative coefficient on Outage of -0.003 ($t=-2.26$). Relative to the sample mean of 0.102, this corresponds to a 2.9% decline in retail herding during ChatGPT outages.

[Insert Table 4 here]

We also examine within-stock coordination from a time-series perspective, since a shared signal may lead to more stable and persistent net order flow across intervals, providing another indication of ChatGPT-related correlated trading.

Building on Bogousslavsky and Collin-Dufresne (2023), we measure the dynamics of net retail order flow using retail order imbalance volatility (ROIV). ROIV is calculated as the standard deviation of five-minute trade-based retail order imbalance (ROI) within each firm-event-day window, requiring at least three valid consecutive five-minute intervals. A higher ROIV therefore indicates greater fluctuations in the direction and magnitude of retail order flow over the event window. To account for mechanical

variation in ROIV arising from retail trading intensity, we also construct an independent-sign benchmark using the retail buy probability estimated from matched control days and the realized level of retail trading activity.¹⁰ Excess ROIV Variance is the observed ROIV variance net of this benchmark variance. We estimate the regression:

$$\ln ROIV_{i,t} = \alpha + \beta_1 Outage_t + \gamma Controls_{i,t} + \gamma_i + \delta_t + \varepsilon_{i,t} \quad (4)$$

Panel A of Table 5 summarizes the results. Both Ln ROIV and Excess ROIV Variance increase significantly during outages (0.034, t=5.20; 0.006, t= 2.21), implying that net retail order flow becomes more volatile when access to ChatGPT is disrupted.

We provide complementary evidence using two measures of first-order serial dependence in ROI. The first is the Pearson correlation between the trade-based ROI in adjacent intervals, calculated within each firm-event-day window and requiring at least three valid adjacent interval pairs as follows.

$$\rho_{i,t}^{Pearson} = Corr(ROI_{i,t,p}, ROI_{i,t,p+1}) \quad (5)$$

The second is the first-order autocorrelation estimator, which is calculated as:

$$\rho_{i,t}^{ACF} = \frac{\sum_{p=1}^P (ROI_{i,t,p+1} - \overline{ROI}_{i,t})(ROI_{i,t,p} - \overline{ROI}_{i,t})}{\sum_{p=1}^P (ROI_{i,t,p} - \overline{ROI}_{i,t})^2} \quad (6)$$

Panel B of Table 5 reports negative coefficients on Outage for both measures (-0.010, t=-2.11; -0.008, t=-1.86), indicating weaker persistence in retail order imbalance during ChatGPT outages.

[Insert Table 5 here]

¹⁰ As in the LSV measure construction, we require at least 10 classified retail trades in the matched-control sample to ensure a reliable benchmark.

All these findings present a consistent pattern: during ChatGPT outages, retail trades are less directionally concentrated, net retail order flow is more volatile, and its persistence weakens. These results are consistent with ChatGPT availability being associated with greater within-stock coordination in retail trading.

4.2.2 Cross-Stock Correlated Trading

ChatGPT may also provide retail investors with common signals about industry- and market-level information, thereby increasing correlated trading across stocks. To examine this, we adapt the high-frequency order-flow approach of Harford and Kaul (2005).

For robustness, we consider three measures of retail order flow. The first is standardized net retail trade flow (STD_NET), defined as the number of retail buy trades minus that of sell trades, standardized using each stock's mean and standard deviation estimated from matched control observations. We also consider trade-based ROI and a standardized directional-flow measure (Z-Flow), which adjusts the observed number of buys relative to the buy probability estimated from matched control observations.

For each stock i in interval p on day t , we construct a leave-one-out industry peer-flow measure as the value-weighted order flow of all other stocks belonging to the same Fama–French 12 industry.¹¹ We estimate:

$$OF_{i,t,p} = \alpha + \beta_1 Outage_t * OF_{Ind,t,p} + \beta_2 Outage_t + \beta_3 OF_{Ind,t,p} + \gamma Controls_{i,t,p} + \gamma_i + \delta_{t,p} + \varepsilon_{i,t,p} \quad (7)$$

The coefficient on OF_{Ind} captures the extent to which a stock's retail order flow comoves with that of its industry peers during control periods, while $Outage \times OF_{Ind}$

¹¹ Our main tests use the Fama–French 12-industry classification and we require the industry-interval to contain at least 100 other stocks. The findings still hold when using the Fama–French 30- and 48-industry classifications.

captures the change in this commonality during ChatGPT outages. Following Harford and Kaul (2005), we control for lagged own-stock order flow and lagged industry order flow, lagged own-stock returns and lagged leave-one-out industry returns.

[Insert Table 6 here]

Panel A of Table 6 summarizes the results. The coefficients on Outage $\times OF_{Ind}$ are consistently negative across the specifications (STD_NET: -0.120, $t=-2.09$; ROI: -0.123, $t=-2.41$; Z-flow: -0.100, $t=-2.60$), indicating that the commonality of retail order flow across stocks in the same industry declines significantly when ChatGPT is unavailable.

We conduct an analogous analysis at the market level, replacing the industry measures with value-weighted leave-one-out measures constructed from all other stocks in the market.¹² Panel B reports negative interaction coefficients across all three measures, although the statistical evidence is weaker (STD_NET: -0.040, $t=-0.63$; ROI: -0.170, $t=-2.37$; Z-flow: -0.086, $t=-1.70$).

Taken together with the within-stock evidence in Section 4.2.1, these results are consistent with ChatGPT availability being associated with greater coordination in retail trading both within and across stocks.

4.3 The Effect of ChatGPT Outages on Retail Trading Quality

The documented increase in ChatGPT-related coordination in retail trading raises an important question for market efficiency: does this coordination reflect improved information processing and contribute to price discovery, or does it instead come at the expense of heterogeneous information and amplify correlated noise?

¹² We require the market-interval contains at least 500 other stocks.

To examine this, following prior work (e.g., Barber et al., 2022; Farrell et al., 2022; Havakhor et al., 2025), we test whether ROI predicts future returns and how this predictability changes with ChatGPT availability. If ChatGPT improves retail investors' information processing, greater coordination may reflect convergence toward informative signals, consistent with a crowding-in effect. In this case, ROI should effectively and persistently predict future returns when ChatGPT is available, with predictability weakening during outages. If, instead, ChatGPT crowds out investors' independent information production and the common signal is less informative, ROI on normal days should predict short-run returns that subsequently reverse, while becoming more informative during outages.

4.3.1 Return Predictability of Retail Order Imbalance

The effect of ChatGPT on trading informativeness may differ between selling and buying because the two decisions face different information constraints. Retail selling is generally limited to securities already held and may therefore depend more on investors' own information production about those positions. Retail buying often faces substantial search and attention constraints when selecting among a large universe of stocks and may possess less prior information about potential purchases (Barber and Odean, 2008; Barber et al., 2022). We allow for different marginal effects of ChatGPT on the two sides by decomposing trade-based ROI into selling and buying components. Specifically, selling ROI retains its original value when ROI is negative and is set to zero otherwise, while buying ROI retains its original value when ROI is positive and is set to zero otherwise (Chordia et al., 2002).¹³

Because our identification is based on intraday outage windows and key trading variables are at the 5-minute frequency, our primary analysis focuses on forward intraday returns over horizons from 5 to 30 minutes. We complement these tests with multi-day returns over one, three, and five trading days using event-level ROI to

¹³ We repeat all tests using volume-based ROI and obtain similar results.

provide suggestive evidence on whether short-run return predictability persists or reverses.¹⁴

Following Boehmer et al. (2021), we adopt a two-stage Fama–MacBeth (1973) approach. For each event and 5-minute interval, we estimate cross-sectional return-predictability regressions separately for the outage day and each matched control day. We then compare the ROI slope during the outage with the average slope across matched control days. The reported Outage $\times$ ROI coefficient therefore captures the outage-minus-control change in return predictability. Standard errors are Newey-West with lags equal to the horizon of the dependent variable. We estimate:

$$\begin{aligned}
 Ret_{i,t,[p+1, p+n]} &= \alpha + \beta_1 Outage_{t,p} * SellROI_{i,t,p} \\
 &+ \beta_2 Outage_{t,p} * BuyROI_{i,t,p} + \beta_3 SellROI_{i,t,p} \\
 &+ \beta_4 BuyROI_{i,t,p} + \beta_5 Outage_{t,p} + \gamma Controls_{i,t} + \varepsilon_{i,t,p} \quad (8)
 \end{aligned}$$

Panel A of Table 7 reports the intraday results. On the sell side, Selling ROI significantly predicts returns at short intraday horizons (0.005, $t = 4.60$ at 5 minutes; 0.004, $t = 2.51$ at 10 minutes; 0.005, $t = 2.18$ at 15 minutes) on control days, while it loses its significance at longer intraday horizons. The pattern changes when ChatGPT is unavailable. The coefficients on Outage $\times$ Selling ROI are positive and statistically significant across all intraday horizons, and the total predictive slope during outages is roughly twice its control-period value at short horizons and remains significant throughout. Thus, retail selling appears relatively more informative during outages.

The buy side exhibits the opposite pattern. On control days, Buying ROI predicts subsequent returns positively across intraday horizons. However, the interaction coefficients are negative and become significant from 15 minutes onward (-0.008, $t = -1.83$ at 15 minutes; -0.011, $t = -2.07$ at 20 minutes; -0.015, $t = -2.47$ at 25 minutes; -

¹⁴ For the intraday analysis, matched controls are drawn from trading days $t-5$ through $t-1$, consistent with our other tests. For the multi-day analysis, we instead use $t-10$ through $t-6$, which prevents the forward return windows of control observations from overlapping the outage date.

0.018, $t = -2.36$ at 30 minutes). Consequently, the total predictive slope during outages becomes insignificant from 20 minutes onward. Therefore, the results suggest that retail buying becomes less informative when ChatGPT is unavailable.

Panel B summarizes the multi-day results. ROI is aggregated over the event window, and future returns are calculated cumulatively from the end of the event window to the end of the first, third, and fifth subsequent trading day. We apply the same two-stage Fama–MacBeth procedure to event-date cross-sections.

Combining the intraday and multi-day results suggests a reversal pattern for retail selling on control days: Selling ROI predicts short-run returns positively, loses significance at longer intraday horizons, and turns significantly negative over the multi-day horizons (-0.002 , $t = -2.15$ at D+1; -0.008 , $t = -3.15$ at D+3; -0.009 , $t = -3.01$ at D+5). Such a reversal is more consistent with temporary price pressure or noise trading pattern (Foucault et al., 2011; Jones et al., 2025) than with permanent price discovery (Barber et al., 2008; Barber et al., 2022). Importantly, this reversal is substantially attenuated when ChatGPT is unavailable: the Outage $\times$ Selling ROI coefficients are positive at D+1 and D+5 (0.007 , $t=6.47$; 0.011 , $t=2.03$), and the reversal observed on control days largely disappears during outages.

The multi-day evidence for retail buying is weaker: Buying ROI predicts future returns at D+1 and D+3 on control days, whereas the interaction coefficients and the total predictive slope during outages are insignificant across all three horizons. Therefore, the evidence that ChatGPT improves the informativeness of retail buying comes primarily from the intraday tests.

[Insert Table 7 here]

Overall, Table 7 documents a sell–buy asymmetry in the effect of ChatGPT on retail trading informativeness. Retail selling exhibits a noise-trading pattern on control days

when ChatGPT is available and becomes relatively more informative during outages. Retail buying, in contrast, is more informative when ChatGPT is available.

This asymmetry may reflect differences in how ChatGPT interacts with investors' existing information sets. On the sell side, where investors may already rely substantially on their own information processing, a low-cost shared signal from ChatGPT may substitute for, or reduce the weight placed on, investors' own costly information, consistent with a crowding-out channel. On the buy side, where investors may possess less prior information about potential purchases, ChatGPT may instead facilitate independent information production, consistent with a crowding-in channel.

We next investigate this explanation by decomposing ROI to identify the underlying mechanism through which ChatGPT affects retail trading informativeness.

4.3.2 Return Predictability of Decomposed ROI

Prior research identifies three main sources of retail order imbalance predictability: persistence in retail order flow (Chordia and Subrahmanyam, 2004), liquidity provision through contrarian trading (Kaniel et al., 2008), and informed trading (Kelley and Tetlock, 2013). If ChatGPT affects return predictability through changes in information processing, the outage effects documented in Table 7 should therefore be concentrated in the information-related component of ROI.

To investigate this, we follow Boehmer et al. (2021) and decompose ROI into three components: $ROI_{persistence}$, a proxy for price pressure; $ROI_{liquidity}$, a proxy for liquidity provision; and ROI_{other} , a residual component that may contain information relevant to future returns.¹⁵ The components are estimated from the following panel regression:

¹⁵ Prior studies use the residual component of ROI as a proxy for informed trading (e.g., Farrell et al., 2022; Goyal et al., 2025). However, the residual may also capture other unobserved trading motives, we interpret ROI_{other} as an information-related residual rather than a pure measure of informed trading

$$ROI_{i,t,p} = \alpha + \beta_1 ROI_{i,t,[p-3,p-1]} + \beta_2 Ret_{i,t,[p-3,p-1]} + \varepsilon_{i,t,p} \quad (9)$$

where $\widehat{ROI}_{persistence} = \hat{\beta}_1 ROI_{i,t,[p-3,p-1]}$, $\widehat{ROI}_{liquidity} = \hat{\beta}_2 Ret_{i,t,[p-3,p-1]}$ and

$$\widehat{ROI}_{other_{i,t,p}} = ROI_{i,t,p} - \hat{\alpha} - \widehat{ROI}_{liquidity_{i,t,p}} - \widehat{ROI}_{persistence_{i,t,p}}.^{16}$$

We then re-estimate the return-predictability regressions separately for selling and buying ROI, interacting each decomposed component with the outage indicator:

$$\begin{aligned} Ret_{i,t,[p+1, p+n]} &= \alpha + \beta_1 Outage_{t,p} * ROI_{other_{i,t,p}} + \beta_2 Outage_{t,p} * ROI_{liquidity_{i,t,p}} \\ &+ \beta_3 Outage_{t,p} * ROI_{persistence_{i,t,p}} + \beta_4 ROI_{other_{i,t,p}} \\ &+ \beta_5 ROI_{liquidity_{i,t,p}} + \beta_6 ROI_{persistence_{i,t,p}} + \beta_7 Outage_{t,p} \\ &+ \gamma Controls_{i,t} + \varepsilon_{i,t,p} \end{aligned} \quad (10)$$

Table 8 shows that the outage effect on return predictability is concentrated in ROI_{other} . In Panel A, the coefficients on Outage x ROI_{other} are positive and statistically significant across horizons on the sell side, indicating that the information-related component of retail selling becomes more predictive when ChatGPT is unavailable. In Panel B, the corresponding coefficients are significantly negative on the buy side, indicating weaker predictability during outages. Meanwhile, the interactions with the other two sources $ROI_{persistence}$ and $ROI_{liquidity}$ are insignificant on both sides.

[Insert Table 8 here]

Taken together, these results suggest that the effect of ChatGPT on retail trading informativeness operates primarily through the information-related component of retail order flow.

¹⁶ We also use $ROI_{i,t,[p-5,p-1]}$ and $Ret_{i,t,[p-5,p-1]}$ for decomposition, and the findings still hold.

4.3.3 Common and Individual Components of Information-Related ROI

Having established that ChatGPT's effect on retail trading informativeness operates primarily through the information-related channel, we next examine how ChatGPT interacts with investors' existing information sets. We further decompose ROI_{other} into a common component ROI_{Other}^{Common} and an individual component $ROI_{Other}^{Individ}$, following the cross-stock commonality approach in Section 4.2.2. ROI_{Other}^{Common} captures the portion of information-related retail order flow that comoves across stocks and therefore serves as a proxy for trading on common signals, whereas $ROI_{Other}^{Individ}$ captures the remaining heterogeneous component of information-related retail order flow.

We interpret this decomposition cautiously. $ROI_{Other}^{Individ}$ is not a direct measure of investors' private signal; it reflects investors' own information acquisition and processing, but also behavioral biases and other idiosyncratic trading motives. Similarly, ROI_{Other}^{Common} need not arise exclusively from ChatGPT. However, the decomposition allows us to examine whether and how ChatGPT changes the composition of information-related retail trading and the informativeness of its common and individual components.

We first examine whether the composition of ROI_{other} changes during ChatGPT outages. If ChatGPT shifts investors' reliance toward low-cost shared signals, then the relative importance of ROI_{Other}^{Common} should be greater when ChatGPT is available and decline during outages. Moreover, if the shared signals substitute for, rather than add to, individual information, a decline in ROI_{Other}^{Common} should be accompanied by an increase in $ROI_{Other}^{Individ}$ while their combined magnitude $ROI_{Other}^{Overall}$ remains broadly unchanged.

Table 9 reports the absolute magnitude of each component, their combined magnitude, and the share attributable to the common component for the full sample (Panel A),

selling side (Panel B), and buying side (Panel C).¹⁷ Panel A shows that, during outages, the absolute ROI_{Other}^{Common} decreases by 0.015 ($t=-3.55$), and its share by 0.019 ($t=-3.22$). Meanwhile, the absolute $ROI_{Other}^{Individ}$ increases by 0.010 ($t=2.03$), and $ROI_{Other}^{Overall}$ has no significant change ($t=-0.74$). This pattern is consistent with a substitution interpretation within information-related retail order flow. Panels B and C show a similar pattern for both the selling and buying sides.

[Insert Table 9 here]

Taken together, Table 9 indicates that ChatGPT outages primarily alter the composition of information-related retail order flow rather than its overall magnitude. When ChatGPT is available, the common component accounts for a relatively larger share and substitutes for the individual component. This evidence complements Section 4.2, which documents greater correlated retail trading when ChatGPT is available. Table 9 further shows that this greater commonality exists in the information-related component of retail order flow, consistent with ChatGPT shifting investors' reliance toward common signals and away from more heterogeneous individual signals.

We then examine whether the informativeness of these two components changes during ChatGPT outages. We re-estimate the return-predictability regressions separately for selling and buying ROI, interacting ROI_{Other}^{Common} , $ROI_{Other}^{Individ}$, $ROI_{persistence}$, and $ROI_{liquidity}$ with the Outage indicator, as in equation (10). The results are shown in Table 10.

¹⁷ We use absolute values because ROI_{Other}^{Common} and $ROI_{Other}^{Individ}$ are signed measures, whereas our interest here is in the relative importance of each component irrespective of trading direction. We measure the combined magnitude as $|ROI_{Other}^{Overall}| = |ROI_{Other}^{Individ}| + |ROI_{Other}^{Common}|$, rather than $|ROI_{Other}|$. If the two components have opposite signs, then $|ROI_{Other}|$ would reflect netting between the components rather than their total magnitude. Therefore, $|ROI_{Other}^{Overall}|$ would provide a more appropriate denominator for assessing how the composition of information-related order flow shifts between the common and individual components.

[Insert Table 10 here]

Table 10 shows that the outage effect on return predictability is concentrated in $ROI_{Other}^{Individ}$ on both sides, whereas ROI_{Other}^{Common} provides little evidence of differential return predictability. On the sell side, the coefficients on Outage x $ROI_{Other}^{Individ}$ are positive across all horizons, indicating that the individual information-related component of retail selling becomes more informative when ChatGPT is unavailable. This result is consistent with the crowding-out interpretation. When ChatGPT is available, the shift toward common inputs documented in Table 9 may reduce the relative weight placed on investors' own heterogeneous signal, which may contain valuable information. When ChatGPT becomes unavailable, the individual component regains importance and retail selling becomes more informative.

The buy side exhibits the opposite pattern. The negative coefficients on Outage x $ROI_{Other}^{Individ}$ indicate that the individual information-related component of retail buying becomes less informative when ChatGPT is unavailable. This result supports the crowding-in interpretation. Retail buyers may have relatively limited prior information about many potential purchases. Thus, although ChatGPT reduces the relative weight of the individual component, it may simultaneously improve the informational quality of that component by facilitating information acquisition, organization, and evaluation. Its absence therefore reduces the informational quality of retail buying.

Importantly, across the specifications, the predictive slopes of ROI_{Other}^{Common} on control and outage days provide little evidence of return predictability, and the coefficients on Outage x ROI_{Other}^{Common} are statistically insignificant. Thus, the common component itself appears to contribute little to return predictability, which is also broadly consistent with the finding in Boehmer et al. (2021) that the predictive power of marketable retail order imbalance is not readily attributable to identifiable public information. This suggests that the greater commonality associated with ChatGPT does not directly translate into more informative retail order flow. Instead, the informativeness effects appear to operate through the individual component.

Overall, Section 4.3 shows that ChatGPT has asymmetric effects on retail trading informativeness through the information-related component of retail order flow. When ChatGPT is available, the common component largely substitutes for the individual component, yet itself provides little evidence of informativeness. The change in predictability is instead concentrated in the individual component, and in opposite directions across the two sides. On the sell side, ChatGPT is associated with a less informative individual component, consistent with a crowding-out effect. On the buy side, it improves the informativeness of that component, consistent with a crowding-in effect. Together, the results suggest that ChatGPT’s informational effect appears to depend on how it interacts with investors’ existing information sets.

4.4 The Effect of ChatGPT Outages on Market Quality

Section 4.2 establishes that ChatGPT availability is associated with greater coordination in retail trading. Prior research shows that correlated trading across securities can generate commonality in liquidity, potentially contributing to systematic liquidity risk (Chordia et al., 2000; Kamara et al., 2008; Karolyi et al., 2012; Koch et al., 2016). We therefore examine whether ChatGPT-induced coordination in retail trading translates into broader liquidity effects.

We first test whether market liquidity changes during ChatGPT outages by estimating:

$$ILLIQ_{i,t,p} = \alpha + \beta_1 Outage_{t,p} + \gamma Controls_{i,t} + \gamma_i + \delta_{t,p} + \varepsilon_{i,t,p} \quad (11)$$

where *ILLIQ* is one of three illiquidity measures: effective spread, realized spread and price impact.

[Insert Table 11 here]

Table 11 shows that liquidity improves when ChatGPT is unavailable, with the effect concentrated primarily in price impact, consistent with the findings in Cheng et al. (2025) and Ecker et al. (2026). Specifically, effective spread declines by 0.010 (t=-

1.91), corresponding to a 6.3% reduction relative to the sample mean, and price impact falls by 12.5% (-0.010, $t=-3.32$). However, realized spread changes little and is statistically insignificant (0.002, $t=1.28$). Thus, the lower trading cost during outages primarily reflects smaller post-trade price movements rather than a change in the realized compensation retained by liquidity providers.

The price impact result connects closely to the trading patterns documented in Section 4.2. ChatGPT makes retail order flow more persistent within stocks, so directional buying or selling pressure is less likely to reverse quickly. As a result, the midpoint stays displaced for longer, leading to the higher price impact observed when ChatGPT is available.

Then we examine whether ChatGPT also affects liquidity commonality. Following Chordia et al. (2000) and Koch et al. (2016), we regress interval changes in stock-level liquidity ($\Delta ILLIQ_{i,t,p} = ILLIQ_{i,t,p} - ILLIQ_{i,t,p-1}$) on contemporaneous changes in leave-one-out industry liquidity and their interaction with the outage indicator.¹⁸ Industry liquidity is constructed as the value-weighted liquidity of all other stocks in the same Fama–French 12 industry:

$$\begin{aligned} \Delta ILLIQ_{i,t,p} = & \alpha + \beta_1 Outage_t * \Delta ILLIQ_IND_{t,p} + \beta_2 \Delta ILLIQ_IND_{t,p} + \beta_3 Outage_t \\ & + \gamma Controls_{i,t,p} + \gamma_i + \delta_{t,p} + \varepsilon_{i,t,p} \end{aligned} \quad (12)$$

The coefficient on the $\Delta ILLIQ_IND$ captures industry-level liquidity commonality during control periods, while its interaction with *Outage* captures the change in commonality when ChatGPT is unavailable. Following the liquidity commonality literature, we control for lead, lag, and contemporaneous industry returns, contemporaneous squared stock returns, and lead and lag changes in industry liquidity.

¹⁸ Following Chordia et al. (2000), we use changes rather than levels of liquidity because our interest is in liquidity comovement and because liquidity levels are more susceptible to time-series econometric concerns, including non-stationarity.

We conduct an analogous analysis using a leave-one-out market liquidity measure, with results reported in Panel B of Table 12.

[Insert Table 12 here]

Table 12 summarizes the results. At the industry level, the coefficients on $\text{Outage} \times \Delta \text{ILLIQ_IND}$ are negative across all three measures (ΔESPD : -0.197, $t=-1.81$; ΔRSPD : -0.223, $t=-2.57$; ΔPI : -0.234, $t=-3.00$). Thus, industry-level liquidity commonality declines when ChatGPT is unavailable, with the evidence particularly strong for realized spread and price impact.

The market-level results point in the same direction. The outage interactions are significantly negative for realized spread and price impact (ΔRSPD : -0.302, $t=-2.42$; ΔPI : -0.285, $t=-2.36$). The effective spread interaction is also negative (-0.386, $t=-1.62$), although statistically insignificant. Overall, ChatGPT availability is associated with greater liquidity commonality both within industries and, to a somewhat weaker extent, across the broader market.¹⁹

Taken together, Tables 11 and 12 show that ChatGPT-related coordination has consequences beyond retail trading behavior. When ChatGPT is available, price impact is higher, and liquidity commonality increases, potentially increasing exposure to systematic liquidity risk.

5 Conclusion

In this paper, we exploit 14 unexpected full ChatGPT outages during 2023–2024 to examine whether access to a common GenAI platform coordinates retail trading and its

¹⁹ Importantly, liquidity commonality is measured as a loading rather than a level (Chordia et al., 2000; Kamara et al., 2008). A reduction in retail trading activity lowers the magnitude of retail liquidity demand without altering its correlation structure, leaving the loading unchanged. The reduction in commonality during outages is therefore unlikely to reflect the decline in retail trading activity documented in Section 4.1.

implications for trading informativeness and market quality. We find that ChatGPT availability is associated with greater correlation in retail trading, asymmetric effects on the informativeness of selling and buying, and higher liquidity commonality.

Our empirical analysis yields three main findings. First, retail trading is more coordinated both within and across stocks when ChatGPT is available. Second, ChatGPT has asymmetric effects on trading informativeness. Further analysis suggests that these effects are concentrated in the information-related component of retail order flow. ChatGPT availability shifts this component toward common inputs and away from individual inputs. However, the common component itself provides little evidence of return predictability, and changes in informativeness are concentrated in the individual component, consistent with a crowding-out effect on retail selling and a crowding-in effect on retail buying. This asymmetry may reflect differences in how a shared, low-cost information tool interacts with investors' existing information sets. Third, ChatGPT availability is associated with higher liquidity commonality, suggesting greater exposure to common liquidity shocks.

Our study contributes to the emerging literature on GenAI and financial markets by extending the focus to the relation among investors' trading decisions, while also investigating the underlying effect mechanism. This perspective provides new empirical evidence on the widely discussed welfare implications of GenAI reliance. The policy implications are also important. Regulators have raised concerns that GenAI-induced herding could become a source of systemic vulnerability, but direct empirical evidence has been lacking. Our results provide evidence consistent with these concerns. As GenAI tools grow in sophistication and retail penetration, and as more trading decisions are delegated to AI-powered systems, the scale of these effects is likely to grow.

Several limitations provide opportunities for future research. Our analysis focuses on ChatGPT over a specific period and is based on a limited set of unexpected outage events, although the high-frequency setting allows us to exploit relatively sharp

disruptions to platform availability. Investors may also substitute toward other GenAI tools during ChatGPT outages, which would tend to attenuate our estimates of the broader effects of GenAI availability. In addition, we do not observe the content of ChatGPT outputs or query-level usage during our sample period, preventing us from directly identifying the specific information through which common reliance affects trading decisions. Nor do we observe account-level trades, which limits our ability to link an individual investor's GenAI use directly to changes in trading behavior. Finally, our analysis emphasizes stocks with meaningful retail trading activity, and the effects may differ in other market segments. Future research using platform-level or account-level data could further identify how investors interact with GenAI and when common AI-generated information complements or substitutes for investors' own information production.

Reference

- Bank for International Settlements. (2024). *Artificial intelligence and the economy: Implications for central banks* (Chapter III, *Annual Economic Report 2024*). <https://www.bis.org/publ/arpdf/ar2024e3.htm>
- Bank of England. (2025). *Financial stability in focus: Artificial intelligence in the financial system*. <https://www.bankofengland.co.uk/financial-stability-in-focus/2025/april-2025>
- Barber, B. M., & Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *The Review of Financial Studies*, 21(2), 785-818.
- Barber, B. M., Odean, T., & Zhu, N. (2008). Do retail trades move markets? *The Review of Financial Studies*, 22(1), 151-186.
- Barber, B. M., Huang, X., Odean, T., & Schwarz, C. (2022). Attention-induced trading and returns: Evidence from Robinhood users. *The Journal of Finance*, 77(6), 3141-3190.
- Barber, B. M., Huang, X., Jorion, P., Odean, T., & Schwarz, C. (2024). A (sub) penny for your thoughts: Tracking retail investor activity in TAQ. *The Journal of Finance*, 79(4), 2403-2427.
- 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*, 101782.
- Blankespoor, E., Croom, J., & Grant, S. M. (2026). Generative AI and investor processing of financial information. *Journal of Accounting and Economics*, 101908.
- Boehmer, E., Jones, C. M., Zhang, X., & Zhang, X. (2021). Tracking retail investor activity. *The Journal of Finance*, 76(5), 2249-2305.
- Bogousslavsky, V., & COLLIN-DUFRESNE, P. I. E. R. R. E. (2023). Liquidity, volume, and order imbalance volatility. *The Journal of Finance*, 78(4), 2189-2232.
- Bradshaw MT, Ma C, Yost BP, Zou Y. Generative AI use by capital market information intermediaries: Evidence from Seeking Alpha. *Journal of Accounting Research*. 2026 Jun;64(3):1233-86.
- Cao, S. S., Chen, W. X., Ma, G., & Srinivasan, S. (2026). Generative AI in capital markets: Information production, dissemination, and processing. *Journal of Accounting Research*, 64(3), 1427-1450.
- Chatterji, A., Cunningham, T., Deming, D. J., Hitzig, Z., Ong, C., Shan, C. Y., & Wadman, K. (2025). *How people use chatgpt* (No. w34255). National Bureau of Economic Research.
- Chang, A. Y., Dong, X., Martin, X., & Zhou, C. (2026). AI democratization and trading inequality. *Journal of Accounting Research*, 64(3), 1287-1331.

- Chelley-Steeley, P., Lambertides, N., & Savva, C. S. (2019). Sentiment, order imbalance, and co-movement: An examination of shocks to retail and institutional trading activity. *European Financial Management*, 25(1), 116-159.
- Cheng, Q., Lin, P., & Zhao, Y. (2025). Does generative AI facilitate investor trading? Early evidence from ChatGPT outages. *Journal of Accounting and Economics*, 101821.
- Choi JH, Hickman KE, Monahan AB, Schwarcz D. ChatGPT goes to law school. *J. Legal Educ.*. 2021;71:387.
- Choi JH, Xie CL. Human+ AI in accounting: Early evidence from the field. *Journal of Accounting Research*. 2026 Jun;64(3):1333-73.
- Chordia, T., Roll, R., & Subrahmanyam, A. (2000). Commonality in liquidity. *Journal of Financial Economics*, 56(1), 3-28.
- Chordia, T., Roll, R., & Subrahmanyam, A. (2002). Order imbalance, liquidity, and market returns. *Journal of Financial Economics*, 65(1), 111-130.
- Chordia, T., & Subrahmanyam, A. (2004). Order imbalance and individual stock returns: Theory and evidence. *Journal of Financial Economics*, 72(3), 485-518.
- Colla, P., & Mele, A. (2010). Information linkages and correlated trading. *The Review of Financial Studies*, 23(1), 203-246.
- Croom J. Interactivity and illusions of ability: How using generative AI affects investor judgments. *Journal of Accounting Research*. 2026 May;64(2):681-719.
- De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of Political Economy*, 98(4), 703-738.
- Deloitte Report (2024). Retail investors may soon rely on generative AI for financial advice. <https://www2.deloitte.com/us/en/insights/industry/financial-services/financial-services-industry-predictions/2024/ai-financial-advisor-for-retail-investment.html?id=us:2el:3dp:wsjspon:awa:WSJCIO:2024:WSJFY25>
- Dorn, D., Huberman, G., & Sengmueller, P. (2008). Correlated trading and returns. *The Journal of Finance*, 63(2), 885-920.
- Dugast, J., & Foucault, T. (2018). Data abundance and asset price informativeness. *Journal of Financial Economics*, 130(2), 367-391.
- 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.
- Ecker, F., Li, X., Li, Y., & Wu, F. (2026). How Stock Market Participants Use Generative Artificial Intelligence: Evidence from User-Platform Interaction Data. *Journal of Accounting Research*, 64(3), 1375-1426.
- eToro. (2023, July 11). Four in ten U.S. retail investors use or are open to using ChatGPT-style AI to help pick and manage investments <https://www.etoro.com/en-us/news-and-analysis/latest-news/press-release/four->

[in-ten-u-s-retail-investors-use-or-are-open-to-using-chat-gpt-style-ai-to-help-pick-and-manage-investments/](#)

- European Securities and Markets Authority. (2023). *The financial stability implications of artificial intelligence*. https://www.esma.europa.eu/sites/default/files/library/ESMA50-164-6247-AI_in_securities_markets.pdf
- Even-Tov, O., Lourie, B., & Nekrasov, A. (2025). The Effect of AI on Retail Investor Behavior: Evidence from Account-Level Trading Data. *Available at SSRN 5778246*.
- Farrell, M., Green, T. C., Jame, R., & Markov, S. (2022). The democratization of investment research and the informativeness of retail investor trading. *Journal of Financial Economics*, 145(2), 616-641.
- Feng, L., & Seasholes, M. S. (2004). Correlated trading and location. *The Journal of Finance*, 59(5), 2117-2144.
- Financial Stability Board. (2024). *The financial stability implications of artificial intelligence*. <https://www.fsb.org/uploads/P14112024.pdf>
- Fisher, A., Martineau, C., & Sheng, J. (2022). Macroeconomic attention and announcement risk premia. *The Review of Financial Studies*, 35(11), 5057-5093.
- Foucault, T., Sraer, D., & Thesmar, D. J. (2011). Individual investors and volatility. *The Journal of Finance*, 66(4), 1369-1406.
- Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American Economic Review*, 70(3), 393-408.
- 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.
- Hansen, A. L., & Lee, S. J. (2025). Financial Stability Implications of Generative AI: Taming the Animal Spirits. *arXiv preprint arXiv:2510.01451*.
- Harford, J., & Kaul, A. (2005). Correlated order flow: Pervasiveness, sources, and pricing effects. *Journal of Financial and Quantitative Analysis*, 40(1), 29-55.
- Havakhori, T., Rahman, M. S., Zhang, T., & Zhu, C. (2025). Tech-enabled financial data access, retail investors, and gambling-like behavior in the stock market. *Management Science*, 71(2), 1646-1670.
- Hirshleifer, D., & Teoh, S. H. (2003). Limited attention, information disclosure, and financial reporting. *Journal of Accounting and Economics*, 36(1-3), 337-386.
- International Monetary Fund. (2024). *GLOBAL FINANCIAL STABILITY REPORT: Steadying the Course: Uncertainty, Artificial Intelligence, and Financial Stability*. <https://www.imf.org/-/media/files/publications/gfsr/2024/october/english/textrevised.pdf>

- Jones, C. M., Shi, D., Zhang, X., & Zhang, X. (2025). Retail trading and return predictability in China. *Journal of Financial and Quantitative Analysis*, 60(1), 68-104.
- Kamara, A., Lou, X., & Sadka, R. (2008). The divergence of liquidity commonality in the cross-section of stocks. *Journal of Financial Economics*, 89(3), 444-466.
- Kaniel, R., Saar, G., & Titman, S. (2008). Individual investor trading and stock returns. *The Journal of Finance*, 63(1), 273-310.
- 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.
- Kelley, E. K., & Tetlock, P. C. (2013). How wise are crowds? Insights from retail orders and stock returns. *The Journal of Finance*, 68(3), 1229-1265.
- Koch, A., Ruenzi, S., & Starks, L. (2016). Commonality in liquidity: a demand-side explanation. *The Review of Financial Studies*, 29(8), 1943-1974.
- Kumar, A., & Lee, C. M. (2006). Retail investor sentiment and return comovements. *The Journal of Finance*, 61(5), 2451-2486.
- Lakonishok, J., Shleifer, A., & Vishny, R. W. (1992). The impact of institutional trading on stock prices. *Journal of Financial Economics*, 32(1), 23-43.
- Lee, C. M., Shleifer, A., & Thaler, R. H. (1991). Investor sentiment and the closed-end fund puzzle. *The Journal of Finance*, 46(1), 75-109.
- Lee, C. M., & Ready, M. J. (1991). Inferring trade direction from intraday data. *The Journal of Finance*, 46(2), 733-746.
- Liu, Yan; Huang, Jingyun; Wang, He. (2025). Who on Earth Is Using Generative AI? Global Trends and Shifts in 2025. Policy Research Working Paper; 11231. © World Bank. <http://hdl.handle.net/10986/43859> License: [CC BY 3.0 IGO](#).
- Lucca, D. O., & Moench, E. (2015). The pre-FOMC announcement drift. *The Journal of Finance*, 70(1), 329-371.
- Morris, S., & Shin, H. S. (2002). Social value of public information. *American Economic Review*, 92(5), 1521-1534.
- Oliver Wyman (2024). How generative AI is transforming business and society: The good, the bad, and everything in between. <https://www.oliverwymanforum.com/content/dam/oliver-wyman/ow-forum/gcs/2023/AI-Report-2024-Davos.pdf>
- Reuters, “ChatGPT sets record for fastest-growing user base - analyst note,” Reuters, February 2 2023.
- Savor, P., & Wilson, M. (2013). How much do investors care about macroeconomic risk? Evidence from scheduled economic announcements. *Journal of Financial and Quantitative Analysis*, 48(2), 343-375.

- Shleifer, A., & Summers, L. H. (1990). The noise trader approach to finance. *Journal of Economic Perspectives*, 4(2), 19-33.
- Veldkamp, L. L. (2006). Information markets and the comovement of asset prices. *The Review of Economic Studies*, 73(3), 823-845.

Table 1: Summary Statistics

	N	Mean	SD	p25	p50	p75
Panel A. 5-Minute Interval-Level Measures						
BEME	3,052,058	0.58	0.56	0.20	0.42	0.77
Effective Spread (%)	3,052,058	0.16	0.32	0.04	0.07	0.15
Ln (Intraday Volatility)	3,052,058	-10.42	1.66	-11.60	-10.58	-9.48
Ln (ME)	3,052,058	8.12	1.91	6.88	8.11	9.36
Ln (Number of Analysts)	3,052,058	2.17	0.74	1.79	2.20	2.71
Ln (Price)	3,052,058	3.52	1.31	2.57	3.59	4.47
LSV	1,855,017	0.10	0.16	-0.01	0.07	0.19
Price Impact (%)	3,008,200	0.08	0.19	0.00	0.03	0.10
Realized Spread (%)	3,008,200	0.06	0.22	-0.02	0.03	0.09
Retail Trade Count	3,052,058	20.59	37.89	3.00	8.00	20.00
Retail Trade Volume (10,000 shares)	3,052,058	0.31	1.92	0.01	0.04	0.14
Return [-1] (%)	3,052,058	-0.26	3.05	-1.70	-0.14	1.19
Return [-5, -2] (%)	3,052,058	-1.43	6.31	-4.55	-1.26	1.51
Return Overnight (%)	3,052,058	-0.01	1.40	-0.64	0.00	0.56
Signed Trade-based ROI	3,052,058	0.01	0.61	-0.39	0.00	0.43
B. Event-Day-Level Measures						
Excess ROIV Variance	84,424	0.09	0.14	0.02	0.07	0.15
ROIV	84,424	0.45	0.21	0.30	0.43	0.57
ROI Pairwise Correlation	74,357	0.03	0.31	-0.15	0.05	0.23
ROI Autocorrelation	74,357	0.03	0.24	-0.12	0.04	0.20

This table reports summary statistics for the main variables used in the analysis. Panel A reports variables constructed at the stock-event-day-interval level, while Panel B reports variables constructed at the stock-event-day level. Continuous variables are winsorized at the 1st and 99th percentiles. Sample sizes vary across variables because some measures are constructed at higher levels of aggregation or require additional data availability. See Appendix B for detailed variable definitions.

Table 2: ChatGPT Outages and Retail Trading Intensity

	Ln Retail Trade Count	Ln Retail Trade Volume
Outage	-0.071** (-2.33)	-0.130*** (-2.98)
Controls	Yes	Yes
Firm FE	Yes	Yes
Event x interval FE	Yes	Yes
Observations	3,052,058	3,052,058
Adj R2	0.669	0.528

The dependent variables are Ln (Retail Trade Count) and Ln (Retail Trade Volume). The key independent variable is Outage, an indicator equal to one if the 5-minute interval overlaps a ChatGPT outage window and zero otherwise. Control variables include Return Overnight, Return [-1], Return [-5, -2], Ln (Intraday Volatility), Ln (ME), BEME, Ln (Price), and Ln (Number of Analysts). Variable definitions are provided in Appendix B. All regressions include firm fixed effects and event-by-interval fixed effects. Standard errors are two-way clustered by firm and event-by-date. t-statistics are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

.

Table 3: Robustness Tests

	Ln Retail Trade Count	Ln Retail Trading Volume
Panel A. Exclude Firm-Specific News Events		
Outage	-0.072** (-2.35)	-0.134*** (-3.01)
Controls	Yes	Yes
Firm FE	Yes	Yes
Event x interval FE	Yes	Yes
Observations	2,658,846	2,658,846
Adj R2	0.659	0.517
Panel B. Exclude Macroeconomic Announcements		
Outage	-0.071** (-2.00)	-0.141*** (-2.79)
Controls	Yes	Yes
Firm FE	Yes	Yes
Event x interval FE	Yes	Yes
Observations	2,639,272	2,639,272
Adj R2	0.669	0.529
Panel C. Exclude Outages Shorter than 30 Minutes		
Outage	-0.074** (-2.37)	-0.132*** (-2.90)
Controls	Yes	Yes
Firm FE	Yes	Yes
Event x interval FE	Yes	Yes
Observations	2,952,403	2,952,403
Adj R2	0.670	0.527
Panel D. Exclude Outages Beginning Before 10:00 AM		
Outage	-0.093** (-2.08)	-0.171*** (-2.70)
Controls	Yes	Yes
Firm FE	Yes	Yes
Event x interval FE	Yes	Yes
Observations	1,859,477	1,859,477
Adj R2	0.679	0.555
Panel E. Pseudo-Outage Test		
Outage	0.045 (1.66)	0.051 (1.15)
Controls	Yes	Yes
Firm FE	Yes	Yes
Event x interval FE	Yes	Yes
Observations	2,608,424	2,608,424
Adj R2	0.687	0.538

This table reports robustness tests for the baseline specification in Table 2. Panel A excludes observations with firm-specific information events. Panel B excludes outage events that overlap major scheduled U.S. macroeconomic announcements (CPI, PPI, and FOMC minutes). Panel C excludes outages shorter than 30 minutes. Panel D excludes outages beginning before 10:00 AM ET. Panel E reports a placebo test that assigns the outage window to day $t-1$ (pseudo-outage). Control variables, fixed effects, clustering, and significance notation follow Table 2. Variable definitions are provided in Appendix B.

Table 4. ChatGPT Outages and Retail Trade-Side Concentration

	LSV
Outage	-0.003** (-2.26)
Controls	Yes
Firm FE	Yes
Event x interval FE	Yes
Observations	1,855,017
Adj R2	0.075

This table reports OLS regression results on the effect of ChatGPT outages on retail trade-side concentration measured by the LSV measure at 5-minute intervals. The LSV construction requires at least five retail trades per interval and at least 10 retail trades in the matched-control benchmark. The key independent variable is Outage, an indicator equal to one if the 5-minute interval overlaps a ChatGPT outage window and zero otherwise. Control variables follow the baseline specification. Variable definitions are provided in Appendix B. The regression includes firm fixed effects and event-by-interval fixed effects. Standard errors are two-way clustered by firm and event-by-date. t-statistics are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 5. ChatGPT Outages and Retail Order Flow Dynamics

Panel A. Retail Order Imbalance Volatility		
	Ln ROIV	Excess ROIV Variance
Outage	0.034*** (5.20)	0.006** (2.21)
Controls	Yes	Yes
Firm FE	Yes	Yes
Event FE	Yes	Yes
Observations	84,424	84,424
Adj R2	0.545	0.209
Panel B. Retail Order Imbalance Persistence		
	ROI Pairwise Correlation	ROI Autocorrelation
Outage	-0.010** (-2.11)	-0.008* (-1.86)
Controls	Yes	Yes
Firm FE	Yes	Yes
Event FE	Yes	Yes
Observations	74,357	74,357
Adj R2	0.082	0.093

This table reports OLS regression results on the effect of ChatGPT outages on the time-series dynamics of retail order flow. Panel A examines retail order imbalance volatility using Ln (ROIV) and Excess ROIV Variance. Panel B examines retail order imbalance persistence using ROI Pairwise Correlation and ROI Autocorrelation. All dependent variables are constructed at the stock-event-day level and require at least three valid adjacent intervals. The key independent variable is Outage, an indicator equal to one for the outage day and zero for matched control days. Control variables follow the baseline specification. Variable definitions are provided in Appendix B. All regressions include firm fixed effects and event fixed effects. Standard errors are two-way clustered by firm and event-by-date. t-statistics are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 6. ChatGPT Outages and Cross-Stock Correlated Retail Trading

	STD_NET	ROI	Z-Flow
Panel A. Industry-Level Correlation			
Outage	0.001 (0.10)	0.003 (0.36)	-0.050 (-1.17)
OF_IND	0.329*** (8.53)	0.669*** (12.16)	0.277*** (6.83)
Outage x OF_IND	-0.120** (-2.09)	-0.123** (-2.41)	-0.100** (-2.60)
Harford-Kaul Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Event x interval FE	Yes	Yes	Yes
Observations	2,119,954	2,119,954	2,119,954
Adj R2	0.028	0.053	0.040
Panel B. Market-Level Correlation			
Outage	0.016 (1.54)	0.011 (1.44)	-0.039 (-0.96)
OF_MKT	0.536*** (13.90)	1.131*** (16.75)	0.372*** (8.09)
Outage x OF_MKT	-0.040 (-0.63)	-0.170** (-2.37)	-0.086* (-1.70)
Harford-Kaul Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Event x interval FE	Yes	Yes	Yes
Observations	2,624,121	2,624,121	2,624,121
Adj R2	0.032	0.056	0.042

This table reports OLS regression results on the effect of ChatGPT outages on correlated retail trading across stocks. The dependent variables are three measures of retail order flow: STD_NET, ROI, and Z-Flow. Panel A examines industry-level correlation using industry peer order flow (OF_IND), and Panel B examines market-level correlation using market-wide order flow (OF_MKT). The key independent variable is the interaction between Outage indicator and peer order flow. OF_IND is a leave-one-out industry peer-flow measure, defined as the value-weighted order flow of all other stocks in the same Fama–French 12 industry, and requires at least 100 other stocks in the industry-interval. OF_MKT is a leave-one-out market peer-flow measure, defined as the value-weighted order flow of all other stocks in the market, and requires at least 500 other stocks in the market-interval. Following Harford and Kaul (2005),

control variables include lagged own-stock order flow, lagged industry or market order flow, lagged own-stock returns, and lagged leave-one-out industry or market returns. Variable definitions are provided in Appendix B. All regressions include firm fixed effects and event-by-interval fixed effects. Standard errors are two-way clustered by firm and event-by-date. t-statistics are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 7. ChatGPT Outages and Return Predictability of Retail Order Imbalance

Panel A. Intraday Return Predictability						
	Return t, [p, p+n]					
	5 min	10 min	15 min	20 min	25 min	30 min
Selling ROI	0.005*** (4.60)	0.004** (2.51)	0.0051** (2.18)	0.0039 (1.33)	0.003 (0.85)	0.001 (0.35)
Buying ROI	0.010*** (10.27)	0.013*** (8.12)	0.014*** (7.18)	0.017*** (7.05)	0.018*** (6.09)	0.018*** (5.49)
Outage	0.029** (2.16)	0.0421 (1.64)	0.0603 (1.55)	0.0810 (1.51)	0.0916 (1.32)	0.092 (1.09)
Outage x Selling ROI	0.006** (2.46)	0.007** (2.11)	0.010* (1.88)	0.014** (2.18)	0.018** (2.35)	0.021** (2.26)
Outage x Buying ROI	-0.004* (-1.87)	-0.005 (-1.47)	-0.008* (-1.83)	-0.011** (-2.07)	-0.015** (-2.47)	-0.018** (-2.36)
<i>Selling ROI: during outage</i>	<i>0.010*** (5.38)</i>	<i>0.012*** (4.07)</i>	<i>0.015*** (3.65)</i>	<i>0.018*** (3.55)</i>	<i>0.021*** (3.49)</i>	<i>0.022*** (3.13)</i>
<i>Buying ROI: during outage</i>	<i>0.007*** (4.12)</i>	<i>0.008*** (3.17)</i>	<i>0.007* (1.93)</i>	<i>0.006 (1.40)</i>	<i>0.003 (0.65)</i>	<i>0.001 (0.12)</i>
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,000,601	2,958,493	2,917,984	2,878,209	2,841,551	2,805,406
Average Adj R2	0.038	0.042	0.045	0.046	0.047	0.048

Panel B. Multi-Day Return Predictability			
	Return [Event End, D+N]		
	End to D+1	End to D+3	End to D+5
Selling ROI	-0.002* (-2.15)	-0.008** (-3.15)	-0.009** (-3.01)
Buying ROI	0.002* (2.49)	0.006* (2.46)	0.005 (1.99)
Outage	0.022** (3.15)	0.019 (1.29)	0.053* (2.30)
Outage x Selling ROI	0.007*** (6.47)	0.006 (0.99)	0.011* (2.03)
Outage x Buying ROI	-0.002 (-1.38)	-0.002 (-1.38)	-0.004 (-1.03)
<i>Selling ROI: during outage</i>	<i>0.004* (2.44)</i>	<i>-0.002 (-0.57)</i>	<i>0.003 (0.75)</i>
<i>Buying ROI: during outage</i>	<i>-0.001 (-0.24)</i>	<i>0.004 (0.77)</i>	<i>0.000 (0.07)</i>
Controls	Yes	Yes	Yes
Observations	91,906	86,082	74,570
Average Adj R2	-0.002*	0.037	0.037

This table reports two-stage Fama–MacBeth regression results examining whether retail order imbalance predicts future returns and how this predictability changes during ChatGPT outages. Panel A examines intraday return predictability at the 5-minute level, with matched controls drawn

from days $t-5$ through $t-1$. The dependent variable, $\text{Return } t, [p, p+n]$, is the intraday return from the end of interval p to n minutes after interval p , where n ranges from 5 to 30 minutes. Panel B examines multi-day return predictability at the event-day level, with matched controls drawn from days $t-10$ through $t-6$ to avoid contamination from the outage event. The dependent variable, $\text{Return} [\text{Event End}, D+N]$, is the cumulative return from the end of the event window to the end of the N th subsequent trading day, where $N=1, 3$ or 5 . The key independent variables are the interactions of the Outage indicator with Selling ROI and Buying ROI. Selling ROI equals ROI when ROI is negative and zero otherwise, while Buying ROI equals ROI when ROI is positive and zero otherwise. The table also reports Selling ROI: during outage and Buying ROI: during outage, which represent the corresponding total predictive slopes during outages. C_{control} variables follow the baseline specification. Variable definitions are provided in Appendix B. Standard errors are Newey–West with lags corresponding to the return horizon. t -statistics are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 8. ChatGPT Outages and Return Predictability of Decomposed Retail Order Imbalance

	Return t, [p, p+n]					
	5 min	10 min	15 min	20 min	25 min	30 min
Panel A. Selling-Side Predictability of ROI Components						
<i>ROI_{persistence}</i>	0.002 (0.50)	0.007 (1.47)	0.013** (2.31)	0.014* (1.91)	0.016* (1.80)	0.010 (0.99)
<i>ROI_{liquidity}</i>	-0.556*** (-2.60)	-1.315*** (-4.70)	-1.540*** (-4.22)	-1.481*** (-3.26)	-1.227** (-2.30)	-0.852 (-1.31)
<i>ROI_{other}</i>	0.004*** (4.30)	0.004*** (2.68)	0.005** (2.37)	0.005** (2.34)	0.006*** (2.63)	0.005* (1.70)
Outage	0.023* (1.74)	0.034 (1.30)	0.062 (1.63)	0.071 (1.36)	0.088 (1.31)	0.099 (1.20)
Outage x <i>ROI_{persistence}</i>	0.004 (0.62)	0.007 (0.88)	0.010 (0.92)	0.006 (0.38)	0.013 (0.79)	0.022 (1.19)
Outage x <i>ROI_{liquidity}</i>	0.272 (0.76)	0.309 (0.57)	0.318 (0.42)	0.280 (0.30)	0.266 (0.24)	-0.173 (-0.14)
Outage x <i>ROI_{other}</i>	0.004** (2.13)	0.005* (1.73)	0.009** (2.23)	0.012*** (2.63)	0.013** (2.49)	0.016** (2.48)
<i>ROI_{persistence}: during outage</i>	0.006 (1.14)	0.014* (1.89)	0.023** (2.50)	0.020* (1.66)	0.029** (2.13)	0.032** (2.21)
<i>ROI_{liquidity}: during outage</i>	-0.284 (-0.96)	-1.006** (-2.08)	-1.222* (-1.80)	-1.201 (-1.39)	-0.962 (-0.91)	-1.026 (-0.87)
<i>ROI_{other}: during outage</i>	0.008*** (5.12)	0.009*** (3.60)	0.014*** (4.10)	0.017*** (4.40)	0.019*** (4.29)	0.021*** (4.04)

Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,217,505	1,198,249	1,179,979	1,161,556	1,144,708	1,128,865
Average Adj R2	0.050	0.055	0.058	0.059	0.059	0.060
Panel B. Buying-Side Predictability of ROI Components						
<i>ROI_{persistence}</i>	0.007** (2.55)	0.010** (2.08)	0.011 (1.65)	0.013 (1.54)	0.008 (0.84)	0.007 (0.60)
<i>ROI_{liquidity}</i>	-0.601*** (-3.37)	-1.247*** (-4.54)	-1.676*** (-5.28)	-1.712*** (-4.17)	-1.616*** (-3.59)	-1.332*** (-2.73)
<i>ROI_{other}</i>	0.008*** (8.02)	0.0010*** (6.35)	0.010*** (5.29)	0.012*** (5.23)	0.013*** (5.01)	0.012*** (4.21)
Outage	0.022 (1.59)	0.037 (1.40)	0.047 (1.17)	0.060 (1.10)	0.063 (0.90)	0.055 (0.64)
Outage x <i>ROI_{persistence}</i>	-0.002 (-0.32)	-0.000 (-0.03)	-0.003 (-0.31)	-0.010 (-0.77)	-0.013 (-0.94)	-0.013 (-0.79)
Outage x <i>ROI_{liquidity}</i>	0.072 (0.21)	0.444 (0.76)	0.550 (0.72)	0.452 (0.47)	0.502 (0.45)	0.299 (0.24)
Outage x <i>ROI_{other}</i>	-0.004** (-2.23)	-0.006** (-2.21)	-0.008* (-1.95)	-0.009* (-1.84)	-0.013** (-2.33)	-0.013** (-2.03)
<i>ROI_{persistence}: during outage</i>	0.006 (1.32)	0.010 (1.39)	0.008 (0.86)	0.003 (0.29)	-0.005 (-0.41)	-0.006 (-0.42)
<i>ROI_{liquidity}: during outage</i>	-0.529* (-1.76)	-0.804 (-1.53)	-1.125 (-1.49)	-1.260 (-1.29)	-1.114 (-0.99)	-1.033 (-0.82)
<i>ROI_{other}: during outage</i>	0.004** (2.32)	0.003 (1.32)	0.003 (0.75)	0.003 (0.80)	0.000 (0.01)	-0.002 (-0.34)
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Observations	1,328,271	1,308,546	1,289,403	1,271,594	1,254,909	1,237,814
Average Adj R2	0.050	0.055	0.057	0.058	0.059	0.060

his table reports two-stage Fama–MacBeth regression results examining how the return predictability of decomposed retail order imbalance changes during ChatGPT outages. Following Boehmer et al. (2021), ROI is decomposed into a Persistence component captured by ROI $[p-3, p-1]$, a Liquidity component captured by Return $[p-3, p-1]$, and an Other component equal to the residual after removing the Persistence and Liquidity components and the intercept. The decomposition coefficients are estimated using matched control observations and then applied to both outage and control observations. Panel A examines the selling-side return predictability of these ROI components, while Panel B examines their buying-side return predictability. The dependent variable, Return $t, [p, p+n]$, is the intraday return from the end of interval p to n minutes after interval p , where n ranges from 5 to 30 minutes. The key independent variables are the interactions between the Outage indicator and each ROI component. The table also reports the corresponding total predictive slopes during outages. Control variables follow the baseline specification. Variable definitions are provided in Appendix B. Standard errors are Newey–West with lags corresponding to the return horizon. t -statistics are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 9: ChatGPT Outages and the Composition of Information-Related ROI

Panel A. Full Sample					
	$ ROI_{Other}^{Common} $	$ ROI_{Other}^{Individ} $	$ ROI_{Other}^{Overall} $	$ ROI_{Other}^{Common} $	Share
Outage	-0.0146*** (-3.55)	0.0103** (2.03)	-0.0043 (-0.74)	-0.0190*** (-3.22)	
Controls	Yes	Yes	Yes	Yes	
Firm FE	Yes	Yes	Yes	Yes	
Event x interval FE	Yes	Yes	Yes	Yes	
Observations	2,843,214	2,843,214	2,843,214	2,843,214	
Adj R2	0.0036	0.0091	0.0052	0.0932	
Panel B. Selling Side					
	$ ROI_{Other}^{Common} $	$ ROI_{Other}^{Individ} $	$ ROI_{Other}^{Overall} $	$ ROI_{Other}^{Common} $	Share
Outage	-0.0103*** (-3.35)	0.0068* (1.80)	-0.0035 (-0.76)	-0.0124*** (-3.20)	
Controls	Yes	Yes	Yes	Yes	
Firm FE	Yes	Yes	Yes	Yes	
Event x interval FE	Yes	Yes	Yes	Yes	
Observations	1,238,132	1,238,132	1,238,132	1,238,132	
Adj R2	0.0349	0.1277	0.0697	0.1286	
Panel C. Buying Side					
	$ ROI_{Other}^{Common} $	$ ROI_{Other}^{Individ} $	$ ROI_{Other}^{Overall} $	$ ROI_{Other}^{Common} $	Share
Outage	-0.0134*** (-3.07)	0.0217*** (2.74)	0.0083 (1.25)	-0.0250*** (-2.85)	
Controls	Yes	Yes	Yes	Yes	

Firm FE	Yes	Yes	Yes	Yes
Event x interval FE	Yes	Yes	Yes	Yes
Observations	1,348,680	1,348,680	1,348,680	1,348,680
Adj R2	0.0267	0.0444	0.0318	0.1393

This table reports OLS regression results examining how the composition of information-related retail order imbalance changes during ChatGPT outages. Building on the decomposition in Table 8, the ROI_{other} is further decomposed following the cross-stock commonality approach in Section 4.2.2, using industry-level commonality based on the Fama–French 12 industry classification. The Common component (ROI_{Other}^{Common}) is the portion of ROI_{other} associated with the value-weighted leave-one-out ROI_{other} of other stocks in the same Fama–French 12 industry, using loadings estimated from matched control observations. The Individual component ($ROI_{Other}^{Individ}$) equals ROI_{other} minus ROI_{Other}^{Common} . The dependent variables are the absolute magnitudes of ROI_{Other}^{Common} , $ROI_{Other}^{Individ}$, their combined magnitude $ROI_{Other}^{Overall}$ ($|ROI_{Other}^{Common}| + |ROI_{Other}^{Individ}|$), together with the share of the ROI_{Other}^{Common} ($ROI_{Other}^{Common} share = |ROI_{Other}^{Common}| / |ROI_{Other}^{Overall}|$). Panel A reports results for the full sample, while Panels B and C report results for the selling and buying sides, respectively. The key independent variable is Outage indicator. Control variables follow the baseline specification. Variable definitions are provided in Appendix B. All regressions include firm fixed effects and event-by-interval fixed effects. Standard errors are two-way clustered by firm and event-by-date. t-statistics are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 10: Return Predictability of Common and Individual Components of Information-Related ROI

	Return t, [p, p+n]					
	5 min	10 min	15 min	20 min	25 min	30 min
Panel A. Selling-Side Predictability of ROI Components						
$ROI_{persistence}$	0.000 (0.06)	0.006 (1.43)	0.013** (2.25)	0.014* (1.90)	0.015* (1.73)	0.009 (0.86)
$ROI_{liquidity}$	-0.638*** (-3.01)	-1.33*** (-4.77)	-1.520*** (-4.20)	-1.354*** (-2.72)	-1.201** (-2.26)	-0.939 (-1.51)
ROI_{Other}^{Common}	0.004 (1.13)	0.009* (1.71)	0.005 (0.79)	-0.001 (-0.16)	-0.004 (-0.40)	-0.002 (-0.24)
$ROI_{Other}^{Individ}$	0.004*** (4.29)	0.004*** (2.75)	0.005** (2.44)	0.005** (2.31)	0.006*** (2.82)	0.005* (1.86)
Outage	0.024* (1.81)	0.035 (1.32)	0.063 (1.63)	0.070 (1.34)	0.089 (1.32)	0.097 (1.18)
Outage x $ROI_{persistence}$	0.005 (0.90)	0.008 (0.90)	0.010 (0.91)	0.005 (0.36)	0.014 (0.82)	0.024 (1.26)
Outage x $ROI_{liquidity}$	0.392 (1.09)	0.366 (0.68)	0.354 (0.47)	0.205 (0.21)	0.308 (0.28)	-0.027 (-0.02)
Outage x ROI_{Other}^{Common}	-0.001 (-0.14)	0.000 (0.04)	0.002 (0.13)	0.004 (0.27)	0.003 (0.15)	0.005 (0.24)
Outage x $ROI_{Other}^{Individ}$	0.004** (2.19)	0.005* (1.68)	0.009** (2.22)	0.013*** (2.71)	0.013** (2.48)	0.016** (2.52)
$ROI_{persistence}$: during outage	0.005 (1.05)	0.014* (1.90)	0.022** (2.46)	0.019 (1.63)	0.028** (2.11)	0.033** (2.21)

$ROI_{liquidity}: \text{during outage}$	-0.246 (-0.83)	-0.959** (-1.98)	-1.1662* (-1.71)	-1.150 (-1.31)	-0.893 (-0.83)	-0.966 (-0.81)
$ROI_{Other}^{Common}: \text{during outage}$	0.003 (0.37)	0.009 (0.90)	0.0070 (0.57)	0.003 (0.23)	-0.001 (-0.07)	0.002 (0.13)
$ROI_{Other}^{Individ}: \text{during outage}$	0.008*** (5.23)	0.009*** (3.61)	0.0138*** (4.17)	0.018*** (4.50)	0.019*** (4.34)	0.021*** (4.10)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,216,451	1,197,205	1,178,960	1,160,539	1,143,699	1,127,884
Average Adj R2	0.052	0.056	0.060	0.060	0.060	0.061

Panel B. Buying-Side Predictability of ROI Components

$ROI_{persistence}$	0.007** (2.40)	0.010** (2.09)	0.011 (1.61)	0.012 (1.48)	0.008 (0.78)	0.007 (0.61)
$ROI_{liquidity}$	-0.579*** (-3.24)	-1.237*** (-4.49)	-1.676*** (-5.25)	-1.696*** (-4.13)	-1.621*** (-3.59)	-1.318*** (-2.70)
ROI_{Other}^{Common}	0.005 (1.39)	0.006 (1.07)	0.004 (0.52)	0.004 (0.46)	-0.002 (-0.26)	-0.003 (-0.32)
$ROI_{Other}^{Individ}$	0.008*** (8.03)	0.009*** (6.28)	0.010*** (5.21)	0.012*** (5.15)	0.013*** (4.88)	0.012*** (4.13)
Outage	0.023 (1.64)	0.038 (1.43)	0.048 (1.20)	0.063 (1.16)	0.066 (0.94)	0.057 (0.66)
Outage x $ROI_{persistence}$	-0.002 (-0.33)	-0.001 (-0.13)	-0.004 (-0.36)	-0.011 (-0.82)	-0.013 (-0.94)	-0.014 (-0.88)
Outage x $ROI_{liquidity}$	0.057 (0.17)	0.440 (0.75)	0.579 (0.76)	0.489 (0.50)	0.568 (0.51)	0.291 (0.23)
Outage x ROI_{Other}^{Common}	0.000	-0.000	-0.005	0.003	-0.010	-0.005

	(0.02)	(-0.02)	(-0.35)	(0.16)	(-0.61)	(-0.31)
Outage x $ROI_{Other}^{Individ}$	-0.004**	-0.006**	-0.008*	-0.009*	-0.013**	-0.014**
	(-2.28)	(-2.18)	(-1.94)	(-1.85)	(-2.32)	(-2.04)
$ROI_{persistence}$: during outage	0.005	0.009	0.007	0.002	-0.005	-0.007
	(1.22)	(1.27)	(0.75)	(0.17)	(-0.46)	(-0.51)
$ROI_{liquidity}$: during outage	-0.522*	-0.797	-1.098	-1.207	-1.054	-1.027
	(-1.74)	(-1.51)	(-1.45)	(-1.23)	(-0.93)	(-0.81)
ROI_{Other}^{Common} : during outage	0.005	0.006	-0.001	0.007	-0.013	-0.008
	(0.76)	(0.55)	(-0.09)	(0.47)	(-0.82)	(-0.49)
$ROI_{Other}^{Individ}$: during outage	0.004**	0.003	0.002	0.003	-0.000	-0.002
	(2.26)	(1.29)	(0.71)	(0.74)	(-0.02)	(-0.36)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,327,111	1,307,424	1,288,295	1,270,522	1,253,843	1,236,770
Average Adj R2	0.051	0.057	0.058	0.060	0.060	0.061

his table reports two-stage Fama–MacBeth regression results examining how the return predictability of the common and individual components of information-related ROI changes during ChatGPT outages. Building on the decomposition in Table 8, the ROI_{Other} is further decomposed following the cross-stock commonality approach in Section 4.2.2, using industry-level commonality based on the Fama–French 12 industry classification. The Common component (ROI_{Other}^{Common}) is the portion of ROI_{Other} associated with the value-weighted leave-one-out ROI_{Other} of other stocks in the same Fama–French 12 industry, using loadings estimated from matched control observations. The Individual component ($ROI_{Other}^{Individ}$) equals ROI_{Other} minus ROI_{Other}^{Common} . Panel A examines the selling-side return predictability of these components, while Panel B

examines their buying-side return predictability. The dependent variable, Return $t, [p, p+n]$, is the intraday return from the end of interval p to n minutes after interval p , where n ranges from 5 to 30 minutes. The key independent variables are the interactions between the Outage indicator and each ROI component. The table also reports the corresponding total predictive slopes during outages. Control variables follow the baseline specification. Variable definitions are provided in Appendix B. Standard errors are Newey–West with lags corresponding to the return horizon. t -statistics are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 11: ChatGPT Outages and Market Liquidity

	Effective Spread	Realized Spread	Price Impact
Outage	-0.010* (-1.91)	0.002 (1.28)	-0.010*** (-3.32)
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Event x interval FE	Yes	Yes	Yes
Observations	3,052,058	3,008,200	3,008,200
Adj R2	0.530	0.273	0.120

This table reports OLS regression results on the effect of ChatGPT outages on market liquidity measured at 5-minute intervals. The dependent variables are Effective Spread, Realized Spread, and Price Impact. The key independent variable is Outage indicator. Control variables follow the baseline specification. Variable definitions are provided in Appendix B. All regressions include firm fixed effects and event-by-interval fixed effects. Standard errors are two-way clustered by firm and event-by-date. t-statistics are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 12: ChatGPT Outages and Liquidity Commonality

	ΔESPD	ΔRSPD	ΔPI
Panel A. Industry-Level Liquidity Commonality			
Outage	0.000 (1.20)	-0.000 (-0.72)	0.000 (0.97)
$\Delta\text{ILLIQ_IND}$	0.262** (2.23)	0.409*** (9.04)	0.392*** (9.39)
Outage x $\Delta\text{ILLIQ_IND}$	-0.197* (-1.81)	-0.223** (-2.57)	-0.234*** (-3.00)
Chordia Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Event x interval FE	Yes	Yes	Yes
Observations	1,991,415	1,957,712	1,957,712
Adj R2	0.008	0.110	0.136
Panel B. Market-Level Liquidity Commonality			
Outage	0.000 (0.78)	-0.000 (-0.72)	0.000 (0.87)
$\Delta\text{ILLIQ_MKT}$	0.554*** (2.88)	0.499*** (8.68)	0.474*** (8.83)
Outage x $\Delta\text{ILLIQ_MKT}$	-0.386 (-1.62)	-0.302** (-2.42)	-0.285** (-2.36)
Chordia Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Event x interval FE	Yes	Yes	Yes
Observations	2,467,835	2,427,694	2,427,694
Adj R2	0.009	0.109	0.135

This table reports OLS regression results on the effect of ChatGPT outages on liquidity commonality measured at 5-minute intervals. The dependent variables are changes in three illiquidity measures (Effective Spread, Realized Spread and Price Impact) calculated as $\Delta\text{ILLIQ}_{i,t,p} = \text{ILLIQ}_{i,t,p} - \text{ILLIQ}_{i,t,p-1}$. Panel A examines industry-level liquidity commonality using $\Delta\text{ILLIQ_IND}$, the contemporaneous change in the value-weighted leave-one-out liquidity measure of all other stocks in the same Fama–French 12 industry, requiring at least 100 industry peers in the same interval. Panel B examines market-level liquidity commonality using $\Delta\text{ILLIQ_MKT}$, requiring at least 500 market peers in the same interval. The key independent variable is the interaction of Outage indicator with $\Delta\text{ILLIQ_IND}$ ($\Delta\text{ILLIQ_MKT}$). Following Chordia et al. (2000), control

variables include lead, lag, and contemporaneous industry or market returns, contemporaneous squared own-stock returns, and lead and lag changes in industry or market liquidity, as appropriate. Variable definitions are provided in Appendix B. All regressions include firm fixed effects and event-by-interval fixed effects. Standard errors are two-way clustered by firm and event-by-date. t-statistics are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Appendices

Appendix A: ChatGPT Outage Events

This appendix lists the ChatGPT outage events in our sample. We collect outage dates and timing from OpenAI's public incident information. For outages that extend beyond regular trading hours (9:30–16:00 ET), we truncate the outage window to trading hours by setting the start time to 9:30 and/or the end time to 16:00. We aggregate trading data into 5-minute intervals and treat any interval that overlaps with the outage window as part of the outage period.

Date	Starting Time (ET)	Ending Time (ET)	Duration (min)	Day of Week
02/14/2023	13:55	14:37	42	Tues
02/27/2023	09:30	13:30	240	Mon
03/15/2023	14:00	14:18	18	Wed
03/20/2023	12:20	16:00	220	Mon
05/24/2023	15:20	16:00	40	Wed
06/30/2023	13:10	15:20	130	Fri
08/29/2023	09:55	11:16	81	Tues
08/31/2023	09:30	12:29	179	Thu
09/13/2023	09:30	11:01	91	Wed
11/08/2023	09:30	10:16	46	Wed
05/21/2024	14:25	15:26	61	Tues
06/04/2024	10:14	13:07	173	Tues
08/21/2024	12:14	12:35	21	Wed
12/26/2024	13:40	16:00	140	Thu
Avg Duration			105.86	
Sum Duration			1482	

Appendix B: Variable Definitions

Variable	Definition
Book-to-Market (BEME)	Ratio of the book value of equity to the market value of equity.
Effective Spread	Dollar-volume-weighted average of the signed effective spread in a 5-minute interval. For each trade, $Effective\ Spread = 2 \times D \times \frac{(P_t - M_t)}{M_t} \times 100$, where D equals 1 for a buy and -1 for a sell, P_t is the trade price, and M_t is the prevailing midquote. Trades are signed using Lee and Ready (1991).
Excess ROIV Variance	Difference between observed ROIV variance and its independent-sign benchmark, calculated as $ROIV_{i,t}^2 - ROIV_{Nul,i,t}^2$. Construction requires at least 3 consecutive valid 5-minute intervals and 10 benchmark trades.
Ln (Intraday Volatility)	Natural logarithm of intraday volatility multiplied by 100.
Ln (ME)	Natural logarithm of market equity, where market equity equals stock price $\times$ shares outstanding.
Ln (Number of Analysts)	Natural logarithm of one plus the number of analysts following the firm.
Ln (Price)	Natural logarithm of one plus the absolute daily stock price.
LSV	Trade-count-based LSV adapted from Lakonishok et al., (1992). For stock i, event e, and 5-minute interval t, $LSV_{i,t,p} = p_{i,t,p} - E[p_{i,t,p}] - E p_{i,t,p} - E[p_{i,t,p}] $, where $p_{i,t,p}$ denotes the proportion of retail buys, $E[p_{i,t,p}]$ represents the expected buy proportion, estimated from matched control intervals. $E p_{i,t,p} - E[p_{i,t,p}] $ is the expected absolute deviation of the realized buy proportion from the expected buy proportion. Construction requires at least 5 retail trades in the interval and 10 benchmark trades.
OF_IND	Value-weighted leave-one-out retail order flow of all other stocks in the same Fama–French 12 industry and 5-minute interval; requires at least 100 industry peers.

OF_MKT	Value-weighted leave-one-out retail order flow of all other stocks in the market and 5-minute interval; requires at least 500 market peers.
Outage	Indicator equal to one for observations in the ChatGPT outage window on the event day and zero for matched control observations.
Price Impact	Equal-weighted average of price impact in a 5-minute interval. For each trade, $Price\ Impact = 2 \times D \times \frac{(M_{t+5} - M_t)}{M_t} \times 100$, where D equals 1 for a buy and -1 for a sell, M_t is the prevailing midquote, and M_{t+5} is the midquote 5 minutes after the trade. Trades are signed using Lee and Ready (1991).
Realized Spread	Equal-weighted average of realized spread in a 5-minute interval. For each trade, $Realized\ Spread = 2 \times D \times \frac{(P_t - M_{t+5})}{M_t} \times 100$, where P_t is the trade price and M_{t+5} is the midquote 5 minutes after the trade. Trades are signed using Lee and Ready (1991).
Retail Order Imbalance (ROI)	Signed imbalance scaled by total retail activity. Trade-based $ROI = (RetailBuyCount - RetailSellCount) / (RetailBuyCount + RetailSellCount)$. Selling $ROI = \min(ROI, 0)$, and Buying $ROI = \max(ROI, 0)$.
Retail Order Imbalance Volatility (ROIV)	Following Bogousslavsky and Collin-Dufresne (2023), ROIV is the standard deviation of 5-minute trade-based retail order imbalance over the event window. Construction requires at least 3 consecutive valid 5-minute intervals.
Retail Trade Count	Number of retail trades in a 5-minute interval.
Retail Trade Volume	Total share volume traded by retail investors in a 5-minute interval, divided by 10,000.
Return [-1]	Stock return on trading day $t-1$.
Return [-5, -2]	Cumulative stock return from trading day $t-5$ through $t-2$.
Return Overnight	Stock return from the closing price on trading day $t-1$ to the opening price on trading day t .
Return t , $[p, p+n]$	Cumulative intraday return from the end of interval p to n minutes after interval.
$ROI_{liquidity}$	Following Boehmer et al. (2021), the component of ROI associated with recent return-driven liquidity provision: $ROI_{liquidity_{i,t,p}} = \widehat{\beta_R} \times Return_{i,t,[p-3,p-1]}$,

	where $\widehat{\beta}_R$ is estimated from matched control observations. Following Boehmer et al. (2021), the component of ROI associated with persistence in recent ROI: $ROI_{persistence_{i,t,p}} = \widehat{\beta}_P \times ROI_{i,t,[p-3,p-1]},$ where $\widehat{\beta}_P$ is estimated from matched control observations.
ROI_{other}	Following Boehmer et al. (2021), the intercept-adjusted component of trade-based ROI unexplained by persistence and liquidity provision: $ROI_{other_{i,t,p}} =$ $ROI_{i,t,p} - \hat{\alpha} - ROI_{liquidity_{i,t,p}} - ROI_{persistence_{i,t,p}}.$
ROI_{Other}^{Common}	The industry-peer-correlated component of $ROI_{other}: ROI_{Other_{i,t,p}}^{Common} = \hat{\lambda}_i \times$ $IND_ROI_{Other_{i,t,p}}, \text{ where } IND_ROI_{Other_{i,t,p}} \text{ is the}$ value-weighted leave-one-out average for other stocks in the same FF12 industry and interval, and $\hat{\lambda}_i$ is estimated from matched controls.
$ROI_{Other}^{Common} \text{ Share}$	The share of the combined absolute magnitude attributable to the Common component: $ROI_{Other}^{Common} \text{ Share} = \frac{ ROI_{Other}^{Common} }{(ROI_{Other}^{Common} + ROI_{Other}^{Individ})}$
$ROI_{Other}^{Individ}$	The stock-specific component of ROI_{Other} remaining after removing the Common component: $ROI_{Other_{i,t,p}}^{Individ} = ROI_{Other_{i,t,p}} - ROI_{Other_{i,t,p}}^{Common}.$
$ROI_{Other}^{Overall}$	The combined absolute magnitude of the Common and Individual components.
ROI Pairwise Correlation	Pearson correlation between ROI in adjacent 5-minute intervals within the event window, calculated as $\rho_{i,t}^{Pearson} = Corr(ROI_{i,t,p}, ROI_{i,t,p+1})$. Construction requires at least 6 scheduled intervals and 3 valid adjacent ROI pairs.
ROI Autocorrelation	Lag-1 autocorrelation of ROI within the event window, calculated as $\rho_{i,t}^{ACF} =$ $\frac{\sum_{p=1}^P (ROI_{i,t,p+1} - \overline{ROI}_{i,t})(ROI_{i,t,p} - \overline{ROI}_{i,t})}{\sum_{p=1}^P (ROI_{i,t,p} - \overline{ROI}_{i,t})^2}.$ Construction

	requires at least 6 scheduled intervals and 3 valid adjacent ROI pairs.
Standardized Net Retail Trading (STD_NET)	Net retail trade count standardized relative to its matched-control distribution: $STD_NET_{i,t} = \frac{NetNum_{i,t,p} - \overline{NetNum_{i,p}^{Ctrl}}}{SD^{Ctrl}(NetNum_{i,p})}, \text{ where}$ NetNum = BuyNum – SellNum, and the control mean and standard deviation are calculated from matched control observations for the same stock, event, and 5-minute interval.
Z-Flow	Retail buy imbalance standardized by its binomial benchmark: $ZFlow_{i,t,p} = \frac{Buy_{i,t,p} - N_{i,t,p}\widehat{P}_{i,p}}{\sqrt{N_{i,t,p}\widehat{P}_{i,p}(1-\widehat{P}_{i,p})}}, \text{ where}$ $N_{i,t,p}$ is the total number of retail trades, and $\widehat{P}_{i,p}$ is the expected retail buy probability estimated from matched control observations for the same stock, event, and 5-minute interval. Construction requires at least 10 retail trades in the matched-control benchmark.
$\Delta ILLIQ$	Interval changes in three illiquidity measures (Effective Spread, Realized Spread and Price Impact) calculated as $\Delta ILLIQ_{i,t,p} = ILLIQ_{i,t,p} - ILLIQ_{i,t,p-1}$.
$\Delta ILLIQ_IND$	Change in the value-weighted leave-one-out $\Delta ILLIQ$ of all other stocks in the same Fama–French 12 industry; requires at least 100 industry peers.
$\Delta ILLIQ_MKT$	Change in the value-weighted leave-one-out $\Delta ILLIQ$ of all other stocks in the market; requires at least 500 market peers.

Appendix C: Supplementary Robustness Tests

This appendix provides supplementary robustness tests for the main analyses in Sections 4.2–4.4. Tables C.1, C.2, and C.3 report robustness tests for Sections 4.2, 4.3, and 4.4, respectively. For each main test, we implement the same five robustness checks used in Table 3. Panel A excludes observations with firm-specific information events. Panel B excludes outage events that overlap major scheduled U.S. macroeconomic announcements, including CPI, PPI, and FOMC minutes. Panel C excludes outages shorter than 30 minutes. Panel D excludes outages beginning before 10:00 AM ET. Panel E conducts a placebo test by assigning the outage window to day $t-1$ (pseudo-outage). Unless otherwise stated, the dependent variables, key independent variables, control variables, fixed effects, standard-error adjustments, and significance notation follow the corresponding main specification. Variable definitions are provided in Appendix B.

Table C.1: Robustness Tests for Section 4.2

	Coeff. on Outage / Interactions with Outage	t-statistic
Panel A. Exclude Firm-Specific News Events		
LSV	-0.004***	(-3.07)
Ln ROIV	0.038***	(5.23)
Excess ROIV Variance	0.008**	(2.50)
ROI Autocorrelation	-0.010*	(-1.86)
Outage x OF_IND	-0.136**	(-2.32)
Outage x OF_MKT	-0.205**	(-2.24)
Panel B. Exclude Macroeconomic Announcements		
LSV	-0.003**	(-2.07)
Ln ROIV	0.033***	(3.94)
Excess ROIV Variance	0.008**	(2.16)
ROI Autocorrelation	-0.006	(-1.13)
Outage x OF_IND	-0.148***	(-2.80)
Outage x OF_MKT	-0.223***	(-3.05)
Panel C. Exclude Outages Shorter than 30 Minutes		
LSV	-0.003**	(-2.41)
Ln ROIV	0.034***	(4.48)
Excess ROIV Variance	0.003	(1.27)
ROI Autocorrelation	-0.010**	(-2.11)
Outage x OF_IND	-0.121**	(-2.31)
Outage x OF_MKT	-0.165**	(-2.18)
Panel D. Exclude Outages Beginning Before 10:00 AM		
LSV	-0.002	(-1.22)
Ln ROIV	0.040***	(4.98)
Excess ROIV Variance	0.008**	(2.41)
ROI Autocorrelation	-0.009	(-1.39)
Outage x OF_IND	-0.130**	(-2.27)
Outage x OF_MKT	-0.219**	(-2.45)
Panel E. Pseudo-Outage Test		
LSV	-0.000	(-0.11)
Ln ROIV	0.004	(0.46)
Excess ROIV Variance	-0.001	(-0.38)
ROI Autocorrelation	0.0001	(0.03)
Outage x OF_IND	0.061	(1.27)
Outage x OF_MKT	0.057	(0.78)

Table C.2: Robustness Tests for Section 4.3

Robustness Tests (A) for Table 7						
	Return t, [p, p+n]					
	5 min	10 min	15 min	20 min	25 min	30 min
Panel A. Exclude Firm-Specific News Events						
Outage x Selling ROI	0.006** (2.54)	0.009** (2.30)	0.010* (1.86)	0.014** (2.17)	0.018** (2.30)	0.021** (2.25)
Outage x Buying ROI	-0.004* (-1.80)	-0.006* (-1.68)	-0.008 (-1.63)	-0.011* (-1.69)	-0.014* (-1.90)	-0.018** (-2.00)
Panel B. Exclude Macroeconomic Announcements						
Outage x Selling ROI	0.005** (2.02)	0.008** (2.03)	0.008 (1.38)	0.011 (1.49)	0.010 (1.40)	0.015 (1.46)
Outage x Buying ROI	-0.003 (-1.33)	-0.006 (-1.42)	-0.006 (-1.05)	-0.007 (-1.06)	-0.009 (-1.15)	-0.013 (-1.34)
Panel C. Exclude Outages Shorter than 30 Minutes						
Outage x Selling ROI	0.006** (2.45)	0.008** (2.07)	0.010* (1.85)	0.014** (2.13)	0.018** (2.24)	0.020** (2.14)
Outage x Buying ROI	-0.004* (-1.91)	-0.005* (-1.68)	-0.009** (-2.08)	-0.012** (-2.29)	-0.016** (-2.59)	-0.019** (-2.46)
Panel D. Exclude Outages Beginning Before 10:00 AM						
Outage x Selling ROI	0.002 (0.71)	0.005 (1.15)	0.006 (1.19)	0.011* (1.66)	0.016* (1.88)	0.018* (1.73)
Outage x Buying ROI	-0.002 (-0.81)	-0.003 (-0.80)	-0.005 (-1.04)	-0.008 (-1.49)	-0.010* (-1.68)	-0.010 (-1.39)

Panel E. Pseudo-Outage Test						
Outage x Selling ROI	0.000 (0.06)	-0.000 (-0.06)	-0.001 (-0.22)	-0.000 (-0.08)	0.004 (0.64)	0.003 (0.41)
Outage x Buying ROI	-0.003 (-1.41)	-0.004 (-1.05)	-0.003 (-0.67)	-0.003 (-0.66)	-0.003 (-0.44)	-0.001 (-0.13)
Robustness Tests (B) for Table 8						
Panel A. Exclude Firm-Specific News Events						
Outage x ROI_{other} (Sell-side)	0.004** (2.18)	0.005 (1.51)	0.008** (1.97)	0.011** (2.38)	0.013** (2.45)	0.016** (2.47)
Outage x ROI_{other} (Buy-side)	-0.005** (-2.47)	-0.0067** (-2.22)	-0.007* (-1.80)	-0.008* (-1.67)	-0.013** (-2.28)	-0.013* (-1.91)
Panel B. Exclude Macroeconomic Announcements						
Outage x ROI_{other} (Sell-side)	0.004* (1.89)	0.006* (1.79)	0.008* (1.91)	0.011** (2.30)	0.010* (1.83)	0.013* (1.83)
Outage x ROI_{other} (Buy-side)	-0.004* (-1.93)	-0.006* (-1.92)	-0.007 (-1.53)	-0.007 (-1.37)	-0.010* (-1.66)	-0.010 (-1.42)
Panel C. Exclude Outages Shorter than 30 Minutes						
Outage x ROI_{other} (Sell-side)	0.004** (0.58)	0.006* (0.84)	0.010** (1.00)	0.013*** (0.40)	0.014** (0.70)	0.017** (0.97)
Outage x ROI_{other} (Buy-side)	-0.004** (-2.19)	-0.007** (-2.36)	-0.009** (-2.15)	-0.010** (-2.01)	-0.014** (-2.40)	-0.014** (-2.00)
Panel D. Exclude Outages Beginning Before 10:00 AM						
Outage x ROI_{other} (Sell-side)	0.002 (0.88)	0.004 (1.07)	0.007 (1.43)	0.010** (1.99)	0.013** (2.04)	0.016* (1.92)

Outage x ROI_{other} (Buy-side)	-0.003 (-1.15)	-0.005 (-1.56)	-0.007 (-1.56)	-0.010* (-1.83)	-0.013** (-2.17)	-0.011 (-1.60)
Panel E. Pseudo-Outage Test						
Outage x ROI_{other} (Sell-side)	-0.001 (-0.47)	-0.002 (-0.65)	-0.006 (-1.48)	-0.006 (-1.33)	-0.005 (-0.86)	-0.007 (-1.14)
Outage x ROI_{other} (Buy-side)	-0.003 (-1.55)	-0.003 (-0.88)	-0.003 (-0.63)	-0.003 (-0.50)	-0.002 (-0.40)	-0.001 (-0.16)
Robustness Tests (C) for Table 10						
Panel A. Exclude Firm-Specific News Events						
Outage x $ROI_{other}^{Individ}$ (Sell-side)	0.004** (2.25)	0.004 (1.45)	0.008** (1.98)	0.011** (2.45)	0.013** (2.45)	0.016** (2.49)
Outage x $ROI_{other}^{Individ}$ (Buy-side)	-0.005** (-2.51)	-0.007** (-2.18)	-0.007* (-1.79)	-0.008* (-1.67)	-0.013** (-2.27)	-0.013* (-1.93)
Panel B. Exclude Macroeconomic Announcements						
Outage x $ROI_{other}^{Individ}$ (Sell-side)	0.004* (1.92)	0.006* (1.73)	0.008* (1.90)	0.011** (2.36)	0.010* (1.82)	0.013* (1.86)
Outage x $ROI_{other}^{Individ}$ (Buy-side)	-0.004** (-2.00)	-0.006* (-1.92)	-0.007 (-1.53)	-0.008 (-1.38)	-0.010* (-1.65)	-0.010 (-1.44)
Panel C. Exclude Outages Shorter than 30 Minutes						
Outage x $ROI_{other}^{Individ}$ (Sell-side)	0.004** (2.33)	0.006* (1.83)	0.010** (2.34)	0.014*** (2.82)	0.014** (2.53)	0.017*** (2.60)
Outage x $ROI_{other}^{Individ}$ (Buy-side)	-0.004** (-2.24)	-0.007** (-2.33)	-0.009** (-2.13)	-0.010** (-2.01)	-0.014** (-2.39)	-0.014** (-2.02)
Panel D. Exclude Outages Beginning Before 10:00 AM						

Outage x $ROI_{Other}^{Individ}$ (Sell-side)	0.002 (0.90)	0.004 (1.00)	0.007 (1.43)	0.011** (2.09)	0.013** (2.05)	0.016* (1.96)
Outage x $ROI_{Other}^{Individ}$ (Buy-side)	-0.003 (-1.29)	-0.005 (-1.55)	-0.007 (-1.56)	-0.010* (-1.83)	-0.012** (-2.15)	-0.011 (-1.62)
Panel E. Pseudo-Outage Test						
Outage x $ROI_{Other}^{Individ}$ (Sell-side)	-0.001 (-0.25)	-0.002 (-0.42)	-0.006 (-1.44)	-0.007 (-1.47)	-0.005 (-0.80)	-0.006 (-1.02)
Outage x $ROI_{Other}^{Individ}$ (Buy-side)	-0.003 (-1.38)	-0.003 (-1.02)	-0.003 (-0.67)	-0.003 (-0.63)	-0.003 (-0.48)	-0.002 (-0.25)

Table C.2: Robustness Tests for Section 4.3 (Continued)

Robustness Tests (D) for Table 9		
	$ ROI_{Other}^{Common} $	Share
Robustness 1. Exclude firm event days		
Outage	-0.0197*** (-3.27)	
Robustness 2. Exclude Marco market events		
Outage	-0.0195*** (-2.90)	
Robustness 3. Exclude outages less than 30 min		
Outage	-0.0198*** (-3.26)	
Robustness 4. Exclude outages that begin before 10:00 AM		
Outage	-0.0218** (-2.57)	
Robustness 5. Pseudo outage		
Outage	0.0014 (0.22)	

Table C.3: Robustness Tests for Section 4.4

	Coeff. on Outage / Interactions with Outage	t-statistic
Panel A. Exclude Firm-Specific News Events		
Effective Spread	-0.011**	(-2.01)
Realized Spread	0.002	(1.08)
Price Impact	-0.011***	(-3.34)
Outage x Δ ESPD_IND	-0.114	(-1.00)
Outage x Δ RSPD_IND	-0.201**	(-2.46)
Outage x Δ PI_IND	-0.230***	(-2.90)
Panel B. Exclude Macroeconomic Announcements		
Effective Spread	-0.011*	(-1.98)
Realized Spread	0.002	(1.16)
Price Impact	-0.012***	(-3.40)
Outage x Δ ESPD_IND	-0.221**	(-2.02)
Outage x Δ RSPD_IND	-0.255**	(-2.64)
Outage x Δ PI_IND	-0.269***	(-3.32)
Panel C. Exclude Outages Shorter than 30 Minutes		
Effective Spread	-0.009*	(-1.80)
Realized Spread	0.002	(1.27)
Price Impact	-0.010***	(-3.17)
Outage x Δ ESPD_IND	-0.198*	(-1.82)
Outage x Δ RSPD_IND	-0.222**	(-2.55)
Outage x Δ PI_IND	-0.234***	(-2.99)
Panel D. Exclude Outages Beginning Before 10:00 AM		
Effective Spread	-0.015**	(-2.02)
Realized Spread	0.001	(0.51)
Price Impact	-0.013***	(-2.85)
Outage x Δ ESPD_IND	-0.286**	(-2.43)
Outage x Δ RSPD_IND	-0.335**	(-2.12)
Outage x Δ PI_IND	-0.280*	(-1.69)
Panel E. Pseudo-Outage Test		
Effective Spread	0.001	(0.33)
Realized Spread	-0.001	(-0.37)
Price Impact	0.001	(0.46)
Outage x Δ ESPD_IND	0.033	(1.26)
Outage x Δ RSPD_IND	-0.024	(-0.16)
Outage x Δ PI_IND	0.002	(0.01)