# When AI Fails:

# A Natural Experiment on Generative AI's Role in Crypto Markets

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

January 2026

*Charles Cao is at the CUHK Business School, Chinese University of Hong Kong, and the Smeal College of Business, Penn 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 School of Management and Economics, Chinese University of Hong Kong, Shenzhen. Email: deliyang@link.cuhk.edu.cn. We are grateful for the comments of Mehmet Canayyaz, Matthew Gustafson, Zongbo Huang, Stephen Lenkey, Tim Simin, Mihail Velikov, Alexey Zhdanov, Burton Hollifield, Jun Pan, Patrick Blonein, Simon Mayer, Hui Chen, William Cong, Mao Ye, Lawrence Jin, Evgeny Lyandres, Bing Han, Marco Pagano, Rui Shen, Pengfei Sui, Eleanor Xu, and seminar participants at the Chinese University of Hong Kong, Penn State University, the 2025 University of Pittsburgh-Carnegie Mellon University-Penn State University Finance Conference, the 2025 CFRC Conference, the 2025 Financial Management Association (FMA) annual meeting, the 2025 New Zealand Finance Meeting, and the 2026 American Economic Association (AEA) annual meeting. All errors are our own.

Acknowledgment: 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.

**When AI Fails:**

**A Natural Experiment on Generative AI's Role in Crypto Markets**


## Abstract

This study examines the impact of ChatGPT on trading volume, return volatility, and market liquidity in cryptocurrency markets. We analyze a key trade-off: while ChatGPT enhances information quality for individual traders, it impairs collective information production by introducing correlated noise. We develop a model and show that ChatGPT adoption enhances market liquidity but reduces trading volume and return volatility. Using ChatGPT outages as exogenous shocks and data from 70 cryptocurrency pairs, we find significant increases in trading volume (24%) and return volatility (19%), alongside a significant deterioration in liquidity—book depth declines by 8% and price impact rises by 12% during outages. These results suggest that ChatGPT adoption reduces disagreement between traders and the market maker and mitigates adverse selection. The absence of significant changes in order imbalance is consistent with a correlated-noise channel and rules out alternative mechanisms that predict a dominant directional trading shift.



Keywords: Generative AI, ChatGPT, Cryptocurrencies, Liquidity, and Market Quality

# 1. Introduction

The advent of Bitcoin in 2009 catalyzed a profound transformation in the global financial landscape and gave rise to a new digital asset class. Since its inception, the cryptocurrency market has experienced exponential growth, with its market capitalization soaring to $3.95 trillion by September 2025. This meteoric rise is underscored by Bitcoin's ascent to become one of the world's top ten most valuable assets, surpassing silver and trailing only gold and a select few technology firms, such as Apple and NVIDIA.

Parallel to the maturation of the cryptocurrency market, the financial sector is undergoing a transformative revolution with the advent of Generative Artificial Intelligence (GenAI), such as ChatGPT. These sophisticated AI systems fundamentally reshape financial decision-making by enhancing the capacity to process, synthesize, and act upon vast quantities of complex information. For example, ChatGPT and other GenAI tools can be deployed to automate and improve a broad spectrum of tasks, including real-time news aggregation and rapid prototyping and back-testing of trading strategies.

There are compelling arguments both for and against the adoption of GenAI in the financial sector. Proponents contend that it augments human intelligence by managing the scale and speed of data processing and allows investors to focus on judgment, creativity, and execution of trading strategies. In the information-saturated environment of financial markets, this synergy of machine speed and human wisdom is seen as a significant competitive advantage. Opponents, however, warn that the widespread adoption of tools like ChatGPT introduces a dangerous combination of correlated decision-making and unproven reliability. They argue it risks creating a fragile financial ecosystem where a single point of failure—an error in a dominant AI model or the coordinated manipulation of its data sources—could have cascading effects, all while potentially eroding the deep, independent analysis that underpins true price discovery.

Despite this debate, there remains a gap in the literature. First, a theoretical framework analyzing the trade-offs of GenAI adoption is lacking. Second, there is little empirical evidence on its impact in the real world. The purpose of this article is to fill in this gap and conduct an empirical study on the impact of ChatGPT on market quality and investor trading in cryptocurrency markets.

To this goal, we first develop a three-period market microstructure model to formalize the effects of ChatGPT on trading volume, return volatility, and market liquidity. The model features traders who adopt ChatGPT in their trading and a competitive market maker who sets prices. We analyze a key trade-off: while ChatGPT adoption enhances an individual trader's information-processing capability and signal precision, it also introduces correlated noise across traders' signals, thereby impairing collective information production. The model indicates that in an environment lacking well-defined fundamental values—such as cryptocurrency markets—this noise-correlation effect dominates. Given the unobservable nature of ChatGPT adoption dynamics, we employ an event-study approach, leveraging unexpected ChatGPT outages as exogenous shocks. We compare market outcomes during outages to those in normal periods. Our hypothesis is that, relative to normal periods, cryptocurrency markets will experience increased trading volume and return volatility and a deterioration in market liquidity during outage windows.

We test these hypotheses by examining changes in trading activity and market liquidity in cryptocurrency markets surrounding unexpected ChatGPT outages. These outages represent sudden, unplanned disruptions in GAI access caused by technical failures. Occurring when ChatGPT experiences significant malfunctions that completely cut off user access, these events are exogenous to market participants and thus serve as a quasi-natural experiment. This approach allows us to isolate the causal effect of GenAI on investor trading and market quality and helps mitigate reverse causality concerns. We compile outage data from OpenAI's official incident reports and analyze changes in trading volume, return volatility, and various

liquidity measures before and during these outage periods.

We focus our empirical tests on cryptocurrency markets because they provide an ideal laboratory for examining our hypotheses. These markets differ fundamentally from traditional equity markets in several dimensions. Unlike stocks, whose prices are anchored in fundamentals such as cash flows, cryptocurrency valuations hinge on more abstract factors like scarcity and user adoption. The high uncertainty surrounding the intrinsic value of cryptocurrencies means that noise plays a prominent role in trading signals generated by models like ChatGPT, making the noise correlation channel more pronounced than in equity markets.

Furthermore, the cryptocurrency market lacks regulations, disclosure, and professional analyst coverage that underpin equity analysis. In this environment, large language models (LLMs) can exert considerable influence by shaping investor beliefs and sentiment and affecting market outcomes. Finally, whereas equity markets are dominated by institutional investors (e.g., banks, insurance companies, mutual funds, and pension funds), cryptocurrency markets are largely driven by retail investors who lack access to sophisticated institutional trading systems. This combination of factors creates a uniquely fragmented information environment, making it an ideal setting in which to examine the trade-off between the precision of trading signals and the correlated noise introduced by widely adopted GenAI tools like ChatGPT.

We assess the impact of ChatGPT outages on investor trading and market quality in cryptocurrency markets by comparing trading volume, return volatility, and liquidity measures during outage periods (the event window) against the same metrics from corresponding time intervals over the preceding five (ten or fifteen) trading days (the control period). Our sample comprises 70 cryptocurrency pairs that include the 10 largest cryptocurrencies by market capitalization, all actively traded on leading exchanges such as Binance, Coinbase, Kraken, and KuCoin. Collectively, the cryptocurrencies in our sample account for about

80% of the total market capitalization of the cryptocurrency market. We also conduct panel regression analyses to control for various confounding factors.

We document that during ChatGPT outages, dollar trading volume increases by 24%, trading frequency rises by 21%, and return volatility increases by 19%. We also find a significant, albeit more modest, deterioration in market liquidity: quoted bid-side (ask-side) book depth declines by 8% (6%), and price impact—a measure of information asymmetry—increases by 12%.

To rule out alternative explanations, we conduct several additional tests. One potential explanation for the increase in trading volume during outages is the comovement of investor attention. ChatGPT outages may redirect investor attention toward cryptocurrencies since they cannot perform their normal job due to the ChatGPT shutdown. We control for investor attention in our baseline regressions, and the results remain unchanged. Another possibility is that investors engage in panic selling in response to technological disruptions such as ChatGPT outages. However, we find that both buy and sell volumes increase during outages, while net order imbalance does not change significantly. This evidence suggests that panic-driven selling is unlikely to be the mechanism underlying our findings.

To provide direct evidence on our core channel—correlated noise—we examine changes in trade imbalance, defined as the absolute value of order imbalance, before and during outages. Both univariate comparisons and regression analyses show that trade imbalance decreases significantly during outages, consistent with the presence of correlated noise. Correlated noise causes investors to trade in the same direction over short intervals, thereby increasing trade imbalance. The observed decline in trade imbalance during outages implies that ChatGPT-induced correlated noise leads to more aligned trading directions.

Our results are robust across a range of tests. Specifically, we (1) employ Poisson regressions to account for integer-valued outcome variables; (2) use median regressions and sample trimming to mitigate the influence of outliers; and (3) incorporate more granular fixed ef-

fects to control for unobserved heterogeneity. We also conduct a placebo test and find no significant effects, alleviating concerns that our results are spurious. In addition, cross-sectional analyses comparing large and small trades reveal no differential impact associated with ChatGPT outages.

We contribute to the emerging literature on the impact of Generative AI in financial markets. To our knowledge, this paper is the first to develop a theoretical framework to analyze the dual impact of ChatGPT on investor trading and market quality. We distinguish between two mechanisms: noise correlation among signals and information production. Our model shows that correlated noise can undermine information production at the aggregate level. This is consistent with the literature of common error (Hassan and Mertens, 2011, 2017). This channel is novel, but difficult to identify empirically because improvements in information production can confound the effects—a challenge our second contribution directly addresses.

Our study is motivated by recent literature that uses outages or malfunctions of information services as exogenous shocks to identify the causal effects of technology on financial markets. For instance, Shkilko and Sokolov (2020) exploit precipitation-induced disruptions in microwave networks to show that speed advantages exacerbate adverse selection and reduce liquidity. Eaton, Green, Roseman, and Wu (2022) document that liquidity improves during outages at Robinhood but deteriorates during disruptions at traditional retail brokers. Most directly related to our work, Cheng, Lin, and Zhao (2025) use ChatGPT outages to provide early evidence on whether investors rely on GenAI to perform professional tasks and its impact on stock price informativeness. They document a significant decline in stock trading volume, short-run price impact, and return volatility during the outage periods for stocks traded on NYSE, AMEX, and Nasdaq, while Bertomeu, Lin, Liu, and Ni (2025) leverage the Italian government's temporary ChatGPT ban to show that its adoption increases analyst coverage and enhances liquidity. Our paper contributes to this stream of research by

providing evidence on the impact of GenAI in the cryptocurrency market—a fast-growing yet under-explored market fundamentally different from the stock market.

Finally, our findings contribute to the policy debate on the benefits and risks of GenAI adoption in financial markets. We document that in cryptocurrency markets, ChatGPT outages lead to a significant increase in trading volume and volatility, and a deterioration in liquidity. This result stands in sharp contrast to evidence from the stock markets, where outages are associated with the opposite effects—reduced volume, lower volatility, and improved liquidity. This dichotomy is critical for regulators and policymakers to consider when evaluating the systemic implications of AI adoption across different market ecosystems.

This paper is organized as follows. Section 2 discusses the institutional background; Section 3 presents our model and empirical predictions; Section 4 describes the data; Section 5 presents the main empirical results; and Section 6 provides results of additional analyses and robustness checks. Concluding remarks are provided in Section 7.

## 2. Institutional Background

ChatGPT, a large language model released by OpenAI in November 2022, is built on the Generative Pre-trained Transformer (GPT) architecture. It generates human-like, conversational text based on user prompts. Grounded in the transformer architecture of Vaswani (2017), GPT models excel at financial tasks like summarizing statements, generating insights, and designing strategies. The premium version relaxes upload limits for diverse data formats (e.g., documents and images) and grants access to more powerful models.

OpenAI's platform allows developers to create tailored GPT-based services for diverse applications. Official statistics from August 2025 indicate that over 15 customized Crypto GPTs had been deployed, facilitating nearly 2 million cryptocurrency-focused consultations. Applications like the popular *Finance & Economics-Stock, Crypto, Trade, Invest* provide

data-driven financial advice, while platforms such as FalconX and Pionex enable brokers and investors to generate personalized investment ideas based on historical activity, portfolio composition, and risk preferences.

FalconX, a leading cryptocurrency prime broker, has developed a chatbot powered by ChatGPT to provide investors with market insights, investment ideas, and trading strategies. The tool aims to democratize access to advanced trading by offering personalized services. Sophisticated cryptocurrency investors—typically large, well-resourced institutions (such as hedge funds)—rely on proprietary algorithmic models and dedicated teams to execute complex strategies like market-neutral and quantitative long/short trades. FalconX's chatbot seeks to level the playing field by providing retail, less sophisticated investors with analytical tools comparable to those used by major quantitative firms.[1]

In March 2023, the trading platform Pionex launched PionexGPT, a feature that integrates its bot service with ChatGPT's analytical capabilities.[2] This interface allows users to solicit feedback on market news or trading ideas. Crucially, it enables the generation of custom trading algorithms based on a user's natural language description, eliminating the need for coding and significantly lowering the technical barrier to sophisticated strategy formulation. These AI-generated strategies can be immediately backtested. Furthermore, Pionex's GPT Marketplace hosts a repository of community-created strategies, allowing users to review historical performance data and engage in copy trading of the most successful ones.

In summary, ChatGPT has emerged as a powerful resource for processing market information and formulating trading strategies across asset classes. Its influence has been further amplified by widely viewed social media tutorials, which highlight its accessibility and rapidly expanding user base.

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

[2]See `https://www.pionex.com/blog/automate-your-trading-ideas-on-pionex-with-pionexgpt-and-tradingview/`

# 3. Token Trading with ChatGPT Outages

We consider a cryptocurrency market with competitive market makers and price-taking traders. Following the literature on token valuation, we assume traders initially purchase tokens for their fundamental utility—either to access platform services or to use as a payment medium (Cong, Li, and Wang, 2021; Sockin and Xiong, 2023; Gryglewicz, Mayer, and Morellec, 2021; Biais, Bisiere, Bouvard, Casamatta, and Menkveld, 2023). Conversely, traders typically sell tokens received as payment in prior transactions to convert them into cash.

## *3.1. Environment*

The market operates across three periods, $t = 0, 1, 2$. In each period, $N$ risk-neutral traders and a representative market maker participate in trading. In $t = 0$, all traders use ChatGPT to generate trading signals. In $t = 1$, an unexpected outage occurs. At this juncture, traders choose between waiting for service restoration or independently generating signals at a fixed cost. In $t = 2$, the ChatGPT service is restored with the probability $\lambda$. Traders discount future payoffs at a constant rate $r$.

In each period $t$, each token provides a utility value $v_t$, which becomes known only at the end of the period. The utility $v_t$ is independently and identically distributed (i.i.d.) across the three periods, with an initial prior distribution $v_0 \sim \mathcal{N}(v_0, \tau_0^{-1})$. At the beginning of each period, traders generate private trading signals about $v_t$ and submit orders to the market maker based on their beliefs. The market maker sets prices and clears the market based on the total order flow. At the end of each period, traders use the tokens to purchase platform services or as a medium of exchange, thereby realizing the utility $v_t$.

### *3.2. Information Production*

With support from ChatGPT, trader $i$ can access and synthesize $\overline{n}$ distinct information sources—ranging from news articles to social media. Each source, after being processed and condensed, yields a single data point, denoted $\{s_{i,t,j}\}_{j=1}^{\overline{n}}$, where $i$ denotes the $i$th trader ($i$ = 1, 2, ....., N) , $j$ is the $j$th information source ($j$ = 1, 2, ....., $\overline{n}$ ), and $t$ is date t. Each data point acts as a noisy but informative signal about the underlying utility value $v_t$:

$$s_{i,t,j} = v_t + \frac{1}{\sqrt{\tau_s}} e_{i,t,j}, \tag{1}$$

where the noise term $e_{i,t,j}$ follows the standard normal distribution and is independent across information sources $j$ and time $t$, but not necessarily across traders (this issue will be elaborated in detail later). $\tau_s$ denotes the precision of a single data point. Given the collection $\{s_{i,t,j}\}_{j=1}^{\overline{n}}$, trader $i$ updates her belief about $v_t$ through Bayesian updating.

**Noise correlation.** While ChatGPT introduces some randomness into its outputs, the same model produces highly consistent answers to identical queries. This suggests a high degree of correlation in the noise terms of the resulting signals. To capture this feature, we assume that, given time $t$ and information source $j$, the noise terms for any two distinct traders $i_1$ and $i_2$ are correlated as follows:

$$\rho = \text{Corr}(e_{i_1,t,j},\, e_{i_2,t,j}) \tag{2}$$

where $\rho \in (0, 1)$. As shown in Equation (2), when ChatGPT refines information source $j$ and generates data points, any deviation from the true value tends to be persistent across traders. If ChatGPT misinterprets information source $j$, similar errors are likely to appear in the data points generated for many traders. This mechanism is analogous to the "common error" problem discussed in Hassan and Mertens (2011) and Hassan and Mertens (2017). In

the extreme case where $\rho = 1$, all traders receive identical signals, implying their noise terms are perfectly correlated.[3]

According to standard Bayesian updating rules, the posterior distribution $v_t \mid \{s_{i,t,j}\}_{j=1}^{\overline{n}}$ depends only on the sample mean $\overline{s}_{i,t} = \frac{1}{\overline{n}} \sum_{j=1}^{\overline{n}} s_{i,t,j}$. For notational simplicity, we define this aggregate signal as $\overline{s}_{i,t}$ for trader $i$ at time $t$:

$$\overline{s}_{i,t} = v_t + \frac{1}{\sqrt{\tau_s}} \overline{e}_{i,t}, \tag{3}$$

where $\overline{e}_{i,t} = \frac{1}{\overline{n}} \sum_{j=1}^{\overline{n}} e_{i,t,j}$. Thus, for trader $i$, observing the full set $\{s_{i,t,j}\}_{j=1}^{\overline{n}}$ is equivalent to observing only the aggregate signal $\overline{s}_{i,t}$.

The correlation in Equation (2) also applies to the aggregate signals $\overline{s}_{i,t}$. It can be verified that the noise term $\frac{1}{\sqrt{\tau_s}}\overline{e}_{i,t}$ has variance $\frac{1}{\overline{n}\tau_s}$, and that for different traders, the noise terms in their aggregate signals also have a correlation coefficient of $\rho$. i.e.,

$$\rho = \text{Corr}\left(\frac{1}{\sqrt{\tau_s}}\overline{e}_{i_1,t}, \frac{1}{\sqrt{\tau_s}}\overline{e}_{i_2,t}\right).$$

Our empirical strategy involves comparing market outcomes with and without access to ChatGPT. We therefore assume an alternative information production technology, which we term independent data generation, exists in the absence of ChatGPT. Under this method, each trader can generate only $\underline{n}$ data points, denoted by $\{s_{i,t,j}\}_{j=1}^{\underline{n}}$, where $\underline{n} < \overline{n}$, reflecting a lower information processing capacity. Because traders conduct research independently, the noise terms in their data points are less correlated. For simplicity, we assume that, in the absence of ChatGPT, for any two different traders $i_1$ and $i_2$, the correlation between their

[3]In reality, noise correlation can vary across investor pairs. We interpret the parameter in Equation (2) as representing the average correlation of the noise terms.

noise terms for a given information source $j$ at time $t$ is zero:

$$\text{Corr}(e_{i_1,t,j},\, e_{i_2,t,j}) = 0. \tag{4}$$

To distinguish this from the case with ChatGPT, we define the aggregate trading signal for trader $i$ as

$$\underline{s}_{i,t} := \frac{1}{\underline{n}} \sum_{j=1}^{\underline{n}} s_{i,t,j},$$

which can be written as $\underline{s}_{i,t} = v_t + \frac{1}{\sqrt{\tau_s}}\, \underline{e}_{i,t}$, where the aggregate noise term is $\underline{e}_{i,t} := \frac{1}{\underline{n}} \sum_{j=1}^{\underline{n}} e_{i,t,j}$. For any two different traders $i_1$ and $i_2$, the aggregate noise terms $\underline{e}_{i_1,t}$ and $\underline{e}_{i_2,t}$ are not correlated. The aggregate noise in the signal, $\frac{1}{\sqrt{\tau_s}}\, \underline{e}_{i,t}$, has a precision of $\underline{n}\tau_s$, which is lower than in the ChatGPT case. Therefore, the absence of ChatGPT creates a clear trade-off: it reduces the average precision of traders' signals but also eliminates the correlation of noise across traders.

### *3.3. Trading Decisions*

We assume tokens materialize their utility and are liquidated in each period, rendering the trading optimization problem static rather than intertemporal. This simplification implies that the optimal number of tokens traded in a given period depends solely on the current period's signals, without influence from past trades or considerations for future periods. Traders maximize their expected trading profits by choosing an order size $x$, conditional on their own trading signals. Under the "agree to disagree" paradigm, trader $i$ in period $t$ solves the following problem:

$$\max_{x_{i,t}} \mathbb{E}\left[ x_{i,t}(v_t - p_t) - \frac{\phi}{2} x_{i,t}^2 \,\middle|\, \overline{s}_{i,t} \right], \tag{5}$$

where $\phi > 0$ represents the transaction cost. The convexity of the cost function ensures that orders from risk-neutral traders remain finite, which is necessary in a setting where traders do not internalize their price impact. This cost structure is empirically grounded, as exchanges commonly charge transaction fees through various fee schedules (Colliard and Foucault, 2012).

The first-order condition for the optimization problem in equation (5) yields the optimal order size as a function of private signal and price:

$$x^*_{i,t}(\overline{s}_{i,t}, p_t) = \frac{1}{\phi}\left(\frac{\tau_0}{\tau_0 + \overline{n}\tau_s} v_0 + \frac{\overline{n}\tau_s}{\tau_0 + \overline{n}\tau_s}\overline{s}_{i,t} - p_t\right) \tag{6}$$

The second-order condition, $-\phi < 0$, is satisfied, indicating that $x^*_{i,t}(\overline{s}_{i,t}, p_t)$ is the unique optimal solution of the profit maximization problem.

The optimal strategy has an intuitive interpretation: the trader's demand is proportional to the difference between their posterior expectation of the value, $\frac{\tau_0}{\tau_0+\overline{n}\tau_s} v_0 + \frac{\overline{n}\tau_s}{\tau_0+\overline{n}\tau_s}\overline{s}_{i,t}$, and the market price, with the sensitivity inversely weighted by the transaction cost.

### *3.4. Equilibrium*

Following Kyle (1985), we assume that the market-making industry is competitive. The representative market maker observes at time $t$ the total order flow $y_t$ ( $= \sum_{i=1}^{N} x_{i,t}$) from $N$ traders, and sets the price according to a zero-profit condition:

$$P(y_t) = E[v_t | y_t] \tag{7}$$

The equilibrium consists of a set of trading strategies, $\{x^*_{i,t}(\overline{s}_{i,t}, p_t)\}_{i=1}^{N}$, where each $x^*_{i,t}$ solves the optimization problem in (5). The market price is determined by $p_t = E[v_t \mid y_t = \sum_{i=1}^{N} x^*_{i,t}]$. Observing the total order flow, the market maker updates her belief about the true

value of $v_t$ by inferring traders' private signals.

In equilibrium, each trader submits an optimal order given by (6). Therefore, the market maker can infer the average of all traders' aggregate signals from the total order flow, that is, $\tilde{s}_t := \frac{1}{N}\sum_{i=1}^{N} \overline{s}_{i,t}$. Consequently, the equilibrium price reflects all the information in the market; specifically, $p_t = E[v_t \mid \tilde{s}_t]$, which is equivalent to $E[v_t \mid \{s_{i,t,j}\}_{i,j}]$.

**Proposition 1.** *In the equilibrium with (resp. without) ChatGPT, the token price is given by $E[v_t \mid \tilde{s}_t]$, where $\tilde{s}_t := \frac{1}{N}\sum_{i=1}^{N} \overline{s}_{i,t}$ (resp. $\tilde{s}_t := \frac{1}{N}\sum_{i=1}^{N} \underline{s}_{i,t}$). For any trader $i$, the optimal order submitted is*

$$\frac{1}{\phi}\left(E[v_t \mid \overline{s}_{i,t}] - E[v_t \mid \tilde{s}_t]\right) \quad \textit{(resp. } \frac{1}{\phi}\left(E[v_t \mid \underline{s}_{i,t}] - E[v_t \mid \tilde{s}_t]\right)\textit{ ).}$$

### *3.4.1. Expected Trading Profits*

Given the optimal order in (6), the expected trading profits for trader $i$ are given by: $\frac{1}{2\phi}\mathbb{E}\left[\left(\mathbb{E}[v_t \mid \overline{s}_{i,t}] - \mathbb{E}[v_t \mid \tilde{s}_t]\right)^2\right]$. Accordingly, before the trading signals are realized, the ex-ante expected trading profits are:

$$\overline{\pi} := \mathbb{E}\left[\frac{1}{2\phi}\mathbb{E}\left[\left(\mathbb{E}[v_t \mid \overline{s}_{i,t}] - \mathbb{E}[v_t \mid \tilde{s}_t]\right)^2\right]\right] = \frac{1}{2\phi}\frac{1}{\tau_0^2}\left[\frac{N\overline{n}\tau_0\tau_s}{\tau_0[1+\rho(N-1)] + \overline{n}\tau_s N} - \frac{\overline{n}\tau_0\tau_s}{\tau_0 + \overline{n}\tau_s}\right], \tag{8}$$

Similarly, we define the ex-ante expected trading profit in the absence of ChatGPT as:

$$\underline{\pi} := \frac{1}{2\phi}\frac{1}{\tau_0^2}\left[\frac{N\underline{n}\tau_0\tau_s}{\tau_0 + \underline{n}\tau_s N} - \frac{\underline{n}\tau_0\tau_s}{\tau_0 + \underline{n}\tau_s}\right], \tag{9}$$

We omit the time index because the expectations are constant across periods. We assume that ChatGPT enhances the ex-ante expected trading profits, i.e., $\overline{\pi} > \underline{\pi}$. As a result, when ChatGPT is available, traders will always utilize it for information production.

**Assumption 1.** *ChatGPT enhances the ex-ante expected trading profits, i.e.,* $\overline{\pi} > \underline{\pi}$*, and when ChatGPT is available, traders prefer to use it.*

### *3.5. ChatGPT Outage*

We model the outage of ChatGPT as an unexpected loss of accessibility at the beginning of period $t = 1$, causing the process of signal generation to halt. Traders must decide how to produce information: one option is to wait for the outage to be resolved; the other is to incur a switching cost, $C$, and set up teams and equipment for independent research.

No information acquisition takes place in period 1 if traders decide to wait. Should ChatGPT restore its service in period 2, information acquisition resumes; otherwise, it remains halted. If traders choose to switch, they acquire information independently (without ChatGPT) in period 1. In period 2, if the ChatGPT service is restored, they revert to ChatGPT-assisted trading, as per Assumption 1. The main tradeoff is between incurring the fixed switching cost and the possibility that ChatGPT will not be restored in period 2. Proposition 2 provides the condition under which traders opt for independent research rather than waiting.

**Proposition 2.** *If* $C < \frac{2+r-\lambda}{1+r}\underline{\pi}$*, then all traders will switch and conduct independent research in period* 1*.*

Proposition 2 indicates that if the switching cost is relatively small compared to the potential trading profits—including the profits foregone in period 1 and in period 2 if ChatGPT service is not restored—the market will reach an equilibrium in which traders conduct independent research.

### *3.6. Empirical Predictions*

We test the model's predictions by comparing changes in trading volume, return volatility, and market depth between periods 0 and 1. Following Wang (1994), we define trading volume as the expected absolute value of orders from all traders in equilibrium, that is, $E\left[\sum_{i=1}^{N}|x_{i,t}^{*}|\right]$. We measure return volatility as the variance of the price gap between the equilibrium price and the prior price. We define returns in this way to capture price fluctuations driven by informed trading and to maintain a consistent measure across all periods. Here, the prior price is given by the mean of the prior distribution of $v_t$, namely, $v_0$. For market depth, we adopt the approach of Kyle (1985) and use the inverse of the price impact. Under the normal distribution assumption, $P_t = E[v_t|y_t]$ has a linear form: $P_t = \mu_t + \eta_t y_t$, where $\eta_t$ represents the price impact coefficient. Market depth is therefore defined as $\frac{1}{\eta_t}$.

Proposition 3 presents the closed-form expressions for these three variables in equilibrium in period 0 and 1. A direct comparison of these solutions yields our main empirical predictions.

**Proposition 3.** *If $C < \frac{2+r-\lambda}{1+r}\underline{\pi}$, the closed-form solutions for trading volume, return volatility, and market depth in equilibrium, in periods 0 and 1, are summarized as follows:*

| **Metric** | **With ChatGPT (Period 0)** | **Without ChatGPT (Period 1)** |
|---|---|---|
| Volume | $\frac{N}{\phi\tau_0}\sqrt{\frac{2}{\pi}}\left[\frac{N\overline{n}\tau_0\tau_s}{\tau_0[1+\rho(N-1)]+\overline{n}\tau_s N}-\frac{\overline{n}\tau_0\tau_s}{\tau_0+\overline{n}\tau_s}\right]^{1/2}$ | $\frac{N}{\phi\tau_0}\sqrt{\frac{2}{\pi}}\left[\frac{N\underline{n}\tau_0\tau_s}{\tau_0+\underline{n}\tau_s N}-\frac{\underline{n}\tau_0\tau_s}{\tau_0+\underline{n}\tau_s}\right]^{1/2}$ |
| Volatility | $\frac{1}{\tau_0}\frac{\overline{n}N\tau_s}{\overline{n}N\tau_s+\tau_0[1+\rho(N-1)]}$ | $\frac{1}{\tau_0}\frac{\underline{n}N\tau_s}{\underline{n}N\tau_s+\tau_0}$ |
| Depth | $\frac{1}{\phi}\frac{\overline{n}\tau_s}{\tau_0+\overline{n}\tau_s}\left\{N+\frac{\tau_0}{\overline{n}\tau_s}[1+\rho(N-1)]\right\}$ | $\frac{1}{\phi}\frac{\underline{n}\tau_s}{\tau_0+\underline{n}\tau_s}\left(N+\frac{\tau_0}{\underline{n}\tau_s}\right)$ |

ChatGPT affects the market through two main channels. First, it enhances information processing capacity and enables traders to generate more data points, i.e., $\overline{n} > \underline{n}$. With greater data availability, traders can form more precise expectations about $v_t$, leading to

increased informativeness. This, in turn, intensifies adverse selection between traders and the market maker.

The second channel operates through noise correlation. By assisting traders in processing information, testing strategies, and generating signals, ChatGPT homogenizes the signals they rely on, leading to more homogeneous beliefs. For example, if multiple traders use ChatGPT to design a momentum strategy—a well-documented and widely followed approach in cryptocurrency markets—they are likely to receive similar recommendations. As shown in Proposition 1, market makers' beliefs equal the average belief of traders, which coincides with the equilibrium price. Consequently, greater investor homogeneity reduces the difference between their beliefs and those of the market maker. This reduction in disagreement—relative to a setting without ChatGPT—lowers trading volume, consistent with the mechanism in Scheinkman and Xiong (2003). A temporary ChatGPT outage forces traders to rely on independent research. In practice, because ChatGPT provides similar outputs for similar queries, this interruption fosters greater disagreement and thus increased trading activity.

At the same time, traders' order flows reflect both fundamental information and noise. When noise components are correlated—as occurs when traders rely on similar LLM-generated signals—the aggregate order flow becomes less informative, as the common noise fails to cancel out. This reduction in informativeness mitigates the adverse selection between traders and the market maker. This mechanism, in which common errors crowd out independent information production, aligns with the common-error literature (Hassan and Mertens, 2011). Specifically, correlated orders create camouflage, obscuring the fundamental information in the order flow. Thus, the price impact of trades declines and market liquidity, as measured by depth, increases.

With these two competing channels, the effects on market metrics depend both on the efficacy of ChatGPT in facilitating information production—which amplifies the first channel—

and the homogeneity of ChatGPT's responses, which reinforces the second channel. Corollary 1 provides a prediction under sufficient conditions, tailored to the empirical setting considered next.

**Corollary 1.** *Suppose the model parameters satisfy the regularity conditions:* $N > \frac{\overline{n}}{\underline{n}} + 1$ *and* $\rho \in \left(\frac{\overline{n}-\underline{n}}{\underline{n}(N-1)}, \frac{\overline{n}-\underline{n}}{\underline{n}}\right)$. *If* $\tau_s$ *and* $C$ *are sufficiently small, compared to period 0 when traders use ChatGPT, trading volume increases, return volatility increases, and market depth decreases during period 1 with outages:*

Corollary 1 characterizes the effects of a ChatGPT outage as a deviation from the baseline equilibrium where ChatGPT is widely adopted. The regularity conditions in Corollary 1 are fairly weak when $N$ is large: as $N \to \infty$, $\frac{\overline{n}-\underline{n}}{\underline{n}(N-1)} \to 0$, while if $\overline{n} \gg \underline{n}$, $\frac{\overline{n}-\underline{n}}{\underline{n}} \to 1$. A small value of $C$ representing low switching costs for traders is a realistic assumption in the retail-dominated cryptocurrency markets. Furthermore, a small $\tau_s$ indicates that each data point carries limited information. This captures a salient feature of crypto markets: the absence of regulated disclosures and unclear fundamentals contribute to generally low information intensity. Thus, even if ChatGPT increases the number of generated data points, the overall improvement in information production remains slight. In this environment, the noise correlation channel emerges as the dominant force shaping market outcomes.

We can interpret these predictions regarding trading volume through the lens of investor disagreement. In equilibrium, a trader's optimal order is proportional to $E[v_t|\overline{s}_{i,t}] - E[v_t|\tilde{s}_t]$, which can be interpreted as the disagreement between the trader and the market maker. Since the market maker aggregates all trading signals, this gap also reflects the average level of disagreement in the market. The noise correlation channel reduces this disagreement, resulting in lower trading volume as market participants increasingly agree on the value of the tokens. With fewer adverse selection issues, price impact declines, market depth increases, and return volatility decreases.

Our model yields an important policy implication: while ChatGPT provides valuable tools for generating trading signals, it also reduces the independence of these signals, leading to common errors. This crowd-out effect reduces information asymmetry and improves market liquidity, but has a negative effect on the aggregate informativeness of prices.

Our framework also generates predictions for scenarios involving high investor switching costs. Institutional investors, for example, face significant fixed costs when a deeply integrated tool like ChatGPT becomes unavailable. These costs include the coordination costs of establishing new research teams. Furthermore, large existing market positions and internal risk budgets can limit their flexibility during outages. Under high switching costs, traders optimally choose to wait rather than generate new signals independently. Consequently, both trading volume and return volatility decline due to the absence of new information and the resulting lack of investor disagreement. In the absence of new signals, any residual order flow is perceived as uninformed noise trading, leading to negligible price impact and a sharp increase in measured market depth. This prediction aligns with the empirical findings of Cheng, Lin, and Zhao (2025), who document reduced trading activity and volatility during ChatGPT outages in equity markets.

# 4. Data and Sample Selection

This study utilizes data from two sources. Data on ChatGPT outage events is obtained from OpenAI's public incident reports. These reports, published whenever a service malfunction occurs, provide the date, severity, and a detailed chronology of each disruption, including its timeline, scope, and root cause. For market data, intraday order book and transaction data for cryptocurrencies are obtained from the Kaiko database.

We manually collected all incident reports related to ChatGPT from its launch on November 30, 2022, through August 31, 2024, to identify service outages. Our selection process

begins with incidents marked as affecting ChatGPT and lasting at least 10 minutes. We then refine this set by focusing on incidents whose titles explicitly indicated a service outage, defined as a period when ChatGPT was functionally inaccessible. We review each incident report to ensure that every outage event is included.

For each qualifying incident, we record the start time, the time when the malfunction began investigation, and the time when the issue was fully resolved. This process yields a final sample of 22 ChatGPT outage events. As summarized in Table A1, the first event occurred on February 21, 2023, and the latest on July 17, 2024. Outage durations range from 14 minutes to 10 hours and 13 minutes, with a mean duration of 1.9 hours and a median of 1 hour 6 minutes.

The intraday cryptocurrency data includes tick-level trade data and full-depth order book snapshots. For each trade, we observe a unique transaction ID, a Unix millisecond timestamp, price, volume, and an indicator for whether the trade was buyer- or seller-initiated. The order book snapshots, recorded every 30 seconds, provide a timestamp, bid and ask prices, and the corresponding quantities (depths). To construct our main variables, we aggregate the high-frequency data into non-overlapping 5-minute intervals.

We measure trading activity using two metrics per coin pair and interval: trading volume (cumulative dollar volume) and trading frequency (number of trades). For liquidity, we calculate time-interval averages for bid-side book depth, ask-side book depth, quoted spread, and five-minute price impact. Following market microstructure convention, we define bid-side book depth as the cumulative dollar value available in the limit order book if the bid price decreases by 0.1% relative to the current best bid price. Similarly, ask-side book depth is the cumulative dollar value available if the ask price increases by 0.1% relative to the current best ask price. Returns are computed as the price difference between the start and end of each interval, calculated separately from trade and quote data. Volatility is measured as the standard deviation of higher-frequency (30-second) returns—calculated from both

trades and quotes—within each 5-minute interval.

Our sample is constructed in two stages. We identify the top 10 largest non-stable cryptocurrencies by market capitalization as of November 30, 2022, coinciding with 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). Then, we select eight major exchanges, including—Binance, Bitstamp, Coinbase, Gemini, Huobi, Kraken, KuCoin, and OkEX, based on their consistent trading of popular cryptocurrency pairs throughout our sample period.

For each exchange and coin, we select the primary trading pair based on the highest trading volume in November 2022. For instance, on Coinbase, the BTC/USD pair was selected over BTC/USDT due to its significantly higher trading volume. When the denominated currency was Tether (USDT), we converted trading volumes to U.S. dollars using the prevailing exchange rate, which fluctuated between \$0.998 and \$1.007 during the sample period.

Our final sample comprises 70 unique coin pairs. Table A2 shows that all ten top cryptocurrencies are traded on Binance, Huobi, KuCoin, and OkEX. The remaining exchanges in our sample list include six to nine of these top coins. Table 1 presents cross-sectional summary statistics for our variables. All variables are winsorized at the 1st and 99th percentiles. The results reveal substantial dispersion across the 70 coin pairs. For instance, the average 5-minute trading frequency is 530 trades (S.D. = 1,807), and the average dollar volume is \$368,740 (S.D. = \$1,568,770). The average quoted spread is \$0.17. Finally, the quotation depth available in the limit order book if the bid (ask) price decreases (increases) by 0.1% relative to the current best bid (ask) price is \$230,830 (\$220,420).

# 5. Empirical Results: ChatGPT's outages

This section empirically tests our hypotheses by examining the impact of ChatGPT outages on trading activity and market quality. Toward this goal, we present results from two approaches: a uniform aggregation procedure and a panel regression that controls for potential confounding factors. Both methods assess changes in trading activity, return volatility, and various liquidity measures from control to outage periods.

Our first test employs a uniform aggregation procedure. For each variable and coin pair, we calculate the ratio of its control-period average to its event-period average. We then take the natural logarithm of this ratio to obtain a measure of proportional change, which also mitigates the effects of the right-skewed cross-sectional distribution. We report the cross-sectional means and medians of these log changes and assess their statistical significance using a t-test, a Wilcoxon signed-rank test (robust to non-normality and outliers), and a bootstrap test (which addresses potential non-normality in the sampling distribution of the mean).

To isolate the effect of ChatGPT outages from intraday seasonality, we define the control period for each outage as the same time interval from the five preceding trading days ([-5, -1]). For example, the outage on February 21, 2023 (7:40 a.m.–12:09 p.m.) is compared to the same window on February 16–20. We also test robustness using control periods from the preceding 10 ([-10, -1]) and 15 ([-15, -1]) trading days.

In the second test, we use a panel regression approach to analyze changes in trading activity and market quality measures between benchmark and outage periods. This method enables us to control for potential confounding factors and to include fixed effects. Using trading volume as an example, we estimate the following specification:

$$\log(\text{Trading Volume})_{i,d,t} = \alpha + \beta \times \text{Outage}_{d,t} + \gamma \times \text{Controls}_{i,d-1} + \text{Fixed Effects}_{d,t} + \epsilon_{i,d,t} \quad (10)$$

where Trading Volume$_{i,d,t}$ is the dollar trading volume for cryptocurrency $i$ during the 5-minute interval $t$ on day $d$. The indicator variable Outage$_{d,t}$ equals one if interval $t$ on day $d$ falls within a ChatGPT outage event window, and zero otherwise. Control variables for cryptocurrency $i$ include: the logarithm of the previous day's market capitalization ($\log(\text{MarketCap}_{i,d-1})$), the previous-day return ($\text{Ret}_{i,d-1}$), the average return from days $d-5$ to $d-2$ ($\text{Ret}_{i,[d-2,d-5]}$), and the logarithm of the previous day's closing price ($\log(\text{Price}_{i,d-1})$).

All specifications include coin pair and time-interval fixed effects. Standard errors are two-way clustered at the coin pair and time levels to account for both serial and cross-sectional correlation. This regression specification is applied analogously to other dependent variables, including trading frequency, return volatility, bid-ask spread, bid-side book depth, ask-side book depth, and price impact. All variables are aggregated into 5-minute intervals.

### *5.1. Trading Volume*

Table 2 presents the summary and test statistics for two measures of trading activity, trading frequency and dollar trading volume, during control and ChatGPT outage periods. We find a statistically and economically significant increase in both measures from the control to the outage periods.

Using [-5, -1] as the control period, Panel A of Table 2 shows that the average trading frequency increases from 630 to 765 trades per five-minute interval. The average increase is 21.64% (median = 22.23%) and significant at the 1% level based on the t-test. Similarly, average dollar volume increases from $511,400 to $624,420, a significant 23.76% increase (median = 24.43%) at the 1% level. These results confirm a substantial increase in market activity during periods when ChatGPT is unavailable.

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. Both the Wilcoxon test and the bootstrapped t-test strongly reject the null hypothesis of no change in trading frequency and dollar volume.

Panel A of Table 2 further reports results using two alternative control periods: [-10, -1] and [-15, -1]. For the [-10, -1] window, the average trading frequency shows a significant increase of 17.48% (median = 19.73%), with all three statistical tests confirming significance. The economic significance is further underscored by the increase in dollar trading volume, which rises from $527,680 to $624,420 per 5-minute interval—a significant 19.98% increase.

Figure 2 plots the 5-minute trading volume for a representative exchange and coin pair—the BTCUSDT pair on Binance—comparing control and outage periods. We find that average trading volume increases noticeably during outage events relative to the control period.

We present the panel regression estimates in Panel B of Table 2. Columns (1) and (2) report the results for trading frequency using the [-5, -1] control period. Column (1) includes coin-pair and time-interval fixed effects, while column (2) adds control variables. The coefficients on the Outage indicator are 0.22 and highly significant (t-statistics = 6.75 and 7.27, respectively). Columns (3) and (4) report the results for dollar trading volume, showing similarly positive and significant Outage coefficients of 0.28 (t = 7.97) and 0.28 (t = 8.10). These results remain positive and statistically significant when we use the alternative [-10, -1] and [-15, -1] control periods. Overall, the panel regression analysis confirms the findings in Panel A, indicating a significant increase in both trading frequency and volume during ChatGPT outage periods.

To assess the economic significance of this result, we compare the magnitude of the increase to the inherent variation in trading activity. As shown in Panel A of Table 2, the average dollar volume increases by $113,020 per five-minute interval (from $511,400 to $624,420) during outages relative to the [-5, -1] control period. This increase represents

7.2% of the variable's sample standard deviation ($1.57 million). This magnitude confirms that the increase in trading volume is not only statistically significant but also economically meaningful.

The significant increase in trading volume provides indirect evidence of ChatGPT's widespread adoption in cryptocurrency markets, as no such change would be expected without substantial reliance on ChatGPT. This finding stands in sharp contrast to equity markets, where Cheng et al. (2025) document a significant 8% decrease in volume during ChatGPT outages for NYSE, AMEX, and Nasdaq stocks.

This divergence can be explained by our theoretical framework. The crypto market's reaction aligns with the prediction in Corollary 1, which applies to settings with low switching costs ($C$) and low information intensity (small $\tau_c$), characteristics typical of retail-driven crypto markets. The equity market decline in trading volume is consistent with a high switching cost ($C$) environment, where institutional traders facing significant switching costs optimally choose to wait and trade less. This leads to reduced trading activity. This stark contrast reveals the dual nature of ChatGPT's impact: its effect on investor trading is not uniform but is determined by the characteristics of the market in which it is deployed.

### *5.2. Return Volatility*

This section examines the impact of ChatGPT outages on return volatility. Following the literature, we measure volatility using realized volatility (RV) computed from high-frequency returns, which provides an efficient, unbiased, and model-free estimator of return variation (Andersen, Bollerslev, Diebold, and Ebens, 2001). We calculate volatility separately using quote-based and trade-based return series.

Quote-based volatility is calculated by using quote revisions from the limit order book. For each five-minute interval, we take 30-second snapshots of the limit-order book and record the corresponding bid-ask midpoint price ($P$). 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:

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

Similarly, trade-based volatility is constructed from transaction prices. The volatility for a given event or control period is then calculated as the average of these five-minute realized volatility measures across the period.

Table 3 presents summary statistics for trade-based and quote-based return volatility, along with their logarithmic change (Log diff in %) during control and ChatGPT outage periods. Using the [-5, -1] control window, Panel A shows that trade-based volatility increases by 19.19% from an average of 5.23 bps in the control period to 6.35 bps during outages. This increase is significant at the 1% level, as confirmed by the t-test, Wilcoxon rank-sum test, and bootstrapped p-value. The results remain consistent when using the alternative [-10, -1] and [-15, -1] control periods.

Panel B reports the results for quote-based volatility. We observe a significant increase during outage periods across all control windows ([-5, -1], [-10, -1], and [-15, -1]). Using the [-5, -1] control period. The average quote-based volatility increases by 18.01%, from 5.51 to 6.64 basis points per five-minute interval, while the median increases by 18.73%, from 5.51 to 6.64 basis points. Both increases are statistically significant at the 1% level based on the t-test and Wilcoxon non-parametric test, a finding further confirmed by a bootstrap test (p = 0.00).

Panel C of Table 3 presents the panel regression results. Columns (1) through (6) report estimates for trade-based volatility, while columns (7) through (12) for quote-based volatility. For trade-based volatility, the coefficients on the Outage dummy are 0.184 (t = 7.74) and 0.192 (t = 7.78) in the specifications without and with control variables, respectively. Both

specifications include fixed effects. We note that both coefficients are highly significant, and their magnitude is consistent with the average percentage change reported in Panel A. These results remain positive and significant when using the alternative [-10, -1] and [-15, -1] control periods. For quote-based volatility, the results in columns (7)-(12) show a consistent and significant increase of between 16% and 19% during outage periods. Collectively, the panel regression analysis reaffirms the findings in Panels A and B and provides evidence of a significant increase in both trade-based and quote-based volatility during ChatGPT outages.

As shown in Section 3, ChatGPT influences investor trading and liquidity through two channels: the noise-correlation channel and the information channel. The former predicts that market makers face less adverse selection from ChatGPT-assisted traders and that they are willing to absorb orders without making price adjustments. This will reduce price impact and volatility. The information channel suggests that improved signal precision intensifies adverse selection. The significant increase in volatility observed during outages provides empirical support for the dominance of the noise-correlation channel in cryptocurrency markets. According to the closed-form solution in Proposition 3, volatility reflects the informativeness of order flow. The rise in volatility during outages therefore indicates a greater degree of informed trading. This result indicates that ChatGPT adoption under normal conditions (non-outage periods) reduces informed trading and dampens volatility. In the following sections, we extend this analysis by examining the impact of outages on price impact and its inverse, market depth.

### *5.3. Market Depth*

Kyle (1985) proposes that the depth of the market, or the amount of cryptocurrencies required to move the price by one unit, corresponds to the inverse of the price impact. Depth is a measure of liquidity; it captures the sensitivity of prices to order flow and reflects the degree of information asymmetry. This measure shares many characteristics with other well-known

liquidity measures (Amihud, 2002).

The Kaiko database provides 30-second snapshots of the limit order book for each coin pair traded on each exchange. Since depth in a limit order book is quoted in the number of coins, we convert it to a dollar value to facilitate comparisons across coin pairs.

Our 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 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 dollar value of coins available if the ask price increases by 0.1% relative to the current best ask. To put book depth in perspective, consider Bitcoin trading at the best bid price of $100,000. The book 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 $99,900. It is commonly referred to as bid-side book depth or ask-side book depth (Bacidore, Battalio, and Jennings, 2003).

Table 4 reports changes in bid-side and ask-side book depths. Using the [-5, -1] control period, Panel A shows that the average bid-side book depth decreases by 6.0%, falling from $256,790 to $247,280 per five-minute interval. This reduction is statistically significant based on both the t-test and bootstrap test, with a median decline of 5.28% confirmed by the Wilcoxon test. These results are robust across the alternative [-10, -1] and [-15, -1] control periods. The ask-side depth also contracts, though to a lesser extent than the bid side. For the [-5, -1] control period, the average decrease is approximately 3.0%, with a median reduction of 4.22%. This decline is statistically significant across all three tests. The results remain consistent with the [-15, -1] control period, though they are weaker when using the [-10, -1] window.

Panel B of Table 4 reports the panel regression estimates for bid-side book depth. The

coefficients on the Outage indicator are -0.066 (t = -3.74) and -0.084 (t = -5.12) in the specifications without and with control variables, respectively. These estimates are economically consistent with the approximately -6% univariate change reported in Panel A. The results remain statistically significant when using the alternative [-10, -1] and [-15, -1] control periods. Panel C presents the corresponding results for ask-side book depth. For the [-5, -1] control period, the Outage coefficients are -0.055 (t = -3.65) and -0.073 (t = -5.41). These estimates suggest a slightly larger effect than the 3% univariate decrease shown in Panel A. Furthermore, columns (4) through (6) confirm a statistically significant reduction in ask-side depth for the [-10, -1] and [-15, -1] control periods.

### *5.4. Bid-Ask Spreads*

In this section, we assess the impact of the ChatGPT outage on another liquidity measure: the bid-ask spread. As a direct reflection of immediate transaction costs for a round-trip trade (Brauneis, Mestel, Riordan, and Theissen, 2021), the bid-ask spread provides a real-time indicator of market liquidity. For our analysis, we convert the spreads for all coin pairs into US dollars to facilitate a standardized comparison across cryptocurrencies. A wide spread signals illiquidity, often resulting from high volatility, or informational asymmetry, which forces market makers to charge a higher premium for the risk of providing liquidity. By examining changes in the bid-ask spread, we can quantify the increase in trading frictions, together with the erosion of market depth discussed in the previous section, directly attributable to the ChatGPT outage events.

The results in Panel A of Table 5 provide evidence that the ChatGPT outage had a significant adverse effect on market liquidity, as measured by the bid-ask spread. The analysis reveals a statistically significant increase in spreads during the outage window across all three control period specifications. For example, the test results using the [-5, -1] control period indicate a significant 4.31% widening of the spread. All three tests—the t-test, the 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 2.11% and 5.28%, respectively.

The panel regression results in Panel B of Table 5 are consistent with findings in Panel A. For the [-5, -1] control period, the coefficients on the Outage indicator are positive and statistically significant in both model specifications—0.063 (t-statistic = 4.92) without control variables and 0.065 (t-statistic = 4.25) with them. These coefficients are slightly larger than the 4.31% average increase reported in Panel A. Thus, after controlling for potential confounding factors and fixed effects, we find that the outage itself had a detrimental impact on liquidity.

In summary, our analysis reveals a significant increase in immediate trading cost during the ChatGPT outage period, with bid-ask spreads rising between 4% and 5%. 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 book depths, deteriorates during the ChatGPT outage period.

### *5.5. Price Impact*

Price impact is a liquidity measure that quantifies the price sensitivity to a trade, specifically the extent to which executing a large order moves the price against the trader. Unlike bid-ask spreads, which reflect the cost of a small, instantaneous trade, price impact captures the hidden cost of liquidity demand for larger volumes. When a trader places a substantial market order to buy, they consume the available limit orders at the best ask and then at progressively worse prices; the opposite occurs for a large sell order. This adverse price movement represents a significant execution cost, as the average execution price is less favorable than the initial quoted price. Thus, true liquidity is not just about the ability to trade immediately, but the ability to do so without materially affecting the market price.

We measure price impact as follows. For a given cryptocurrency $i$ and a trade executed at timestamp $t_s$, let $Midpoint_{t_s}$ denote the prevailing bid-ask midpoint immediately before the trade, and $Midpoint_{t_s+5min}$ the midpoint five minutes after the trade. The trade direction indicator D equals 1 if it is a purchase and -1 if it is a sale. The price impact for a single trade is then calculated as:

$$PriceImpact_{i,t_s} = 2 \times D \times \frac{Midpoint_{t_{s+5minute}} - Midpoint_{t_s}}{Midpoint_{t_s}} \tag{12}$$

This measures the price move (in percent) five minutes after a trade. For a given 5-minute interval t, we calculate the aggregate price impact as the equally weighted average of the absolute values of all individual price impacts, $|PriceImpact_{t_s}|$, for trades within that interval.

Table 6 presents the results for price impact. Panel A provides summary statistics and shows a significant increase in price impact during ChatGPT outages. Using the [-5, -1] control period, price impact rises by 12.25%, from 12.41 to 13.96 basis points. This increase is statistically significant across all three tests (t-test, Wilcoxon, and bootstrap). The results are robust, with increases of 10.61% and 12.65% for the [-10, -1] and [-15, -1] control periods, respectively. Panel B reports findings with panel regressions. The coefficients on the Outage dummy are 0.144 (t-statistic = 4.82) and 0.147 (t-statistic = 5.00) in the specifications without and with controls, respectively. These estimates are highly significant and consistent with the 12.25% increase observed in the univariate tests.

Our model posits that ChatGPT affects investor trading and market quality through two channels: the information production channel and the noise-correlation channel. Price impact—the sensitivity of prices to order flow—directly reflects the market maker's assessment of the informational content embedded in trades. Under the information channel, ChatGPT adoption enhances traders' information-processing capacity and increases the informational content of order flow regarding fundamentals. This leads market makers to

update prices more aggressively, resulting in higher price impact. On the other hand, under the noise-correlation channel, ChatGPT introduces correlated noise across traders' signals, reducing the overall informativeness of order flow. Perceiving a greater proportion of trading as non-informative, market makers are willing to absorb more orders without significant price adjustments, The significant increase in price impact we document during ChatGPT outages provides empirical evidence that the noise-correlation channel dominates in cryptocurrency markets. The removal of ChatGPT reduces correlated noise, thereby increasing the informational content of order flow and raising price impact.

# 6. Additional Tests

## *6.1. Define Outage Windows using Data from StatusGator*

In our primary tests, we rely on official OpenAI records to identify ChatGPT outages. A potential concern is that these records may not perfectly align with the exact times users experience service disruptions. To address this concern, we validate our outage events using data from StatusGator, a third-party monitoring service. StatusGator aggregates data from proprietary probes and user reports to track service status, often detecting outages before official provider acknowledgment.

Specifically, we cross-reference outage events between the official OpenAI records and StatusGator's enterprise database. For events recorded by both sources, the discrepancy in start times is consistently within a five-minute window. This confirms the reliability of our primary data source. For these overlapping events, we redefine the outage window using StatusGator's timestamps and re-estimate all panel regression models in Tables 2-6. The results, reported in Table 7, are very close to our baseline findings. The coefficients on the Outage indicator across all specifications are quantitatively similar to those in our main

tables and retain their statistical significance. We conclude that our main findings are not sensitive to the specific source of the outage data.

### *6.2. Alternative Mechanism: Investor Attentions*

We have documented a significant increase in trading activity during ChatGPT outages. An alternative explanation for this result is **attention-driven trading volume.** According to this hypothesis, investor attention may co-move with outage events: when access to generative AI tools is disrupted, investors may be unable to continue their regular tasks and instead divert their attention to trading cryptocurrencies, thereby generating a confounding increase in trading volume.

To test this alternative explanation, we construct measures of investor attention using the RavenPack News Analytics database, which aggregates data from over 40,000 news and social media sources. RavenPack is particularly suitable for our high-frequency analysis for two reasons. First, it provides millisecond-level timestamps, enabling precise alignment of attention metrics with market activity. Second, and more importantly, its coverage extends beyond traditional media to include investor-driven social media platforms. This is critical, as a surge in posting volume captures heightened attention driven not only by passive information consumption but also by active investor contribution and engagement.

For cryptocurrencies in our sample, we extract all postings mentioning them from RavenPack and compute the total number of new postings per five-minute interval. As an alternative, we construct a second attention measure using only the top 20 sources by posting frequency—such as Ticker Report, CryptoPanic, Investing.com via Web, Yahoo! Finance, Benzinga, Dow Jones Newswires, Business Insider, and Nasdaq. We then control for atten-

tion measures and estimate the following regression model:

$$\log(\text{Trading Volume})_{i,d,t} = \alpha + \beta \times \text{Outage}_{d,t} + \theta \times \text{Attention}_{i,d,t} + \gamma \times \text{Controls}_{i,d-1} + \text{Fixed Effect}_i + \epsilon_{i,d,t} \quad (13)$$

where $\text{Attention}_{i,d,t}$ is the total number of postings from all sources or from the top 20 sources for cryptocurrency $i$ during the 5-minute interval $t$ on day $d$. All other variables are defined in previous specifications,

As shown in Table 8, the estimated effect of outages ($\beta$) remains statistically and economically significant after controlling for investor attention. When controlling for the all-source posting-based attention measure, trading frequency and dollar volume increase by 20.8% and 21.6%, respectively—closely mirroring the baseline increases of 22.6% and 23.8% in Table 2. These results are robust to using alternative attention proxies, including the daily Google search index for Bitcoin (unreported). Thus, we rule out the attention-driven trading channel as an explanation for the observed increase in trading activity.

### *6.3. Alternative Mechanism: Panic Selling*

An alternative mechanism that could explain increased trading volume is panic selling, whereby the technological disruption itself erodes market confidence. An outage of a critical tool like ChatGPT might generate broad negative sentiment, prompting a net sell-off as investors reduce exposure. This hypothesis predicts a significant skew toward sell orders during outages.

To test this, we analyze signed trading volume and net order imbalance. Panel A of Table 9 reports univariate differences between outage and control periods, and Panel B presents panel regression estimates. The results show that both buyer- and seller-initiated volumes increase by a similar magnitude during outages. The change in net order imbalance is small

and statistically insignificant. This symmetric increase in buying and selling is inconsistent with the panic-selling hypothesis.

### *6.4. Evidence of Correlation Noises*

This subsection provides evidence for the existence of correlated noise. A key feature of correlated noise is that it induces traders to submit orders in the same direction within a short interval, even though the direction itself remains random over time. If correlated noise plays an important role, traders are more likely to trade on the same side over short intervals, resulting in larger absolute order imbalances during periods when ChatGPT is available. Thus, we expect the trade imbalance to decline during its outages. To test this hypothesis, we examine changes in trade imbalance, defined as the absolute value of order imbalance over five-minute intervals, from control to outage periods.

Table 9 reports the results. Panel A shows that the average trade imbalance decreases by 7.65% during outage periods relative to the baseline control window $[-5, -1]$. This reduction is statistically significant based on a t-test, the nonparametric Wilcoxon signed-rank test, and bootstrap simulations. The result is robust to alternative control windows ($[-10, -1]$ and $[-15, -1]$); these untabulated estimates are qualitatively similar. Panel B presents regression estimates that control for coin and interval fixed effects. The coefficient on the outage dummy is $-0.073$ (without controls) and $-0.075$ (with controls). Both estimates are statistically significant and quantitatively align with the 7.65% decrease shown in Panel A. These findings support the view that the adoption of ChatGPT increases correlated noise in order flow.

### *6.5. Differential Impact on Large and Small Trades*

When ChatGPT fails, its impact on investors' trading may vary across trade-size groups. To explore this heterogeneity, we partition our control sample into ten subsamples based on trade size: below \$100, (\$100, \$200], (\$200, \$300], (\$300, \$500], (\$500, \$1000], (\$1000, \$2000], (\$2000, \$3000], (\$3000, \$5000], (\$5000, \$10,000], and above \$10,000. For each group, we report the total trading frequency and total dollar volume during the control period and outage period, along with their respective changes from the control to the outage period. Our focus is on the difference in these changes between the smallest trade group and the largest trade group.

Table 10 summarizes these results. During the control period [-5, -1], trading frequency varies dramatically by trade-size groups (Column 2), with over 11 million total trades in the smallest group (below \$100 per trade) versus 260,000 trades in the largest (above \$10,000 per trade). Despite this disparity in levels, Column 3 reveals that the outage-induced percentage changes in trade frequency fall within a range of 21.59% to 29.74% for all groups. The difference in these changes between the smallest and largest trade-size groups is statistically insignificant (difference = -0.06%, p-value = 0.99 using a bootstrap test).

The results for total dollar volume (Columns 5-7) reveal a pattern opposite to that of trading frequency. While small trades dominate in count, large trades dominate in dollar volume. The smallest group contributes approximately \$330 million in total volume, compared to \$6.54 billion from the largest group. Nevertheless, the effect of the outage on dollar volume is similarly homogeneous—changes again fall within the 21–30% range across the size spectrum. The difference in volume changes between the extreme groups is also indistinguishable from zero (difference = -0.65%, p-value = 0.92). We conclude that the impact of ChatGPT outage on small- and large-trade size groups is similar.

### *6.6. Robustness Checks*

This subsection presents a series of tests to establish the robustness of our primary findings. First, we address the integer nature of our dependent variable, trading frequency. In the spirit of Cohn, Liu, and Wardlaw (2022), we estimate a Poisson pseudo-maximum-likelihood regression model with fixed effects using the raw (untransformed) trading frequency. The results, presented in Panel A of Table 11, show that the coefficients on *Outage* remain positive and significant at the 1% level.

Second, to mitigate the influence of outliers. We employ a median regression with standard errors clustered by coin pair and time interval, following Machado and Silva (2019), and also trim the top and bottom 2.5% of outliers. Panel B of Table 11 confirms that our results are robust. We also verify that the shortest service disruptions do not drive our findings. Excluding the shortest or the four shortest outage events (each lasting less than 20 minutes) produces results similar to our baseline findings, as reported in Panel C of Table 8.

Third, we vary the granularity of our fixed effects to control for unobserved heterogeneity. We sequentially include outage–coin-pair and outage–interval fixed effects to absorb any event-specific confounding factors, and the more stringent outage–coin-pair–interval fixed effects. Across all specifications, the *Outage* coefficients remain statistically significant with only minor changes in magnitude. These results are reported in Panel D of Table 11.

Next, we conduct a placebo test by randomly assigning a pseudo-outage date from the ten-day window preceding each actual outage. We then re-estimate our main specifications using data from these pseudo-outage events (over the same time period as the actual outage) and their time-matched control windows during [-5, -1]. The insignificant coefficients on these pseudo-outage terms across all outcomes support a causal interpretation of our primary findings. These results are reported in Panel E of Table 11.

As a final robustness check, we graphically analyze six measures of trading activity and

market liquidity around outages to assess our results more directly. Figure 1 plots the gap between the logged value of each measure in a 5-minute interval and its time-matched control value from a fifteen-day baseline period, averaged across all outage events (requiring at least ten events). We demean each series by its average level in the 30 minutes preceding the outage onset to focus on the dynamic response relative to the immediate pre-outage baseline. For brevity, we display only a 60-minute window during each outage, along with the 30 minutes before and after, which reveals a clear pattern: trading volume and return volatility spike during the disruption before reverting, while market depth and bid–ask spreads move consistently with our regression estimates. We obtain qualitatively similar dynamics when plotting the raw log differences without demeaning or when using alternative pre-event windows.

These visual results have two implications. First, because the log difference between outage and control values accounts for intraday seasonality, no systematic changes in the variables are present prior to the onset of the outage. Second, the effects vanish after the outage, with variables reverting toward their levels observed in the 30 minutes preceding the disruption. The absence of pre-event movements and the sharp post-event dynamics are consistent with outages representing unanticipated shocks rather than the continuation of pre-existing trends.

Figure 2 compares the density distributions of the variables used in our main regressions between control and outage periods. Each panel reports the natural logarithm of trading volume, volatility, or liquidity measures, standardized at the exchange–coin level. Relative to control periods, outage periods exhibit distributional shifts across a range of percentiles. The difference between the two distributions indicates that outages affect both the central tendency and the dispersion of trading activity and liquidity.

# 7. Conclusion

This paper examines the impact of generative AI on investor trading and market quality. We first develop a theoretical model that formalizes how GenAI influences market activity and liquidity. Our framework identifies two channels through which ChatGPT operates: the information production channel and the noise-correlation channel. A novel friction emerges from the latter: correlated noise in AI-generated outputs homogenizes trader signals, leading to more aligned beliefs. While this reduces disagreement, it simultaneously impairs collective information production by making aggregated signals less informative. This, in turn, mitigates the adverse selection problem faced by market makers. Conversely, the information channel posits that ChatGPT adoption enhances traders' information-processing capacity, increasing the fundamental informational content of order flow and thus intensifying adverse selection. Our model yields a testable prediction: if the noise-correlation channel dominates, ChatGPT outages will lead to increased trading volume, higher return volatility, and deteriorated market liquidity.

We use unexpected ChatGPT outages as a natural experiment to test our model's predictions in cryptocurrency markets—a setting dominated by retail investors who historically lacked access to sophisticated analytical tools. Using high-frequency data for 70 major coin pairs traded on the top 10 cryptocurrency exchanges, we document that, during ChatGPT outage periods, there is a significant increase in trading frequency, dollar volume, and return volatility. Concurrently, we observe a significant deterioration in liquidity, evidenced by a decline in limit-order book depth and an increase in price impact. These changes in trading activity and liquidity are both statistically and economically significant, suggesting the dominance of the noise-correlation channel in this market.

Our model and empirical analysis yields two insights. First, investors in cryptocurrency markets heavily rely on generative AI tools like ChatGPT to formulate their trading

strategies. Second, while this AI-assisted trading enhances market liquidity under normal conditions, it does so through a nuanced mechanism: the synchronization of noise across traders, which reduces adverse selection.

The findings of this paper reveal the dual nature of GenAI's impact in financial markets: it democratizes access to sophisticated analysis while simultaneously creating a novel, systemic risk through correlated signal noise. Whereas existing literature primarily emphasizes GenAI's role in enhancing private information precision, our study provides a critical counterpoint by demonstrating how its widespread adoption can homogenize investor behavior and alter the fundamental properties of market quality.

# References


Amihud, Y., 2002. Illiquidity and stock returns: cross-section and time-series effects. Journal of Financial Markets 5, 31–56.

Andersen, T. G., Bollerslev, T., Diebold, F. X., Ebens, H., 2001. The distribution of realized stock return volatility. Journal of Financial Economics 61, 43–76.

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

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 pp. 1–23.

Biais, B., Bisiere, C., Bouvard, M., Casamatta, C., Menkveld, A. J., 2023. Equilibrium bitcoin pricing. Journal of Finance 78, 967–1014.

Brauneis, A., Mestel, R., Riordan, R., Theissen, E., 2021. How to measure the liquidity of cryptocurrency markets? Journal of Banking & Finance 124.

Cheng, Q., Lin, P., Zhao, Y., 2025. Does generative ai facilitate investor trading? early evidence from chatgpt outages. Journal of Accounting and Economics .

Cohn, J. B., Liu, Z., Wardlaw, M. I., 2022. Count (and count-like) data in finance. Journal of Financial Economics 146, 529–551.

Colliard, J.-E., Foucault, T., 2012. Trading fees and efficiency in limit order markets. Review of Financial Studies 25, 3389–3421.

Cong, L. W., Li, Y., Wang, N., 2021. Tokenomics: Dynamic adoption and valuation. Review of Financial Studies 34, 1105–1155.

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, 502–528.

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

Gryglewicz, S., Mayer, S., Morellec, E., 2021. Optimal financing with tokens. Journal of Financial Economics 142, 1038–1067.

Hassan, T. A., Mertens, T. M., 2011. Market sentiment: A tragedy of the commons. American Economic Review 101, 402–405.

Hassan, T. A., Mertens, T. M., 2017. The social cost of near-rational investment. American Economic Review 107, 1059–1103.

Kyle, A. S., 1985. Continuous auctions and insider trading. Econometrica pp. 1315–1335.

Machado, J. A., Silva, J. S., 2019. Quantiles via moments. Journal of Econometrics 213, 145–173.

Scheinkman, J. A., Xiong, W., 2003. Overconfidence and speculative bubbles. Journal of Political Economy 111, 1183–1220.

Shkilko, A., Sokolov, K., 2020. Every cloud has a silver lining: Fast trading, microwave connectivity, and trading costs. Journal of Finance 75, 2899–2927.

Sockin, M., Xiong, W., 2023. A model of cryptocurrencies. Management Science 69, 6684–6707.

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

Wang, J., 1994. A model of competitive stock trading volume. Journal of Political Economy 102, 127–168.

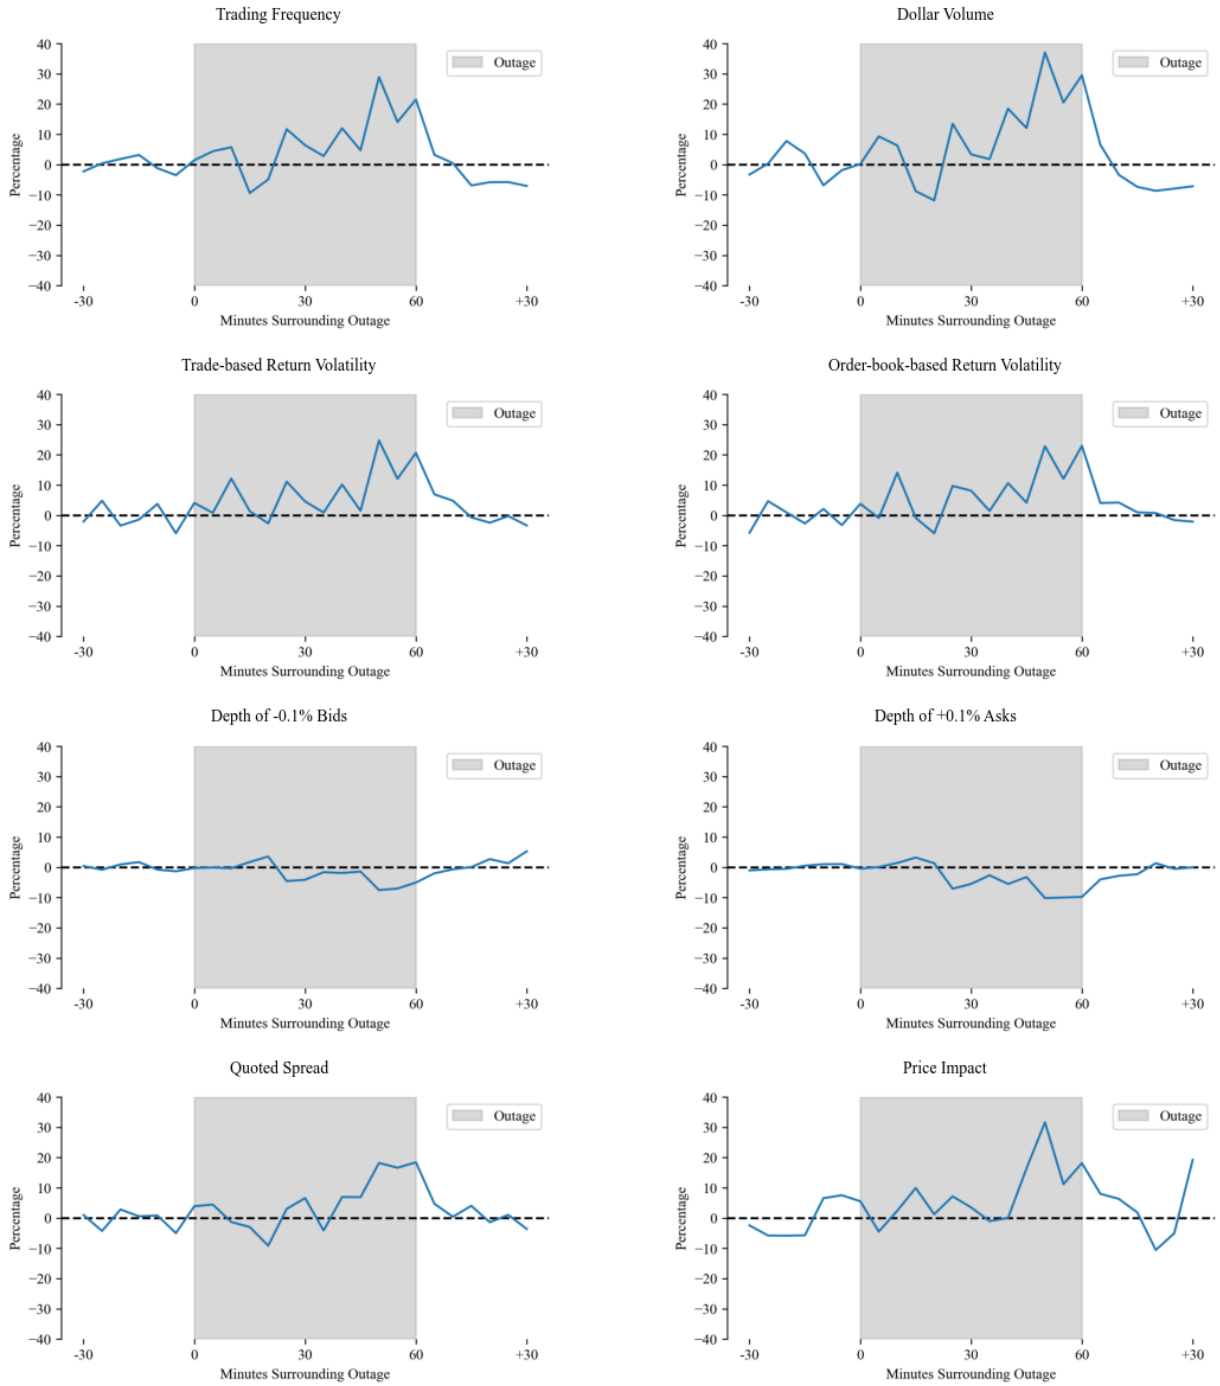



Fig. 1: **Market Quality Surrounding ChatGPT outages**

This figure shows changes in trading volume and market liquidity surrounding ChatGPT outages. Each panel reports the log-differences between outage days and and its time-matched control value from a fifteen-day baseline period, averaged across all outage events (requiring at least ten events). The plots use one-hour windows for outages because at least ten outage events last longer than one hour. We demean each series by its average level in the 30 minutes preceding the outage onset.

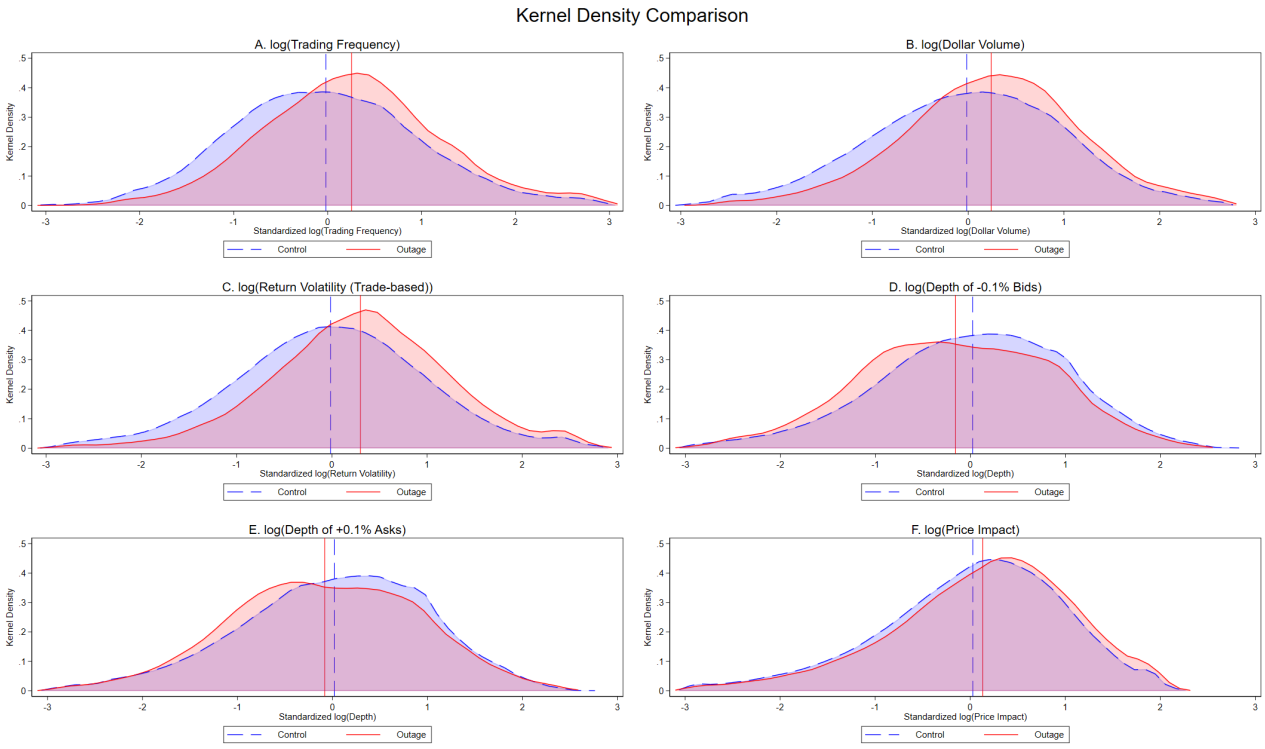



Fig. 2: **Market Quality Distribution in Control and Outage Periods**

This figure shows distributions in trading volume and market liquidity in control and outage periods. Each panel reports the natural logarithm of trading volume or liquidity measures. We standardized all logarithm measures for each exchange-coin pair. The outage periods are defined as the 5-minute intervals during ChatGPT outages. The control periods are defined as the same 5-minute interval in fifteen days preceding the outages. The solid curve (vertical line) shows the distribution (mean) of trading volume or liquidity measures in outage periods. The dash curve (vertical line) shows the distribution (mean) of trading volume or liquidity measures in control periods.

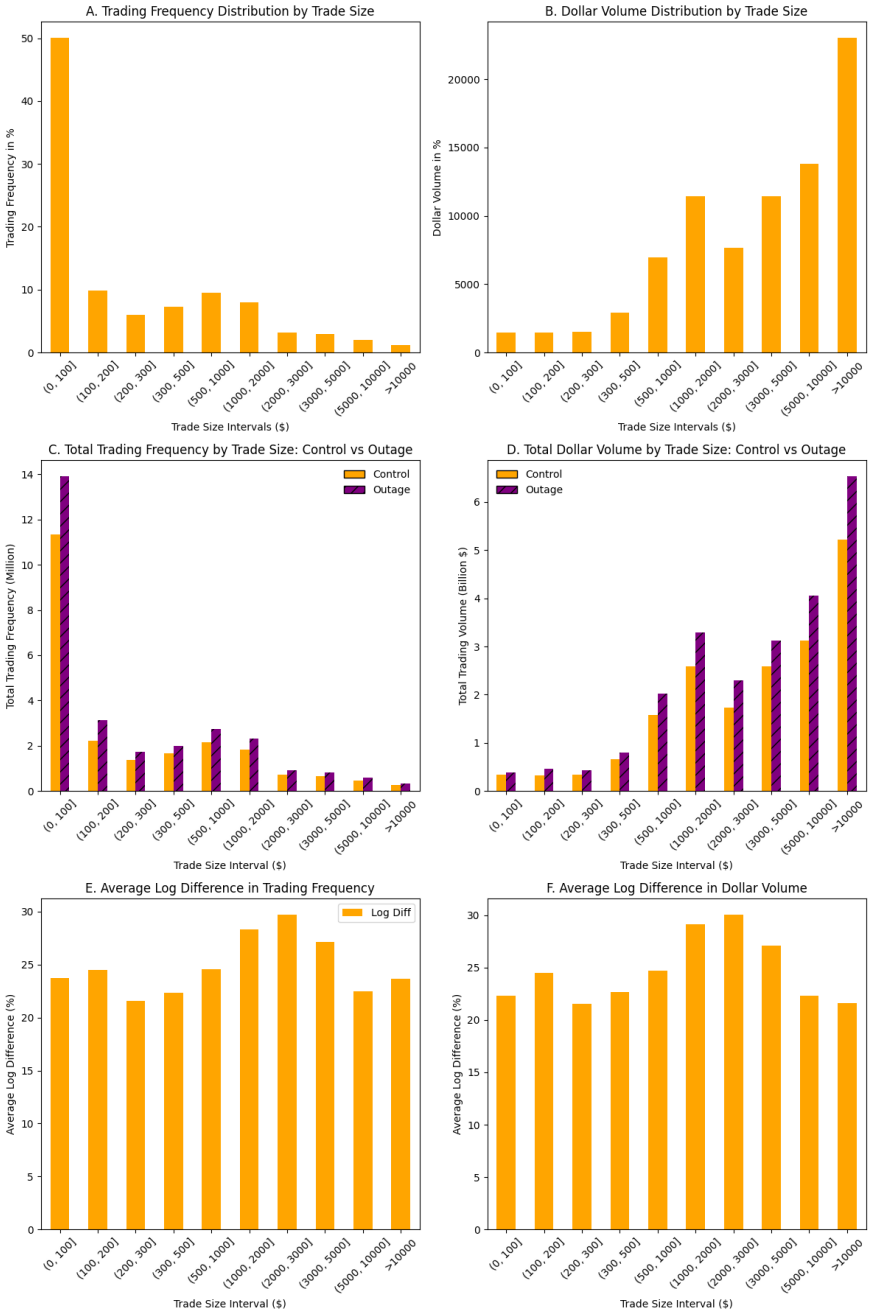



Fig. 3: **Distribution by Order Size**

This figure shows the distributions and changes of trading frequency and dollar volume among ten transaction size groups. Panel A (B) presents the fraction of trading frequency (dollar volume) for each transaction size group during control periods. Panel C (D) presents the total trading frequency (trading volume) for each transaction size group during the control and outage periods. Panel E (F) presents the logarithmic difference (in %) in total trading frequency (dollar volume) for each transaction size group between the outage and control periods.

**Table 1. Characteristics of Sample Cryptocurrency Pairs**

This table presents cross-sectional summary statistics for the 70 cryptocurrency pairs in our sample. Intraday data are aggregated at the 5-minute frequency. For each coin pair and variable, we compute its time-series average over the sample period and report the cross-sectional mean, standard deviation (S.D.), and the 25th, 50th, and 75th percentiles. Trading frequency is the number of transactions per 5-minute interval. Dollar volume is the total dollar value traded per 5-minute interval (in $1,000). Return is the 5-minute return calculated from the log difference of transaction prices. Trade-based (quote-based) return volatility is the standard deviation of 5-minute returns, calculated from the log difference of transaction prices (bid-ask midpoints). Quoted spread is the difference between the best bid and ask prices. Depth at -0.1% Bid (+0.1% Ask) is the total dollar value (in $1,000) available in the limit order book if the bid (ask) price moves -0.1% (+0.1%) from the current best bid (ask). Price impact is the change in the bid-ask mid-point (in bps) five minutes after a trade. Buy volume is the total dollar volume purchased per 5-minute interval (in $1,000). Sell volume is the total dollar volume sold per 5-minute interval (in $1,000). Order imbalance is the difference between buy and sell volume scaled by total dollar volume. Trade imbalance is the absolute value of order imbalance. The sample period is November 2022 to December 2022, a two-month window preceding the first ChatGPT outage event. The panel comprises 1,545,679 observations.

| | Mean | S.D. | P25 | P50 | P75 |
|---|---|---|---|---|---|
| Trading Frequency | 529.64 | 1807.34 | 44.15 | 186.37 | 372.59 |
| Dollar Volume ($1,000) | 367.74 | 1568.77 | 17.98 | 50.98 | 149.82 |
| Return (bps) | -0.09 | 0.42 | -0.14 | -0.02 | 0.05 |
| Trade-based Return Volatility (bps) | 4.85 | 1.43 | 3.62 | 4.86 | 5.79 |
| Quote-based Return Volatility (bps) | 4.95 | 1.30 | 3.76 | 5.09 | 5.78 |
| Quoted Spread (cents) | 17.43 | 70.00 | 0.01 | 0.39 | 4.80 |
| Depth of -0.1% Bids ($1,000) | 230.83 | 380.55 | 35.59 | 80.48 | 248.86 |
| Depth of +0.1% Asks ($1,000) | 220.42 | 362.78 | 33.00 | 81.34 | 238.68 |
| Price Impact (bps) | 13.59 | 12.86 | 6.50 | 8.84 | 13.42 |
| Buy Volume ($1,000) | 190.77 | 815.75 | 9.31 | 26.40 | 86.93 |
| Sell Volume ($1,000) | 192.63 | 825.05 | 9.56 | 27.82 | 77.55 |
| Order Imbalance (%) | 3.80 | 8.77 | -1.81 | 0.41 | 7.43 |
| Trade Imbalance (%) | 42.41 | 17.65 | 30.49 | 37.57 | 58.67 |

## Table 2. ChatGPT Outages and Trading Activity

Panel A presents the cross-sectional means and medians of trading frequency and dollar volume during the control and outage periods, along with their logarithmic changes (Log Diff in %). Intraday data for all variables are aggregated at the 5-minute level. For each outage event, we identify its actual time window and define the same time window from 1-5 days before the outage date (e.g., [−5,−1]) as the control period. Two alternative control periods are [−10,−1] and [−15,−1]. We compute the average trading frequency (dollar volume) across all outage event (control) periods for each cryptocurrency pair, and then obtain the cross-sectional summary statistics. The significance of the change is assessed using a t-test, the nonparametric Wilcoxon signed-rank test, and bootstrap simulations. Panel B presents estimates from the following panel regression:

$$\log(\text{Trading Activity})_{i,d,t} = \alpha + \beta \times \text{Outage}_{d,t} + \gamma \times \text{Controls}_{i,d-1} + \text{Fixed Effects}_{d,t} + \epsilon_{i,d,t}$$

where the dependent variable is log trading frequency or log dollar volume for cryptocurrency $i$ during the 5-minute interval $t$ on day $d$. The indicator $\text{Outage}_{d,t}$ equals one if interval $t$ on day $d$ falls within a ChatGPT outage event window, and zero otherwise. Control variables include the logarithm of the previous day's market capitalization and return, the average return from days $d-5$ to $d-2$, and the logarithm of the previous day's closing price. Each specification includes coin and interval fixed effects. Standard errors are clustered at the coin and interval levels. The sample period spans from November 2022 to August 2024. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively.

Panel A. Univariate Analysis

| | | Trading Frequency | | | Dollar Volume ($1,000) | | |
|---|---|---|---|---|---|---|---|
| | | Control | Outage | Log Diff. (%) | Control | Outage | Log Diff. (%) |
| [-5,-1] | Mean | 629.65 | 764.61 | 21.64*** | 511.40 | 624.42 | 23.76*** |
| | Median | 169.41 | 234.64 | 22.23 | 62.13 | 83.48 | 24.43 |
| | Wilcoxon p-value | | | (0.00) | | | (0.00) |
| | Bootstrap p-value | | | (0.00) | | | (0.00) |
| [-10,-1] | Mean | 644.54 | 764.61 | 17.48*** | 527.68 | 624.42 | 19.98*** |
| | Median | 176.87 | 234.64 | 19.73 | 65.95 | 83.48 | 20.39 |
| | Wilcoxon p-value | | | (0.00) | | | (0.00) |
| | Bootstrap p-value | | | (0.00) | | | (0.00) |
| [-15,-1] | Mean | 606.85 | 764.61 | 21.07*** | 491.50 | 624.42 | 22.77*** |
| | Median | 173.71 | 234.64 | 24.22 | 66.34 | 83.48 | 22.92 |
| | Wilcoxon p-value | | | (0.00) | | | (0.00) |
| | Bootstrap p-value | | | (0.00) | | | (0.00) |

Panel B. Panel Regressions of Trading Activity

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) | (12) |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | [-5, -1] | | | | [-10, -1] | | | | [-15, -1] | | | |
| | Trading Frequency | | Dollar Volume | | Trading Frequency | | Dollar Volume | | Trading Frequency | | Dollar Volume | |
| $Outage_{d,t}$ | 0.217*** | 0.221*** | 0.281*** | 0.278*** | 0.191*** | 0.166*** | 0.243*** | 0.197*** | 0.236*** | 0.196*** | 0.288*** | 0.224*** |
| | (6.75) | (7.27) | (7.97) | (8.10) | (6.13) | (5.41) | (7.33) | (5.97) | (6.98) | (5.92) | (8.73) | (6.75) |
| $log(MarketCap_{d-1}$ | | -0.555 | | 0.661 | | -0.540 | | 0.406 | | -0.208 | | 0.706 |
| | | (-0.69) | | (0.90) | | (-0.57) | | (0.43) | | (-0.25) | | (0.78) |
| $Ret_{d-1}$ | | 1.664*** | | 2.062*** | | 1.099*** | | 1.368*** | | 0.841** | | 0.970** |
| | | (4.62) | | (4.61) | | (2.90) | | (2.96) | | (2.59) | | (2.51) |
| $Ret_{[d-5,d-2]}$ | | 1.733*** | | 2.120*** | | 0.495* | | 0.535* | | 0.913*** | | 0.995*** |
| | | (5.11) | | (5.43) | | (2.05) | | (1.97) | | (3.64) | | (3.48) |
| $Price_{d-1}$ | | 0.933 | | 0.197 | | 0.858 | | 0.394 | | 0.534 | | 0.115 |
| | | (1.08) | | (0.24) | | (0.85) | | (0.39) | | (0.61) | | (0.12) |
| Observations | 178,444 | 178,444 | 178,444 | 178,444 | 327,234 | 327,234 | 327,234 | 327,234 | 473,403 | 473,403 | 473,403 | 473,403 |
| Coin Pair FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Interval FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Adj. R-squared | 0.307 | 0.319 | 0.318 | 0.339 | 0.307 | 0.313 | 0.316 | 0.332 | 0.291 | 0.298 | 0.306 | 0.324 |

## Table 3. ChatGPT Outages and Return Volatility

Panels A and B report the cross-sectional means and medians of trade-based and quote-based return volatility during the control and outage periods, along with their logarithmic changes (Log Diff in %). Intraday data for all variables are aggregated at the 5-minute level. For each outage event, we identify its actual time window and define the same time window from 1-5 days before the outage date (e.g., [−5,−1]) as the control period. Two alternative control periods are [−10,−1] and [−15,−1]. Trade-based (quote-based) return volatility is the standard deviation of 5-minute returns, calculated from the log difference of transaction prices (bid-ask midpoints). We compute the average volatility across all outage event (control) periods for each cryptocurrency pair, and then obtain the cross-sectional summary statistics. The significance of the change is assessed using a t-test, the nonparametric Wilcoxon signed-rank test, and bootstrap simulations. Panel B presents estimates from the following panel regression:

$$\log(\text{Volatility})_{i,d,t} = \alpha + \beta \times \text{Outage}_{d,t} + \gamma \times \text{Controls}_{i,d-1} + \text{Fixed Effects}_{d,t} + \epsilon_{i,d,t}$$

where the dependent variable is log trade-based or log quote-based return volatility for cryptocurrency $i$ during the 5-minute interval $t$ on day $d$. The indicator $\text{Outage}_{d,t}$ equals one if interval $t$ on day $d$ falls within a ChatGPT outage event window, and zero otherwise. Control variables include the logarithm of the previous day's market capitalization and return, the average return from days $d-5$ to $d-2$, and the logarithm of the previous day's closing price. Each specification includes coin and interval fixed effects. Standard errors are clustered at the coin and interval level. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively.

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

| | [−5,−1] | | | [−10,−1] | | | [−15,−1] | | |
|---|---|---|---|---|---|---|---|---|---|
| | Control | Outage | Log Diff. (%) | Control | Outage | Log Diff. (%) | Control | Outage | Log Diff. (%) |
| Mean | 5.23 | 6.35 | 19.19*** | 5.37 | 6.35 | 16.55*** | 5.18 | 6.35 | 20.39*** |
| Median | 5.10 | 6.28 | 19.01 | 5.13 | 6.28 | 17.11 | 5.07 | 6.28 | 21.03 |
| Wilcoxon p-value | | | (0.00) | | | (0.00) | | | (0.00) |
| Bootstrap p-value | | | (0.00) | | | (0.00) | | | (0.00) |

Panel B. Univariate Analysis: quote-based Return Volatility (in bps)

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Mean | 5.51 | 6.64 | 18.01*** | 5.63 | 6.64 | 15.97*** | 5.44 | 6.64 | 19.60*** |
| Median | 5.61 | 6.78 | 18.73 | 5.75 | 6.78 | 17.39 | 5.54 | 6.78 | 20.21 |
| Wilcoxon p-value | | | (0.00) | | | (0.00) | | | (0.00) |
| Bootstrap p-value | | | (0.00) | | | (0.00) | | | (0.00) |

Panel C. Panel Regressions of Return Volatility

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) | (12) |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Trade-based Return Volatility | | | | | | Order-book-based Return Volatility | | | | | |
| | [-5, -1] | | [-10, -1] | | [-15, -1] | | [-5, -1] | | [-10, -1] | | [-15, -1] | |
| $Outage_{d,t}$ | 0.184*** | 0.192*** | 0.158*** | 0.149*** | 0.201*** | 0.181*** | 0.191*** | 0.198*** | 0.164*** | 0.155*** | 0.205*** | 0.187*** |
| | (7.74) | (7.78) | (7.00) | (6.76) | (8.92) | (8.36) | (8.21) | (8.28) | (7.39) | (7.08) | (9.40) | (8.72) |
| $log(MarketCap)_{d-1}$ | | -0.804* | | -0.906** | | -0.569 | | -0.768* | | -0.834* | | -0.522 |
| | | (-1.97) | | (-2.11) | | (-1.66) | | (-1.99) | | (-2.03) | | (-1.61) |
| $Ret_{d-1}$ | | 1.274*** | | 0.717** | | 0.362 | | 1.304*** | | 0.640** | | 0.271 |
| | | (4.31) | | (2.51) | | (1.48) | | (4.36) | | (2.29) | | (1.13) |
| $Ret_{d-2,d-5}$ | | 1.330*** | | 0.372* | | 0.707*** | | 1.303*** | | 0.350* | | 0.659*** |
| | | (4.81) | | (1.91) | | (3.75) | | (4.99) | | (1.82) | | (3.55) |
| $Price_{d-1}$ | | 0.853* | | 0.933* | | 0.618 | | 0.834* | | 0.878* | | 0.592 |
| | | (1.99) | | (2.02) | | (1.71) | | (2.04) | | (1.98) | | (1.73) |
| Observations | 178,376 | 178,376 | 327,100 | 327,100 | 473,224 | 473,224 | 178,194 | 178,194 | 326,808 | 326,808 | 472,752 | 472,752 |
| Coin Pair FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Interval FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Adj. R-squared | 0.218 | 0.238 | 0.207 | 0.211 | 0.204 | 0.210 | 0.230 | 0.251 | 0.216 | 0.221 | 0.212 | 0.219 |

## Table 4. ChatGPT Outages and Book Depths

Panel A presents the cross-sectional means and medians of two book depth measures during the control and outage periods, along with their logarithmic changes (Log Diff in %). Depth at -0.1% Bid (+0.1% Ask) is the total dollar value (in $1,000) available in the limit order book if the bid (ask) price moves -0.1% (+0.1%) from the current best bid (ask). Intraday data for all variables are aggregated at the 5-minute level. For each outage event, we identify its actual time window and define the same time window from 1-5 days before the outage date (e.g., [−5,−1]) as the control period. Two alternative control periods are [−10,−1] and [−15,−1]. We compute the average of each depth measure across all outage event (control) periods for each cryptocurrency pair, and then obtain the cross-sectional summary statistics. The significance of the change is assessed using a t-test, the nonparametric Wilcoxon signed-rank test, and bootstrap simulations. Panel B presents estimates from the following panel regression:

$$\log(\text{Depth})_{i,d,t} = \alpha + \beta \times \text{Outage}_{d,t} + \gamma \times \text{Controls}_{i,d-1} + \text{Fixed Effects}_{d,t} + \epsilon_{i,d,t}$$

where the dependent variable is log Depth at -0.1% Bid or log depth at +0.1% Ask for cryptocurrency $i$ during the 5-minute interval $t$ on day $d$. The indicator $\text{Outage}_{d,t}$ equals one if interval $t$ on day $d$ falls within a ChatGPT outage event window, and zero otherwise. Control variables include the logarithm of the previous day's market capitalization and return, the average return from days $d-5$ to $d-2$, and the logarithm of the previous day's closing price. Each specification includes coin and interval fixed effects. Standard errors are clustered at the coin and interval level. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively.

Panel A. Univariate Analysis

| | | Depth of -0.1% Bids ($1,000) | | | Depth of +0.1% Asks ($1,000) | | |
|---|---|---|---|---|---|---|---|
| | | Control | Outage | Log Diff. (%) | Control | Outage | Log Diff. (%) |
| [-5,-1] | Mean | 256.79 | 247.28 | -6.00*** | 269.63 | 262.88 | -3.09** |
| | Median | 85.88 | 78.05 | -5.28 | 88.50 | 87.47 | -4.22 |
| | Wilcoxon p-value | | | (0.00) | | | (0.00) |
| | Bootstrap p-value | | | (0.00) | | | (0.01) |
| [-10,-1] | Mean | 255.29 | 247.28 | -4.69*** | 264.31 | 262.88 | -0.55 |
| | Median | 83.80 | 78.05 | -3.90 | 88.93 | 87.47 | -0.69 |
| | Wilcoxon p-value | | | (0.00) | | | (0.22) |
| | Bootstrap p-value | | | (0.00) | | | (0.66) |
| [-15,-1] | Mean | 265.58 | 247.28 | -7.39*** | 275.99 | 262.88 | -3.15** |
| | Median | 85.87 | 78.05 | -8.33 | 90.98 | 87.47 | -3.74 |
| | Wilcoxon p-value | | | (0.00) | | | (0.00) |
| | Bootstrap p-value | | | (0.00) | | | (0.02) |

Panel B. Panel Regressions of Book Depth of -0.1% Bid

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | [-5, -1] | | [-10, -1] | | [-15, -1] | |
| $Outage_{d,t}$ | -0.066*** | -0.084*** | -0.040** | -0.053*** | -0.083*** | -0.080*** |
| | (-3.74) | (-5.12) | (-2.32) | (-2.98) | (-5.16) | (-4.67) |
| $log(MarketCap)_{d-1}$ | | 2.077*** | | 2.181*** | | 1.796** |
| | | (3.33) | | (2.92) | | (2.81) |
| $Ret_{d-1}$ | | -1.147*** | | -1.342*** | | -1.320*** |
| | | (-4.39) | | (-4.41) | | (-4.83) |
| $Ret_{d-2,d-5}$ | | -1.156*** | | -0.247 | | -0.580*** |
| | | (-4.80) | | (-1.44) | | (-4.23) |
| $Price_{d-1}$ | | -1.567** | | -1.619* | | -1.250* |
| | | (-2.40) | | (-2.05) | | (-1.84) |
| Observations | 178,444 | 178,444 | 327,234 | 327,234 | 473,403 | 473,403 |
| Coin Pair FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Interval FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Adj. R-squared | 0.579 | 0.598 | 0.564 | 0.584 | 0.569 | 0.587 |

Panel C. Panel Regressions of Book Depth of 0.1% Ask

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | [-5, -1] | | [-10, -1] | | [-15, -1] | |
| $Outage_{d,t}$ | -0.055*** | -0.073*** | -0.012 | -0.032** | -0.053*** | -0.058*** |
| | (-3.65) | (-5.41) | (-0.80) | (-2.20) | (-3.85) | (-4.16) |
| $log(MarketCap)_{d-1}$ | | 1.498*** | | 1.839*** | | 1.464** |
| | | (3.31) | | (2.98) | | (2.80) |
| $Ret_{d-1}$ | | -1.159*** | | -0.990*** | | -1.060*** |
| | | (-4.09) | | (-2.88) | | (-3.75) |
| $Ret_{d-2,d-5}$ | | -1.062*** | | 0.011 | | -0.382** |
| | | (-4.77) | | (0.06) | | (-2.56) |
| $Price_{d-1}$ | | -0.913* | | -1.232* | | -0.889 |
| | | (-1.94) | | (-1.91) | | (-1.62) |
| Observations | 178,444 | 178,444 | 327,234 | 327,234 | 473,403 | 473,403 |
| Coin Pair FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Interval FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Adj. R-squared | 0.584 | 0.603 | 0.567 | 0.587 | 0.574 | 0.591 |

## Table 5. The Impact of ChatGPT Outages on Quoted Spread

Panel A presents the cross-sectional means and medians of quoted spread during the control and outage periods, along with their logarithmic changes (Log Diff in %). Intraday data for all variables are aggregated at the 5-minute level. For each outage event, we identify its actual time window and define the same time window from 1-5 days before the outage date (e.g., [−5,−1]) as the control period. Two alternative control periods are [−10,−1] and [−15,−1]. We compute the average quoted spread across all outage event (control) periods for each cryptocurrency pair, and then obtain the cross-sectional summary statistics. The significance of the change is assessed using a t-test, the nonparametric Wilcoxon signed-rank test, and bootstrap simulations. Panel B presents estimates from the following panel regression:

$$\log(\text{Quoted Spread})_{i,d,t} = \alpha + \beta \times \text{Outage}_{d,t} + \gamma \times \text{Controls}_{i,d-1} + \text{Fixed Effects}_{d,t} + \epsilon_{i,d,t}$$

where the dependent variable is log quoted spread for cryptocurrency $i$ during the 5-minute interval $t$ on day $d$. The indicator $\text{Outage}_{d,t}$ equals one if interval $t$ on day $d$ falls within a ChatGPT outage event window, and zero otherwise. Control variables include the logarithm of the previous day's market capitalization and return, the average return from days $d-5$ to $d-2$, and the logarithm of the previous day's closing price. Each specification includes coin and interval fixed effects. Standard errors are clustered at the cryptocurrency and day level. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively.

Panel A. Univariate Analysis

| | [−5,−1] | | | [−10,−1] | | | [−15,−1] | | |
|---|---|---|---|---|---|---|---|---|---|
| | Control | Outage | Log Diff. (%) | Control | Outage | Log Diff. (%) | Control | Outage | Log Diff. (%) |
| Mean | 21.19 | 21.98 | 4.31*** | 22.33 | 21.98 | 2.11* | 20.40 | 21.98 | 5.28*** |
| Median | 0.50 | 0.56 | 2.28 | 0.50 | 0.56 | 2.11 | 0.49 | 0.56 | 4.13 |
| Wilcoxon p-value | | | (0.00) | | | (0.01) | | | (0.00) |
| Bootstrap p-value | | | (0.00) | | | (0.06) | | | (0.00) |

Panel B. Panel Regressions of Quoted Spread

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | [-5, -1] | | [-10, -1] | | [-15, -1] | |
| $Outage_{d,t}$ | 0.063*** | 0.065*** | 0.050*** | 0.028** | 0.082*** | 0.048*** |
| | (4.92) | (4.25) | (3.73) | (2.15) | (5.13) | (3.76) |
| $log(MarketCap)_{d-1}$ | | -0.116 | | -0.276 | | -0.327 |
| | | (-0.17) | | (-0.38) | | (-0.49) |
| $Ret_{d-1}$ | | 1.242*** | | 1.046** | | 0.855** |
| | | (3.17) | | (2.63) | | (2.36) |
| $Ret_{d-2,d-5}$ | | 1.273*** | | 0.529* | | 0.701** |
| | | (2.92) | | (1.87) | | (2.60) |
| $Price_{d-1}$ | | 0.400 | | 0.533 | | 0.607 |
| | | (0.54) | | (0.67) | | (0.84) |
| Observations | 178,444 | 178,444 | 327,234 | 327,234 | 473,403 | 473,403 |
| Coin Pair FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Interval FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Adj. R-squared | 0.926 | 0.928 | 0.925 | 0.926 | 0.923 | 0.924 |

## Table 6. The Impact of ChatGPT Outages on Price Impact

Panel A presents the cross-sectional means and medians of price impact during the control and outage periods, along with their logarithmic changes (Log Diff in %). Intraday data for all variables are aggregated at the 5-minute level. For each outage event, we identify its actual time window and define the same time window from 1-5 days before the outage date (e.g., [−5,−1]) as the control period. Two alternative control periods are [−10,−1] and [−15,−1]. We compute the average price impact across all outage event (control) periods for each cryptocurrency pair, and then obtain the cross-sectional summary statistics. The significance of the change is assessed using a t-test, the nonparametric Wilcoxon signed-rank test, and bootstrap simulations. Panel B presents estimates from the following panel regression:

$$\log(\text{Price Impact})_{i,d,t} = \alpha + \beta \times \text{Outage}_{d,t} + \gamma \times \text{Controls}_{i,d-1} + \text{Fixed Effects}_{d,t} + \epsilon_{i,d,t}$$

where the dependent variable is price impact for cryptocurrency $i$ during the 5-minute interval $t$ on day $d$. The indicator $\text{Outage}_{d,t}$ equals one if interval $t$ on day $d$ falls within a ChatGPT outage event window, and zero otherwise. Control variables include the logarithm of the previous day's market capitalization and return, the average return from days $d-5$ to $d-2$, and the logarithm of the previous day's closing price. Each specification includes coin and interval fixed effects. Standard errors are clustered at the coin and internal level. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively.

Panel A. Univariate Analysis

| | [−5,−1] | | | [−10,−1] | | | [−15,−1] | | |
|---|---|---|---|---|---|---|---|---|---|
| | Control | Outage | Log Diff. (%) | Control | Outage | Log Diff. (%) | Control | Outage | Log Diff. (%) |
| Mean | 12.41 | 13.96 | 12.25*** | 12.64 | 13.96 | 10.61*** | 12.37 | 13.96 | 12.65*** |
| Median | 8.87 | 10.84 | 13.28 | 9.02 | 10.84 | 10.93 | 8.88 | 10.84 | 12.60 |
| Wilcoxon p-value | | | (0.00) | | | (0.00) | | | (0.00) |
| Bootstrap p-value | | | (0.00) | | | (0.00) | | | (0.00) |

Panel B. Panel Regressions of Price Impact

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | [-5, -1] | | [-10, -1] | | [-15, -1] | |
| $Outage_{d,t}$ | 0.144*** | 0.147*** | 0.127*** | 0.125*** | 0.150*** | 0.142*** |
| | (4.82) | (5.00) | (4.49) | (4.40) | (5.43) | (5.22) |
| $log(MarketCap)_{d-1}$ | | -1.921*** | | -1.890*** | | -1.500*** |
| | | (-4.46) | | (-4.33) | | (-3.78) |
| $Ret_{d-1}$ | | 1.052*** | | 0.421 | | 0.228 |
| | | (3.40) | | (1.71) | | (1.08) |
| $Ret_{d-2,d-5}$ | | 0.755*** | | 0.119 | | 0.360** |
| | | (3.85) | | (0.76) | | (2.68) |
| $Price_{d-1}$ | | 1.927*** | | 1.903*** | | 1.528*** |
| | | (4.31) | | (4.14) | | (3.72) |
| Observations | 178,363 | 178,363 | 327,106 | 327,106 | 473,214 | 473,214 |
| Coin Pair FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Interval FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Adj. R-squared | 0.091 | 0.095 | 0.086 | 0.088 | 0.084 | 0.086 |

**Table 7. Test of the Impact of ChatGPT Outages using Records from StatusGator**

The table reports the logarithmic change (Log Diff in %) of trading activity, return volatility, book depths, quoted spread, and price impact from the control to the outage periods. Outage events are reported by the third-party website detector, StatusGator, instead of OpenAI. Intraday data for all variables are aggregated at the 5-minute level. For each outage event, we identify its actual time window and define the same time window from 1-5 days before the outage date (e.g., [−5,−1]) as the control period. Two alternative control periods are [−10,−1] and [−15,−1]. We compute the average of each variable across all outage event (control) periods for each cryptocurrency pair, and then obtain the cross-sectional summary statistics. The significance of the change is assessed using a t-test, the nonparametric Wilcoxon signed-rank test, and bootstrap simulations. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively.

| | | Trading Volume | | Return Volatility | | 0.1% Depth | | Quoted Spread | Price Impact |
|---|---|---|---|---|---|---|---|---|---|
| | | Freqency | Dollar | Trade-based | Order-book-based | Bid | Ask | | |
| | | Log Diff. (%) | Log Diff. (%) | Log Diff. (%) | Log Diff. (%) | Log Diff. (%) | Log Diff. (%) | Log Diff. (%) | Log Diff. (%) |
| [-5,-1] | Mean | 19.24*** | 25.49*** | 16.99*** | 15.47*** | -5.88*** | -4.06*** | 6.27*** | 11.45*** |
| | Median | 25.67 | 35.28 | 16.39 | 15.34 | -5.30 | -4.63 | 2.99 | 14.65 |
| | Wilcoxon p-value | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) |
| | Bootstrap p-value | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.01) | (0.00) |
| [-10,-1] | Mean | 20.37*** | 28.13*** | 21.90*** | 20.04*** | -4.59*** | -1.50 | 7.77*** | 14.51*** |
| | Median | 24.18 | 29.56 | 25.17 | 21.11 | -3.92 | -0.93 | 3.95 | 17.00 |
| | Wilcoxon p-value | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.10) | (0.00) | (0.00) |
| | Bootstrap p-value | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.12) | (0.00) | (0.00) |
| [-15,-1] | Mean | 22.58*** | 30.74*** | 23.75*** | 21.88*** | -7.30*** | -4.11*** | 8.70*** | 15.29*** |
| | Median | 24.71 | 30.68 | 24.88 | 22.28 | -8.20 | -3.81 | 6.01 | 17.67 |
| | Wilcoxon p-value | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) |
| | Bootstrap p-value | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) | (0.00) |

## Table 8. ChatGPT Outages and Investors' Attention

This table reports the effects of ChatGPT outages on trading activity, controlling for investor attention. We present estimates from the following panel regression:

$$\log(\text{Trading Activity})_{i,d,t} = \alpha + \beta \times \text{Outage}_{d,t} + \theta \times \text{Attention}_{i,d,t} + \gamma \times \text{Controls}_{i,d-1} + \text{Fixed Effect}_i + \epsilon_{i,d,t}$$

where the dependent variable is the log trading frequency or log dollar volume for cryptocurrency $i$ during the 5-minute interval $t$ on day $d$. The indicator $\text{Outage}_{d,t}$ equals one if interval $t$ on day $d$ falls within a ChatGPT outage event window, and zero otherwise. Attention is the total number of postings from Ravenpack for cryptocurrency $i$ during the 5-minute interval $t$ on day $d$. We construct two attention measures: (1) postings from all sources, and (2) postings from only the top 20 sources by posting frequency in our sample. Control variables include the logarithm of the previous day's market capitalization and return, the average return from days $d-5$ to $d-2$, and the logarithm of the previous day's closing price. Each specification includes coin and interval fixed effects. Standard errors are clustered at the coin and internal levels. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively.

| | Postings from All Sources | | | | | Postings from the Top 20 Sources | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | $\beta$ | t-stat | $\theta$ | t-stat | Adj. $R^2$ | $\beta$ | t-stat | $\theta$ | t-stat | Adj. $R^2$ | Obs |
| Trading Frequency | 0.208*** | (6.61) | 0.017*** | (6.27) | 0.317 | 0.215*** | (6.82) | 0.014*** | (3.75) | 0.316 | 159,722 |
| Dollar Volume | 0.261*** | (7.35) | 0.028*** | (7.83) | 0.347 | 0.272*** | (7.56) | 0.027*** | (5.72) | 0.346 | 159,722 |

## Table 9. ChatGPT Outages and Changes in Order Imbalance and Trade Imbalance

Panel A presents the cross-sectional means and medians of buy volume, sell volume, order imbalance, and trade imbalance during the control and outage periods, along with their logarithmic changes (Log Diff in %). Intraday data for all variables are aggregated at the 5-minute level. For each outage event, we identify its actual time window and define the same time window from 1-5 days before the outage date (e.g., [−5,−1]) as the control period. The results for alternative control periods, [−10,−1] and [−15,−1], are not tabulated for brevity. Order imbalance is the difference between buyer-initiated and seller-initiated volume scaled by total volume. Trade imbalance is the absolute value of order imbalance. We compute the average of each variable across all outage event (control) periods for each cryptocurrency pair, and then obtain the cross-sectional summary statistics. The significance of the change is assessed using a t-test, the nonparametric Wilcoxon signed-rank test, and bootstrap simulations. Panel B presents estimates from the following panel regression:

$$\log(\text{Dep. Variable})_{i,d,t} = \alpha + \beta \times \text{Outage}_{d,t} + \gamma \times \text{Controls}_{i,d-1} + \text{Fixed Effects}_{d,t} + \epsilon_{i,d,t}$$

where the dependent variable is buy volume, sell volume, order imbalance, or trade imbalance for cryptocurrency $i$ during the 5-minute interval $t$ on day $d$. The indicator $\text{Outage}_{d,t}$ equals one if interval $t$ on day $d$ falls within a ChatGPT outage event window, and zero otherwise. Control variables include the logarithm of the previous day's market capitalization and return, the average return from days $d-5$ to $d-2$, and the logarithm of the previous day's closing price. Each specification includes coin and interval fixed effects. Standard errors are clustered at the coin and internal level. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively.

Panel A. Univariate Analysis

| | Buy Volume (n $1,000) | | | Sell Volume (in $1,000) | | | Order Imbalance (in %) | | | Trade Imbalance (in %) | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Before | Outage | Log Diff. (%) | Before | Outage | Log Diff. (%) | Before | Outage | Diff. | Before | Outage | Log Diff. (%) |
| Mean | 255.76 | 308.66 | 22.76*** | 254.67 | 315.16 | 25.30*** | 1.23 | 1.71 | 0.48 | 40.03 | 37.76 | -7.65*** |
| S.E. | (141.48) | (161.41) | (2.24) | (140.47) | (164.06) | (2.23) | (0.83) | (0.81) | (0.63) | (2.36) | (2.40) | (0.80) |
| Median | 30.55 | 42.09 | 23.20 | 31.07 | 43.40 | 26.19 | -0.13 | 0.24 | 0.00 | 33.81 | 31.14 | -7.27 |
| Wilcoxon p-value | | | (0.00) | | | (0.00) | | | (0.85) | | | (0.00) |
| Bootstrap p-value | | | (0.00) | | | (0.00) | | | (0.45) | | | (0.00) |

Panel B. Panel Regression

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
|---|---|---|---|---|---|---|---|---|
| | | | | | [-5, -1] | | | |
| | Buy Volume | | Sell Volume | | Order Imbalance | | Trade Imbalance | |
| $Outage_{d,t}$ | 0.293*** | 0.293*** | 0.282*** | 0.282*** | 0.278 | 0.313 | -0.073*** | -0.078*** |
| | (7.71) | (7.90) | (7.49) | (7.64) | (0.41) | (0.45) | (-4.23) | (-4.35) |
| $log(MarketCap)_{d-1}$ | | 0.648 | | 0.560 | | -3.549 | | -0.609 |
| | | (0.82) | | (0.65) | | (-0.31) | | (-1.15) |
| $Ret_{d-1}$ | | 2.002*** | | 2.319*** | | -11.780* | | -0.334* |
| | | (4.26) | | (4.51) | | (-1.93) | | (-1.97) |
| $Ret_{d-2,d-5}$ | | 2.253*** | | 2.361*** | | -1.991 | | -0.660*** |
| | | (5.29) | | (5.40) | | (-0.43) | | (-3.88) |
| $Price_{d-1}$ | | 0.214 | | 0.288 | | 3.893 | | 0.558 |
| | | (0.25) | | (0.31) | | (0.32) | | (1.03) |
| Observations | 177,097 | 177,097 | 175,271 | 175,271 | 178,444 | 178,444 | 178,444 | 178,444 |
| Coin Pair FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Interval FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Adj. R-squared | 0.281 | 0.301 | 0.276 | 0.295 | 0.011 | 0.011 | 0.084 | 0.086 |

**Table 10. Changes in Trading Volume by Trade Size Group**

This table presents the cross-sectional means of total trading frequency and total dollar volume by trade-size group during control and outage periods, along with their logarithmic changes (Log Diff in %). Intraday data for all variables are aggregated at the 5-minute level. For each outage event, we identify its actual time window and define the same time window from 1-5 days before the outage date (e.g., [−5,−1]) as the control period. The results for alternative control periods, [−10,−1] and [−15,−1], are not tabulated for brevity. For each cryptocurrency pair and trade-size group, we compute total trading frequency (in million trades) and total dollar volume (in billion dollars) across all outage (control) periods, and then obtain the cross-sectional means. The last two rows test the significance of the difference in volume change between the largest and smallest trade-size groups using a t-test and bootstrap simulations; the corresponding p-values are reported.

| | Trading Frequency (in million trades) | | | Dollar Volume (in billion dollars) | | |
|---|---|---|---|---|---|---|
| Trade Size Group | Control | Outage | Log Diff (%) | Control | Outage | Log Diff (%) |
| (0, $100] (Small) | 11.34 | 13.93 | 23.74 | 0.33 | 0.39 | 22.29 |
| ($100, $200] | 2.23 | 3.12 | 24.50 | 0.33 | 0.45 | 24.48 |
| ($200, $300] | 1.36 | 1.72 | 21.59 | 0.34 | 0.43 | 21.57 |
| ($300, $500] | 1.66 | 2.00 | 22.37 | 0.66 | 0.79 | 22.70 |
| ($500, $1,000] | 2.14 | 2.75 | 24.59 | 1.57 | 2.02 | 24.69 |
| ($1,000, $2,000] | 1.82 | 2.30 | 28.34 | 2.59 | 3.29 | 29.13 |
| ($2,000, $3,000] | 0.71 | 0.93 | 29.74 | 1.74 | 2.29 | 30.08 |
| ($3,000, $5,000] | 0.67 | 0.80 | 27.16 | 2.59 | 3.12 | 27.08 |
| ($5,000, $10,000] | 0.45 | 0.59 | 22.46 | 3.12 | 4.06 | 22.34 |
| >$10000 (Large) | 0.26 | 0.31 | 23.68 | 5.21 | 6.54 | 21.64 |
| Large - Small | | | -0.06 | | | -0.65 |
| T-test (p-value) | | | (0.99) | | | (0.92) |
| Bootstrap-test (p-value) | | | (0.99) | | | (0.92) |

**Table 11. Robustness Tests**

This table assesses the robustness of our results. The baseline regression refers to the specification (10). For brevity, we only report the results using the control period of [-5, -1] and coefficients on *Outage*. Panel A reports the Poisson regression results using trading frequency as the dependent variable to address the integer nature of the variable. Panel B presents the results using a median regression or using the subsample excluding 2.5% bottom/top residuals. We follow Machado and Silva (2019) to compute standard errors clustered at coin pair and time intervals. Panel C reports the regression results excluding the shortest outage or the four shortest outages. Panel D presents the regression results using (1) Outage events × Exchange-Coin pair and Outage events × time interval fixed effects, and (2) Outage events × Exchange-coin pair × time interval fixed effects. Panel E displays the placebo test results. We randomly assign one day in the control period of [-5, -1] as the outage event day for each event and then redo all tests. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively.

Panel A. Poisson Regression (Trading Frequency)

| | Outage | T-stat | FE |
|---|---|---|---|
| [-5,-1] | 0.233*** | (5.28) | Coin Pair + Interval |
| [-10,-1] | 0.153*** | (3.61) | Coin Pair + Interval |
| [-15,-1] | 0.160*** | (3.87) | Coin Pair + Interval |

Panel B. Sensitivity to Outliers

| | Median Regression | | Excluding 2.5% | |
|---|---|---|---|---|
| | Outage | T-stat | Outage | T-stat |
| Trading Frequency | 0.222*** | (29.05) | 0.229*** | (9.84) |
| Dollar Volume | 0.277*** | (27.98) | 0.270*** | (10.25) |
| Return Volatility (trade) | 0.190*** | (31.09) | 0.177*** | (9.13) |
| Return Volatility (OB) | 0.199*** | (32.81) | 0.183*** | (9.45) |
| Depth of -0.1% Bid | -0.089*** | (-23.67) | -0.092*** | (-7.17) |
| Depth of -0.1% Ask | -0.077*** | (-20.91) | -0.080*** | (-6.20) |
| Quoted Spread | 0.065*** | (16.39) | 0.058*** | (4.91) |
| Price Impact | 0.150*** | (14.17) | 0.149*** | (6.34) |

Panel C. Excluding Short Events

| | Exclude the shortest event | | Exclude four shortest events | |
|---|---|---|---|---|
| | Outage | T-stat | Outage | T-stat |
| Trading Frequency | 0.222*** | (7.25) | 0.235*** | (7.50) |
| Dollar Volume | 0.276*** | (8.06) | 0.296*** | (8.53) |
| Return Volatility (trade) | 0.190*** | (7.76) | 0.198*** | (8.00) |
| Return Volatility (OB) | 0.198*** | (8.22) | 0.206*** | (8.42) |
| Depth of -0.1% Bid | -0.087*** | (-5.13) | -0.090*** | (-5.29) |
| Depth of +0.1% Ask | -0.075*** | (-5.39) | -0.077*** | (-5.47) |
| Quoted Spread | 0.065*** | (4.31) | 0.068*** | (4.32) |
| Price Impact | 0.146*** | (4.99) | 0.160*** | (5.46) |

Panel D. Alternative Fixed Effects

| | Event * Coin Pair+Event * Interval | | Event * Coin Pair * Interval | |
|---|---|---|---|---|
| | Outage | T-stat | Outage | T-stat |
| Trading Frequency | 0.191*** | (6.25) | 0.191*** | (6.24) |
| Dollar Volume | 0.255*** | (7.29) | 0.253*** | (7.30) |
| Return Volatility (trade) | 0.155*** | (6.29) | 0.154*** | (6.27) |
| Return Volatility (OB) | 0.161*** | (6.65) | 0.160*** | (6.62) |
| Depth of -0.1% Bid | -0.043** | (-2.18) | -0.043** | (-2.17) |
| Depth of +0.1% Ask | -0.038** | (-2.32) | -0.038** | (-2.28) |
| Quoted Spread | 0.030** | (2.42) | 0.029** | (2.36) |
| Price Impact | 0.131*** | (4.28) | 0.130*** | (4.25) |

Panel E. Pseudo Outages

| | Pseudo Outage | T-stat | FE |
|---|---|---|---|
| Trading Frequency | -0.009 | (-0.26) | Coin Pair + Interval |
| Dollar Volume | -0.016 | (-0.39) | Coin Pair + Interval |
| Return Volatility (trade) | 0.025 | (0.94) | Coin Pair + Interval |
| Return Volatility (OB) | 0.028 | (1.00) | Coin Pair + Interval |
| Depth of -0.1% Bid | 0.032 | (1.66) | Coin Pair + Interval |
| Depth of +0.1% Ask | 0.038 | (1.58) | Coin Pair + Interval |
| Quoted Spread | -0.000 | (-0.02) | Coin Pair + Interval |
| Price Impact | 0.032 | (1.03) | Coin Pair + Interval |

# Appendix A

### Table A1. 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 August 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 | 2023/2/27 15:00 | 2023/2/27 18:30 | 210 mins |
| ChatGPT Outage | 2023/3/15 17:55 | 2023/3/15 18:18 | 23 mins |
| Outage on chat.openai.com | 2023/3/20 9:18 | 2023/3/20 10:17 | 59 mins |
| ChatGPT Web Interface Incident | 2023/3/20 16:41 | 2023/3/21 2:54 | 613 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 |
| ChatGPT web & mobile UIs unavailable | 2023/8/29 14:04 | 2023/8/29 15:16 | 72 mins |
| Degraded performance on ChatGPT | 2023/8/31 13:52 | 2023/8/31 18:03 | 251 mins |
| ChatGPT severely degraded | 2023/9/1 4:36 | 2023/9/1 5:11 | 35 mins |
| Major Outage across ChatGPT and API | 2023/11/8 13:54 | 2023/11/8 15:46 | 112 mins |
| Labs, API, ChatGPT, and Platform.openai.com are inaccessible | 2023/11/12 5:57 | 2023/11/12 6:44 | 47 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 |
| Partial outage for GPT-4o free users | 2024/5/21 18:40 | 2024/5/21 19:48 | 68 mins |
| Outage on ChatGPT and API Platform | 2024/7/5 22:23 | 2024/7/5 22:42 | 19 mins |
| Outage on ChatGPT and Assistants API | 2024/7/12 10:03 | 2024/7/12 10:52 | 49 mins |
| ChatGPT outage | 2024/7/17 8:18 | 2024/7/17 10:18 | 120 mins |

## Table A2. 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 August 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 |

# Appendix B

**Proof of Proposition 2**

If traders wait, they receive no signal and there is no trading in period 1. In period 2, with probability $\lambda$, ChatGPT is restored, and trading yields an expected profit of $\overline{\pi}$. The expected payoff from waiting is $\frac{1}{1+r}\lambda\overline{\pi}$.

If traders switch, they incur a cost $C$ and earn expected trading profits of $\underline{\pi}$ in period 1. In period 2, the expected profit is $\lambda\overline{\pi} + (1-\lambda)\underline{\pi}$. Thus, the total discounted payoff is

$$\underline{\pi} + \frac{1}{1+r}\left[\lambda\overline{\pi} + (1-\lambda)\underline{\pi}\right] - C.$$

All traders will choose to switch if the total expected discounted payoff from switching exceeds that from waiting, i.e.,

$$\underline{\pi} + \frac{1}{1+r}\left[\lambda\overline{\pi} + (1-\lambda)\underline{\pi}\right] - C > \frac{1}{1+r}\lambda\overline{\pi}.$$

Simplifying yields the condition $C < \frac{2+r-\lambda}{1+r}\underline{\pi}$, as stated.

**Proof of Proposition 3** The trading volume in equilibrium is given by

$$E\left[\sum_{i=1}^{N}|x_{i,t}^*|\right] = \sum_{i=1}^{N}E\left[|x_{i,t}^*|\right] = NE\left[|x_{i,t}^*|\right] = N\sqrt{\frac{2}{\pi}}\sqrt{\mathbb{V}ar(x_{i,t}^*)}, \tag{14}$$

where the last equality uses the property that if $x \sim N(0,\sigma^2)$, then $E[|x|] = \sqrt{2/\pi}\,\sigma$.

From Proposition 1, in equilibrium

$$x_{i,t}^* = \frac{1}{\phi}(E[v_t|\overline{s}_{i,t}] - E[v_t|\tilde{s}_t]),$$

where the conditional means can be calculated using the following lemma (see also Grossman and Stiglitz (1980)):

**Lemma 1.** *If $X^*$ and $Y^*$ are jointly normally distributed,*

$$E[X^*|Y^* = Y] = EX^* + \frac{Cov(X^*, Y^*)}{Var(Y^*)}(Y - EY^*), \tag{15}$$

$$Var[X^*|Y^* = Y] = Var(X^*) - \frac{(Cov(X^*, Y^*))^2}{Var(Y^*)}. \tag{16}$$

Applying Lemma 1, we get

$$E[v_t|\overline{s}_{i,t}] = v_0 + \frac{\mathrm{Cov}(\overline{s}_{i,t}, v_t)}{\mathbb{V}ar(\overline{s}_{i,t})}(\overline{s}_{i,t} - v_0), \tag{17}$$

$$E[v_t|\tilde{s}_t] = v_0 + \frac{\mathrm{Cov}(\tilde{s}_t, v_t)}{\mathbb{V}ar(\tilde{s}_t)}(\tilde{s}_t - v_0). \tag{18}$$

With $\tilde{s}_t = \frac{1}{N}\sum_{i=1}^{N} \overline{s}_{i,t}$ and $E[\tilde{s}_t] = E[\overline{s}_{i,t}] = v_0$, we have $\mathrm{Cov}(\overline{s}_{i,t}, v_t) = \mathrm{Cov}(\tilde{s}_t, v_t) = 1/\tau_0$, and

$$\mathbb{V}ar(\overline{s}_{i,t}) = \frac{1}{\tau_0} + \frac{1}{\overline{n}\tau_s}, \tag{19}$$

$$\mathbb{V}ar(\tilde{s}_t) = \frac{1}{\tau_0} + \frac{1}{\overline{n}N\tau_s}\big[1 + \rho(N-1)\big]. \tag{20}$$

Since $\tilde{s}_t - v_0 = \frac{1}{N}\sum_{i=1}^{N}(\overline{s}_{i,t} - v_0)$, it follows that

$$\mathrm{Cov}(\overline{s}_{i,t} - v_0, \tilde{s}_t - v_0) = \mathbb{V}ar(\tilde{s}_t) \tag{21}$$

by homogeneity. Therefore, the variance of $x^*_{i,t}$ is

$$\mathbb{V}ar(x^*_{i,t}) = \frac{1}{\phi^2\tau_0^2}\left[\frac{1}{\mathbb{V}ar(\overline{s}_{i,t})} + \frac{1}{\mathbb{V}ar(\tilde{s}_t)} - 2\frac{\mathbb{V}ar(\tilde{s}_t)}{\mathbb{V}ar(\overline{s}_{i,t})\mathbb{V}ar(\tilde{s}_t)}\right] = \frac{1}{\phi^2\tau_0^2}\left[\frac{1}{\mathbb{V}ar(\tilde{s}_t)} - \frac{1}{\mathbb{V}ar(\overline{s}_{i,t})}\right]. \tag{22}$$

Plugging in the expressions,

$$\mathbb{V}ar(x^*_{i,t}) = \frac{1}{\phi^2\tau_0^2}\left[\frac{N\overline{n}\tau_0\tau_s}{\tau_0[1 + \rho(N-1)] + \overline{n}\tau_s N} - \frac{\overline{n}\tau_0\tau_s}{\tau_0 + \overline{n}\tau_s}\right]. \tag{23}$$

Thus, the trading volume in equilibrium with ChatGPT is

$$\frac{N}{\phi\tau_0}\sqrt{\frac{2}{\pi}}\left[\frac{N\overline{n}\tau_0\tau_s}{\tau_0[1+\rho(N-1)]+\overline{n}\tau_s N}-\frac{\overline{n}\tau_0\tau_s}{\tau_0+\overline{n}\tau_s}\right]^{1/2}. \tag{24}$$

Without ChatGPT (independent signal generation), replace $\overline{n}$ by $\underline{n}$ and $\rho = 0$:

$$\frac{N}{\phi\tau_0}\sqrt{\frac{2}{\pi}}\left[\frac{N\underline{n}\tau_0\tau_s}{\tau_0+\underline{n}\tau_s N}-\frac{\underline{n}\tau_0\tau_s}{\tau_0+\underline{n}\tau_s}\right]^{1/2}. \tag{25}$$

The return volatility is, by definition,

$$\mathbb{V}ar(p_t - v_0) = \mathbb{V}ar(\mathbb{E}[v_t|\tilde{s}_t]) = \frac{1}{\tau_0^2}\frac{1}{\mathbb{V}ar(\tilde{s}_t)} = \frac{1}{\tau_0}\frac{\overline{n}N\tau_s}{\overline{n}N\tau_s+\tau_0[1+\rho(N-1)]}, \tag{26}$$

and in the independent case,

$$\frac{1}{\tau_0}\frac{\underline{n}N\tau_s}{\underline{n}N\tau_s+\tau_0}. \tag{27}$$

For market depth, with the linear pricing rule for the market maker, Lemma 1 yields

$$E[v_t|y_t] = v_0 + \frac{\mathrm{Cov}(v_t, y_t)}{\mathbb{V}ar(y_t)}(y_t - E[y_t]), \tag{28}$$

with

$$\begin{aligned} y_t &= \frac{1}{\phi}\left(\sum_{i=1}^{N} E[v_t|\overline{s}_{i,t}] - Np_t\right) \\ &= \frac{1}{\phi}\left(\frac{N\tau_0 v_0}{\tau_0+\overline{n}\tau_s}+\frac{\overline{n}\tau_s}{\tau_0+\overline{n}\tau_s}\sum_{i=1}^{N}\overline{s}_{i,t} - Np_t\right). \end{aligned} \tag{29}$$

We then have

$$\mathrm{Cov}(v_t, y_t) = \frac{1}{\phi\tau_0}\frac{N\overline{n}\tau_s}{\tau_0 + \overline{n}\tau_s}, \tag{30}$$

$$\mathbb{V}ar(y_t) = \frac{1}{\phi^2}\left(\frac{\overline{n}\tau_s}{\tau_0 + \overline{n}\tau_s}\right)^2 \left\{\frac{N^2}{\tau_0} + \frac{N}{\overline{n}\tau_s}[1 + \rho(N-1)]\right\}. \tag{31}$$

Hence, market depth, as the inverse of price impact, is

$$\frac{\mathbb{V}ar(y_t)}{\mathrm{Cov}(v_t, y_t)} = \frac{1}{\phi}\frac{\overline{n}\tau_s}{\tau_0 + \overline{n}\tau_s}\left\{N + \frac{\tau_0}{\overline{n}\tau_s}[1 + \rho(N-1)]\right\}. \tag{32}$$

In the case of independent data generation:

$$\frac{1}{\phi}\frac{\underline{n}\tau_s}{\tau_0 + \underline{n}\tau_s}\left(N + \frac{\tau_0}{\underline{n}\tau_s}\right). \tag{33}$$

**Proof of Corollary 1**

**Trading Volume:** To compare trading volumes, define

$$F_1(\tau_s) := \frac{N\overline{n}\tau_0\tau_s}{\tau_0[1 + \rho(N-1)] + \overline{n}\tau_s N} - \frac{\overline{n}\tau_0\tau_s}{\tau_0 + \overline{n}\tau_s} - \left[\frac{N\underline{n}\tau_0\tau_s}{\tau_0 + \underline{n}\tau_s N} - \frac{\underline{n}\tau_0\tau_s}{\tau_0 + \underline{n}\tau_s}\right]. \tag{34}$$

To determine the sign near $\tau_s = 0$, calculate the derivative:

$$\begin{aligned} F_1'(\tau_s) = {} & \frac{N\overline{n}\tau_0^2[1 + \rho(N-1)]}{[\tau_0[1 + \rho(N-1)] + \overline{n}\tau_s N]^2} - \frac{\overline{n}\tau_0^2}{(\tau_0 + \overline{n}\tau_s)^2} \\ & - \left[\frac{N\underline{n}\tau_0^2}{(\tau_0 + \underline{n}\tau_s N)^2} - \frac{\underline{n}\tau_0^2}{(\tau_0 + \underline{n}\tau_s)^2}\right]. \end{aligned} \tag{35}$$

At $\tau_s = 0$,

$$F_1'(0) = (N-1)\left[\frac{1-\rho}{1 + \rho(N-1)}\overline{n} - \underline{n}\right]. \tag{36}$$

This is negative under the regularity conditions, so for small $\tau_s$,

$$F_1(\tau_s) \approx F_1'(0)\tau_s < 0. \tag{37}$$

**Return Volatility:** Similarly, define the difference

$$F_2(\tau_s) = \frac{1}{\tau_0}\frac{\overline{n}N\tau_s}{\overline{n}N\tau_s + \tau_0[1+\rho(N-1)]} - \frac{1}{\tau_0}\frac{\underline{n}N\tau_s}{\underline{n}N\tau_s + \tau_0}. \tag{38}$$

With derivative

$$F_2'(\tau_s) = \frac{\overline{n}N\tau_0[1+\rho(N-1)]}{(\overline{n}N\tau_s + \tau_0[1+\rho(N-1)])^2} - \frac{\underline{n}N\tau_0}{(\underline{n}N\tau_s + \tau_0)^2}, \tag{39}$$

evaluated at $\tau_s = 0$:

$$F_2'(0) = \frac{N}{\tau_0}\left(\frac{\overline{n}}{1+\rho(N-1)} - \underline{n}\right). \tag{40}$$

This is negative under the regularity conditions, so for small $\tau_s$,

$$F_2(\tau_s) \approx F_2'(0)\tau_s < 0. \tag{41}$$

**Market Depth:** Define the difference

$$F_3(\tau_s) = \frac{1}{\phi}\frac{\overline{n}\tau_s}{\tau_0 + \overline{n}\tau_s}\left\{N + \frac{\tau_0}{\overline{n}\tau_s}[1+\rho(N-1)]\right\} - \frac{1}{\phi}\frac{\underline{n}\tau_s}{\tau_0 + \underline{n}\tau_s}\left(N + \frac{\tau_0}{\underline{n}\tau_s}\right). \tag{42}$$

Its derivative is

$$F_3'(\tau_s) = \frac{1}{\phi}\frac{\overline{n}\tau_0(1-\rho)(N-1)}{(\tau_0 + \overline{n}\tau_s)^2} - \frac{1}{\phi}\frac{\underline{n}\tau_0(N-1)}{(\tau_0 + \underline{n}\tau_s)^2}, \tag{43}$$

so at $\tau_s = 0$,

$$F_3'(0) = \frac{(N-1)}{\phi\tau_0}\left[\overline{n}(1-\rho) - \underline{n}\right]. \tag{44}$$

This is positive under the regularity conditions, so for small $\tau_s$,

$$F_3(\tau_s) \approx F_3'(0)\tau_s > 0. \tag{45}$$