Does ChatGPT Enhance Market Liquidity? Evidence from Its Introduction and Outages

Charles Cao, Kenneth (Wei) Wang, and Deli Yang*

March 2025

*Charles Cao is at the Smeal College of Business, Pennsylvania State University. Email: qxc2@psu.edu. Kenneth (Wei) Wang is at the School of Management and Economics, Chinese University of Hong Kong, Shenzhen. Email: kennethwang@cuhk.edu.cn. Deli Yang is at the HKUST Business School, Hong Kong University of Science and Technology. Email: acdeliyang@ust.hk. We are grateful for the comments of Mehmet Canayyaz, Matthew Gustafson, Zongbo Huang, Stephen Lenkey, Tim Simin, Mihail Velikov, Alexey Zhdanov, and seminar and conference participants at the Chinese University of Hong Kong, Shenzhen, and the Pitt-CMU-PSU Finance Conference. All errors are our own.

Acknowledgement: We confirm that we have not received any research support and have no financial relationships or other potential conflicts of interest that apply to any of us.

Does ChatGPT Enhance Market Liquidity? Evidence from Its Introduction and Outages

Abstract

We examine how ChatGPT reshapes market liquidity by improving information quality and increasing noise correlation. We extend the Kyle model to capture this trade-off: while ChatGPT improves signal precision, it also introduces correlated noise among traders. Our model predicts that when noise correlation is sufficiently high, ChatGPT paradoxically enhances market liquidity while reducing volume and return volatility. To empirically test these predictions, we examine significant events related to ChatGPT - its initial release and subsequent outages as exogenous shocks. During ChatGPT outages, we document a 22% increase in trading volume, a 16% rise in volatility, and significant liquidity deterioration, evidenced by an 8% decline in market depth and a 15% widening of bid-ask spreads in cryptocurrency markets. In contrast, the launch of ChatGPT on November 30, 2022, had the opposite effect. These findings suggest that, following the introduction of ChatGPT, traders' information sets become more correlated due to the common noise it introduced. In response, traders strategically submit smaller orders to reduce price impact. Market makers find it more challenging to learn from the total order flow and attribute more orders to informed trades. This dynamic leads to lower price impact and mitigates adverse selection. Our results underscore ChatGPT's underappreciated role in improving market liquidity and stabilizing markets through more synchronized trading, especially in the cryptocurrency market characterized by retail dominance and the absence of fundamentals.

I Introduction

Generative artificial intelligence (GAI) models, such as ChatGPT, have introduced powerful tools for processing and generating information, sparking widespread interest in their potential impact on financial markets. GAIs enable users to perform complex tasks, including textual analysis, data processing, portfolio construction, and back-testing of trading strategies, with unprecedented ease. Investors who previously lacked the resources to interpret financial news, analyze earnings reports, or employ sophisticated models can now design and implement advanced trading strategies. Recent research has examined the impact and performance of GAI models in equity markets. These studies focus on earnings forecasts, portfolio construction, and institutional investor behavior (Lopez-Lira and Tang, 2023; Beckmann et al., 2024; Cheng et al., 2024; Chen et al., 2025). However, less attention has been paid to GAIs' impact on the cryptocurrency markets—an equally important domain characterized by retail dominance and the absence of clear fundamentals.

In this paper, we examine the impact of ChatGPT on cryptocurrency market liquidity by analyzing changes in liquidity, trading volume, and return volatility surrounding two important events: the launch of ChatGPT on November 30, 2022, and its subsequent outages between December 2022 and October 2024. Our analysis begins with a theoretical model that provides a framework for understanding how the introduction and temporary outages of ChatGPT affect market liquidity and trading volume.

Our model is motivated by two considerations. First, while GAIs enhance investors' ability to process information, their outputs are often homogenized. Trained on large datasets, GAIs tend to generate correlated responses across users, effectively injecting common noise into private signals. This creates a fundamental tension: although ChatGPT increases the precision of private signals, it also introduces noise correlation among traders and leads to more synchronized trading behavior. Second, despite growing anecdotal evidence on the

role of GAIs in stock trading, there is a notable lack of theoretical models that examine the effects of ChatGPT's introduction and outages on market liquidity and trading dynamics. This paper addresses these gaps in the literature.

To formalize the ambiguous effects of GAIs on market liquidity, we extend the Kyle (1985) model by introducing a key friction: correlated noise resulting from homogenized AI-generated outputs. Unlike traditional frameworks, our model emphasizes how GAI-induced synchronization of trades generates correlated noise across investors, even as it reduces information asymmetry. By deriving equilibrium outcomes with and without access to GAIs, we develop testable hypotheses suggesting that the introduction of ChatGPT improves market liquidity, while its service outages impair it.

Our model yields three hypotheses regarding market behavior when GAIs suddenly become unavailable, such as during ChatGPT outages. First, trading volume increases. Without access to precise signals, traders submit more orders that are less tied to asset fundamentals. In the absence of correlated noise from homogenized GAI outputs, traders are less concerned about synchronized trades and consequently trade more aggressively. Second, return volatility rises through two channels: (1) noisier signals generate more erratic trades, increasing volatility; and (2) the abrupt loss of GAIs pushes traders to rely more heavily on private signals, which are incorporated into prices more quickly, resulting in larger price movements. Finally, market liquidity decreases. While noisier signals alleviate market makers' concerns about adverse selection, the disappearance of noise correlation makes order flow more informative. If the pre-outage noise correlation is sufficiently high, the outage negatively impacts market liquidity and leads to a reduction in market depth and an increase in bid-ask spreads.

In the first set of tests, we use ChatGPT outages as exogenous shocks to analyze their impact on cryptocurrency market's liquidity. These events represent sudden and unexpected disruptions to access GAI. They occur when ChatGPT experienced significant malfunctions

and completely cut off user access to ChatGPT. We compile outage histories from OpenAI's official incident reports and investigate changes in trading volume, return volatility, and various liquidity measures in the periods before and during the ChatGPT outages. Importantly, outages are unpredicted by market participants (for instance, triggered by technical failures rather than scheduled maintenance) and serve as a quasi-natural experiment to isolate the causal effects of GAIs on market liquidity. This setup reduces concerns about reverse causality or the impact of confounding factors, as traders cannot predict the timing or duration of these disruptions.

To assess the impact of ChatGPT outages on trading volumes, we compare average volumes during the outage period (the event period) to those in the same time intervals over the preceding five trading days (the control period) for the 70 coin pairs in our sample. Using intraday order book data and transaction data from the Kaiko database, we document a significant increase in dollar trading volume of 26% and in trading frequency of 22% during the ChatGPT outage period. These results in trading volume support the prediction of our model: GAIs suppress noise-driven trading by improving signal precision. Meanwhile, it also increases the correlation of private signals. Investors will take into account the potential synchronized trades and reduce their order sizes. When access to ChatGPT is disrupted, investors lose the guidance of GAI-generated signals. As a result, they trade based on noisier signals and submit orders aggressively.

We next examine changes in return volatility using two measures: trade-based and order book-based realized volatility. During ChatGPT outages, we find a significant increase in volatility of 16% and 20% for trade-based and order-book-based measures. These results support the predictions of our model: the absence of ChatGPT amplifies noise in trading signals and removes the stabilizing effects of synchronized trading. When ChatGPT is accessible, GAIs homogenize signals, prompting traders to submit smaller orders to avoid crowded trades, while market makers attribute order flow to correlated noise rather than informed

trading. Both effects suppress volatility. However, when ChatGPT become inaccessible, this coordination breaks down. With lower correlation of information sets, traders trade more aggressively on noisier signals. The noises are independent and cancel out in the total order flow, making it more informative. Therefore, market makers become more sensitive to order flow, leading to higher price impact. The surge in volatility during outages underscores ChatGPT's underappreciated role in stabilizing markets through synchronized trading.

Our model predicts that losing access to GAIs leads to a reduction in market liquidity. To test this hypothesis, we examine changes in three liquidity measures: depth behind the best-bid and best-ask prices, bid-side book depth and ask-side book depth, and bid-ask spreads. Our findings indicate a significant decline in all liquidity measures during outage periods. Consistent with the previous analysis, the noise correlation from wide adoption of ChatGPT reduces price impact and mitigates adverse selection. Thus, given the order flow of the same size, the price changes less and the market depth is higher. The consistent patterns across liquidity measures—reduced depth and widening spreads—reinforce our conclusion that ChatGPT improves cryptocurrency markets' liquidity.

In our second set of tests, we examine the impact of another significant shock to GAI accessibility: the public launch of ChatGPT on November 30, 2022. Based on our model's predictions, we hypothesize that trading volume will decrease, return volatility will decrease, market depth will increase, and the bid-ask spread will decrease following the introduction of ChatGPT. This shock represents a more impactful event than other shocks. In contrast to the impact of ChatGPT on market liquidity during its outages, we find that the introduction of ChatGPT leads to a reduction in trading volume and return volatility while improving market liquidity.

This paper contributes to the literature on the impact of generative artificial intelligence on financial markets in several important ways. First, we provide evidence that ChatGPT's introduction and its subsequent outages significantly affect market liquidity and trading

activity. In cryptocurrency markets, we observe a notable improvement in liquidity following ChatGPT's introduction, as reflected in a narrower bid-ask spread and increased market depth, along with a decline in trading volume and return volatility within the event window of [-7, +7] days. Moreover, during periods of ChatGPT outages, we find a significant reduction in market liquidity, evidenced by a wider bid-ask spread and reduced depth, while trading volume and volatility increase. Second, our findings stand in sharp contrast to those of Cheng et al. (2024), who examined the impact of ChatGPT outages on the stock market. During the ChatGPT's outages, they document a significant improvement in market liquidity, a decrease in trading volume and return volatility for stocks traded on the NYSE, a pattern opposite to that for the cryptocurrency markets. The differing effects suggest that GAI's impact on financial markets is not uniform. In comparison to the stock market, the cryptocurrency market is characterized as a market that lacks fundamentals dominated by retail investors, and that investors' sentiment plays an important role in determining returns.

Our paper builds on recent studies examining whether ChatGPT assists investors in processing information and enhancing information acquisition. Cheng et al. (2024) analyze the impact of ChatGPT outages on stock market trading and price informativeness. Bertomeu et al. (2023) investigate the effect of ChatGPT's ban in Italy on analysts' forecasting accuracy. Sheng et al. (2024) explore hedge funds' adoption of generative AI using ChatGPT outages as a natural experiment. These studies focus on the impact of ChatGPT on the stock market, where institutional investors dominate. In contrast, our study examines the effects of ChatGPT's introduction and outages in the cryptocurrency market, where retail investors play a central role (Makarov and Schoar, 2020; Kogan, et al., 2024).

Finally, this study contributes to the broader literature on AI in capital markets. While general-purpose AIs like ChatGPT reduce the costs of acquiring and processing information, they also democratize access to these capabilities for a broader group of investors, potentially driving behavioral shifts across financial markets. Cao et al. (2023, 2024) document how

advancements in textual analysis influence corporate disclosure strategies and stock analyses. Rossi and Utkus (2024), D'Acunto et al. (2019) and Loos et al. (2020) examine the role of robo-advisors in improving household financial decision-making. Expanding on this literature, our study provides empirical evidence of GAIs' impact on cryptocurrency market liquidity. By extending the analysis to this distinct market segment, our paper complements existing research and offers new insights into their role in retail-dominated markets.

This paper is organized as follows: Section II discusses the institutional background; Section III presents our model and its predictions; Section IV describes the data and sample selection. In Sections V and VI, we present the empirical results, using ChatGPT's outages and introduction as two types of events, respectively. Section VII concludes.

II Institutional Background

ChatGPT is a large-scale language model developed by OpenAI and released to the public on November 30, 2022. It is built on the Generative Pre-trained Transformer (GPT) architecture, a state-of-the-art framework for natural language processing. ChatGPT is designed to generate human-like responses based on user prompts. It can understand and generate text in a conversational manner and allow users to interact with it. Its model was trained on an extensive corpus of publicly available text data. Based on the deep neural network architecture introduced by Vaswani et al. (2017), GPT has demonstrated exceptional performance in various language-processing tasks and has been widely used by institutional and retail investors. Its pre-training process enables the model to encode semantic meanings into vector representations, which can then be fine-tuned for specific applications.

ChatGPT's advanced conversational capabilities support various applications in the financial sector, including summarizing news articles and financial statements, analyzing historical data, generating investment insights, and designing trading strategies. Furthermore,

the premium version allows investors to input diverse data sources, such as documents, images, and audio, and enables advanced tasks like technical and sentiment analysis.

Since the introduction of ChatGPT, financial institutions have increasingly integrated Generative AI into their operations. For instance, Bloomberg developed BloombergGPT, which is a 50-billion-parameter large language model for the financial sector. It is designed for natural language processing tasks such as sentiment analysis and is trained on a vast dataset comprising 363 billion tokens of financial data. It can process and extract insights from financial reports, earnings transcripts, and news articles that might influence trading decisions. It can also assist investors in generating trading strategy ideas and formatting trading rules. Institutional and retail investors use ChatGPT to analyze historical data, review regulatory filings, and develop creative trading strategies. Empirical evidence indicates widespread adoption among investors. A survey by Ebert (2023) found that half of the surveyed fund managers employed ChatGPT for regulatory filing analysis and other tasks.

In addition, cryptocurrency brokers and investors leverage ChatGPT to generate investment ideas based on investors' historical trading activity, portfolio composition, and risk preferences. For example, FalconX, a prime broker in cryptocurrencies, developed a chatbot powered by ChatGPT's technology that provides investors with market insights, investment ideas, and trading strategies. This chatbot aims to democratize access to advanced trading tools by offering personalized services.

In the cryptocurrency market, sophisticated investors rely on algorithmic models to trade. These investors, typically well-resourced, hire teams to develop sophisticated trading platforms and strategies. They implement a range of strategies, including market-neutral and quantitative long/short strategies. FalconX's chatbot aims to level the playing field by providing retail, less sophisticated investors with tools comparable to those used by large

quantitative firms.[1]

As an illustrated example, FalconX's chatbot enhances investor decision-making in three ways. First, it aggregates and analyzes relevant news and information from both traditional and social media to provide tailored briefings based on an investor's interests (Canayaz et al., 2025). For example, it can answer questions such as, "Which cryptocurrencies benefited the most after President Trump signed an executive order on March 2, 2025, to establish a U.S. digital asset stockpile?" or "Which concentrated liquidity pools generated annual percentage yield (APY) exceeding 100% over the past 12 months?" Second, investors can test trading strategies by posing service-based queries like, "On a risk-adjusted basis, what should be the optimal approach to providing liquidity within Uniswap pools to achieve an expected yield of 100% (APY) over an extended period of 12 months?" Finally, the chatbot features integrated buy/sell buttons within the platform, enabling users to execute strategies seamlessly in real-time.

III A Model of Informed Trading with ChatGPT

We extend the Kyle (1985) model to examine the impact of ChatGPT on market liquidity, focusing on its dual role in information production. In our adaptation of the Kyle framework, we model ChatGPT as a signal-generating technology. On one hand, it improves investors' ability to process vast amounts of information and complex data, thereby increasing the accuracy of private signals. On the other hand, due to the homogeneity of the model across users, investors receive similar trading signals when consulting ChatGPT. As a result, even the noise terms in investors' private signals become highly correlated. This interdependence alters traders' strategies and market makers' pricing strategies.

We derive a rational-expectation equilibrium in which the strategies of market partici-

[1]See https://www.forbes.com/sites/stevenehrlich/2023/04/13/a-new-chatgpt-powered-bot-named-satoshi-will-soon-help-crypto-traders/.

pants reflect both the informational efficiency gains from ChatGPT and the impact of correlated trades. By comparing equilibria with and without ChatGPT, we examine how its dual role - improving information quality while introducing noise correlation - shapes market liquidity, trading volume, and return volatility.

III.1 Model Setup

Our model is single-period. At the beginning, a risky asset with a random liquidation value v is traded. Investors cannot directly observe v and rely on private signals and a common prior belief, which follows a normal distribution with the mean v_0 and the variance σ_0^2. For simplicity and analytical convenience, we define the precision of this distribution, τ_0, as $\tau_0 := 1/\sigma_0^2$. In our one-period setting, we interpret the liquidation value as the period-end market price of the asset.

There are three types of traders: N expert traders, M non-expert traders, and a continuum of noise traders. Expert traders receive perfect signals about v, meaning that for trader i, $s_i = v$ (no noise). Non-expert traders receive noisy signals $s_i = v + e_i$ where the noise term $e_i \sim N(0, \sigma_e^2)$, with precision $\tau_e = 1/\sigma_e^2$. Noise traders submit aggregate order flow $u \sim N(0, \sigma_u^2)$, which is independent of v and e_i terms, with precision $\tau_u = 1/\sigma_u^2$. They trade for reasons unrelated to the fundamental value of the asset, introducing randomness into the order flow.

We collectively refer to expert and non-expert traders as "informed traders"($I := N+M$), where i denotes the ith informed trader. Informed traders have market power and trade strategically, taking into account the strategies of one another. The difference between expert and non-expert traders captures the differential impact of ChatGPT on them. Historically, expert traders, such as quantitative fund managers, had exclusive access to sophisticated tools, while others faced barriers to effectively leveraging these technologies due to their technical complexity. ChatGPT democratizes access, enabling non-experts to generate

trading signals using similarly advanced tools.

Each informed trader i submits an order $\tilde{x}_i(s_i)$ based on her own signal s_i. A representative market maker observes the total order flow, given by $y = \sum_{i=1}^{I} \tilde{x}_i + u$, and sets the price p of the risky asset. We denote the pricing rule as $p = P(y)$. Once the price is set, trades are executed. At the end of the period, the risky asset is liquidated, its true value is realized, and traders receive their payoffs.

III.1.1 Information Production and ChatGPT

Our model examines the impact of ChatGPT on market liquidity through two different channels. First, non-expert traders may benefit from ChatGPT due to its ability to facilitate information processing, thereby improving the quality of their private signals. Second, since traders rely on a single model to generate private signals, the correlation between signals, including the noise components, can be high. To capture this dual mechanism, Assumption 1 specifies the information structure, characterizing the precision and correlation of signals in scenarios with and without ChatGPT.

Assumption 1 (Information Structure). For any two different non-expert traders, i and j, the correlation between the noise terms (e.g., e_i and e_j) associated with their private signals s_i and s_j is a constant, ρ, i.e.,

$$Cov(e_i, e_j) = \rho \frac{1}{\tau_e} \tag{1}$$

where τ_e is the precision of the noise term. We assume that the precision and correlation of the noise term are equal to 0 and τ_θ, respectively, without ChatGPT. When ChatGPT is available, the precision and correlation of the noise term change to $\tau_s > \tau_\theta$ and $\rho = \rho^* > 0$.

We use this assumption to capture ChatGPT's dual potential: improving the quality of information while simultaneously increasing the correlation of noise terms. For simplicity, we assume that, without ChatGPT, the noise terms in private signals are independent of

each other. It is important to note that independent noise terms do not imply independent signals. Investors' private signals can still be correlated due to common components related to fundamental values. As we will demonstrate, noise correlation plays a critical role in deriving our empirical implications.

III.2 Market Equilibrium

Using our notations, the trading profits of trader i, denoted as $\tilde{\pi}_i$, are calculated as $\tilde{\pi}_i(\tilde{x}_i) = (v - P)\tilde{x}_i$ where P is the end-period price of the asset and $\tilde{x}_i$ is the submitted order size. We assume that the market-making industry is competitive. Next, we define the rational expectation equilibrium following Kyle (1985, 1989).

Definition 1. An equilibrium is a set of $\{\tilde{x}_1, \tilde{x}_2, \cdots, \tilde{x}_I, P\}$ such that:

1. For all $i = 1, \cdots, I$, $\tilde{x}_i$ dominates any alternative strategy $\tilde{x}_i'$ for any private signal s_i given other traders' strategies fixed, i.e.,:

$$E[\tilde{\pi}_i(\tilde{x}_i)|s_i] \geq E[\tilde{\pi}_i(\tilde{x}_i')|s_i] \tag{2}$$

2. With a competitive market-making industry, the pricing function $P(y)$ satisfies

$$P(y) = E[v|y] \tag{3}$$

Equation (2) represents the standard Nash equilibrium condition, where each agent takes her best response to others' strategies. Equation (3) captures the competitive pricing condition in the market-making industry, implying zero ex-ante expected profits. We focus on the linear equilibrium in which traders' strategies are linear in their signals, and the pricing function is linear in total order flow (Kyle, 1985, 1989; Foster et al, 1996). Proposition 1 solves the equilibrium in Definition 1.

Proposition 1: In a linear rational expectation equilibrium, traders follow linear strategies:

$\tilde{x}_i = \beta_i(s_i - v_0)$, where $\beta_i = \overline{\beta}$ for expert traders and $\beta_i = \underline{\beta}$ for non-expert traders. The competitive market maker sets the price by $P(y) = v_0 + \lambda * y$, where v_0 is the mean of the prior belief on v. The parameters $\overline{\beta}$, $\underline{\beta}$, and λ are solved from the following equation system:

$$\overline{\beta} = \frac{1}{2\lambda} - \frac{N-1}{2}\overline{\beta} - \frac{M}{2}\underline{\beta} \tag{4}$$

$$\underline{\beta} = \frac{1}{2\lambda}\frac{\tau_e}{\tau_0+\tau_e} - \frac{\overline{\beta}N}{2}\frac{\tau_e}{\tau_0+\tau_e} - \frac{M-1}{2}\underline{\beta}\frac{\rho\tau_0+\tau_e}{\tau_0+\tau_e} \tag{5}$$

$$\lambda = \frac{(N\overline{\beta}+M\underline{\beta})\frac{1}{\tau_0}}{(N\overline{\beta}+M\underline{\beta})^2\frac{1}{\tau_0} + \underline{\beta}^2[M+\rho M(M-1)]\frac{1}{\tau_e}+\frac{1}{\tau_u}} \tag{6}$$

.

Appendix A provides the proof of Proposition 1. In the equation system characterizing the equilibrium, Equations (4) and (5) represent the first-order conditions of the optimization problem (2) for expert and non-expert traders, respectively. Equation (6) corresponds to the competitive pricing condition of market makers. These equations illustrate how the strategies of market participants are interdependent and shaped by the information structure. For instance, from the first-order condition of expert traders, Equation (4), the best response (the left-hand side) increases with the market maker's order-flow sensitivity, proxied by λ, but decreases with the aggressiveness of other expert and non-expert traders, proxied by $\overline{\beta}$ and $\underline{\beta}$, which appear on the right-hand side. These relationships are intuitive: if market makers are sensitive to the total order flow, traders prefer to place smaller orders to reduce the price impact. If other traders already placed large orders, a trader will reduce her order size accordingly to minimize the price impact.

Equation (5) is quite similar except for the information structure. Different from expert traders, non-expert traders do not have perfect signals. This feature means that they have to "guess" other traders' order sizes based on their own signals. The posterior belief depends on how precise the private signal is, τ_e, and the correlation of noise terms, ρ, which matters

when they extrapolate on other non-expert traders' order sizes.

For Equation (6), the numerator is the covariance between the total order flow y and the fundamental value v. The denominator is the variance of the total order flow. The pricing condition can be interpreted as the ratio of informed trades to the total trades, which captures adverse selection. The market maker is more sensitive to the order flow if the adverse selection is more severe.

III.2.1 Equilibrium Properties

To examine the impact of GAIs, we focus on trading volume, return volatility, and market depth, defined as $\frac{1}{\lambda}$ in (Kyle, 1985). We compare the three variables in the case with/without GAIs. Our definition of trading volume follows Wang (1994), $V := E[\sum_i |x_i| + |u|]$. To calculate the return, we assume that before the model starts, the price is equal to v_0. The return volatility is defined as $Vol = Var(p - v_0)$. Proposition 2 gives the explicit expressions of trading volume and return volatility in equilibrium. Both variables depend on the number of traders, information quality, and traders' strategies. Note that the noise correlation is implicitly included in $\overline{\beta}$ and $\underline{\beta}$.

Proposition 2: Defining the trading volume as $V := E[\sum_i |x_i| + |u|]$, we have

$$V = \sqrt{\frac{2}{\pi}} \left(N\overline{\beta}\sqrt{1/\tau_0} + M\underline{\beta}\sqrt{1/\tau_0 + 1/\tau_e} + \sqrt{1/\tau_u} \right) \tag{7}$$

Defining the return volatility as $Vol = Var(p - v_0)$, we have

$$Vol = \frac{(N\overline{\beta} + M\underline{\beta})^2 \frac{1}{\tau_0^2}}{(N\overline{\beta} + M\underline{\beta})^2 \frac{1}{\tau_0} + \underline{\beta}^2[M + M(M-1)\rho]\frac{1}{\tau_e} + \frac{1}{\tau_u}} \tag{8}$$

The parameters of interest include the precision of noise terms in signals, τ_e, and the correlation between noise terms, ρ. Since these parameters enter the trading strategies of

expert and non-expert traders, i.e., $\overline{\beta}$ and $\underline{\beta}$, trading volumes and return volatility, as shown in Proposition 2, are non-linear functions of τ_e and ρ. To better understand the mechanism, we start with a hypothetical setting where traders follow fixed trading strategies. Proposition 3 characterizes the equilibrium properties.

Proposition 3: Suppose traders follow fixed strategies, i.e., $\overline{\beta}$ and $\underline{\beta}$ are constant. The model predicts:

- Trading volume is lower with ChatGPT compared to without ChatGPT if, and only if, ChatGPT improves information quality, i.e.,$\tau_s > \tau_\theta$.
- Market depth is higher (and volatility lower) with ChatGPT compared to without ChatGPT if, and only if, the noise correlation introduced by ChatGPT, ρ^*, previously exceeded the critical threshold:

$$\rho^* > \frac{1}{M-1}\left(\frac{\tau_s}{\tau_\theta} - 1\right) \tag{9}$$

  where M is the number of non-expert traders, and ρ^* is the noise correlation in signals with ChatGPT.

Proposition 3, while in a simplified setting, yields tractable insights. First, if ChatGPT improves information quality ($\tau_s > \tau_\theta$), it will help traders trade on less noisy signal (τ_θ), which reduces volume. Second, equation (9) shows correlated noise, when properly anticipated by market makers, mitigates adverse selection. From the market maker's perspective, idiosyncratic noise (uncorrelated errors across traders) enables the aggregate order flow to convey meaningful information about v. When noise terms are independent, the Law of Large Numbers ensures they partially cancel out in the total order flow, leaving a clearer signal about v. Market makers can then disentangle informed trading from random fluctuations, setting prices that closely reflect fundamentals. In contrast, correlated noise

(systematic errors driven by GAIs) distorts this mechanism. When $\rho^* > 0$, noise terms no longer offset-they compound, injecting persistent bias into the order flow. This obscures the signal about v, as market makers cannot distinguish whether elevated trading activity stems from informed speculation or trades induced by noise trades. Recognizing the correlation structure, market makers attribute synchronized trades to GAI-induced noise rather than informed speculation. This reduces perceived adverse selection, enabling them to set smaller price impacts (improved depth).

A central contribution of our analysis is identifying noise correlation as a critical friction of ChatGPT. Unlike traditional models where improved signal precision unambiguously worsens liquidity (via heightened adverse selection), ChatGPT introduces a countervailing force. The noise correlation creates an empirical regularity absent in conventional settings—liquidity can improve even as signal quality rises, provided ρ^* is sufficiently large. For example, when $\rho^* = 0$ (no ChatGPT-induced correlation), ChatGPT's adoption cannot generate the opposing movements in volume and liquidity we predict (e.g., volume decreasing while liquidity improves). Instead, ChatGPT would either uniformly harm liquidity (if $\frac{\tau_s}{\tau_\theta} > 1$) or improve it (if $\frac{\tau_s}{\tau_\theta} < 1$), depending on precision gains. Only with $\rho^* > 0$ can these effects diverge, revealing how ChatGPT reshapes markets through both precision and correlation.

III.2.2 Comparative Statics

A constraint of Proposition 3 lies in the assumption that trading strategies remain fixed. We now perform a comparative static analysis on the full model. Although an analytical characterization of the parametric regimes is not feasible, we can utilize numerical solutions to visualize these regimes. The numerical results closely resemble the implications of Proposition 3.

We numerically solve the full model and compare three key market measures—trading volume, return volatility, and market depth—under alternative parametric assumptions. Our

focus centers on the regime where GAI outages (loss of access) increase trading volume and return volatility while decreasing market depth. Figure 1 visualizes the parametric regime of the scenario of interest. In Figure 1, the parameters are set as follows: $N = 3$, $M = 4$, $\tau_0 = 1$, and $\tau_u = 3$. Panel (a) assumes low noise correlation $\rho^* = 0.1$, while Panel (b) adopts high noise correlation $\rho_1 = 0.8$. The relevant regimes are color-coded in green. Similar to Proposition 3, an increase in noise correlation ρ^* expands the green area. Crucially, in simulations with $\rho^* = 0$ (unreported), the green regime vanishes entirely, confirming that GAI-induced noise correlation is indispensable for generating these outcomes.

However, unlike Proposition 3, the numerical results demonstrate the strategic channel via which ChatGPT imposes impacts. Different investors' trades will be similar (in size and direction) due to the high correlation of signals, the synchronized trades altogether are likely to generate large price impacts. Investors also want to reduce the size of their trades to avoid the big price impact, as shown in Figure 2.

Our model identifies two channels through which GAI accessibility impacts markets. First, traders can trade on high-precision GAI signals ($\tau_s > \tau_\theta$). This reduces trading volume (as agents overreact to imprecise signals) and return volatility. Second, with GAI, the noise correlation is introduced. Investors' signals cannot average out and become less informative about the fundamental value, reducing equilibrium price impact.

Empirically, we predict that if the noise correlation induced by ChatGPT is high enough, after the ChatGPT launch (during ChatGPT outages), the markets will exhibit: (1) decreased (increased) trading volume; (2) lower (heightened) return volatility; (3) improved (reduced) market liquidity. By exploiting launches and outages of ChatGPT as exogenous shocks, we isolate how GAIs shape market liquidity.

IV Data and Sample Selection

The data used in this study comes from two sources. Information on ChatGPT outrage events is obtained from OpenAI, while intraday order book data and transaction data are from the Kaiko database.

OpenAI maintains a public record that tracks the operational status of its products, including ChatGPT and models such as DALL-E. The official website publishes incident reports whenever one or more services experience malfunctions. These reports document the date and severity of each event and provide detailed information about the timeline, extent of service disruption, and underlying causes.[2].

We manually collect all incident events and the corresponding reports related to ChatGPT to identify outages from its release on November 30, 2022, until the end of our cryptocurrency dataset on October 31, 2024. We begin with incidents marked as affecting ChatGPT and lasting at least 10 minutes, then narrow our selection to those whose titles explicitly indicate outages during which the service is inaccessible. We review each incident report to ensure that every outage event is included. For each incident, we record the start time, the beginning of the investigation of the malfunction, and the end time when the issue is fully resolved. To construct a clean control period for each event, we only include outages that have no other outage occurring within the preceding 15 days. In total, we identify 12 ChatGPT outage events, as presented in Table 1. The first event occurred on February 21, 2023, and the latest on July 5, 2024. The duration of the outages range from 14 minutes to 4 hours and 29 minutes, with a mean of 1.5 hours and a median of 52 minutes.

Incident reports only partially disclose the causes of both outage and non-outage events. They indicate that outages are not primarily due to database instability, program misconfigurations, or software bugs. ChatGPT is a general-purpose generative AI that serves a

[2]See https://status.openai.com/history

wide range of users beyond investors. Thus, its outages are unlikely to be directly tied to economy-wide disruptions or solely the result of usage surges during market frenzies.

The intraday cryptocurrency transaction data is provided by Kaiko Inc., a leading cryptocurrency data provider. Kaiko supplies tick-level trade data and full-depth order book snapshots across a large number of cryptocurrency exchanges and coin pairs. For each coin pair, the trade data includes a unique transaction ID, a Unix millisecond timestamp, price, trading volume, and a trade direction indicator, which identifies whether the trade was initiated by the buyer or seller. Similarly, order book snapshots (taken every 30 seconds) include a Unix millisecond timestamp, price, traded currency amount, and a bid-ask type indicator. For variables such as trading frequency, trading volume, bid-ask spreads, and market depth, we aggregate the intraday data into 5-minute intervals. For example, within each 5-minute interval for a given coin pair, trading volume (or trading frequency) is defined as the total volume (or total number of trades) during that period. The bid-ask spread (or depth) is computed as the average quoted spread (or average depth) over the interval. Returns are calculated using the prices at beginng and the end of each 5-minute interval, while volatility is the standard deviation of returns using 30-second snapshots with each 5-minute interval.

We first select eight exchanges based on their consistent trading of popular coin pairs throughout our sample period. These exchanges include Binance, Bitstamp, Coinbase, Gemini, Huobi, Kraken, KuCoin and OkEX. Daily ranking of cryptocurrency exchanges is provided by CoinMarketCap.com. Then we select the top 10 largest non-stable cryptocurrencies by market capitalization as of November 30, 2022 (which is around the release of ChatGPT). The selected coins are Bitcoin (BTC), Ethereum (ETH), Binance Coin (BNB), Ripple (XRP), Dogecoin (DOGE), Cardano (ADA), Polygon (MATIC), Polkadot (DOT), Litecoin (LTC), and Solana (SOL).

For each exchange and each coin, we select the primary coin pair with the highest trading volume in November 2022 when multiple pairs were available. For example, on the exchange

of Coinbase, there are two Bitcoin-related coin pairs, one is quoted in dollars (BTC/USD) and the other one in Tether (BTC/USDT), a stable coin pegged at the dollars. Both BTC/USD and BTC/USDT are actively traded; however, the trading volume of BTC/USD is 17 times larger than that of BTC/USDT in October 2022, so we include BTC/USD in our sample. During our sample period, the exchange rate between USD and Tether was between $0.998 and $1.007. When Tether is a denominated currency, we convert trading volume to dollar volume.

Overall, our sample comprises 70 coin pairs, and spans from December 2022 to October 2024. We provide a summary of the cryptocurrency exchanges and coin pairs in Table 2. Notably, all top 10 coins are traded on Binance, Huobi, KuCoin and OkEX. On the other exchanges, the number of top 10 coins included in our sample ranges from six to nine.

Table 3 reports cross-sectional summary statistics—mean, standard deviation (S.D.), and the 25th, 50th, and 75th percentiles for several variables: trading frequency, dollar volume, 5-minute return, price, trade-based volatility, order book-based volatility, quoted spread, proportional spread, best-bid and best-ask depth, depth of -0.5% bid, and depth of +0.5% ask (definitions follow). The table shows that, across 70 coin pairs, the average trading frequency is 494 trades per 5-minute interval, with an S.D. of 1,747 trades, while the average dollar volume is $352,000 per 5-minute interval, with an S.D. of $1,524,000. The average coin price is $2,850, and the 25th and 75th percentiles are $0.44 and $287, respectively. On average, the quoted spread is $0.16, and the proportional spread is 3.20%. These results highlight a large dispersion in coin pair characteristics.

V Empirical Results: ChatGPT's outages

In this section, we seek to answer two questions. First, to what extent do crypto market investors rely on generative AI (GAI) for trading cryptocurrencies? Second, how does GAI-

assisted trading affect market trading activity and liquidity? We conduct an analysis to investigate whether there is a significant change in trading activity and liquidity measures from the period of before ChatGPT outages to the period of during outages. We also provide supplementary evidence through panel regressions that control for other confounding factors.

To assess the impact of ChatGPT outage events on trading activity and liquidity in cryptocurrency markets, we employ a standardized aggregation procedure. For each coin pair and variable of interest, we compute the time-series average for the period preceding the outage (control period) and the period during the ChatGPT outage (See Ahn, et. al., 1996; Ahn, et. al., 1998). Given the skewed distribution of dollar trading volume across coin pairs, we apply dollar volume weights from the control period to each coin pair and then calculate the corresponding cross-sectional mean, standard error, median, and test results for the variable. For example, the average bid-ask spread across the 70 coin pairs is the volume-weighted average spread. By doing so, we give a larger weight to a coin pair with higher dollar trading volume and the economic interpretation of our result is straightforward. For simplicity, we omit the term "volume-weighted" when discussing cross-sectional summary statistics.

Because changes in these variables might be influenced by other confounding factors, it is important to control for these effects when comparing the control and outage periods. For this reason, we define the control period as follows. For each outage event, we designate the period from one to five days before the outage date (e.g., [–5, –1]) as the primary control period. For instance, the first outage occurred on February 21, 2023, from 7:40 a.m. to 12:09 p.m. Thus, the outage period is 7:40 a.m.–12:09 p.m. on February 21, while the control period consists of the same time interval on February 16, 17, 18, 19, and 20. We also consider alternative control periods using the same time intervals from the previous 10 ([–10, –1]) and 15 ([–15, –1]) trading days for robustness checks. Note that we only include outage events where no other outages occurred during the corresponding control periods.

This approach has two advantages: first, selecting a control period immediately preceding an outage minimizes changes in the coin's characteristics; and second, by ensuring that no other outages occur during the control period, any observed changes in trading activity and liquidity can be attributed to the ChatGPT outage event.

Before conducting statistical tests, we first select one exchange and one coin pair to visualize the 5-minute trading volume during both control and outage periods. Figure 2 displays the average trading volume for the Binance BTCUSDT pair during the control period of [-5, -1] and the outage periods. The figure clearly shows an increase in trading volume from the control period to the outage events. We now turn to formal statistical tests.

V.1 Trading Volume

Table 4 presents summary statistics of two different measures of trading volume, trading frequency and dollar trading volume, in the control period and the period of ChatGPT outage events. As shown in Table 3, the distribution of each volume measure is skewed in the cross section. Thus, for each volume measure, we compute the ratio of control-period volume to the event-period volume, and then take the natural log to obtain the change in trading volume.

We find a significant increase in both trading volume measures from the control period to the event period. Using [-5, -1] as the control period, Table 4 shows that the average trading frequency across the 70 coin pairs increases by 1,655 trades. Specifically, the average trading frequency is 5,157 trades per five-minute interval during the control period, increasing to 6,812 trades during the outage period. The logarithmic difference (Log. Diff.) in trading frequency is 22.80% and is significant at the 5% level based on the t-test. The increase in dollar trading volume is slightly larger, at 26.12% and significant, from $4,914,830 to $6,806,720 per five-minute interval. To address concerns about the non-normal distribution of volume changes, we employ the non-parametric Wilcoxon test and report the corresponding

p-value. Additionally, to account for potential cross-sectional dependence, we use a bootstrap method to recover the empirical distribution of the t-statistic and report the bootstrapped p-value for the t-test. Both the Wilcoxon test and the bootstrapped t-test strongly reject the null hypothesis of no change in trading frequency.

Table 4 also report results using two alternative control periods: [-10, -1] and [-15, -1]. For the first alternative control period, [-10, -1], the average trading frequency increases by 28.01%, while the median trading frequency rises by 35.71%. All three tests confirm a significant increase in trading frequency. A similar conclusion holds for dollar trading volume.

In summary, the empirical evidence from the control and event periods indicates a significant increase in trading frequency and dollar trading volume during the ChatGPT outage. When ChatGPT is unavailable, investors trade more frequently and in larger volumes.

V.2 Return Volatility

We now turn to the analysis of return volatility using realized volatility as our metric. Prior literature has demonstrated that realized volatility (RV) is an efficient, unbiased, and model-free estimator of return volatility when computed from high-frequency returns (Andersen et al., 2001). We construct realized volatility using two approaches.

The first measure is order book–based volatility. Specifically, for each five-minute interval, we take 30-second snapshots of the limit-order book from Kaiko Limit Order Book Data and record the corresponding bid-ask midpoint. Each snapshot captures the quantities and prices of the best bids and best asks that lie within 10% of the mid-price at the time of the snapshot. We then compute return volatility using 11 midpoints obtained within each five-minute interval. The realized volatility $RV_{i,t}$ for coin pair i and time interval t is calculated

as follows:

$$RV_{i,t} = \sqrt{\sum_{s=1}^{11} [\log(P_{i,t_{s+1}}) - \log(P_{i,t_s})]^2 / 10} \tag{10}$$

This measure is defined as the order book-based volatility. Our second volatility measure is constructed using Kaiko's Trade Data. We identify a transaction price that is the closest to the end of a given 30-second interval. Then, we obtain return volatility using the 11 prices within each 5-minute interval. This measure is defined as trade-based volatility. For both volatility measures, the volatility during a given time period is the average of 5-minute interval volatility over the period.

Table 5 presents summary statistics of trade-based return volatility, order book-based return volatility, and the logarithmic change (Log diff in %) of these variables for the periods before and during the ChatGPT outages. Panel A shows that trade-based volatility increases by 16.22% during the ChatGPT outage period when using the control period of [-5, -1]. The averages of trade-based volatility are 3.88 bps and 4.48 bps for the control period and outage period, respectively. Further, the t-test, the Wilcoxon test, and the bootstrapped t-statistic strongly reject the null hypothesis that there is no change in traded-based volatility. When using [-10, -1] and [-15, -1] as control periods, we obtain consistent results.

Panel B of Table 5 shows changes in order book-based volatility. Once again, we see a significant increase in volatility from the control period to the outage period, regardless of whether we use [-5, -1], [-10, -1] or [-15, -1] as a control period. When using [-5, -1] as the control period, the increase in order book-based volatility is 20.35%, on average, and the increase in median volatility is 25.94%. Both increases are significant at the 5% level.

Collectively, these results indicate that when ChatGPT is unavailable, cryptocurrency returns become more volatile, as investor and trader behavior may be influenced by behavioral biases. These findings support our model's predictions. In the next section, we examine

the impact of ChatGPT outages on liquidity.

V.3 Market Depth

The depth, or the quantity of coins behind the bid and ask prices, corresponds to the inverse of λ in our model and is a measure of liquidity. It captures the sensitivity of cryptocurrency prices to order flow and reflects the degree of information asymmetry. This measure shares many characteristics with other well-known liquidity measures (Amihud, 2002; Chordia et al., 2009).

The Kaiko database provides 30-second snapshots of the limit order book for each coin pair traded on each exchange. In our analysis, we employ two depth measures. The first is the depth in dollar value behind the best bid or best ask price. Since depth in a limit order book is typically quoted in the number of coins for a given coin pair, we convert it to a dollar value to facilitate meaningful cross-sectional comparisons across coin pairs. This measure is commonly referred to as the best-bid depth or best-ask depth.

Table 6 presents summary statistics for best-bid depth, best-ask depth, and their logarithmic changes (Log diff in %). Using the period [-5, -1] as a control, we observe a significant decline of 7.56% in best-bid depth during the outage period, which is statistically significant at the 5% level. The best-bid depth decreases from $70,070 in the control period to $65,090 in the outage period. Both the median-based Wilcoxon test and bootstrapped t-statistic confirm this significant change. Similar results hold when using alternative control periods.

We also report the results for best-ask depth in Table 6. The findings show that both best-bid and best-ask depths decline, though the reduction is more pronounced on the bid side. Specifically, the best-ask depth decreases by 1.58%, 6.91%, and 8.84% between the control and outage periods when using control periods of [-5, -1], [-10, -1], and [-15, -1], respectively. The reduction in best-ask depth is statistically significant for the [-10, -1] and [-15, -1] control periods, but not for [-5, -1]. Collectively, these results indicate a deterioration

in liquidity on both sides, with a more substantial decline in market depth on the best-bid side.

Our second depth measure accounts for the multi-level structure of the limit order book. A large sell order may need to traverse multiple price levels if the market is illiquid, consuming liquidity at various bid prices along the way. This measure is defined as follows: Depth at -0.1% Bid represents the total dollar value of coins available in the limit order book if the bid price decreases by 0.1% relative to the current best bid. Similarly, Depth at +0.1% Ask denotes the total dollar value of coins available if the ask price increases by 0.1% relative to the current best ask. To put this in perspective, consider Bitcoin trading at the best bid price of $20,000. This alternative depth measure quantifies the dollar value of Bitcoin that a large sell order can be filled if a seller places a limit order at a price of $19,900. It is commonly referred to as bid-side book depth or ask-side book depth (Bacidore et al., 2003).

In Table 7, we report changes in the bid-side and ask-side book depths. Using [-5, -1] as the control period, the average bid-side book depth decreases by 3.81%, dropping from $5,476,000 to $5,229,000 in 5-minute intervals. This reduction is statistically significant based on both the t-test and bootstrap test. Although the median reduction is 5.51%, the Wilcoxon test indicates that this change is not significant. Similar quantitative results are observed when using [-10, -1] and [-15, -1] as alternative control periods.

The results for ask-side book depth reveal a reduction magnitude similar to that observed on the bid side. Specifically, when using control periods of [-5, -1], [-10, -1], and [-15, -1], the ask-side book depth decreases by 4.17%, 3.57%, and 4.13%, respectively. The median-based analysis shows a comparable magnitude of reduction. Furthermore, both the t-test and bootstrap test confirm that the reduction in ask-side book depth is statistically significant at the 5% level.

Collectively, these findings indicate a deterioration in liquidity on both the best-bid and best-ask sides during the ChatGPT outages, with a more pronounced decline in market depth

on the best-bid side. In contrast, the reductions in liquidity for the bid-side and ask-side book depths are quite similar.

V.4 Bid-Ask Spreads

In this section, we assess the impact of the ChatGPT outage on another liquidity measure: the bid-ask spread, which serves as a proxy for trading costs. Given the wide dispersion in average prices across the 70 coin pairs in our sample (see Table 3), we employ two spread measures—the quoted spread and the proportional spread. The quoted spread is defined as the difference between the ask and bid prices, while the proportional spread is calculated by dividing the quoted spread by the bid-ask midpoint.

Using the quoted bid-ask spread may cause our analysis to be dominated by coin pairs with high prices—for instance, Bitcoin, which had an average price of $23, during our control period of [-5, -1]. In contrast, using the proportional spread could lead to findings primarily influenced by coin pairs with lower prices, such as DOGEcoin, with an average price of $0.09. To capture changes in trading costs across the entire distribution of 70 coin pairs, we employ both spread measures.

Table 8 presents the results for both quoted and proportional spreads during the control and ChatGPT outage periods. When using [-5, -1] as the control period, the quoted spread increases significantly by 14.87%, with the average bid-ask spread rising from $0.21 to $0.24. All three tests—the t-test, the non-parametric Wilcoxon test, and the bootstrap test—reject the null hypothesis of no change in the bid-ask spread. Moreover, when using [-10, -1] and [-15, -1] as control periods, the bid-ask spread increases by 18.29% and 20.08%, respectively. The results for the proportional spread show a similar magnitude of increase.

In summary, our analysis reveals a significant increase in trading costs during the ChatGPT outage period, with bid-ask spreads rising between 15% and 20%. Although our model does not have a direct implication for the bid-ask spread, these findings are consistent with

our earlier results, which indicate that market liquidity, as measured by market depths, deteriorates during the ChatGPT outage period.

VI Empirical Results: ChatGPT's Introduction

In this section, we assess the impact of ChatGPT's introduction on market liquidity and trading activity. ChatGPT was officially launched on November 30, 2022, which we define as the event day (t = 0). We adopt a [-7, +7] event window, where the seven days preceding November 30, 2022, constitute the pre-event period, and the seven days following form the post-event period. This event window is chosen to ensure that no overlapping events affect our analysis. To ensure a clean event window, we conducted an Internet search and confirmed that no major confounding events occurred during the [-7, +7] window. Due to the absence of two coin pairs during this period, our final sample consists of 68 coin pairs instead of 70.

We summarize the test results in Table 9. Panel A shows that following the launch of ChatGPT, both trading frequency and dollar volume decline significantly. Specifically, trading frequency decreases by 16.10%, while dollar trading volume drops by 11.47%. For instance, the average dollar volume per 5-minute interval declines from $7,042,000 in the pre-event period to $6,761,000 in the post-event period, a reduction that is statistically significant at the 1% level using the t-test, the non-parametric Wilcoxon test, and the bootstrapped-value. Our model predicts that ChatGPT reduces information noise, thereby diminishing the noisy component of total trades. Additionally, increased signal correlation may raise concerns about synchronized trading behavior, leading to more cautious trading overall. The results in Panel A provides supporting evidence for our model's predictions but contrast with the intuitive expectation that ChatGPT's trading advice for retail investors would increase trading activity.

Panel B of Table 9 presents results for changes in return volatility. Trade-based and

quote-based return volatility decrease by 17.07% and 23.34%, respectively, from the pre-event to the post-event period. Panel C shows a reduction in the quoted bid-ask spread by 11.33%. The bid-ask spreads decline from 0.33 bps in the pre-event period to 0.23 bps in the post-event period, a reduction that is statistically significant at the 1% level across all three tests. Panel D presents results for the best-bid and best-ask depths. Both measures show a significant increase in the depths behind the best-bid and best-ask prices. For instance, the best-bid depth increases from $7,170 to $8,970, reflecting a 12.98% increase, which is significant at the 1% level. Similarly, bid-side and ask-side book depths increase significantly by 13.47% and 14.57%, respectively (see Panel E). These findings are consistent with our model's prediction that access to ChatGPT enhances market liquidity.

Overall, the results reported in this section, namely, a significant decrease in trading frequency, dollar volume, and return volatility and a significant increase in market liquidity, as reflected in narrower bid-ask spreads and larger depths, are consistent with the predictions of our model.

VII Conclusion

In this paper, we develop a theoretical model to examine how generative AI (GAI) reshapes market liquidity through a novel friction: correlated noise arising from homogenized AI-generated outputs. When traders receive similar responses from AI, their order flows become synchronized due to correlated signals. On one hand, traders recognize the risk of crowded positions and adjust their trading aggressiveness accordingly. On the other hand, the overall order flow becomes less informative for market makers seeking to infer fundamental value. Our model predicts that access to ChatGPT results in lower trading volume, reduced return volatility, and improved market liquidity.

Our study uses unexpected ChatGPT outages as natural experiments to empirically test

our model's predictions in cryptocurrency markets, which are largely dominated by retail investors with limited access to traditional analytical tools. Using high-frequency trade and quote data for 70 major coin pairs traded on the 10 most active cryptocurrency exchanges, we find that ChatGPT outages lead to a significant increase in trading frequency, dollar volume, and return volatility, alongside a notable decline in quoted depths and a widening of the bid-ask spread. These results indicate a deterioration in market liquidity during ChatGPT outage periods

We also examine the impact of ChatGPT's launch on market liquidity and trading activity, treating this event as an exogenous shock to GAI accessibility. In contrast to the effects observed during ChatGPT outages, we find a significant decrease in trading frequency, dollar volume, and return volatility following its introduction. Additionally, bid-ask spreads narrow, and depths behind the best-bid and best-ask prices increase, indicating improved market liquidity in the seven-day period after ChatGPT's official launch. The results from these two contrasting scenarios, the official launch of ChatGPT and its subsequent unexpected outages, provide strong empirical support for our model's predictions. They highlight the impact of both gaining and losing access to GAI tools on cryptocurrency market liquidity

Our theoretical model and empirical tests yield two insights. Investors of the cryptocurrency market heavily rely on generative AI (GAI) tools like ChatGPT to formulate their trading strategies. Next, GAI-assisted trading improves market liquidity. Beyond testing our model's prediction that noise correlation plays a significant role, our findings underscore the dual impact of GAI in cryptocurrency markets: democratizing access to advanced analytical tools while simultaneously synchronizing noise. While most research on GAI emphasizes its ability to enhance the precision of private signals, our study provides a novel perspective on how AI influences market quality.

References

Ahn, Hee-Joon; Cao, Charles Q.; and Choe, Hyuk (1996). Tick size, spread, and volume. Journal of Financial Intermediation 5(1): 2-22.

Ahn, Hee-Joon; Cao, Charles Q.; and Choe, Hyuk (1998). Decimalization and competition among stock markets: Evidence from the Toronto Stock Exchange cross-listed securities. Journal of Financial Markets 1(1) : 51-87.

Amihud, Yakov (2002). Illiquidity and stock returns: cross-section and time-series effects. Journal of Financial Markets 5(1): 31-56.

Andersen, Torben G.; Bollerslev, Tim; Diebold, Francis X; Labys, Paul (2001). The distribution of realized exchange rate volatility. Journal of the American statistical association, 96 (453), 42-55.

Bacidore, Jeffrey; Battalio, Robert H; and Jennings, Robert H (2003). Order submission strategies, liquidity supply, and trading in pennies on the New York Stock Exchange. Journal of Financial Markets, 6, 337–362.

Chen, Jian; Tang, Guohao; Zhou, Guofu; Zhu, Wu (2025). New ChatGPT and DeepSeek: Can They Predict the Stock Market and Macroeconomy? Working Paper

Beckmann, Lars; Beckmeyer, Heiner; Filippou, Ilias; Menze, Stefan; and Zhou, Guofu (2024). Unusual Financial Communication: ChatGPT, Earnings Calls, and Financial Markets. Working Paper

Bertomeu, Jeremy; Lin, Yupeng; Liu, Yibin; and Ni, Zhenghui (2023). Capital market consequences of generative AI: Early evidence from the ban of ChatGPT in Italy. Working Paper, Washington University in St. Louis, National University of Singapore and Remin University of China.

Canayaz, Mehmet; Cao, Charles; Nguyen, Giang; Wang, Qiang and Huang, Zongbo (2025). An Anatomy of Cryptocurrency Sentiment. Working Paper, Penn State University and Chinese University of Hong Kong.

Cao, Sean; Jiang, Wei; Wang, Junbo; and Yang, Baozhong (2024). From man vs. machine to man+ machine: The art and AI of stock analyses. Journal of Financial Economics, 160, 103910.

Cao, Sean; Jiang, Wei; Yang, Baozhong; and Zhang, Alan L (2023). How to talk when a machine is listening: Corporate disclosure in the age of AI. Review of Financial Studies, 36(9), 3603–3642.

Cheng, Qiang; Lin, Pengkai; and Zhao, Yue (2024). Does Generative AI Facilitate Investor Trading? Evidence from ChatGPT Outages. Working Paper, Singapore Management University.

Chordia, Tarun; Huh, Sahn-Wook, and Subrahmanyam, Avanidhar (2009). Theory-based illiquidity and asset pricing. The Review of Financial Studies 22(9): 3629-3668.

D'Acunto, Francesco; Prabhala, Nagpurnanand; and Rossi, Alberto G (2019). The promises and pitfalls of robo-advising. Review of Financial Studies, 32, 1983–2020.

Ebert, Lauren (2023). The Rise of AI Assistants: Hedge Fund Managers and ChatGPT. Retrieved from https://globalmarkets.cib.bnpparibas/the-rise-of-ai-assistants-hedge-fund-managers-and-chatgpt/.

Foster, F Douglas and Viswanathan, S (1996). Strategic trading when agents forecast the forecasts of others. Journal of Finance, 51, 1437–1478.

Goyenko, Ruslan Y.; Holden, Craig W.; and Trzcinka, Charles A. (2009). Do liquidity measures measure liquidity?. Journal of financial Economics 92(2): 153-181.

Grossman, Sanford J and Stiglitz, Joseph E (1980). On the impossibility of informationally efficient markets. American Economic Review, 70, 393–408.

Kogan, Shimon; Makarov, Igor; Niessner, Marina; and Schoar, Antoinette (2024). Are cryptos different? Evidence from retail trading. Journal of Financial Economics, 159, 103897.

Kyle, Albert (1985). Continuous auctions and insider trading. Econometrica, 53(6), 1315–1335.

Kyle, Albert (1989). Informed speculation with imperfect competition. Review of Economic Studies, 56, 317–355.

Loos, Benjamin; Previtero, Alessandro; Scheurle, Sebastian; and Hackethal, Andreas (2020). Robo-advisers and investor behavior. Working Paper, University of Florida.

Lopez-Lira, Alejandro and Tang, Yuehua (2023). Can chatgpt forecast stock price movements? Return predictability and large language models. Working Paper, University of Florida.

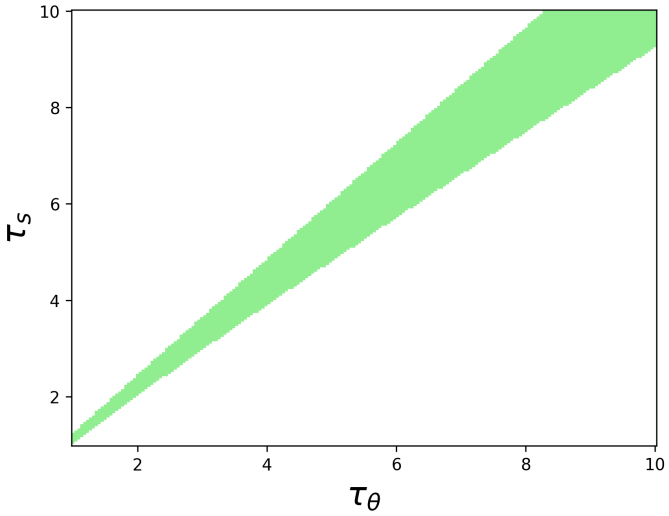
Makarov, Igor and Schoar, Antoinette (2020). Trading and arbitrage in cryptocurrency markets. Journal of Financial Economics, 135, 293–319.

Rossi, Alberto G and Utkus, Stephen (2024). The diversification and welfare effects of roboadvising. Journal of Financial Economics, 157, 103869.

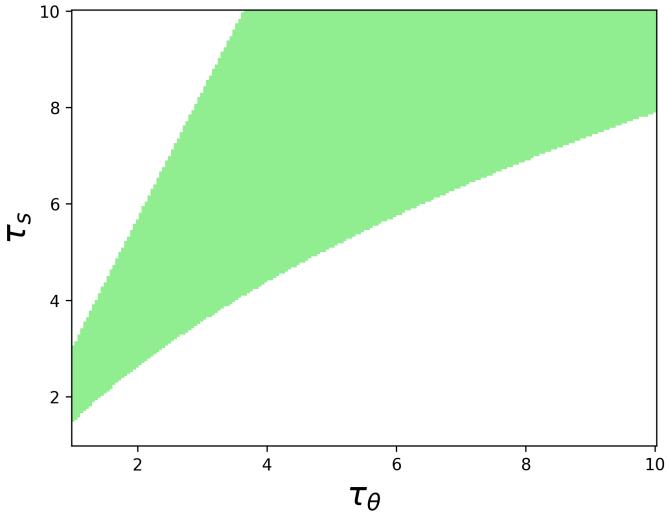
Sheng, Jinfei; Sun, Zheng; Yang, Baozhong; and Zhang, Alan L (2024). Generative AI and Asset Management. Working Paper, University of California, Irvine, Georgia State University and Florida International University.

Vaswani, A (2017). Attention is all you need. Advances in Neural Information Processing Systems.

Wang, Jiang (1994). A model of competitive stock trading volume. Journal of Political Economy, 102(1), 127–168. The University of Chicago Press.



(a) $\rho_1 = 0.1$



(b) $\rho_1 = 0.8$

Figure 1: **Market Dynamics in (τ_θ, τ_s) Regimes after ChatGPT Adoption**

This figure depicts the parametric regimes of (τ_θ, τ_s) for which, without the access to ChatGPT, trading volumes and return volatility are higher while market depth is lower. These regimes are color-coded in green. The parameters are set as follows: $N = 3$, $M = 4$, $\tau_0 = 1$, and $\tau_u = 3$. In Panel (a), $\rho^* = 0.1$, while in Panel (b), $\rho^* = 0.8$.

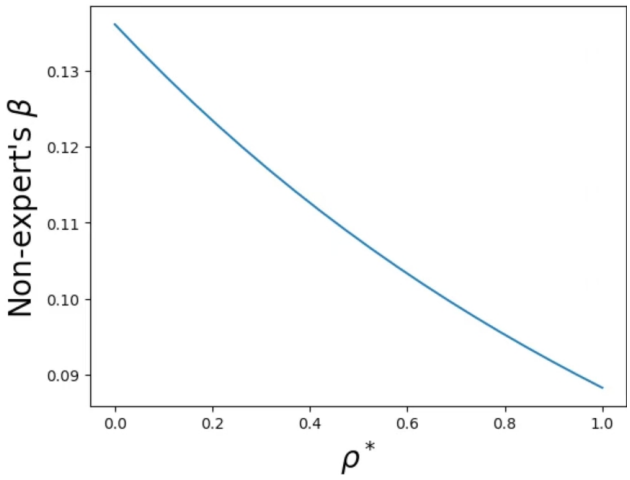
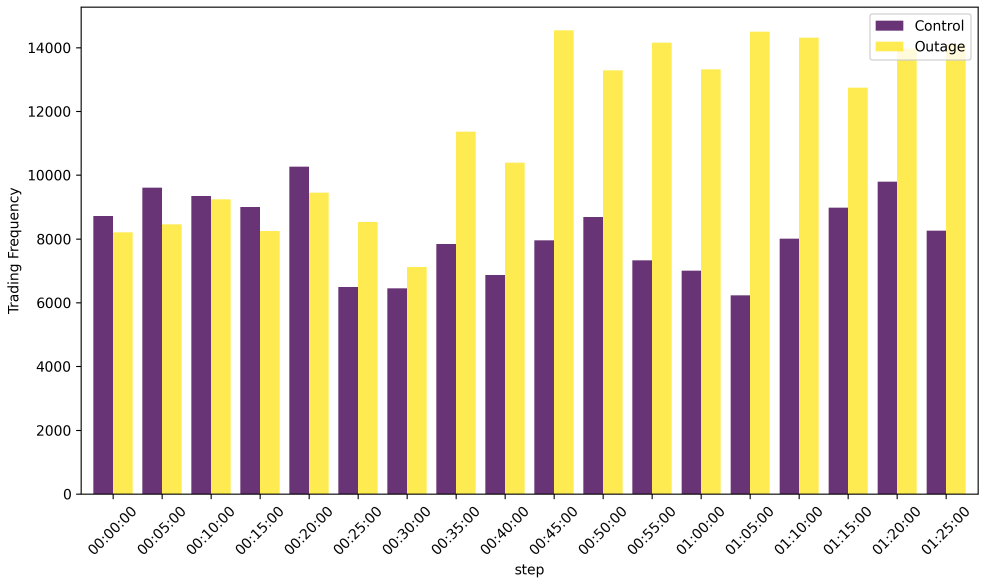
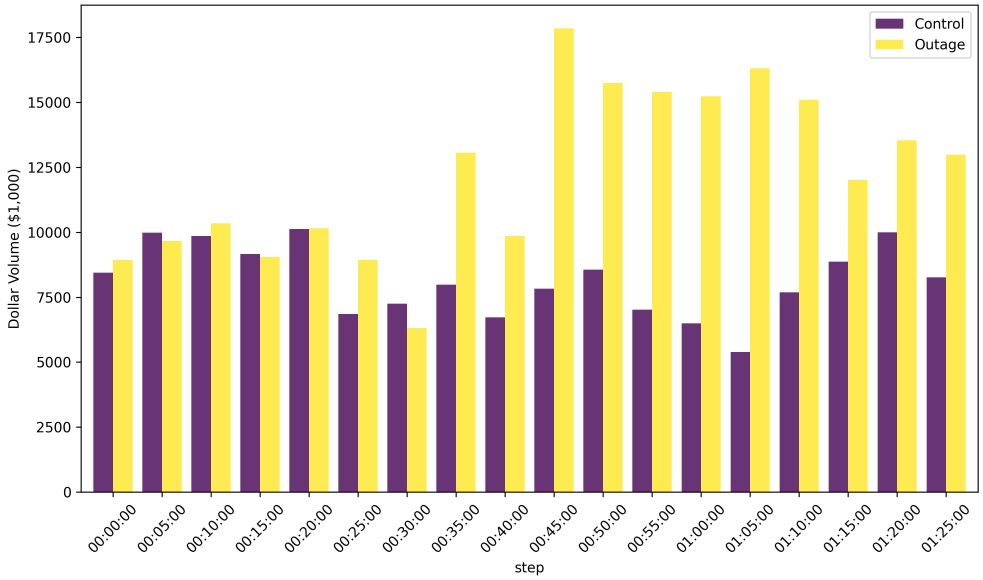


Figure 2: **ChatGPT's Impact on Non-Expert Traders' Strategy**

This figure depicts the impact of ChatGPT on non-expert traders' strategies. The parameters are set as follows: $N = 3$, $M = 4$, $\tau_0 = 1$, and $\tau_u = 3$. ρ^* is the noise correlation induced by ChatGPT.



(a) Trading Frequency



(b) Dollar Volume

Figure 3: **Binance-BTCUSDT Trading Frequency and Dollar Volume Before and During Outage Periods**

This figure illustrates the 5-minute average trading frequency (dollar volume) of Binance-BTCUSDT before and during outages. For each outage event, we define the same period from one to five days before the outage date (e.g., [-5,-1]) as the control period. We compute the average trading frequency (dollar volume) from the first 5-minute interval ("step 00:00:00") to the last interval for both control and outage periods.

Table 1. ChatGPT Outage Events

This table presents the ChatGPT outage events included in our sample. We obtain information on these events and their severity levels from OpenAI's weblogs. Outage incidents lasting less than 10 minutes are excluded. The sample period spans from November 2022 to October 2024.

Title	Start Time (UTC)	End Time (UTC)	Duration
Outage on All Models	2023/2/21 7:40	2023/2/21 12:09	269 mins
ChatGPT Outage	2023/3/15 17:55	2023/3/15 18:18	23 mins
ChatGPT Outage	2023/4/24 3:23	2023/4/24 3:37	14 mins
ChatGPT Site Outage	2023/5/25 11:23	2023/5/25 12:27	64 mins
Customer Logins to ChatGPT Are Not Available	2023/6/30 18:41	2023/6/30 19:20	39 mins
ChatGPT Severe Degradation	2023/8/3 23:25	2023/8/3 23:57	32 mins
ChatGPT Web UI Unavailable	2023/8/24 0:33	2023/8/24 0:55	22 mins
Major Outage across ChatGPT and API	2023/11/8 13:54	2023/11/8 15:46	112 mins
Major Outage on ChatGPT on Web	2023/12/14 1:52	2023/12/14 3:15	83 mins
Service Degradation in ChatGPT	2024/1/11 4:41	2024/1/11 8:19	218 mins
Partial Outage for ChatGPT	2024/5/9 23:52	2024/5/10 1:18	86 mins
Outage on ChatGPT and API Platform	2024/7/5 22:23	2024/7/5 22:42	19 mins

Table 2. Selection of Cryptocurrency Exchanges and Coin Pairs

This table summarizes the cryptocurrency exchanges and coin pairs used in this study. We select eight exchanges based on their consistent trading of popular coin pairs throughout our sample period. We select 10 non-stable coins with the largest market capitalization as of the end of November 2022: Bitcoin (BTC), Ethereum (ETH), Binance Coin (BNB), Ripple (XRP), Dogecoin (DOGE), Cardano (ADA), Polygon (MATIC), Polkadot (DOT), Litecoin (LTC), and Solana (SOL). For each exchange and each coin, we select the primary coin pair with the highest trading volume from November 2022 to January 2023 when multiple pairs are available. For example, on Coinbase, both BTC/USD and BTC/USDT are actively traded. However, the trading volume of BTC/USD is 17 times larger than that of BTC/USDT from November 2022 to January 2023, so we include BTC/USD in our sample. There are 70 coin-pairs in total in our sample. The sample period spans from November 2022 to October 2024.

	Exchanges							
Coin-pair	Bitstamp	Coinbase	Gemini	Kraken	Binance	Huobi	KuCoin	OkEX
Bitcoin	BTCUSD	BTCUSD	BTCUSD	BTCUSD	BTCUSDT	BTCUSDT	BTCUSDT	BTCUSDT
Ethereum	ETHUSD	ETHUSD	ETHUSD	ETHUSD	ETHUSDT	ETHUSDT	ETHUSDT	ETHUSDT
Binance Coin					BNBUSDT	BNBUSDT	BNBUSDT	BNBUSDT
Ripple	XRPUSD	XRPUSD		XRPUSD	XRPUSDT	XRPUSDT	XRPUSDT	XRPUSDT
Doge Coin		DOGEUSD	DOGEUSD		DOGEUSDT	DOGEUSDT	DOGEUSDT	DOGEUSDT
Cardano	ADAUSD	ADAUSD		ADAUSD	ADAUSDT	ADAUSDT	ADAUSDT	ADAUSDT
Polygon	MATICUSD	MATICUSD	MATICUSD	MATICUSD	MATICUSDT	MATICUSDT	MATICUSDT	MATICUSDT
Polkadot		DOTUSD	DOTUSD	DOTUSD	DOTUSDT	DOTUSDT	DOTUSDT	DOTUSDT
Litecoin	LTCUSD	LTCUSD	LTCUSD	LTCUSD	LTCUSDT	LTCUSDT	LTCUSDT	LTCUSDT
Solana		SOLUSD	SOLUSD	SOLUSD	SOLUSDT	SOLUSDT	SOLUSDT	SOLUSDT

Table 3. Characteristics of Sample Coin Pairs

This table presents summary statistics for the 70 coin pairs in our sample. Intraday data for each variable is aggregated at the 5-minute level. For each variable, we first compute its time-series average and then report cross-sectional summary statistics, including the mean, standard deviation (S.D.), and the 25th, 50th, and 75th percentiles. Trading frequency is the number of transactions within each 5-minute interval. Dollar volume is the value of coins traded in each 5-minute interval (in $1,000). Trade-based (order book-based) return volatility is the standard deviation of 5-minute returns, where returns are based on the log difference of transaction prices (the midpoint of the bid and ask prices). Quoted spread is the difference between the best bid and ask prices, while proportional quoted spread is the quoted spread divided by the midpoint of the bid and ask prices. Best bid (ask) depth is the number of coins (in $1,000) available at the best bid (ask) price. Depth at -0.1% Bid is the number of coins (in $1,000) available in the limit order book if the bid price decreases by 0.1% relative to the current best bid price, whereas Depth at +0.1% Ask is the number of coins (in $1,000) available if the ask price increases by 0.1% relative to the current best ask price. For each coin pair and variable, we first calculate its time series average, then report its cross-sectional summary statistics, including mean, standard deviation (S.D.), and the 25th, 50th, and 75th percentiles. The sample comprises 1,545,679 panel observations across all coin pairs, covering the period from November 2022 to December 2022—a two-month window preceding the first ChatGPT outage event.

	Mean	S.D.	P25	P50	P75
Trading Frequency	512.36	1763.38	44.22	192.35	367.44
Dollar Volume ($1,000)	358.76	1530.72	16.98	48.46	147.97
Return (bps)	-0.12	0.42	-0.17	-0.05	0.05
Price ($)	2828.20	7287.06	0.44	12.39	288.15
Trade-based Return Volatility (bps)	4.41	1.34	3.27	4.36	5.24
Order Book-based Return Volatility (bps)	4.92	1.29	3.75	5.04	5.76
Quoted Spread	0.17	0.70	0.00	0.00	0.05
Proportional Quoted Spread	3.32	3.82	1.28	2.42	3.55
Best-Bid Depth ($1,000)	9.36	15.58	0.97	2.34	9.18
Best-Ask Depth ($1,000)	9.69	16.37	0.91	2.37	9.73
Depth of -0.1% Bids ($1,000)	169.77	316.85	22.81	53.10	160.85
Depth of +0.1% Asks ($1,000)	176.77	321.54	23.13	54.26	139.26

Table 4. Trading Frequency and Dollar Volume for the Periods before and during the ChatGPT Outages

This table presents summary statistics of trading frequency, dollar volume, and the logarithmic change (Log diff in %) of these variables for the periods before and during the ChatGPT outages. For each outage event, we define the same period from one to five days before the outage date (e.g., [−5,−1]) as the primary control period. To ensure robustness, we also use alternative control periods [−10,−1] and [−15,−1]. For each coin pair, we compute the average trading frequency (dollar volume) across all outage events and define it as the average trading frequency (dollar volume) during the outage period. The corresponding values during the pre-event period are defined similarly. The table reports the cross-sectional mean, standard error (S.E.), median, results from the nonparametric Wilcoxon test, and empirical p-values from bootstrap simulations. Intraday data for each variable is aggregated at the 5-minute level. The sample period spans from November 2022 to October 2024. ***, **, and * indicate the statistical significance of a t-test at 1%, 5%, and 10% levels, respectively.

		Trading Frequency			Dollar Volume ($1,000)		
Control Period		Control	Outage	Log Diff. (%)	Control	Outage	Log Diff. (%)
[−5,−1]	Mean	5157.74	6812.45	22.80***	4914.83	6806.72	26.14***
	S.E.	(662.77)	(887.89)	(1.89)	(630.73)	(883.44)	(2.34)
	Median	2353.78	3063.51	28.60	2481.66	3504.09	33.14
	Wilcoxon p-value			(0.00)			(0.00)
	Bootstrap p-value			(0.00)			(0.00)
[−10,−1]	Mean	4648.49	6686.75	28.01***	4669.84	6686.84	27.69***
	S.E.	(595.25)	(880.65)	(2.25)	(598.28)	(875.77)	(2.37)
	Median	2269.22	3063.51	35.71	2494.98	3504.09	37.11
	Wilcoxon p-value			(0.00)			(0.00)
	Bootstrap p-value			(0.00)			(0.00)
[−15,−1]	Mean	4554.90	6681.10	29.95***	4529.55	6679.27	30.54***
	S.E.	(583.24)	(880.32)	(2.23)	(580.91)	(875.66)	(2.36)
	Median	2233.88	3063.51	39.70	2420.82	3504.09	40.05
	Wilcoxon p-value			(0.00)			(0.00)
	Bootstrap p-value			(0.00)			(0.00)

Table 5. Return Volatility for the Periods before and during the ChatGPT Outages

This table presents summary statistics of trade-based return volatility, order book-based return volatility, and the logarithmic change (Log diff in %) of these variables for the periods before and during the ChatGPT outages. For each outage event, we define the same period from one to five days before the outage date (e.g., [−5,−1]) as the primary control period. To ensure robustness, we also use alternative control periods [−10,−1] and [−15,−1]. Trade-based (order book-based) return volatility is the standard deviation of 5-minute returns, where returns are based on the log difference of transaction prices (the midpoint of the bid and ask prices). For each coin pair and each volatility measure, we compute the average volatility across all outage events and define it as the average volatility during the outage period. The corresponding values during the pre-event period are defined similarly. The table reports the cross-sectional mean, standard error (S.E.), median, results from the nonparametric Wilcoxon test, and empirical p-values from bootstrap simulations. Intraday data for each variable is aggregated at the 5-minute level. The sample period spans from November 2022 to October 2024. ***, **, and * indicate the statistical significance of a t-test at 1%, 5%, and 10% levels, respectively.

Panel A. Trade-based Return Volatility (in bps)

Control Period	[−5,−1]			[−10,−1]			[−15,−1]		
	Control	Outage	Log Diff. (%)	Control	Outage	Log Diff. (%)	Control	Outage	Log Diff. (%)
Mean	3.38	4.05	20.35***	3.32	4.04	22.08***	3.25	4.04	24.50***
S.E.	(0.17)	(0.15)	(1.23)	(0.17)	(0.15)	(1.20)	(0.17)	(0.15)	(1.29)
Median	2.68	3.48	25.94	2.64	3.48	27.64	2.55	3.48	30.99
Wilcoxon p-value			(0.00)			(0.00)			(0.00)
Bootstrap p-value			(0.00)			(0.00)			(0.00)

Panel B. Order Book-based Return Volatility (in bps)

Control Period	[−5,−1]			[−10,−1]			[−15,−1]		
	Control	Outage	Log Diff. (%)	Control	Outage	Log Diff. (%)	Control	Outage	Log Diff. (%)
Mean	3.88	4.48	16.22***	3.76	4.47	19.10***	3.68	4.48	21.73***
S.E.	(0.19)	(0.17)	(1.10)	(0.18)	(0.17)	(1.11)	(0.18)	(0.17)	(1.22)
Median	3.09	3.81	20.90	3.02	3.81	23.23	2.91	3.81	26.80
Wilcoxon p-value			(0.00)			(0.00)			(0.00)
Bootstrap p-value			(0.00)			(0.00)			(0.00)

Table 6. Best-bid and Best-ask Depths for the Periods before and during the ChatGPT Outages

This table reports summary statistics of best-bid depth, best-ask depth and the logarithmic change (Log diff in %) of these variables for the periods before and during the ChatGPT outages. For each outage event, we define the same period from one to five days before the outage date (e.g., [−5,−1]) as the primary control period. To ensure robustness, we also use alternative control periods [−10,−1] and [−15,−1]. For each coin pair, we compute the average of the best-bid (best-ask) depth across all outage events and define it as the average best-bid (best-ask) depth during the outage period. The corresponding values during the pre-event period are defined similarly. The table reports the cross-sectional mean, standard error (S.E.), median, results from the nonparametric Wilcoxon test, and empirical p-values from bootstrap simulations. Intraday data for each variable is aggregated at the 5-minute level. The sample period spans from November 2022 to October 2024. ***, **, and * indicate the statistical significance of a t-test at 1%, 5%, and 10% levels, respectively.

		Best-bid Depth (in $1,000)			Best-ask Depth (in $1,000)		
Control Period		Control	Outage	Log Diff. (%)	Control	Outage	Log Diff. (%)
[−5,−1]	Mean	70.07	65.09	-7.56***	71.56	69.36	-4.02***
	S.E.	(6.80)	(6.55)	(1.41)	(7.02)	(7.07)	(1.37)
	Median	80.13	77.28	-9.11	83.86	72.33	-1.54
	Wilcoxon p-value			(0.03)			(0.21)
	Bootstrap p-value			(0.00)			(0.01)
[−10,−1]	Mean	73.98	64.50	-12.35***	74.85	68.66	-9.27***
	S.E.	(7.30)	(6.53)	(1.57)	(7.39)	(7.04)	(1.24)
	Median	84.48	77.28	-15.83	87.31	72.33	-7.51
	Wilcoxon p-value			(0.00)			(0.01)
	Bootstrap p-value			(0.00)			(0.00)
[−15,−1]	Mean	74.85	64.36	-13.64***	75.60	68.50	-10.36***
	S.E.	(7.41)	(6.53)	(1.64)	(7.48)	(7.03)	(1.26)
	Median	85.09	77.28	-18.28	86.93	72.33	-8.88
	Wilcoxon p-value			(0.00)			(0.00)
	Bootstrap p-value			(0.00)			(0.00)

Table 7. Bid-side and Ask-side Book Depths for the Periods before and during the ChatGPT Outages

This table reports summary statistics of two alternative depth measures and the logarithmic change (Log diff in %) of these variables for the periods before and during the ChatGPT outages. Depth at -0.1% Bid is the number of coins (in $1,000) available in the limit order book if the bid price decreases by 0.1% relative to the current best bid price, whereas Depth at +0.1% Ask is the number of coins (in $1,000) available if the ask price increases by 0.1% relative to the current best ask price. For each outage event, we define the same period from one to five days before the outage date (e.g., [−5,−1]) as the primary control period. To ensure robustness, we also use alternative control periods [−10,−1] and [−15,−1]. For each coin pair, we compute the average of Depth at -0.1% Bid (Depth at +0.1% Ask) across all outage events and define it as the average Depth at -0.1% Bid (Depth at +0.1% Ask) during the outage period. The corresponding values during the pre-event period are defined similarly. The table reports the cross-sectional mean, standard error (S.E.), median, results from the nonparametric Wilcoxon test, and empirical p-values from bootstrap simulations. Intraday data for each variable is aggregated at the 5-minute level. The sample period spans from November 2022 to October 2024. ***, **, and * indicate the statistical significance of a t-test at 1%, 5%, and 10% levels, respectively.

		Depth of -0.1% Bids ($1,000)			Depth of +0.1% Asks ($1,000)		
Control Period		Control	Outage	Log Diff. (%)	Control	Outage	Log Diff. (%)
[−5,−1]	Mean	1500.42	1459.20	-3.49***	1633.33	1470.06	-8.03***
	S.E.	(129.27)	(126.15)	(0.78)	(142.35)	(122.28)	(1.15)
	Median	1464.71	1371.72	-2.64	1607.16	1432.44	-13.71
	Wilcoxon p-value			(0.06)			(0.36)
	Bootstrap p-value			(0.00)			(0.00)
[−10,−1]	Mean	1567.28	1450.12	-7.11***	1644.21	1461.87	-8.65***
	S.E.	(137.36)	(124.76)	(0.87)	(143.87)	(120.86)	(1.18)
	Median	1485.87	1371.66	-9.07	1597.19	1432.44	-14.21
	Wilcoxon p-value			(0.00)			(0.40)
	Bootstrap p-value			(0.00)			(0.00)
[−15,−1]	Mean	1576.92	1448.69	-7.90***	1657.35	1459.07	-9.36***
	S.E.	(138.72)	(125.00)	(0.92)	(145.74)	(120.90)	(1.21)
	Median	1495.26	1371.66	-9.88	1598.21	1432.44	-12.43
	Wilcoxon p-value			(0.00)			(0.47)
	Bootstrap p-value			(0.00)			(0.00)

Table 8. Quoted Spread for the Periods before and during the ChatGPT Outages

This table reports summary statistics of Quoted Spread, and the logarithmic change (Log diff in %) of these variables for the periods before and during the ChatGPT outages. For each outage event, we define the same period from one to five days before the outage date (e.g., [−5,−1]) as the primary control period. To ensure robustness, we also use alternative control periods [−10,−1] and [−15,−1]. For each coin pair, we compute the average quoted spread across all outage events and define it as the average quoted spread during the outage period. The corresponding values during the pre-event period are defined similarly. The table reports the cross-sectional mean, standard error (S.E.), median, results from the nonparametric Wilcoxon test, and empirical p-values from bootstrap simulations. Intraday data for each variable is aggregated at the 5-minute level. The sample period spans from November 2022 to October 2024. ***, **, and * indicate the statistical significance of a t-test at 1%, 5%, and 10% levels, respectively.

Control Period	[−5,−1]			[−10,−1]			[−15,−1]		
	Control	Outage	Log Diff. (%)	Control	Outage	Log Diff. (%)	Control	Outage	Log Diff. (%)
Mean	0.21	0.24	14.87***	0.20	0.24	18.29***	0.20	0.24	20.08***
S.E.	(0.07)	(0.07)	(2.20)	(0.07)	(0.07)	(2.57)	(0.06)	(0.07)	(2.78)
Median	0.11	0.14	10.56	0.11	0.14	11.75	0.11	0.14	15.83
Wilcoxon p-value			(0.00)			(0.00)			(0.00)
Bootstrap p-value			(0.00)			(0.00)			(0.00)

Table 9. The Impact of ChatGPT Introduction

This table presents summary statistics of variables on trading volume, return volatility, market liquidity, and the logarithmic change (Log diff in %) of these variables for the periods before and after the introduction of ChatGPT on Nov 30, 2022. For control and treated period, we use seven days before the launch date as the primary control period and seven days after as the primary treated period. The table reports the cross-sectional mean, standard error (S.E.), median, results from the nonparametric Wilcoxon test, and empirical p-values from bootstrap simulations. Intraday data for each variable is aggregated at the 5-minute level. The sample period spans from November 24, 2022 to December 7, 2022. ***, **, and * indicate the statistical significance of a t-test at 1%, 5%, and 10% levels, respectively.

Panel A. Trading Frequency and Dollar Volume

	Trading Frequency			Dollar Volume ($1,000)		
	Before	Post	Log Diff. (%)	Before	Post	Log Diff. (%)
Mean	9406.92	8425.88	-16.10***	7042.05	6761.63	-11.47***
S.E.	(1123.53)	(1016.37)	(2.42)	(844.44)	(821.60)	(2.42)
Median	16681.06	15006.61	-10.58	12504.46	12078.80	-3.46
Wilcoxon p-value			(0.00)			(0.00)
Bootstrap p-value			(0.00)			(0.00)

Panel B. Return Volatility

	Trade-based Return Volatility (bps)			Order Book-based Return Volatility (bps)		
	Before	Post	Log Diff. (%)	Before	Post	Log Diff. (%)
Mean	0.53	0.47	-17.07***	3.55	2.75	-23.34***
S.E.	(0.14)	(0.13)	(1.33)	(0.28)	(0.18)	(0.97)
Median	0.15	0.12	-17.59	2.52	2.03	-21.23
Wilcoxon p-value			(0.00)			(0.00)
Bootstrap p-value			(0.00)			(0.00)

Panel C. Quoted Spread

	Quoted Spread (bps)				
	Mean	S.E.	Median	Wilcoxon p-value	Bootstrap p-value
Before	0.33	(0.07)	0.40		
Post	0.28	(0.04)	0.36		
Log Diff.	-11.33***	(2.11)	-12.47	(0.00)	(0.00)

Panel D. Best-bid and Best-ask Depths

	Best Bid Depth ($1,000)			Best Ask Depth ($1,000)		
	Before	Post	Log Diff. (%)	Before	Post	Log Diff. (%)
Mean	7.17	8.97	12.98***	7.45	9.50	31.63***
S.E.	(1.76)	(2.24)	(2.85)	(1.83)	(2.32)	(3.32)
Median	1.01	1.18	15.13	0.96	1.54	46.76
Wilcoxon p-value			(0.01)			(0.00)
Bootstrap p-value			(0.00)			(0.00)

Panel E. Bid-side and Ask-side Book Depths

	Depth of -0.1% Bids ($1,000)			Depth of +0.1% Asks ($1,000))		
	Before	Post	Log Diff. (%)	Before	Post	Log Diff. (%)
Mean	1483.40	1648.88	13.47***	1382.91	1535.90	14.57***
S.E.	(143.55)	(158.45)	(1.69)	(128.23)	(138.28)	(1.73)
Median	2381.53	2642.99	10.42	2176.05	2387.99	9.29
Wilcoxon p-value			(0.00)			(0.00)
Bootstrap p-value			(0.00)			(0.00)

Appendix A

Proof of Proposition 1

Our proof needs a property of Gaussian distribution to facilitate the calculation of coniditional expectations. See Equation $(A1)$ and $(A2)$ in Grossman and Stiglitz (1980)

Lemma 1: If X^* and Y^* are jointly normally distributed then

$$E[X^*|Y^* = Y] = EX^* + \frac{Cov(X^*, Y^*)}{Var(Y^*)}\{Y - EY^*\} \tag{11}$$

$$Var[X^*|Y^* = Y] = Var(X^*) - \frac{[Cov(X^*, Y^*)]^2}{Var(Y^*)} \tag{12}$$

In a linear rational expectation equilibrium, the demand of trader i is a linear function of her signal. Within non-expert traders or expert traders, everyone is the same. Due to this within-group symmetry, we assume that the expert trader's strategy is denoted as:

$$x(s) = \overline{\beta}(s - v_0) \tag{13}$$

And the strategy of non-expert traders is:

$$x(s) = \underline{\beta}(s_i - v_0) \tag{14}$$

The pricing function $P(y)$ is also linear:

$$P(y) = v_0 + \lambda y \tag{15}$$

Given these function, the total order flow, y, is normally distributed

$$y = \sum_{i=1}^{N} \overline{\beta}(s_i - v_0) + \sum_{i=N+1}^{N+M} \underline{\beta}(s_i - v_0) + u \tag{16}$$

The competitive market-making assumption implies the pricing condition, $P(y) = E[v|y]$. Using the Lemma 1, we have

$$\lambda = \frac{N\overline{\beta} + M\underline{\beta}}{(N\overline{\beta} + M\underline{\beta})^2 + \underline{\beta}^2[M + \rho M(M-1)]\frac{\tau_0}{\tau_s} + \frac{\tau_0}{\tau_u}} \tag{17}$$

The optimization problem of an expert trader is

$$\max_{x_i} \quad E[(v - P(y))x_i|v] \tag{18}$$

The first-order condition is

$$v - E[P(y)|v] - \lambda x_i = 0 \tag{19}$$

which can be simplified to

$$\overline{\beta} = \frac{1}{2\lambda} - \frac{N-1}{2}\overline{\beta} - \frac{M}{2}\underline{\beta} \tag{20}$$

The optimization problem of a non-expert trader is

$$\max_{x_i} \quad E[(v - P(y))x_i|s_i] \tag{21}$$

The first-order condition is

$$v - E[P(y)|s_i] - \lambda x_i = 0 \tag{22}$$

Using Lemma 1, the FOC can be simplified to

$$\underline{\beta} = (\frac{1}{2\lambda} - \frac{\overline{\beta}N}{2})\frac{\tau_s}{\tau_0 + \tau_s} - \frac{M-1}{2}\underline{\beta}\frac{\tau_s + \rho\tau_0}{\tau_s + \tau_0} \tag{23}$$

According to Definition 1, the equilibrium can be solved jointly from equation (17), (20), and (23).

Proof of Proposition 2

We need to use the property of a mean-zero normally distributed variable: If $x \sim N(0, \sigma^2)$, then $E[|x|] = \sqrt{\frac{2}{\pi}}\sigma$. Using the property, we immediately get the expression of the trading volume.

For the return volatility, using Lemma 1

$$Vol = Var(P(y) - v_0) = \lambda^2 * Var(y) = \left(\frac{Cov(v,y)}{Var(y)}\right)^2 * Var(y) = \frac{Cov(v,y)^2}{Var(y)} \tag{24}$$

We can get the expression of return volatility by calculating $Cov(v, y)$ and $Var(y)$.

Proof of Proposition 3

Given $\overline{\beta}$ and $\underline{\beta}$ fixed and Assumption 1, we can calculate the trading volume, return volatility, and market depth using Equation (6), (7), and (8) for the case with/without GAIs. We can immediately arrive at the result.