

Azar, Pablo; Olivas, Sergio; Sinha, Nish D.

Working Paper

The price of processing: Information frictions and market efficiency in DeFi

Staff Reports, No. 1153

Provided in Cooperation with:

Federal Reserve Bank of New York

Suggested Citation: Azar, Pablo; Olivas, Sergio; Sinha, Nish D. (2025) : The price of processing: Information frictions and market efficiency in DeFi, Staff Reports, No. 1153, Federal Reserve Bank of New York, New York, NY,
<https://doi.org/10.59576/sr.1153>

This Version is available at:

<https://hdl.handle.net/10419/324782>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

NO. 1153
APRIL 2025

The Price of Processing: Information Frictions and Market Efficiency in DeFi

Pablo D. Azar | Sergio Olivas | Nish D. Sinha

The Price of Processing: Information Frictions and Market Efficiency in DeFi

Pablo D. Azar, Sergio Olivas, and Nish D. Sinha

Federal Reserve Bank of New York Staff Reports, no. 1153

April 2025

<https://doi.org/10.59576/sr.1153>

Abstract

This paper investigates the speed of price discovery when information becomes publicly available but requires costly processing to become common knowledge. We exploit the unique institutional setting of hacks on decentralized finance (DeFi) protocols. Public blockchain data provides the precise time a hack's transactions are recorded—becoming public information—while subsequent social media disclosures mark the transition to common knowledge. This empirical design allows us to isolate the price impact occurring during the interval characterized by information asymmetry driven purely by differential processing capabilities. Our central empirical finding is that substantial price discovery precedes common knowledge: approximately 36 percent of the total 24-hour price decline (~27 percent) materializes *before* the public announcement. This evidence suggests sophisticated traders rapidly exploit their ability to process complex, publicly available on-chain data, capturing informational rents. We develop a theoretical model of informed trading under processing costs which predicts strategic, slow information revelation, consistent with our empirical findings. Our results quantify the limits imposed by information processing costs on market efficiency, demonstrating that transparency alone does not guarantee immediate information incorporation into prices.

JEL classification: G12, G14, G18, G23, L86

Key words: information asymmetry, price discovery, common knowledge, information processing costs, market microstructure, event study, high-frequency data, cryptocurrency, defi, cybersecurity hacks, market efficiency

Azar, Sinha: Federal Reserve Bank of New York (email: pablo.azar@ny.frb.org, nish.sinha@ny.frb.org).
Olivas: University of Texas at Austin (email: solivas@utexas.edu). The authors thank Owen Engbretson for helpful comments and discussions.

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the author(s) and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the author(s).

To view the authors' disclosure statements, visit
https://www.newyorkfed.org/research/staff_reports/sr1153.html.

1 Introduction

A central question in financial economics concerns *price discovery*—the process by which new information becomes embedded in asset prices. While canonical models often depict information arriving continuously or through predictable channels, many real-world events manifest as discrete, initially opaque shocks. Understanding how prices respond during the interval between an event’s hidden occurrence and its eventual public revelation is crucial for assessing market efficiency. Yet, empirically isolating this interval has proven challenging due to the difficulty in precisely observing both the timing of private information generation and its transition into the public domain.

This paper studies this question within the context of hacks in decentralized finance (DeFi). DeFi protocols aim to provide financial services, such as lending or automated trading, directly on a blockchain, bypassing traditional intermediaries. Many protocols issue equity-like governance tokens whose market value reflects the protocol’s perceived health, user base, and expected future earnings, akin to traditional equity claims. However, these protocols operate by holding user-deposited assets (collateral) within smart contracts, making them targets for cybersecurity breaches, or hacks. During a hack, attackers typically exploit vulnerabilities in the protocol’s code to illicitly drain substantial amounts of this collateral, representing a direct and often significant loss to the protocol’s value and undermining user confidence. These hacks constitute sudden, adverse shocks whose informational impact is the focus of our study.

The informational environment surrounding these hacks provides a valuable research setting. Most DeFi protocols operate on public blockchains like Ethereum, where all transactions are recorded on an immutable, publicly accessible ledger. This radical transparency means the raw data concerning a hack becomes *public information* almost instantaneously when malicious transactions are confirmed on the chain. A naive interpretation might suggest this implies strong-form market efficiency, where prices immediately reflect all available information (Fama, 1970).

We show that there’s a critical distinction between information being publicly available and it being *common knowledge*. Interpreting raw blockchain data—identifying a specific sequence of transactions as a value-destroying hack amidst millions of legitimate ones—requires significant technical expertise, continuous monitoring infrastructure, and substantial computational effort. These *information processing costs* create a barrier: while the *data* is public, understanding its economic significance is initially limited to sophisticated actors (e.g., specialized crypto hedge funds, security firms, the hackers themselves) who have invested in the necessary resources. For the broader market, the information content remains opaque until it is distilled into a simpler, easily digestible signal.

In contrast, information disseminated through widely followed channels like official X¹ an-

¹X is the social media website formerly known as Twitter.

nouncements rapidly overcomes these processing hurdles and becomes *common knowledge*. This creates a distinct timeline and a period of *information asymmetry* driven not by access, but by processing capability: the initial on-chain hack transaction (at time t_0) marks the arrival of complex *public information*, while the subsequent public announcement (typically via X, at time $t_{common} > t_0$) marks the transition of this information into easily processed *common knowledge*.




Figure 1: Information Flow and Price Discovery During a DeFi Hack. This figure illustrates the stylized timeline of information arrival and price reaction following a DeFi hack. At time t_0 , the malicious transaction occurs on-chain, making the event technically *public information* but requiring significant processing cost (expertise, monitoring) to interpret. At time t_{common} , a public announcement (e.g., via social media) makes the information easily accessible and transforms it into *common knowledge* (low processing cost). The shaded area represents the information asymmetry window created by these processing costs. During this window, sophisticated traders who can process the complex on-chain data begin trading, leading to significant pre-announcement price discovery (empirically estimated at ~57% of the total impact) before the broader market reacts to the common knowledge signal.

This unique informational structure allows us to ask: how much of the price impact occurs *after* information becomes public (t_0) but *before* it becomes common knowledge (t_{common})? As illustrated in Figure 1, DeFi hacks offer a rare controlled setting where we know the exact timing of a significant information event and the timing of its public disclosure, contrasting sharply with traditional finance where such precise timing is often elusive. This setting lets us isolate the ‘hidden’ information window—where information is public but not common knowledge—in a way

not possible with most corporate news or macro announcements, making it a powerful empirical testbed for theories of information asymmetry and price discovery under processing costs.

We perform a high-frequency event study to show that a substantial portion of the price adjustment occurs before the public announcement makes the information common knowledge. While the total negative impact following the initial hack averages approximately -27% over 24 hours for the protocol’s governance token, examining cumulative abnormal returns relative to the announcement time shows significant negative returns hours prior. For instance, cumulative abnormal returns reach -9.5% twelve hours *before* the announcement. This indicates that sophisticated actors, by expending resources to process the complex on-chain data, impound approximately 36% of the total price impact before the news becomes easily digestible for the broader public. This rapid pre-announcement price discovery underscores the economic value of specialized processing capabilities and provides quantitative evidence on the limits of market efficiency imposed by information processing costs, even in transparent blockchain environments.

1.1 Broader Financial Stability Concerns

The findings in this paper extend beyond the specialized domain of digital assets, speaking directly to emerging financial stability concerns. The lines between decentralized and traditional finance are increasingly blurring. Regulated investment vehicles like spot Bitcoin exchange-traded funds now provide broad investor access. Major financial institutions are actively tokenizing real-world assets on blockchains, exemplified by initiatives like BlackRock’s BUIDL fund. Furthermore, comprehensive regulatory frameworks for stablecoins and crypto-assets, such as Europe’s MiCA, are integrating these markets more formally into the conventional financial structure.

Against this backdrop of increasing integration, the cybersecurity vulnerabilities inherent in DeFi protocols—manifesting as frequent, large-scale hacks—represent a primary source of instability. Because hacks can cause large, rapid price shifts in governance tokens, understanding the precise market mechanics during these crises—how quickly adverse information is processed, who processes it, and how prices react *before* events become common knowledge—is crucial not only for crypto participants but also for policymakers, regulators, and traditional financial institutions monitoring potential systemic risks emanating from this evolving technological frontier. The growing linkages mean shocks originating in DeFi could increasingly spill over into the traditional system. This paper offers a crucial lens into these dynamics by measuring the speed of information incorporation and highlighting the real-world impact of processing costs on market efficiency.

1.2 Related Work

A rich literature in finance and economics examines how information asymmetries affect price discovery. In this review, we organize the discussion thematically, covering foundational theories, modern strategic trading models, the distinction between public information and common knowledge, empirical evidence on informed trading with hard-to-process public signals, and recent insights from cryptocurrency and decentralized finance (DeFi) markets.

Classic Foundations of Price Discovery and Asymmetric Information Early work established the mechanisms by which private information becomes reflected in asset prices. Grossman and Stiglitz (1980) argue that if prices fully reveal all private information, informed traders would earn no compensation, thereby precluding an informationally efficient equilibrium. Building on this insight, Kyle (1985) develops a dynamic insider trading model in which a monopolistic informed trader optimally discloses information over time, while market makers set prices based on the total order flow. Similarly, Glosten and Milgrom (1985) present a sequential trade model in which a specialist infers private information from order flows, resulting in endogenous bid-ask spreads. Both models emphasize that private information is *gradually* incorporated into prices through strategic trading. While these models provide the bedrock for understanding information asymmetry, our setting allows for an unusually precise empirical examination of the speed of this process when information is technically public but costly to process, transitioning later to low-cost common knowledge.

Dynamic and Strategic Trading: Modern Extensions Subsequent research has enriched these early models by incorporating multiple trading periods and strategic considerations. Extensions of Kyle (1985) (see, e.g., Foster and Viswanathan (1996) and Back (1992)) consider multiple informed traders and intertemporal information extraction, demonstrating that competition among informed agents can accelerate price discovery while simultaneously compressing individual profits. Meanwhile, Duffie et al. (2005) develop an over-the-counter market framework where search frictions and bilateral trading delay the aggregation of private information. In parallel, the rise of high-frequency trading has spurred studies by Budish et al. (2015) and Brogaard et al. (2014), who show that technological advancements in trading speed can both enhance and, paradoxically, impede efficient price discovery.

Public Information vs. Common Knowledge A crucial nuance in the literature is the distinction between *public information* and *common knowledge*. Although public information is available to all, it becomes common knowledge only after it is widely recognized and acknowledged through

iterative confirmation. Morris and Shin (2002) formalize this difference, showing that the transition from public information to common knowledge can significantly influence market dynamics. Further work by Allen et al. (2006) further demonstrates that when information remains public but not common, market participants who process it earlier can exploit the delay, leading to gradual price adjustments until a formal announcement is made. Our paper directly leverages this distinction, using the precisely timed on-chain hack as the arrival of complex public information and the subsequent announcement as the transition to common knowledge.

Informed Trading on Hard-to-Process Public Information In many cases, public information is available in principle but requires sophisticated analysis to extract its economic significance. Tetlock (2007) provides evidence that nuanced media sentiment predicts short-term price movements, implying that only those with the ability to process complex public signals can trade profitably. Laboratory experiments by Bloomfield and O’Hara (1999) confirm that even with complete transparency, differences in processing capabilities can create substantial informational advantages. In the era of big data, studies such as Gross et al. (2011) show that machine-readable news feeds enable a subset of traders to react almost instantaneously to new information, capturing profits before the broader market is informed. In spirit, our analysis relates to studies of how private information is traded on before public news, similar to classic insider trading studies, except here the ‘insider’ information is derived legally from skill in parsing public data.

Market Frictions and Price Discovery in Cryptocurrency and DeFi Recent advances in blockchain technology have opened a new avenue to study information asymmetry and price discovery. Despite the inherent transparency of blockchain ledgers, research by Makarov and Schoar (2020) and Liu et al. (2022) documents significant arbitrage opportunities and pricing inefficiencies across cryptocurrency exchanges, suggesting frictions prevent immediate price convergence. Liu et al. (2022) also find that standard asset pricing factors explain crypto returns, implying some market structure parallels, but highlight potential mispricings.

The unique nature of decentralized finance means that on-chain data, while publicly available, must be processed with specialized tools and expertise. Azar et al. (2024) argue that information advantages persist even with ledger transparency due to complexities in the intermediation chain and inherent processing costs. Others explore how blockchain technology and smart contracts alter information diffusion pathways (Cong and He, 2019; Saleh, 2021) or how phenomena like Miner Extractable Value (MEV) create informational rents (e.g., Azar et al., 2024; Capponi et al., 2022; Adams et al., 2024) by allowing certain traders to act on visible, yet complex, order flow information.

Our research also relates directly to the nascent literature examining the economic consequences

of security breaches in digital asset markets. Chen et al. (2023), for example, provide valuable early analysis, documenting significant negative returns on bitcoin and heightened volatility following security incidents, often focusing on events like cryptocurrency exchange hacks using daily frequency data. While demonstrating the detrimental impact of such events, the use of daily data limits the ability to pinpoint intraday price discovery. Our work advances this line of inquiry substantially by employing *high-frequency* (minute-level) market data for DeFi *governance tokens* and, crucially, by precisely separating the timing of the *initial on-chain exploit transaction* from the *subsequent public announcement*. This granular approach allows us to dissect the intraday price discovery process and isolate the market reaction occurring *before* the hack becomes common knowledge, providing direct empirical measurement of the value of early, costly information processing in the specific context of DeFi security breaches and quantifying the extent to which sophisticated actors anticipate public news by overcoming these processing hurdles.

1.3 Roadmap

The remainder of the paper is organized as follows. Section 2 provides institutional details on DeFi protocols and Hacks. Section 3 describes the data and empirical methodology. Section 4 presents the main empirical findings. Section 5 develops our theoretical framework. Section 6 concludes.

2 Institutional Background: DeFi Protocols and Hacks

Decentralized Finance (DeFi) represents a rapidly growing ecosystem aiming to replicate and innovate upon traditional financial services using blockchain technology. Understanding the institutional context, particularly the nature of DeFi protocols and the prevalence of hacks, is crucial for interpreting our empirical findings on information processing.

At its core, DeFi relies on public, permissionless blockchains—distributed ledgers maintained by a network of computers through a consensus mechanism. Blockchains like Ethereum provide not only a means of transferring value (via native cryptocurrencies like Ether) but also a platform for executing *smart contracts*. These are self-executing contracts which take digital assets as inputs and have the terms of financial agreements directly written into code. They automatically perform actions when predefined conditions are met, without the need for traditional intermediaries like banks or brokers.

DeFi protocols leverage smart contracts to offer a wide array of financial services. These include decentralized exchanges (DEXs) that allow peer-to-peer token trading (e.g., Uniswap, Curve), lending and borrowing platforms where users can supply assets to earn interest or borrow against collateral (e.g., Aave, Compound), derivatives platforms, asset management services, and

more. Users interact directly with these protocols' smart contracts, typically depositing digital assets (cryptocurrencies or stablecoins) as collateral or liquidity. Many protocols issue governance tokens, which grant holders voting rights on protocol upgrades and parameter changes, and often accrue a share of protocol revenues, making their market value sensitive to the protocol's performance and security, analogous to equity in a traditional firm.

A defining characteristic of DeFi is its reliance on these smart contracts to custody user funds, often referred to as the protocol's Total Value Locked (TVL). While proponents highlight the transparency and efficiency gains from removing intermediaries, concentrating potentially billions of dollars worth of assets into complex, automated code creates significant security challenges. These challenges manifest most starkly as "DeFi hacks" or "exploits."

A DeFi hack occurs when an attacker finds and leverages a vulnerability to illegitimately withdraw assets from a protocol's smart contracts. These vulnerabilities can arise from various sources. *Smart contract bugs* are flaws in the code itself—logical errors, incorrect assumptions, or unforeseen interactions between different functions—that allow attackers to bypass intended constraints. Given the complexity and novelty of smart contract programming languages (like Solidity on Ethereum) and the financial stakes involved, even audited code can contain subtle, exploitable bugs.

Beyond direct code flaws, attackers often target the broader infrastructure supporting DeFi protocols. *Oracle manipulation* involves exploiting the external data feeds (oracles) that smart contracts rely on for real-world information, such as asset prices. By manipulating the price reported by an oracle, an attacker might trick a lending protocol into issuing loans against undervalued collateral or liquidating positions unfairly. *Cross-chain bridge exploits* target the protocols designed to transfer assets between different blockchains; vulnerabilities in these bridges have led to some of the largest DeFi hacks, as attackers drain the assets locked on one chain that back the wrapped tokens issued on another. Flash loan attacks involve borrowing extremely large sums with no upfront collateral, using the borrowed funds to manipulate market prices or exploit other vulnerabilities within a single atomic transaction, and then repaying the loan, pocketing the difference. Governance attacks involve malicious actors acquiring enough governance tokens to pass proposals that allow them to steal funds.

The scale and frequency of these hacks are substantial. Since the rise of DeFi in 2020, hundreds of significant exploits have occurred, with cumulative losses estimated in the billions of dollars. Prominent examples illustrate the magnitude and variety: the Ronin hack in March 2022 saw over \$600 million stolen due to compromised private keys; the Poly Network exploit in August 2021 involved a similar amount taken via a cross-chain vulnerability (though much was later returned); the Wormhole bridge was exploited for over \$320 million in February 2022; lending protocols like Euler Finance lost nearly \$200 million in March 2023 due to a flash loan and logic vulnerability;

and even established protocols like Curve Finance suffered significant losses in July 2023 due to a vulnerability in a specific version of its programming language compiler.

From an economic perspective, these hacks occur because the potential rewards for attackers are immense, while the costs of securing complex, rapidly evolving systems are high. The "composability" or interconnectedness of DeFi protocols, often lauded as a feature, can also create systemic risks, where a vulnerability in one protocol can cascade to others that rely on it. The fast pace of innovation often means protocols are launched quickly, sometimes deprioritizing exhaustive security audits. Furthermore, the open-source nature of the code, while promoting transparency, also allows potential attackers to scrutinize it for weaknesses.

Each successful hack represents a direct, negative economic shock to the affected protocol. It typically results in a loss of user funds, damages the protocol's reputation, reduces its TVL and fee-generating potential, and consequently leads to a sharp decline in the market value of its governance token. It is the market's reaction to this specific type of negative shock, and particularly the information flow surrounding it, that forms the basis of our empirical investigation.

3 Data and Methodology

Our empirical analysis relies on a novel, hand-collected dataset that combines information on cryptocurrency hack events with high-frequency market data for the affected crypto assets. This section details the data sources, variable construction, and matching procedures.

3.1 Hack Event Data

We compile a comprehensive list of DeFi hack events occurring between June 2016 and January 2024. To identify potential hacks, we draw upon multiple sources, including the public dashboard provided by DeFi Llama (<https://defillama.com/hacks>),² systematic monitoring of reputable blockchain security news outlets and blogs (such as Rekt News, PeckShield Alerts, CertiK Skynet Alerts, and SlowMist), and reporting from major financial news aggregators and crypto-focused media outlets. Each potential event undergoes a detailed verification process, involving cross-referencing information across sources and consulting primary on-chain data via blockchain explorers (e.g., Etherscan, BSCScan). We include only confirmed malicious exploits resulting in quantifiable asset loss in our sample.

For each confirmed hack event i , we collect several key pieces of information. We identify the name of the targeted DeFi protocol and the ticker symbol and contract address of its primary governance token, focusing on protocols with publicly traded tokens. A critical variable is the

²We thank DeFi Llama for making this data publicly available.

precise hack timestamp ($t_{0,i}$), defined as the UTC timestamp (typically at minute or second-level precision based on block times) of the *first confirmed malicious transaction* associated with the hack sequence, as recorded on the blockchain. Identifying $t_{0,i}$ requires careful analysis of transaction logs, often guided by post-mortem security reports and news articles. Equally important is the public disclosure timestamp ($t_{common,i}$), which captures when information about the hack likely became common knowledge. We define $t_{common,i}$ as the UTC timestamp (minute-level precision) of the *first verifiable public announcement* regarding the specific hack, typically via the protocol’s official X account or a reputable security firm’s alert.³ This identification relies on monitoring the official X account of the hacked protocol, tweets from widely followed blockchain security firms or researchers, and initial reports from major crypto news outlets. We use the earliest of these verifiable disclosures as $t_{common,i}$. Furthermore, we record the estimated value in USD of the assets stolen, based on prices around $t_{0,i}$ as reported by DeFi Llama or security firms. We also identify the market capitalization of each protocol as of 2 days prior to $t_{0,i}$.

Our final sample consists of 49 distinct hack events satisfying our criteria during the sample period September 2020 to March 2024. Table 1 presents summary statistics for the key characteristics of these events, detailing the distribution of stolen values and, crucially, the time lag between the hack occurrence and its public disclosure.

	Amount Stolen _{<i>i</i>} (\$M)	Lagged Protocol Market Cap _{<i>i</i>} (\$M)	Amount Stolen / Lagged Protocol Market Cap _{<i>i</i>}	$(t_{common} - t_{0,i})$	$\sum_{t=t_{start}}^{t_{end}} \log(R_{it})$	Std(log(R_{it}))
Count	49	49	49	49	49	49
Mean	49.60	260.74	0.51	4 days 23:54:42	-0.17	0.01
Std	104.98	639.54	0.74	32 days 20:48:18	0.21	0.01
Min	0.50	0.99	0.00	0 days 00:00:59	-0.75	0.00
25%	3.00	20.99	0.08	0 days 00:43:02	-0.22	0.00
50%	7.90	65.63	0.22	0 days 01:50:49	-0.11	0.00
75%	25.00	148.27	0.51	0 days 05:45:12	-0.03	0.01
Max	624.00	3846.27	3.17	230 days 06:55:55	0.19	0.04

Table 1: Summary Statistics of Hack Events

3.2 Market Data

To analyze the market impact of these hacks, we collect high-frequency trading data for the identified protocol governance tokens and the major collateral tokens stolen during the exploits. Our primary source for this market data is CryptoCompare. Specifically, we utilize the CryptoCompare Aggregated Index (CADLI) for each relevant token pair (e.g., TOKEN-USD). The CADLI methodology provides a robust, aggregated measure of market activity by calculating a 24-hour rolling volume-weighted average price across a vetted set of constituent exchanges for each trading pair. This approach yields a comprehensive market price and volume feed, mitigating concerns

³We detail the precise method for extracting timestamps from tweet URLs in Section 3.2.

related to relying on a single exchange’s data, which might be subject to manipulation, downtime, or idiosyncratic liquidity conditions.

From the CADLI feed, we obtain data at minute-level frequency covering our sample period. For each relevant token j associated with a hack event i , we compile time series for the aggregated price ($P_{j,t}$) in USD and the total aggregated trading volume ($Vol_{j,t}$) during interval t . We also gather market capitalization data from CryptoCompare to serve as control variables and facilitate cross-sectional analysis.

To determine the public disclosure timestamp with high precision when it originates from X, we utilize the tweet’s unique Snowflake ID.⁴ This method allows for highly accurate identification of t_{common} when relying on X as the source of initial public disclosure.

3.3 Data Matching

The final dataset is constructed by merging the hack event information with the high-frequency aggregated market data from CryptoCompare. For each hack event i , we match the identified protocol governance token and relevant stolen collateral tokens to their corresponding CADLI price and volume time series using standard ticker symbols (e.g., ETH, WBTC, CRV) paired with USD.

The precise event timestamps, $t_{0,i}$ and $t_{common,i}$, are central to our analysis. We align these timestamps with the corresponding intervals in our minute-level CADLI market data, enabling a granular measurement of aggregated market dynamics around both the on-chain exploit and the subsequent public information release. Rigorous verification of token identities is performed throughout the matching process to ensure the correct CADLI series is used for each token involved in a hack event.

3.4 Event Definition

For each hack event i in our sample, we define the event window spanning 48 hours, from $T_{start} = -1440$ minutes to $T_{end} = 1440$ minutes relative to the event time ($t = 0$). We focus primarily on the price dynamics of the affected protocol’s main governance token.

⁴The web address (URL) of any specific tweet contains a unique numerical identifier (Snowflake ID). This identifier encodes the tweet’s creation timestamp with millisecond precision. By extracting this ID, converting it to an integer, performing a bitwise right-shift operation (by 22 bits), and adding the X epoch offset (1288834974657 milliseconds since the Unix epoch, corresponding to Nov 04 2010 01:42:54 UTC), we recover the precise UTC timestamp of the tweet’s publication.

3.5 Event Study Regression

To estimate the average return profile around the events across our sample, while controlling for event-specific heterogeneity, we use the following panel regression model:

$$AR_{it} = \alpha_i + \sum_{\tau=T_{start}}^{T_{end}} \beta_{\tau} D_{i\tau}(t) + \epsilon_{it} \quad (1)$$

Here, AR_{it} is the return for the governance token of event i at minute t relative to the event time ($t = 0$). α_i represents event-specific fixed effects, absorbing any time-invariant differences across hack events. $D_{i\tau}$ are dummy variables equal to one if the observation (i, t) corresponds to relative event time τ (minute τ relative to the event at $\tau = 0$), and zero otherwise. The coefficients of interest are β_{τ} , which measure the average abnormal return across all events at minute τ relative to the event time, compared to a baseline period (implicitly the period just before T_{start}). This regression is estimated separately using t_0 as the event time and using t_{common} as the event time. Standard errors are clustered at the event level, to account for autocorrelation and heteroskedasticity.

3.6 Cumulative Returns and Presentation

From the estimated $\hat{\beta}_{\tau}$ coefficients, we calculate the Cumulative Abnormal Return (CAR) up to time T within the event window as the sum of the average minute-level abnormal returns:

$$CAR_T = \sum_{\tau=T_{start}}^T \hat{\beta}_{\tau} \quad (2)$$

We present the results of our event study graphically by plotting the estimated CAR sequence (CAR_T for T from -24 hours to $+24$ hours). Additionally, we report the estimated CARs at key time points relative to the event ($\tau = -24h, -12h, -6h, -3h, 0h, +3h, +6h, +12h, +24h$) in tables, including their respective t-statistics. This allows for a detailed examination of the timing and statistical significance of the market reaction to both the hack itself and its public disclosure.

4 Empirical Results

This section presents the main empirical findings from our high-frequency event study analysis. We examine the evolution of cumulative abnormal returns for protocol governance tokens around the time of the initial on-chain hack transaction and the time of the first public announcement.

4.1 Event Study: Price Dynamics Around Hacks and Disclosures

Figures 2 and 3 graphically depict the average CARs for affected protocol tokens in the hours surrounding the two key event times. Both figures show results from regressions weighted differently (by assets stolen divided by lagged market cap in Figure 2, unweighted in Figure 3), yielding qualitatively similar patterns. Panel (a) in each figure centers the event time ($t = 0$) on the first malicious on-chain transaction associated with the hack (t_0), while Panel (b) centers the event time ($t = 0$) on the first identified public announcement (t_{common}). Table 2 provides the corresponding CAR estimates and their statistical significance at selected horizons.

Reaction to the Initial Hack Transaction. Panel (a) of Figures 2 and 3, together with the first column of Table 2, decisively illustrates the market’s reaction at the moment the hack is initiated on-chain. Although the cumulative abnormal returns (CARs) curve remains largely flat—reflecting the unexpected nature of these events—a modest yet telling decline of approximately 1.55% from $t = -6$ hours to $t = t_0$ emerges. This subtle downturn provides suggestive evidence that some hackers may exploit their advanced knowledge to trade ahead of the cyberattack, capturing a fleeting informational edge before the event fully unfolds.


Following the initiation of the hack at $t = 0$, prices begin a sharp and statistically significant decline. As shown in Table 2, the CAR becomes significantly negative within the first few hours, reaching -17.2 percent by hour 3 and -19.1 percent by hour 6. The decline continues over the subsequent hours, stabilizing around -26 percent between 9 and 24 hours after the initial transaction. This pattern clearly indicates that negative information begins to impact prices shortly after the hack occurs on the blockchain, as soon as it becomes public information accessible to those equipped to process it.

Reaction Relative to Social Media Disclosures. Panel (b) of Figures 2 and 3, along with the third column of Table 2, presents the price dynamics centered on the time of the first social media announcement of the hack, when the information becomes common knowledge. The results reveal a striking pattern of pre-announcement price discovery. Cumulative abnormal returns are already substantially negative and statistically significant many hours *before* the announcement occurs at $t = 0$. For instance, Table 2 shows that the CAR averages -9.5 percent twelve hours prior to the announcement and -7 percent three hours prior. This demonstrates that a significant portion of the negative information associated with the hack is impounded into prices well before it becomes common knowledge through easily accessible public channels like X. This suggests that a subset of market participants—likely those with the capability and incentive to incur the costs of monitoring and interpreting blockchain data in real time—capitalize on the news before it hits the broader market.


Around the time of the announcement, the price decline appears to continue, although potentially at a decelerated rate initially. The CAR reaches its minimum point and stabilizes several hours after

the announcement, for example, around -20.8 percent at 15 hours post-announcement.

Comparing the results from the two event windows allows us to quantify the extent of price discovery occurring before public disclosure makes the information common knowledge. The total average price impact over 24 hours following the initial hack is approximately -27 percent. The CAR measured just prior to the public announcement represents the portion of this total impact realized due to costly information processing before widespread public awareness. This implies that *more than one-third* (approximately $0.095/0.266 \approx 36\%$ using the 12-hour pre-announcement value) of the total price discovery occurs within the window between the on-chain event and the social media announcement. This underscores the significant role of sophisticated traders who can interpret complex, costly-to-process public information and trade on it before the news becomes easily accessible common knowledge.



(a) $t = t_0$ (Hack Time)




(b) $t = t_{common}$ (Announcement Time)


Figure 2: Event Study: Cumulative Abnormal Returns Around Hack Events (Weighted by Value Stolen over Lagged Market Capitalization). Panel (a) shows time relative to the initial hack transaction time ($t = t_0$). Panel (b) shows time relative to the first public announcement time ($t = t_{common}$). The solid line shows the average cumulative abnormal return (CAR), weighted by the total USD value stolen in the hack. Time is measured in hours relative to the event.

5 A Simple Model of Strategic Information Revelation

All information on the blockchain is technically public information, yet our empirical results show that prices don't automatically adjust when a hack happens, but rather decline slowly over time. Furthermore, it seems that informed investors capture more than half of the returns before the hack becomes common knowledge. To understand the microfoundations driving this empirical result, we develop a simple theoretical model incorporating information processing considerations. This model illustrates how the presence of future trading opportunities, coupled with uncertainty about when costly-to-process public information will become low-cost common knowledge, incentivizes informed traders (those who have incurred the processing cost) to manage the release of their infor-



(a) $t = t_0$ (Hack Time)



(b) $t = t_{common}$ (Announcement Time)

Figure 3: Event Study: Cumulative Abnormal Returns Around Hack Events (Unweighted). Panel (a) shows time relative to the initial hack transaction time ($t = t_0$). Panel (b) shows time relative to the first public announcement time ($t = t_{common}$). The solid line shows the average cumulative abnormal return (CAR), weighted by the protocol market capitalization two days before the hack. Time is measured in hours relative to the event.

t	$\hat{\beta}: t_0$	$t\text{-stat}$	$\hat{\beta}: t_{common}$	$t\text{-stat}$
-21	0.026	(1.13)	0.099	(1.54)
-18	0.028	(1.33)	0.096	(1.61)
-15	0.022	(1.19)	0.097*	(1.77)
-12	0.029*	(1.75)	0.094*	(1.90)
-9	0.022	(1.57)	0.087**	(2.00)
-6	0.016*	(1.67)	0.081**	(2.31)
-3	0.002	(0.46)	0.070***	(2.73)
0	0	(.)	0	(.)
3	-0.172***	(-3.45)	-0.099**	(-2.03)
6	-0.191***	(-2.64)	-0.150*	(-1.91)
9	-0.263***	(-2.73)	-0.201*	(-1.92)
12	-0.255**	(-2.26)	-0.206*	(-1.66)
15	-0.268**	(-2.11)	-0.208	(-1.48)
18	-0.274*	(-1.96)	-0.198	(-1.27)
21	-0.262*	(-1.74)	-0.192	(-1.13)
24	-0.266*	(-1.65)	-0.194	(-1.05)

Table 2: Cumulative Abnormal Returns Around Hack Events

This table reports cumulative abnormal return (CAR) estimates from the event study regression (Equation 1) at selected horizons relative to the event time ($t = 0$). Columns (1) and (2) use the initial on-chain hack transaction time (t_0) as the event time. Columns (3) and (4) use the first public announcement time (t_{common}) as the event time. CARs are calculated as the sum of estimated minute-level average abnormal returns ($\hat{\beta}_\tau$) up to the specified hour. Regressions are weighted by value stolen over lagged market capitalization. t -statistics are reported in parentheses. The symbols *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively. The number of events is 49.

mation strategically over time. It provides a rational basis for observing gradual, pre-announcement price discovery even when informed agents possess perfect information about an event's impact from the outset.

We adapt the sequential trading framework of Kyle (1985) to a three-period setting ($t = 0, 1, 2$). At $t = 0$, a potential value-altering event occurs with probability ρ_{event} , changing the asset value from V_0 to either a low value V_L or a high value V_H (with $E[V] = V_0$). A single risk-neutral informed trader (who has incurred the cost to process the on-chain information) learns the true final value $V \in \{V_0, V_L, V_H\}$ immediately at $t = 0$. Trading occurs at $t = 1$ and $t = 2$. Competitive, risk-neutral market makers observe the noisy total order flow $Y_t = x_t + u_t$ (where x_t is the informed trade and u_t is noise) and set prices P_t based on conditional expectations, using a linear rule where price impact is denoted by λ_t . Critically, after trading at $t = 1$, the true value V becomes common knowledge (e.g., via a public announcement) with probability ρ_c ; otherwise, asymmetric information (due to differential processing) persists into $t = 2$.

The core economic mechanism arises from the informed trader's dynamic optimization problem at $t = 1$. Knowing the true value V , the trader chooses their initial trade size x_1 . A larger trade potentially yields higher immediate profits by exploiting the current mispricing ($V - P_0$). However, a larger trade also reveals more information to the market maker, causing the period 1 price P_1 to move closer to V . This has two adverse effects on potential future profits: it directly reduces the mispricing available to exploit at $t = 2$, ($V - P_1$), and it reduces the market maker's residual uncertainty (Σ_1), which in turn lowers the market depth (increases price impact λ_2) in the second period. Therefore, the informed trader faces an intertemporal trade-off: maximize current profits versus preserving future profit opportunities.

This trade-off is moderated by the probability of public disclosure (transition to common knowledge), ρ_c . If disclosure is very likely after period 1 (ρ_c is high), future profit opportunities are improbable, incentivizing the trader to trade more aggressively in period 1 to capitalize immediately. Conversely, if disclosure is unlikely (ρ_c is low), the prospect of significant period 2 profits looms larger, incentivizing the trader to moderate their period 1 trades (*i.e.*, trade more inconspicuously) to avoid revealing too much information too soon. The equilibrium reflects the informed trader optimally balancing these considerations.

Theorem 1 (Equilibrium with Strategic Moderation). *In the 3-period model, a unique linear Bayesian Nash Equilibrium exists. The informed trader's strategy involves trading $x_1 = \beta_1 V$ in period 1 and $x_2 = \beta_2 (V - P_1)$ in period 2 (if reached). Market makers set prices $P_1 = \lambda_1 Y_1$ and $P_2 = P_1 + \lambda_2 Y_2$. Crucially, the informed trader's initial trading intensity, β_1 , is increasing in the probability of public disclosure ρ_c .*

Proof. See Appendix A for the derivation of equilibrium parameters $\lambda_1, \beta_1, \lambda_2, \beta_2$ and the comparative static with respect to ρ_c . ■

This simple model provides a clear economic rationale for the gradual price discovery observed empirically between the initial hack and the public announcement . The theorem’s key result—that informed traders strategically moderate their initial trades when future trading opportunities exist and disclosure is uncertain ($\rho_c < 1$)—explains why prices do not immediately jump to the full information value at the moment sophisticated traders detect the hack , even though the information is technically public. Instead, information is impounded incrementally through period 1 trading as informed agents trade on their processed information, with the speed influenced by factors like disclosure risk (ρ_c). The model thus microfound our empirical finding that a significant fraction, but not all, of the price adjustment occurs before the news becomes common knowledge. It generates testable implications, explored further in the appendix, regarding how the speed of initial price discovery should relate to characteristics that influence disclosure probability or the value of future trading.

6 Conclusion

How quickly do markets process bad news, especially when the initial event is technically public but costly to interpret? This paper studies this fundamental question using hacks on decentralized finance platforms. Our key advantage is the ability to precisely time both when the hack begins on the public blockchain record —making the information publicly available but not yet common knowledge—and when the news is later announced plainly on social media , making it common knowledge. This allows us to measure price discovery driven by costly information processing.

Our central finding is clear: prices move significantly before the news becomes common knowledge. We document that while the total price drop after a hack averages around 27 percent over a day, *more than one-third* (approximately 36 percent) of this drop occurs before the first public announcement. This shows that sophisticated traders, by incurring the costs to process complex but publicly available data, anticipate the news hours in advance and trade on information that is not yet widely understood. This result provides sharp evidence that even transparent markets are not strong-form efficient when significant processing costs prevent public information from immediately becoming common knowledge. It highlights that price discovery unfolds at different speeds, depending on information complexity and the resources required to interpret it.

These findings matter beyond digital assets. As blockchain-based finance becomes increasingly linked to the traditional financial system through new products and regulations, vulnerabilities like DeFi hacks present potential systemic risks. The rapid, yet initially opaque, price adjustments we document illustrate how shocks originating in crypto could create unexpected losses or liquidity strains in connected markets before the full picture is clear to regulators or traditional institutions relying on standard news cycles or common knowledge signals. Understanding these information

dynamics, particularly the lag induced by processing costs, is therefore essential for monitoring and managing financial stability in an increasingly interconnected world.

References

- ADAMS, A., M. IBERT, AND G. LIAO (2024): “What drives crypto asset prices?” Working Paper, May 2024.
- ALLEN, F., S. MORRIS, AND H. SHIN (2006): “Beauty Contests and Iterated Expectations in Asset Markets,” *Review of Financial Studies*, 19, 975–1005.
- AZAR, P. D., A. CASILLAS, AND M. FARBOODI (2024): “Information and Market Power in DeFi Intermediation,” *Federal Reserve Bank of New York Staff Report*.
- BACK, K. (1992): “Insider Trading in Continuous Time,” *The Review of Financial Studies*, 5, 387–409.
- BLOOMFIELD, R. AND M. O’HARA (1999): “Market transparency: Who wins and who loses?” *Review of Financial Studies*, 12, 5–35.
- BROGAARD, J., T. HENDERSHOTT, AND R. RIORDAN (2014): “High-frequency trading and price discovery,” *Review of Financial Studies*, 27, 2267–2306.
- BUDISH, E., P. CRAMTON, AND J. SHIM (2015): “The high-frequency trading arms race: Frequent batch auctions as a market design response,” *Quarterly Journal of Economics*, 130, 1547–1622.
- CAPPONI, A., J. JIA, AND R. SHAMAUGH (2022): “Price discovery on decentralized exchanges,” Working Paper.
- CHEN, Y.-L., Y.-T. CHANG, AND J. J. YANG (2023): “Cryptocurrency Hacking Incidents and the Price Dynamics of Bitcoin Spot and Futures,” *Finance Research Letters*, 55, 103955.
- CONG, L. W. AND Z. HE (2019): “Blockchain Disruption and Smart Contracts,” *Review of Financial Studies*, 32, 1754–1797.
- DUFFIE, D., N. GÂRLEANU, AND L. H. PEDERSEN (2005): “Over-the-counter markets,” *Econometrica*, 73, 1815–1847.
- FAMA, E. F. (1970): “Efficient Capital Markets: A Review of Theory and Empirical Work,” *Journal of Finance*, 25, 383–417.
- FOSTER, F. D. AND S. VISWANATHAN (1996): “Strategic Trading When Agents Forecast the Forecasts of Others,” *Journal of Finance*, 51, 1437–1478.
- GLOSTEN, L. R. AND P. MILGROM (1985): “Bid, ask and transaction prices in a specialist market with heterogeneously informed traders,” *Journal of Financial Economics*, 14, 71–100.

- GROSS, W., E. KLUSSMANN, AND N. HAUTSCH (2011): “When machines read the news: Using automated text analytics to quantify high frequency news-implied market reactions,” *Journal of Empirical Finance*, 18, 1056–1075.
- GROSSMAN, S. J. AND J. E. STIGLITZ (1980): “On the impossibility of informationally efficient markets,” *American Economic Review*, 70, 393–408.
- KYLE, A. S. (1985): “Continuous Auctions and Insider Trading,” *Econometrica*, 53, 1315–1335.
- LIU, Y., A. TSYVINSKI, AND X. WU (2022): “Common risk factors in cryptocurrency,” *The Journal of Finance*, 77, 1133–1177.
- MAKAROV, I. AND A. SCHOAR (2020): “Trading and arbitrage in cryptocurrency markets,” *Journal of Financial Economics*, 135, 293–319.
- MORRIS, S. AND H. S. SHIN (2002): “Social value of public information,” *American Economic Review*, 92, 1521–1534.
- SALEH, F. (2021): “Blockchain without waste: Proof-of-stake,” *Review of Financial Studies*, 34, 1156–1190.
- TETLOCK, P. C. (2007): “Giving content to investor sentiment: The role of media in the stock market,” *Journal of Finance*, 62, 1139–1168.

A Detailed Derivations for the 3-Period Model

This appendix provides a step-by-step derivation of the linear Bayesian Nash Equilibrium for the 3-period trading model. Our goal is to understand how the informed trader chooses their trades and how market makers set prices, ultimately showing how the probability of information disclosure (ρ_c) affects initial trading intensity (β_1).

Model Setup Recap:

- Trading occurs at $t = 1$ and $t = 2$.
- The true value V is realized at $t = 0$, drawn from a distribution with mean $E[V] = V_0 = 0$ and variance $\text{Var}(V) = \Sigma_0$.
- One informed trader (IT) knows V from $t = 0$.
- Market makers (MMs) are competitive and risk-neutral; they set prices.
- Noise traders submit random orders $u_t \sim N(0, \sigma_u^2)$, independent of V and each other.
- Total order flow at time t is $Y_t = x_t + u_t$, where x_t is the IT's trade.
- After trading at $t = 1$, the true value V becomes public with probability ρ_c . If not, trading continues at $t = 2$.

A.1 Analysis of Period $t = 2$

We first analyze the decisions made in the final period, $t = 2$, assuming this period is reached (which happens with probability $1 - \rho_c$). At the start of $t = 2$, the information available to the market makers is the history of order flow from period 1, $\mathcal{F}_1 = \{Y_1\}$. The price set at the end of period 1 is $P_1 = E[V|\mathcal{F}_1]$. The remaining uncertainty about the asset's value, from the market makers' perspective, is the conditional variance $\Sigma_1 = \text{Var}(V|\mathcal{F}_1)$.

Informed Trader's Optimal Trade x_2^* : The IT knows the true value V and the price P_1 . They choose their trade size x_2 to maximize their expected profit in period 2. The profit is the quantity traded (x_2) times the difference between the true value (V) and the execution price (P_2).

$$\text{Maximize} \quad E[x_2(V - P_2)|V, P_1]$$

The IT anticipates that the market makers will set the price P_2 based on the total order flow $Y_2 = x_2 + u_2$. We assume MMs use a linear rule: $P_2 = P_1 + \lambda_2 Y_2 = P_1 + \lambda_2(x_2 + u_2)$. Here, λ_2

represents the price impact of the order flow in period 2. Substituting this into the IT's objective:

$$\text{Maximize } E[x_2(V - (P_1 + \lambda_2(x_2 + u_2)))|V, P_1]$$

The IT knows V , P_1 , and anticipates the value of λ_2 . The only uncertainty from the IT's perspective is the noise trade u_2 . Since u_2 has an expected value of zero ($E[u_2] = 0$) and is independent of the IT's choice x_2 , the expectation simplifies:

$$E[x_2(V - P_1 - \lambda_2 x_2 - \lambda_2 u_2)|V, P_1] = x_2(V - P_1 - \lambda_2 x_2) - \lambda_2 x_2 E[u_2] = x_2(V - P_1 - \lambda_2 x_2)$$

So, the IT's problem is to choose x_2 to maximize a simple quadratic function, let's call it the Period 2 Profit Objective:

$$\Pi_2(x_2) = x_2(V - P_1 - \lambda_2 x_2)$$

To find the maximum, we take the derivative with respect to x_2 and set it to zero (First-Order Condition, FOC):

$$\frac{d\Pi_2}{dx_2} = V - P_1 - 2\lambda_2 x_2 = 0$$

Solving for x_2 gives the IT's optimal trade in period 2:

$$x_2^*(V, P_1) = \frac{V - P_1}{2\lambda_2} \quad (3)$$

This makes intuitive sense: the IT trades more aggressively (larger x_2^*) when the perceived mispricing ($V - P_1$) is larger, and less aggressively when the price impact λ_2 is higher. We can define the trading intensity parameter $\beta_2 = 1/(2\lambda_2)$, so the strategy is linear: $x_2^* = \beta_2(V - P_1)$.

Market Maker's Optimal Pricing λ_2 : Competitive market makers set the price equal to the expected value of the asset given all publicly available information. At the end of period 2, this information is $\mathcal{F}_2 = \{Y_1, Y_2\}$. So, $P_2 = E[V|Y_1, Y_2]$. The linear rule $P_2 = P_1 + \lambda_2 Y_2$ is consistent with this if λ_2 correctly reflects how much information about V is contained in the new order flow Y_2 , given what was already known from Y_1 . Specifically, λ_2 is the coefficient from projecting V onto Y_2 after accounting for $P_1 = E[V|Y_1]$. This projection coefficient is given by the conditional covariance divided by the conditional variance:

$$\lambda_2 = \frac{\text{Cov}(V, Y_2|\mathcal{F}_1)}{\text{Var}(Y_2|\mathcal{F}_1)}$$

The notation " $|\mathcal{F}_1$ " means "conditional on knowing the information from period 1 (i.e., Y_1 and thus P_1 and Σ_1 are known)". Let's calculate the numerator and denominator.

Numerator (Conditional Covariance): Substitute $Y_2 = x_2^* + u_2 = \beta_2(V - P_1) + u_2$:

$$\begin{aligned}\text{Cov}(V, Y_2 | \mathcal{F}_1) &= \text{Cov}(V, \beta_2(V - P_1) + u_2 | \mathcal{F}_1) \\ &= \text{Cov}(V, \beta_2 V | \mathcal{F}_1) - \text{Cov}(V, \beta_2 P_1 | \mathcal{F}_1) + \text{Cov}(V, u_2 | \mathcal{F}_1) \quad (\text{Linearity of Covariance}) \\ &= \beta_2 \text{Cov}(V, V | \mathcal{F}_1) - \beta_2 \text{Cov}(V, P_1 | \mathcal{F}_1) + 0 \quad (\text{Pull out constant } \beta_2; V, u_2 \text{ independent})\end{aligned}$$

Since P_1 is known given $\mathcal{F}_1$, it's treated as a constant in the conditional covariance, so $\text{Cov}(V, P_1 | \mathcal{F}_1) = 0$. Also, $\text{Cov}(V, V | \mathcal{F}_1) = \text{Var}(V | \mathcal{F}_1) = \Sigma_1$.

$$\text{Cov}(V, Y_2 | \mathcal{F}_1) = \beta_2 \Sigma_1 - 0 + 0 = \beta_2 \Sigma_1$$

Denominator (Conditional Variance):

$$\begin{aligned}\text{Var}(Y_2 | \mathcal{F}_1) &= \text{Var}(\beta_2(V - P_1) + u_2 | \mathcal{F}_1) \\ &= \text{Var}(\beta_2 V - \beta_2 P_1 + u_2 | \mathcal{F}_1) \\ &= \text{Var}(\beta_2 V | \mathcal{F}_1) + \text{Var}(u_2 | \mathcal{F}_1) + 2\text{Cov}(\beta_2 V - \beta_2 P_1, u_2 | \mathcal{F}_1) \quad (\text{Variance of sum}) \\ &= \beta_2^2 \text{Var}(V | \mathcal{F}_1) + \text{Var}(u_2) + 0 \quad (\text{Pull out constant } \beta_2; V, u_2 \text{ independent; } P_1 \text{ known}) \\ &= \beta_2^2 \Sigma_1 + \sigma_u^2\end{aligned}$$

Solving for λ_2 and β_2 : Substitute the numerator and denominator back into the formula for λ_2 :

$$\lambda_2 = \frac{\beta_2 \Sigma_1}{\beta_2^2 \Sigma_1 + \sigma_u^2}$$

We have two equations relating λ_2 and β_2 : this one, and $\beta_2 = 1/(2\lambda_2)$ from the IT's optimization. Substitute $\beta_2 = 1/(2\lambda_2)$ into the equation for λ_2 :

$$\lambda_2 = \frac{(1/2\lambda_2)\Sigma_1}{(1/(2\lambda_2))^2 \Sigma_1 + \sigma_u^2} = \frac{\Sigma_1/(2\lambda_2)}{(\Sigma_1/(4\lambda_2^2)) + \sigma_u^2}$$

Multiply both sides by the denominator and by $2\lambda_2$ (assuming $\lambda_2 \neq 0$):

$$\begin{aligned}
\lambda_2 \left(\frac{\Sigma_1}{4\lambda_2^2} + \sigma_u^2 \right) &= \frac{\Sigma_1}{2\lambda_2} \\
\frac{\Sigma_1}{4\lambda_2} + \lambda_2 \sigma_u^2 &= \frac{\Sigma_1}{2\lambda_2} \quad (\text{Multiply by } 2\lambda_2) \\
\frac{\Sigma_1}{2} + 2\lambda_2^2 \sigma_u^2 &= \Sigma_1 \\
2\lambda_2^2 \sigma_u^2 &= \Sigma_1 - \frac{\Sigma_1}{2} = \frac{\Sigma_1}{2} \\
\lambda_2^2 &= \frac{\Sigma_1}{4\sigma_u^2}
\end{aligned}$$

Price impact λ_2 must be positive if there's any uncertainty left ($\Sigma_1 > 0$). So we take the positive square root:

$$\lambda_2 = \frac{\sqrt{\Sigma_1}}{2\sigma_u} \quad (4)$$

Now find the corresponding β_2 :

$$\beta_2 = \frac{1}{2\lambda_2} = \frac{1}{2(\sqrt{\Sigma_1}/2\sigma_u)} = \frac{\sigma_u}{\sqrt{\Sigma_1}} \quad (5)$$

These give the equilibrium price impact and trading intensity in period 2, depending on the residual uncertainty Σ_1 and noise trader variance σ_u^2 .

Informed Trader's Expected Profit in $t = 2$: What is the value of playing the period 2 game for the IT? We substitute the optimal trade x_2^* back into the IT's objective function IT_2 's Objective(x_2) = $x_2(V - P_1 - \lambda_2 x_2)$.

$$\begin{aligned}
\Pi_2^*(V, P_1) &= x_2^*(V - P_1 - \lambda_2 x_2^*) \\
&= \frac{V - P_1}{2\lambda_2} \left(V - P_1 - \lambda_2 \frac{V - P_1}{2\lambda_2} \right) \\
&= \frac{V - P_1}{2\lambda_2} \left(V - P_1 - \frac{V - P_1}{2} \right) \\
&= \frac{V - P_1}{2\lambda_2} \left(\frac{V - P_1}{2} \right) = \frac{(V - P_1)^2}{4\lambda_2}
\end{aligned}$$

Now substitute the equilibrium value $\lambda_2 = \frac{\sqrt{\Sigma_1}}{2\sigma_u}$:

$$\Pi_2^*(V, P_1) = \frac{(V - P_1)^2}{4(\sqrt{\Sigma_1}/2\sigma_u)} = \frac{(V - P_1)^2}{2\sqrt{\Sigma_1}/\sigma_u} = \frac{\sigma_u(V - P_1)^2}{2\sqrt{\Sigma_1}} \quad (6)$$

This is the profit the IT expects to make in period 2, given the true value V and the price P_1 set after period 1. It depends on the squared mispricing $(V - P_1)^2$ and the market conditions (σ_u, Σ_1) .

Residual Variance Σ_1 (Bayesian Updating): How much uncertainty about V remains after observing the period 1 order flow Y_1 ? This is $\Sigma_1 = \text{Var}(V|Y_1)$. This is a standard Bayesian updating problem. When the prior distribution of V is Normal and the signal $Y_1 = \beta_1 V + u_1$ is linear with Normal noise u_1 , the posterior variance is given by a specific formula. Even without assuming Normality, this formula gives the variance of the best linear predictor of V given Y_1 . The formula is:

$$\text{Var}(V|Y_1) = \text{Var}(V) - \frac{[\text{Cov}(V, Y_1)]^2}{\text{Var}(Y_1)}$$

Let's calculate the terms:

- $\text{Var}(V) = \Sigma_0$ (Prior variance before any trading)
- $\text{Cov}(V, Y_1) = \text{Cov}(V, \beta_1 V + u_1) = \text{Cov}(V, \beta_1 V) + \text{Cov}(V, u_1) = \beta_1 \text{Var}(V) + 0 = \beta_1 \Sigma_0$
- $\text{Var}(Y_1) = \text{Var}(\beta_1 V + u_1) = \text{Var}(\beta_1 V) + \text{Var}(u_1) + 2\text{Cov}(\beta_1 V, u_1) = \beta_1^2 \text{Var}(V) + \sigma_u^2 + 0 = \beta_1^2 \Sigma_0 + \sigma_u^2$

Substituting these into the formula for Σ_1 :

$$\Sigma_1 = \Sigma_0 - \frac{(\beta_1 \Sigma_0)^2}{\beta_1^2 \Sigma_0 + \sigma_u^2}$$

We can rewrite this using the definition of the period 1 price impact parameter, λ_1 . By the same logic used for λ_2 , the market maker sets $P_1 = \lambda_1 Y_1$ where:

$$\lambda_1 = \frac{\text{Cov}(V, Y_1)}{\text{Var}(Y_1)} = \frac{\beta_1 \Sigma_0}{\beta_1^2 \Sigma_0 + \sigma_u^2}$$

Now notice that the fraction in the expression for Σ_1 is related to λ_1 :

$$\frac{(\beta_1 \Sigma_0)^2}{\beta_1^2 \Sigma_0 + \sigma_u^2} = \beta_1 \Sigma_0 \left(\frac{\beta_1 \Sigma_0}{\beta_1^2 \Sigma_0 + \sigma_u^2} \right) = \beta_1 \Sigma_0 \lambda_1$$

So, we can write the updated variance very neatly:

$$\Sigma_1 = \Sigma_0 - \lambda_1 \beta_1 \Sigma_0 = \Sigma_0 (1 - \lambda_1 \beta_1) \quad (7)$$

This shows how the initial uncertainty Σ_0 is reduced by the trading in period 1. The reduction depends on the product of the price impact λ_1 and the trading intensity β_1 . More informative trading (higher $\lambda_1 \beta_1$) leads to lower residual uncertainty Σ_1 .

A.2 Analysis of Period $t = 1$

Now we move back to period 1. The IT knows V and needs to choose the trade size x_1 .

Informed Trader's Problem (t=1): The IT chooses x_1 to maximize their total expected profit across both periods. The profit consists of:

1. Profit from period 1 trading: $x_1(V - P_1)$
2. Expected profit from period 2 trading: This is only realized if the information does not become public after period 1 (which happens with probability $1 - \rho_c$). The value of this future profit is $\Pi_2^*(V, P_1)$ as calculated in Eq. 6.

The IT's total objective is to maximize the expectation of the sum, conditional on knowing V :

$$\text{Maximize } E [x_1(V - P_1) + (1 - \rho_c)\Pi_2^*(V, P_1) \mid V]$$

The IT anticipates the pricing rule $P_1 = \lambda_1 Y_1 = \lambda_1(x_1 + u_1)$. The future profit Π_2^* depends on P_1 and Σ_1 . The residual variance $\Sigma_1 = \Sigma_0(1 - \lambda_1\beta_1)$ depends on the anticipated equilibrium intensity β_1 (which determines λ_1). When choosing x_1 , the IT treats λ_1 and Σ_1 as fixed parameters determined by the overall equilibrium, but knows that P_1 will depend on their chosen x_1 and the random noise u_1 .

Substitute P_1 and Π_2^* into the objective:

$$\text{Maximize } E \left[x_1(V - \lambda_1(x_1 + u_1)) + (1 - \rho_c) \frac{\sigma_u(V - \lambda_1(x_1 + u_1))^2}{2\sqrt{\Sigma_1}} \mid V \right]$$

Now we need to take the expectation over the random noise u_1 . Let $A = V - \lambda_1 x_1$ and $B = \lambda_1$. The expression involves terms like $E[A - Bu_1]$ and $E[(A - Bu_1)^2]$.

- $E[x_1(V - \lambda_1 x_1 - \lambda_1 u_1) \mid V] = x_1(V - \lambda_1 x_1) - x_1 \lambda_1 E[u_1] = x_1(V - \lambda_1 x_1)$
- $E[(V - \lambda_1 x_1 - \lambda_1 u_1)^2 \mid V] = E[(A - Bu_1)^2 \mid V] = E[A^2 - 2ABu_1 + B^2 u_1^2 \mid V] = A^2 - 2ABE[u_1] + B^2 E[u_1^2] = A^2 + B^2 \text{Var}(u_1) = (V - \lambda_1 x_1)^2 + \lambda_1^2 \sigma_u^2$

So the IT's objective function, let's call it the Total Expected Profit, becomes:

$$\Pi_1(x_1) = x_1(V - \lambda_1 x_1) + (1 - \rho_c) \frac{\sigma_u}{2\sqrt{\Sigma_1}} [(V - \lambda_1 x_1)^2 + \lambda_1^2 \sigma_u^2]$$

To maximize this with respect to x_1 , we take the derivative and set it to zero:

$$\begin{aligned}
\frac{d\Pi_1}{dx_1} &= \frac{d}{dx_1} [x_1 V - \lambda_1 x_1^2] + (1 - \rho_c) \frac{\sigma_u}{2\sqrt{\Sigma_1}} \frac{d}{dx_1} [(V - \lambda_1 x_1)^2 + \lambda_1^2 \sigma_u^2] \\
&= (V - 2\lambda_1 x_1) + (1 - \rho_c) \frac{\sigma_u}{2\sqrt{\Sigma_1}} [2(V - \lambda_1 x_1)(-\lambda_1) + 0] \\
&= V - 2\lambda_1 x_1 - (1 - \rho_c) \frac{\sigma_u \lambda_1}{\sqrt{\Sigma_1}} (V - \lambda_1 x_1) = 0
\end{aligned}$$

This equation must hold for the optimal x_1 . We are looking for a linear equilibrium where the IT chooses $x_1 = \beta_1 V$. Substitute this into the FOC:

$$V - 2\lambda_1(\beta_1 V) - (1 - \rho_c) \frac{\sigma_u \lambda_1}{\sqrt{\Sigma_1}} (V - \lambda_1(\beta_1 V)) = 0$$

Assuming $V \neq 0$, we can divide the entire equation by V :

$$1 - 2\lambda_1\beta_1 - (1 - \rho_c) \frac{\sigma_u \lambda_1}{\sqrt{\Sigma_1}} (1 - \lambda_1\beta_1) = 0$$

Now substitute the expression for $\Sigma_1 = \Sigma_0(1 - \lambda_1\beta_1)$ (from Eq. 7):

$$\begin{aligned}
1 - 2\lambda_1\beta_1 - (1 - \rho_c) \frac{\sigma_u \lambda_1 (1 - \lambda_1\beta_1)}{\sqrt{\Sigma_0(1 - \lambda_1\beta_1)}} &= 0 \\
1 - 2\lambda_1\beta_1 - (1 - \rho_c) \frac{\sigma_u \lambda_1 \sqrt{1 - \lambda_1\beta_1}}{\sqrt{\Sigma_0}} &= 0 \tag{8}
\end{aligned}$$

This is the crucial equation resulting from the IT's optimization in period 1. It provides a relationship that must hold between the equilibrium price impact λ_1 and trading intensity β_1 .

Market Maker's Problem (t=1): As mentioned before, the market maker sets λ_1 such that $P_1 = E[V|Y_1]$. This leads to the standard Kyle model formula for price impact:

$$\lambda_1 = \frac{\text{Cov}(V, Y_1)}{\text{Var}(Y_1)} = \frac{\beta_1 \Sigma_0}{\beta_1^2 \Sigma_0 + \sigma_u^2} \tag{9}$$

Equilibrium Solution (β_1, λ_1): The equilibrium values of β_1 and λ_1 are found by solving the two key equations simultaneously:

1. IT's Optimality (FOC): Eq. 8
2. MM's Pricing Rule: Eq. 9

Solving this system explicitly for β_1 and λ_1 is algebraically complex. A common technique is to

introduce the variable $X = \lambda_1 \beta_1$, which represents the fraction of variance explained by the first period's trade (since $\Sigma_1 = \Sigma_0(1 - X)$). For Σ_1 to be positive, we need $X < 1$. Also, since $\lambda_1, \beta_1 > 0$, we need $X > 0$.

We can express β_1 and λ_1 in terms of X using the MM's pricing rule (Eq. 9):

- From $\lambda_1 = \frac{\beta_1 \Sigma_0}{\beta_1^2 \Sigma_0 + \sigma_u^2}$ and $\lambda_1 = X/\beta_1$, we get:

$$\beta_1 = \sigma_u \sqrt{\frac{X}{\Sigma_0(1 - X)}} \quad (10)$$

- And substituting back $\lambda_1 = X/\beta_1$:

$$\lambda_1 = \frac{X}{\beta_1} = \frac{X}{\sigma_u} \sqrt{\frac{\Sigma_0(1 - X)}{X}} = \frac{1}{\sigma_u} \sqrt{X \Sigma_0(1 - X)} \quad (11)$$

Now, substitute these expressions for λ_1 and β_1 (only in terms of X) into the IT's FOC (Eq. 8):

$$\begin{aligned} 1 - 2X - (1 - \rho_c) \frac{\sigma_u}{\sqrt{\Sigma_0}} \left(\frac{1}{\sigma_u} \sqrt{X \Sigma_0(1 - X)} \right) \sqrt{1 - X} &= 0 \\ 1 - 2X - (1 - \rho_c) \frac{\sqrt{X \Sigma_0(1 - X)}}{\sqrt{\Sigma_0}} \sqrt{1 - X} &= 0 \\ 1 - 2X - (1 - \rho_c) \sqrt{X(1 - X)} \sqrt{1 - X} &= 0 \\ 1 - 2X - (1 - \rho_c) \sqrt{X}(1 - X) &= 0 \end{aligned} \quad (*)$$

This single equation implicitly defines the equilibrium value of $X = \lambda_1 \beta_1$. Once we solve this equation for X (likely numerically), we can plug the value of X back into Eqs. 10 and 11 to find the equilibrium β_1 and λ_1 .

A.3 Proof of Strategic Moderation (How β_1 changes with ρ_c)

Our main goal is to show that β_1 increases when ρ_c increases ($\partial \beta_1 / \partial \rho_c > 0$). We use the equilibrium condition derived above (Eq. *) and the relationship between β_1 and X .

The proof proceeds in two steps: 1. Show that X increases with ρ_c ($\partial X / \partial \rho_c > 0$). 2. Show that β_1 increases with X ($\partial \beta_1 / \partial X > 0$). If both are true, then by the chain rule, β_1 must increase with ρ_c .

Step 1: Show $\partial X / \partial \rho_c > 0$ The equilibrium value of X is defined implicitly by the equation $1 - 2X - (1 - \rho_c) \sqrt{X}(1 - X) = 0$. Let's define a function $G(X, \rho_c) = 1 - 2X - (1 - \rho_c) \sqrt{X}(1 - X)$. The equilibrium condition is $G(X, \rho_c) = 0$. We want to find how X changes when ρ_c changes,

which is the derivative $dX/d\rho_c$. We can use the implicit function theorem, which states that if $G(X, \rho_c) = 0$, then:

$$\frac{dX}{d\rho_c} = -\frac{\partial G/\partial \rho_c}{\partial G/\partial X}$$

We need to calculate the two partial derivatives.

Partial derivative with respect to ρ_c : Treat X as constant.

$$\frac{\partial G}{\partial \rho_c} = \frac{\partial}{\partial \rho_c} [1 - 2X - (1 - \rho_c)\sqrt{X}(1 - X)] = 0 - 0 - (-\sqrt{X}(1 - X)) = \sqrt{X}(1 - X)$$

Since X must be between 0 and 1 in equilibrium ($0 < X < 1$), both $\sqrt{X}$ and $(1 - X)$ are positive. Thus, $\frac{\partial G}{\partial \rho_c} > 0$.

Partial derivative with respect to X : Treat ρ_c as constant. We need the product rule for the term $\sqrt{X}(1 - X)$.

$$\begin{aligned} \frac{\partial G}{\partial X} &= \frac{\partial}{\partial X} [1 - 2X - (1 - \rho_c)(\sqrt{X} - X^{3/2})] \\ &= 0 - 2 - (1 - \rho_c) \left[\frac{d}{dX} (\sqrt{X} - X^{3/2}) \right] \\ &= -2 - (1 - \rho_c) \left[\frac{1}{2\sqrt{X}} - \frac{3}{2}X^{1/2} \right] \\ &= -2 - (1 - \rho_c) \left[\frac{1}{2\sqrt{X}} - \frac{3\sqrt{X}}{2} \right] = -2 - (1 - \rho_c) \left[\frac{1 - 3X}{2\sqrt{X}} \right] \end{aligned}$$

This expression looks complicated, but we can simplify it by using the equilibrium condition $G(X, \rho_c) = 0$ itself, which implies $1 - 2X = (1 - \rho_c)\sqrt{X}(1 - X)$. From this, we can isolate

$(1 - \rho_c) = \frac{1-2X}{\sqrt{X}(1-X)}$. Substitute this into the expression for $\partial G/\partial X$:

$$\begin{aligned}
\frac{\partial G}{\partial X} &= -2 - \left(\frac{1-2X}{\sqrt{X}(1-X)} \right) \left[\frac{1-3X}{2\sqrt{X}} \right] \\
&= -2 - \frac{(1-2X)(1-3X)}{2X(1-X)} \quad (\text{Combine terms}) \\
&= \frac{-2 \times 2X(1-X) - (1-2X)(1-3X)}{2X(1-X)} \quad (\text{Common denominator}) \\
&= \frac{-4X(1-X) - (1-3X-2X+6X^2)}{2X(1-X)} \\
&= \frac{-4X+4X^2 - (1-5X+6X^2)}{2X(1-X)} \\
&= \frac{-4X+4X^2 - 1+5X-6X^2}{2X(1-X)} \\
&= \frac{-2X^2+X-1}{2X(1-X)}
\end{aligned}$$

Now we need the sign of this derivative. The denominator $2X(1-X)$ is positive since $0 < X < 1$. Consider the numerator: $-2X^2 + X - 1$. This is a downward-opening parabola. To check if it can be positive, we look at its discriminant: $\Delta = b^2 - 4ac = (1)^2 - 4(-2)(-1) = 1 - 8 = -7$. Since the discriminant is negative and the parabola opens downward, the numerator is *always negative* for any real X . Therefore, $\frac{\partial G}{\partial X} = \frac{\text{Negative}}{\text{Positive}} = \text{Negative}$.

Combining the results:

$$\frac{dX}{d\rho_c} = -\frac{\partial G/\partial \rho_c}{\partial G/\partial X} = -\frac{(+)}{(-)} = (+)$$

We have successfully shown that $X = \lambda_1 \beta_1$ increases as ρ_c increases.

Step 2: Show $\partial \beta_1/\partial X > 0$ We know the relationship between β_1 and X from Eq. 10:

$$\beta_1 = \sigma_u \sqrt{\frac{X}{\Sigma_0(1-X)}} = \sqrt{\frac{\sigma_u^2}{\Sigma_0}} \sqrt{\frac{X}{1-X}}$$

Since σ_u^2 and Σ_0 are positive constants, β_1 is proportional to $\sqrt{h(X)}$ where $h(X) = \frac{X}{1-X}$. To see how β_1 changes with X , we just need to see how $h(X)$ changes with X . Let's find the derivative of $h(X)$:

$$h'(X) = \frac{d}{dX} \left(\frac{X}{1-X} \right) = \frac{(1-X)(1) - X(-1)}{(1-X)^2} = \frac{1-X+X}{(1-X)^2} = \frac{1}{(1-X)^2}$$

Since $(1-X)^2$ is always positive for $X \neq 1$, $h'(X)$ is positive for $X \in (0, 1)$. This means $h(X)$ is strictly increasing in X . Because β_1 is proportional to the square root of an increasing function of X (and the square root function is itself increasing for positive arguments), β_1 must be strictly

increasing in X . That is, $\partial\beta_1/\partial X > 0$.

Conclusion: We showed that X increases with ρ_c (Step 1), and β_1 increases with X (Step 2). Using the chain rule for derivatives:

$$\frac{\partial\beta_1}{\partial\rho_c} = \underbrace{\frac{\partial\beta_1}{\partial X}}_{(+)} \times \underbrace{\frac{\partial X}{\partial\rho_c}}_{(+)} = (+) > 0$$

This completes the proof. It confirms the theorem's statement and the economic intuition: when the informed trader perceives a higher probability (ρ_c) that their information advantage will disappear soon (due to public disclosure), they have a stronger incentive to trade more aggressively in the first period to capitalize on the information before it becomes common knowledge. This leads to a higher initial trading intensity β_1 .