

Resisting Manipulative Bots in Meme Coin Copy Trading: A Multi-Agent Approach with Chain-of-Thought Reasoning

Yichen Luo

UCL, Centre for Blockchain Technologies
London, United Kingdom
The University of Hong Kong, FinTech Academy
Hong Kong SAR, China
yichen.luo.22@ucl.ac.uk

Jiahua Xu

UCL, Centre for Blockchain Technologies
DLT Science Foundation
London, United Kingdom
jiahua.xu@ucl.ac.uk

Yebo Feng*

Nanyang Technological University
Singapore
yebo.feng@ntu.edu.sg

Yang Liu

Nanyang Technological University
Singapore
yangliu@ntu.edu.sg

Abstract

Copy trading has become the dominant entry strategy in meme coin markets. However, due to the market's extremely illiquid and volatile nature, the strategy exposes an exploitable attack surface: adversaries deploy manipulative bots to front-run trades, conceal positions, and fabricate sentiment, systematically extracting value from naïve copiers at scale. Despite its prevalence, bot-driven manipulation remains largely unexplored, and no robust defensive framework exists. We propose a manipulation-resistant copy-trading system based on a multi-agent architecture powered by a multi-modal large language model (LLM) and chain-of-thought (CoT) reasoning. Our approach outperforms zero-shot and most statistic-driven baselines in prediction accuracy as well as all baselines in economic performance, achieving an average copier return of 3% per meme coin investment under realistic market frictions. Overall, our results demonstrate the effectiveness of agent-based defenses and predictability of trader profitability in adversarial meme coin markets, providing a practical foundation for robust copy trading.

CCS Concepts

• **Computing methodologies** → **Artificial intelligence**; • **Applied computing** → **Economics**.

Keywords

Multi-Agent System, Chain of Thought, Copy Trading, Meme Coin

ACM Reference Format:

Yichen Luo, Yebo Feng, Jiahua Xu, and Yang Liu. 2026. Resisting Manipulative Bots in Meme Coin Copy Trading: A Multi-Agent Approach with Chain-of-Thought Reasoning. In *Proceedings of the ACM Web Conference 2026 (WWW '26)*, April 13–17, 2026, Dubai, United Arab Emirates. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3774904.3792635>

*Yebo Feng is the corresponding author.



This work is licensed under a Creative Commons Attribution 4.0 International License. *WWW '26, Dubai, United Arab Emirates*
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2307-0/2026/04
<https://doi.org/10.1145/3774904.3792635>

1 Introduction

While manipulative tactics such as rat trading are well-trodden ruses in traditional financial markets, these old-fashioned ploys have now become bot-driven and prey on naïve copy traders in a new hunting ground: the meme coin market. The launch of the \$TRUMP meme coin on January 17, 2025 ignited a surge of speculation in the meme coin market, drawing millions of traders. However, many of these entrants lack prior trading experience. To lower the barrier to participation, meme coin tracking platforms such as GMGN introduced **copy trading**, an automated, one-click feature that allows users to replicate the trades of other wallets. In practice, copiers attempt to identify so-called smart money—wallets perceived to possess insider knowledge, trading expertise, and consistent profitability—and replicate their trading activity [2, 30].

However, the characteristics of the meme coin market, most notably extreme illiquidity and high volatility, combined with the naïveté of copy traders, create an ideal environment for exploitation. Sophisticated adversaries deploy automated trading bots (e.g., **PumpKing**) to systematically manipulate prices and induce predictable copy-trading behavior. In particular, some key opinion leaders (KOLs) featured on tracking platforms may abuse their perceived credibility by coordinating multiple bots to accumulate positions early at low prices, artificially inflate prices while concealing their true exposure, and engage in wash trading and fabricated on-chain or social activity to create an illusion of demand [6, 12, 25]. These optics of profitability attract copiers, inflate prices, and allow manipulators to gain by using exit liquidity from copy traders. Despite the prevalence of these attacks, neither industry nor academia has established a copy-trading framework robust to manipulative bots, hence our motivation to design such a framework.

Recent advances in large language models (LLMs) offer a promising foundation for our design. Leveraging their strong capabilities in multi-modal understanding and reasoning [15, 17, 33, 43–45], LLM-powered multi-agent systems (MASs) are well suited for bot-resilient copy trading. First, the on-chain meme coin market can be encoded as a machine-readable environment for sequential copy-trading decisions, in which agents interact with multi-modal on-chain and off-chain signals while reasoning over latent factors such

as market sentiment and sophisticated manipulation. This formulation aligns well with the multi-modal inference and reasoning strengths of LLM-powered agent systems. Moreover, investment decisions in traditional financial institutions are typically made by teams rather than individuals; MASs mirror this collective human cognitive decision-making process by enabling multiple specialized agents to collaborate, deliberate, and reach more robust copy-trading decisions, thereby improving resilience against manipulative bots and adversarial trading behavior.

However, due to limited domain knowledge of meme coin markets, zero-shot LLMs often fail to identify subtle manipulative patterns or reliably distinguish genuine smart money from adversarial wallets. Chain-of-thought (CoT) reasoning, in contrast, injects domain knowledge into agent deliberation, enabling agents to internalize copy-trading knowledge and reason in a manner closer to experienced human copy traders.

Accordingly, we propose an LLM-powered MAS for manipulation-resistant meme coin copy trading with CoT reasoning. Our framework decomposes copy trading into three core subtasks—coin evaluation, wallet selection, and timing assessment—each handled by a specialized agent guided by few-shot CoT prompts. We evaluate our framework on historical wallet-level data from 6,000 meme coins and show that the proposed multi-agent system outperforms the zero-shot and most statistic-driven baselines in prediction accuracy. Economically, the smart-money wallets identified by our framework achieve an average return of 14%, while the corresponding estimated copier return is 3% per meme coin investment under realistic market frictions. Our contributions are multifold:

- We propose an LLM-based MAS for meme coin copy trading. To the best of our knowledge, this is the *first* application of a multi-agent system to copy trading.
- We provide a comprehensive analysis of manipulative bots in meme coin markets under copy-trading settings. We formalize manipulation heuristics, develop detection algorithms, and empirically quantify bot activity in the meme coin market.
- We design a structured feature space tailored to meme coin markets that captures manipulative bot presence as well as traders' historical performance and experience. Using these features and proposed prediction models, we show that trader profitability is predictable, providing a foundation for meme coin copy trading.

2 Related Works

We next review related work on crypto market manipulation and LLM-powered agents in blockchain to highlight the gaps our work addresses in adversarial copy-trading settings.

Crypto Market Manipulation. Manipulation in cryptocurrency markets has been widely documented, including pump-and-dump schemes [5, 35, 41], rug pull [24, 26, 31], wash trading [1, 3, 28], metric distortion [21, 29], and wallet drainers [7]. In the context of meme coin markets, Mongardini and Mei [25] and Li et al. [16] document the prevalence of systematic manipulation. Ding et al. [6] empirically quantify the causal effects of distinct meme coin manipulation strategies on project performance and participant profitability. Xiang et al. [36] develop a framework to quantify the fragility in the meme coin market. This line of work primarily focuses on measuring manipulation, whereas defensive

mechanisms for copy trading under adversarial conditions remain underexplored.

LLM-powered Agents in Blockchain. Recent advances in LLM-powered autonomous agents have begun to transform the blockchain ecosystem. Prior work has explored the use of LLM-based agents in cryptocurrency markets and decentralized finance (DeFi), particularly in quantitative trading [14, 22, 32], smart contract auditing [34], and user intent understanding [23]. Building on evidence that CoT prompting substantially enhances the reasoning capabilities of LLMs [33, 38, 39], recent studies have increasingly incorporated CoT-based reasoning to improve decision-making in financial agents [9, 27, 37, 46]. However, the application of multi-agent LLM systems to adversarially robust copy trading in meme coin markets has received limited attention.

3 Background

3.1 Solana and Decentralized Finance

Solana is a high-throughput permissionless blockchain that supports DeFi applications via smart contracts and a block-based execution model. It features sub-second block times (approximately 400 ms) and processes thousands of transactions per block, enabling rapid on-chain state updates and dense trading activity. At the time of writing, Solana has processed over 500 billion transactions. Leveraging its high throughput, Solana does not rely on a public mempool in which pending transactions remain visible for an extended period prior to block inclusion. Instead, transactions are forwarded directly to validators and incorporated into blocks with minimal delay, leaving little time for the general public to observe pending transactions and react to emerging trading activity. As a result, the ability to respond strategically, such as through front-running or coordination, is concentrated among well-positioned or infrastructure-level participants.

DeFi activity on Solana consists of repeated interactions between wallets and the evolving on-chain state. At each block, wallets observe prices, balances, and transactions, and submit actions such as token creation, buys, sells, and transfers. State updates follow deterministic smart-contract logic [40], but these trading decisions are subject to uncertainty from transaction ordering, concurrent wallet activity, and latent coordination.

3.2 Meme Coin Markets and Copy Trading

Meme coins are cryptocurrencies originating from internet memes [18, 19]. They are typically characterized by low liquidity and extreme volatility, with their prices driven by community hype, social media activity, and celebrity endorsements. Pump.fun is the largest on-chain meme coin launch platform on Solana, enabling users to rapidly create and trade meme coins. Traders purchase meme coins using SOL, Solana's native cryptocurrency, paying a 1% transaction fee and following a bonding curve mechanism. The bonding curve induces a strictly increasing and convex price path, so prices rise deterministically with cumulative demand, yielding diminishing tokens received per unit of SOL deposited. As a result, late entry is mechanically penalized along the curve (see §A.1). Besides the financial market, pump.fun also provides a social platform. Any user with a Solana wallet can comment on any coin and interact with other users by referencing specific comments via unique comment



Figure 1: Excerpts of transaction history and token holding concentration on pump.fun of meme coin Oracle. The label “dev” denotes the meme coin creator, whereas “liquidity pool” denotes unpurchased coins.

IDs (e.g., “@144840307 it is wild”). Each trade bumps the coin’s name to the front page with a brief jiggle. As illustrated in Fig. 1, pump.fun publicly displays both transaction histories and token holding distributions, with the meme coin creator explicitly labeled.

In parallel, meme coin tracking platforms such as GMGN aggregate on-chain trading activity and highlight so-called KOL wallets that are ranked based on returns or number of trades, providing copy-trading tools to replicate trades from selected addresses.

4 Modeling and Detecting Manipulations

This section describes manipulations in detail. We introduce our threat model, discuss ways to carry out manipulation, propose detection metrics and algorithms, and present detection results.

4.1 Threat Model

We assume a strategic adversary capable of rapidly creating and controlling a large number of on-chain wallets. The adversary can coordinate these wallets to create, buy, and sell meme coins with precise timing and high execution speed, thereby ensuring transaction inclusion in targeted blocks.

The adversary’s ultimate objective is to extract economic value from copy traders. Specifically, the adversary seeks to (i) select a historically profitable wallet to attract victims on meme coin tracking websites or social media, (ii) acquire meme coins at very early stages, either by launching the meme coins themselves (as *bundle*) or by entering immediately after meme coin creation (as *sniper*), (iii) conceal true positions by distributing purchases across

multiple controlled wallets, (iv) fake trading volume and social-media sentiment through wash trading and coordinated comment bots, thereby increasing copier confidence, and (v) strategically exit positions by coordinating sales across controlled wallets, thereby further strengthening the appearance of consistent profitability and sustaining the manipulation cycle.

We model the interaction between the adversary and copy traders as a near zero-sum game. The adversary’s profits come directly from the losses of copy traders who follow manipulated on-chain and off-chain signals.

4.2 Manipulative Bot Behaviors

To operationalize at scale, the adversary deploys automated bots that programmatically control wallets and execute transactions. We categorize these bots into two broad classes based on their primary role: *position manipulation bots* and *attention manipulation bots*.

4.2.1 Position Manipulation Bots. *Position manipulation bots*, including *bundle bots* and *sniper bots*, exploit temporal advantages at or near meme coin creation to secure (and conceal) positions before copy traders can react. Beyond early acquisition, these bots coordinate wallet splitting and synchronized execution to obfuscate true positions and evade detection. We further illustrate the underlying mechanisms and heterogeneity of frontrunning and concealment strategies. Fig. 2 provides evolution and heuristics of position manipulation bots in copy trading. Fig. 2a illustrates an unmanipulated copy-trading process. The meme coin creator launches a meme coin at block b_1 (①). A KOL wallet identifies the meme coin’s potential and purchases it at block b_3 (②). A copier wallet imitates the KOL’s action and buys the meme coin at block b_4 (③). From block b_4 to b_{4+i} , additional retail traders enter the market, trade the meme coin, and drive the price upward. The KOL sells the meme coin at a profit at block b_{6+i} (④). The copier wallet also sells the meme coin at a profit at block b_{7+i} (⑤).

While the process above reflects an unmanipulated copy-trading scenario, it assumes that the KOL acts honestly. To maximize profit, a KOL may instead employ a strategy known as a *naïve bundle*, as illustrated in Fig. 2b. Unlike unmanipulated copy trading, the KOL in this strategy creates the meme coin and purchases it himself within the same block b_1 (①). When a copier wallet subsequently

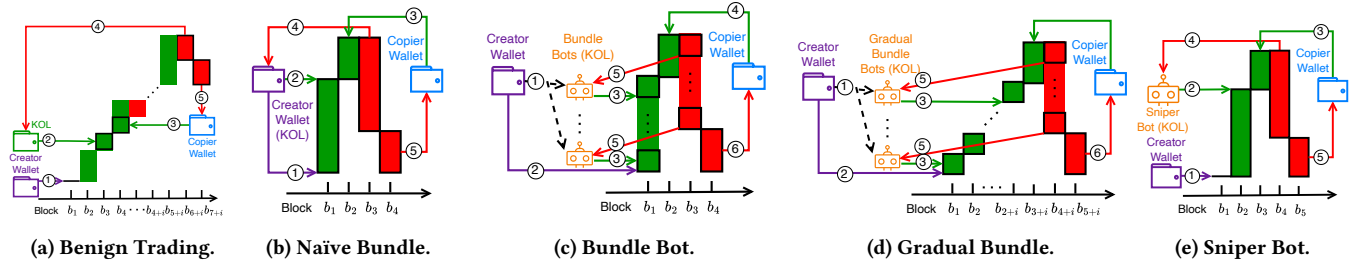


Figure 2: Evolution and heuristics of position manipulation bots in copy trading, as detailed in §4.2. → denotes buys, ↗ denotes sells, ↘ denotes coin creation, and -> denotes coin transfer. □ denotes price impact of each buy or sell action. The horizontal axis shows block number, and circled numbers indicate action order. A bot is a wallet controlled by an automated script. For each candlestick, the bottom and top of the body indicate the price at the beginning and end of the block, respectively: a green candle indicates that the end-of-block price is higher than the start-of-block price, while a red candle indicates the opposite.

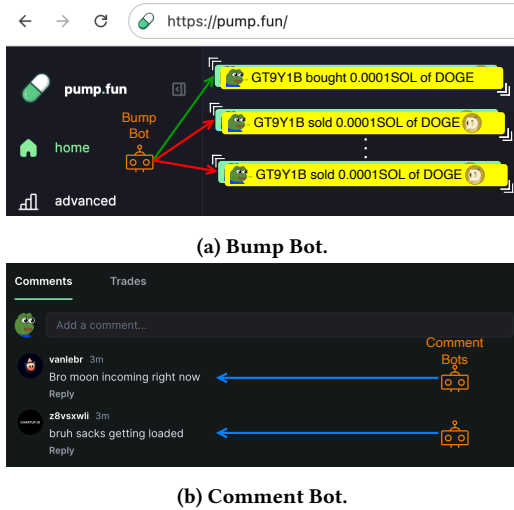


Figure 3: Heuristics of attention manipulation bots. → denotes buys, → denotes sells, and → denotes comments. A bot represents a wallet controlled by an automated script.

purchases the meme coin (2), the KOL (also the creator) immediately sells the meme coin for a profit (3), leaving the copiers with a loss (4).

However, this naïve strategy is easy to detect, since pump.fun and other meme coin tracking platforms publicly expose transaction histories and token holding distributions, with the meme coin creator explicitly identified (see §3.2 and Fig. 1). As a result, copiers can readily identify the creator’s concentrated ownership and avoid such meme coins. In response, adversaries adopt more sophisticated strategies, namely, the *bundle bot* (Fig. 2c) and the *gradual bundle bot* (Fig. 2d), to decouple meme coin creation from accumulation. In these strategies, the KOL first creates and funds multiple fresh wallets (1), launches the meme coin (2), and then uses automated scripts to coordinate purchases across these wallets, either within a single block (Fig. 2c) or gradually over multiple blocks (Fig. 2d) (3). Because the purchasing wallets appear independent, copiers are less likely to associate them with the meme coin creator and are more inclined to follow the observed buying activity (4). Finally, once sufficient copy-trading demand has accumulated, the KOL coordinates these wallets to exit their positions and dump the meme coin (5), leaving the copier with a loss (6).

In addition to bundle-based strategies, adversaries may employ a *sniper bot* to exploit copy trading without participating in meme coin creation, as shown in Fig. 2e. A sniper bot monitors newly launched meme coins and identifies those that have not yet been bundled by their creators. Upon detecting such a meme coin (1), the bot executes purchases within a very short time window following meme coin creation, typically within the first one to five blocks (approximately 0.4 to 2 seconds on Solana, see §3.1) (2). This early purchase subsequently attracts copier wallets, which replicate the observed buying activity and enter the market (3), driving up the meme coin price. The sniper bot then exits its position for a profit (4), leaving copier wallets exposed to losses (5).

4.2.2 Attention Manipulation Bots. Attention manipulation bots generate artificial trading activity and coordinated social signals to capture and amplify attention around specific meme coins. By faking demand and legitimacy, these bots steer copy traders’ attention toward targeted meme coins, reinforce their confidence in following observed trading behavior, and attract additional uninformed retail participants. We further decompose attention manipulation bots into two representative categories: *bump bots* and *comment bots*. Fig. 3 illustrates the heuristics of these bots, highlighting how on-chain activity and off-chain signals jointly amplify attention.

Bump bots exploit the bumping mechanism in pump.fun by repeatedly buying and selling identical amounts of a meme coin, thereby inflating short-term trading activity and visibility while maintaining minimal net exposure, as illustrated in Fig. 3a. Comment bots, in contrast, consist of wallets that automatically generate coordinated hype comments to simulate community engagement and amplify perceived interest, as shown in Fig. 3b¹.

4.3 Detection Metrics and Algorithms

We next explain how we detect manipulative bots and the results of the detection.

4.3.1 Bundle Bot. Based on the heuristic illustrated in Fig. 2c, a key characteristic of bundle bots is that, within the meme coin creation block, the creator wallet not only deploys the token but is immediately followed by near-simultaneous purchases from additional controlled wallets. Given Solana’s lack of a public mempool and its very short block interval (see §3.1), such same-block coordination is unlikely to be achieved by independent frontrunners and instead indicates pre-coordinated control across wallets. Accordingly, we detect bundle bots by checking whether non-creator wallets execute purchases within the meme coin creation block, as formalized in Alg. 1.

4.3.2 Sniper Bot. Similarly, sniper bots exhibit salient temporal characteristics, as illustrated in Fig. 2e. They typically execute purchases within the first one to five blocks after meme coin creation, a timescale well below human reaction time. We present the sniper bot detection algorithm in Alg. 2.

4.3.3 Bump Bot. Unlike bundle bots and sniper bots, bump bots do not exhibit salient temporal characteristics. Instead, their behavior is characterized by repetitive trading patterns. As illustrated in Fig. 3a, bump bots repeatedly buy and sell identical amounts of a meme coin. We define a *flip* as a pair of consecutive transactions with opposite directions and identical meme coin quantities to capture this repetition. Because bump bots are controlled by automated scripts with limited strategic complexity, their net positions are expected to remain largely unchanged over time, aside from minor variations due to normal market frictions. Accordingly, our metric rewards repetitive flip behavior while penalizing significant net position changes. Specifically, for meme coin i and wallet j , we define the bump bot score as the ratio of the number of flips to the net position change:

$$\alpha_{i,j} = \frac{F_{i,j}}{\Delta P_{i,j} + \epsilon}, \tag{1}$$

¹“Moon” is crypto jargon for a sharp price increase, while “sacks” is crypto jargon for accumulated capital.

Algorithm 1: Bundle Bot Detection.

Input: Meme coin address A , on-chain transactions \mathcal{T} for A
Output: $\mathbb{1}_{\text{Bundle Bot}} \in \{0, 1\}$
 Extract creator wallet w_c and launch block b_{launch} for A ;
 $\mathcal{T}_{\text{launch}} \leftarrow \{tx \in \mathcal{T} \mid tx.\text{block} = b_{\text{launch}}\}$;
if $\exists tx \in \mathcal{T}_{\text{launch}}$ s.t. $tx.\text{type} = \text{buy} \wedge tx.\text{trader} \neq w_c$ **then**
 | **return** 1; // simultaneous non-creator buys in launch block
return 0;

Algorithm 2: Sniper Bot Detection.

Input: Meme coin address A , on-chain transactions \mathcal{T} for A , block window K (default $K=5$)
Output: $\mathbb{1}_{\text{Sniper Bot}} \in \{0, 1\}$
 Extract launch block b_{launch} and creator wallet w_c for A ;
 $\mathcal{T}_{\text{early}} \leftarrow \{tx \in \mathcal{T} \mid b_{\text{launch}} < tx.\text{block} \leq b_{\text{launch}}+K\}$;
if $\exists tx \in \mathcal{T}_{\text{early}}$ s.t. $tx.\text{type} = \text{buy} \wedge tx.\text{trader} \neq w_c$ **then**
 | **return** 1; // non-creator buy within first K blocks
return 0;

Algorithm 3: Bump Bot Detection.

Input: Meme coin address A , on-chain transactions \mathcal{T} for A , threshold ξ (default $\xi = 50$), small constant ϵ (default $\epsilon = 1$)
Output: $\mathbb{1}_{\text{Bump Bot}} \in \{0, 1\}$
 Group \mathcal{T} by trader wallet to obtain $\{\mathcal{T}_w\}$;
foreach wallet w with transaction list \mathcal{T}_w **do**
 | Sort \mathcal{T}_w by block height and index within block;
 | $F \leftarrow 0$; // number of flips
 | $\Delta P \leftarrow 0$; // net position change
 | **for** $i \leftarrow 1$ **to** $|\mathcal{T}_w| - 1$ **do**
 | | $\Delta P \leftarrow \Delta P + \mathcal{T}_w[i]$;
 | | **if** $\mathcal{T}_w[i] + \mathcal{T}_w[i+1] = 0$ **then**
 | | | $F \leftarrow F + 1$; // a flip
 | | $\Delta P \leftarrow |\Delta P|$;
 | | **if** $\frac{F}{\Delta P + \epsilon} \geq \xi$ **then**
 | | | **return** 1; // wallet w is classified as a bump bot
return 0;

where $F_{i,j}$ is the number of flips, $\Delta P_{i,j}$ is the net position change (in meme coin units), and ϵ is a small constant to avoid a division-by-zero error (default $\epsilon=1$). We classify wallet j as a bump bot for meme coin i if $\alpha_{i,j} \geq \xi$, where ξ is an empirical threshold (default $\xi=50$). The corresponding detection procedure is given in Alg. 3.

4.3.4 Comment Bot. The most prominent characteristic of comment bots is their use of contextless, hype-driven, and slogan-like comments (e.g., “Bro moon incoming right now”), as illustrated in Fig. 3b. In addition, bot-generated comments typically do not reference other wallets (see §3.2). Such patterns are well suited to detection by LLMs. Accordingly, we manually curate a set of human-generated and bot-generated comments as few-shot examples and provide them in the LLM context using Prompt 1 in §A.2. We then use the LLM to classify comments associated with each meme coin. If more than one comment is classified as bot-generated, we label the meme coin as exhibiting comment bot activity.

4.3.5 Gradual Bundle. Gradual bundles represent a more sophisticated variant of bundle bots. As illustrated in Fig. 2d, they decouple token creation from purchases over time and may involve multi-layer funding networks. Combined with the scale of the Solana

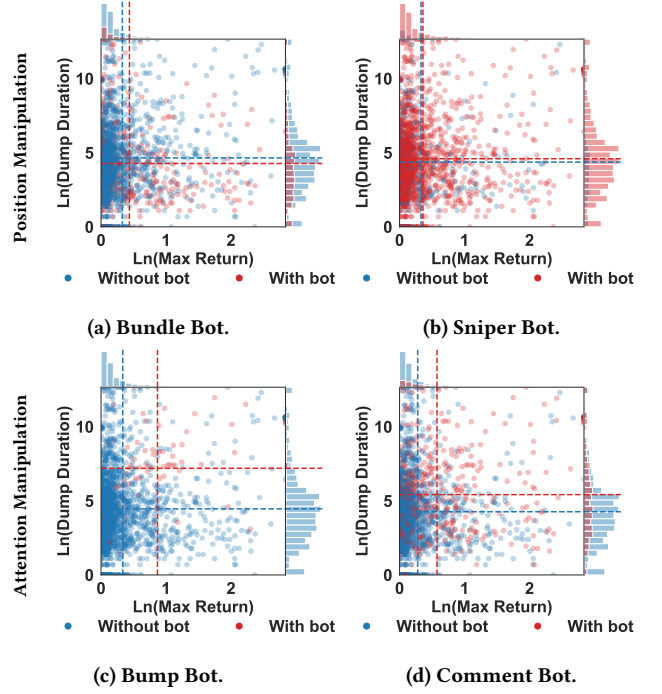


Figure 4: Distribution of performance metrics for meme coins with and without specific bots, grouped into position manipulation (first row) and attention manipulation bots (second row). The dashed line denotes the mean of the distribution.

blockchain (see §3.1), these properties obscure temporal signals and attribution, making gradual bundles much harder to detect accurately. In practice, experienced meme coin traders may rely on candlestick patterns to infer and avoid such coins, as the candlestick dynamics of gradual bundles differ systematically from those of non-manipulative tokens: the former tend to exhibit mechanically regular patterns, whereas the latter display more organic patterns.

4.4 Detection Results

For on-chain data, we use account-level transaction and transfer records obtained from Flipside, a blockchain data provider that maintains archive nodes with fully indexed, historical access to the Solana ledger. We additionally collect off-chain comment data from pump.fun. Our final dataset comprises 6,000 meme coin projects, including all associated creators and traders, along with their complete historical trading records.

Fig. 4 presents the distributions of meme coin projects with and without four types of bots, grouped into position manipulation and attention manipulation bots, and their heterogeneous effects on two performance metrics: maximum log return and log dump duration. The maximum log return is defined as the peak price appreciation during the lift phase, $\text{Ln}(\text{Max Return}) = \ln(P_{\text{max}}/P_{\text{launch}})$, where P_{max} and P_{launch} denote the peak and launch prices, respectively. The $\text{Ln}(\text{Dump Duration})$ is defined as the natural logarithm of seconds from the peak price until liquidity declines to 10% of its peak level. For position manipulation bots, the performance distributions

with and without bots largely overlap, indicating limited influence on either return or duration. Bundle bots appear in roughly one quarter of projects and are weakly associated with lower returns and shorter dump durations (Fig. 4a), while sniper bots are widespread but show minimal observable performance effect (Fig. 4b). This pattern suggests that the primary role of position manipulation bots is not to alter overall project performance, but to redistribute profits within the participant pool, shifting gains from slower retail traders to early or better-positioned adversaries. In contrast, attention manipulation bots, namely bump bots (Fig. 4c) and comment bots (Fig. 4d), are associated with improved performance metrics, although bump bots occur less frequently than comment bots, likely due to their higher operational cost.

5 Bot-Resilient Copy Trading

Based on the analysis in the previous section, we find that manipulative bots distort market efficiency, which in turn affects copy traders' profitability. This motivates the design of bot-resilient copy-trading frameworks that help copy traders follow profitable wallets in a given meme coin while avoiding harmful bot interactions and, where appropriate, leveraging attention manipulation bots.

We therefore focus on a fundamental profitability prediction task: given a wallet i that initiates its first position in meme coin j , we aim to predict whether the position of the wallet will ultimately be profitable (TRUE) or not (FALSE). Formally, we model this prediction as $profitable_{i,j} = \mathcal{M}(feature_1, feature_2, \dots)$, where $\mathcal{M}(\cdot)$ denotes our proposed **statistic-driven** and **MAS-based** models.

5.1 Settings

5.1.1 Sample Specification. We first describe the construction of our training samples. Using the wallet-level data described in §4.4, we chronologically order each trader's first transaction within a given meme coin to prevent **information leakage**. Based on this temporal ordering, we partition the data into training, validation, and test sets using a 70%–15%–15% split. Finally, we define the ground-truth label as the sign of the wallet's realized profit in the corresponding meme coin.

5.1.2 Feature Selection. We next describe the features selected in our models. A key assumption underlying copy trading is that selected KOL wallets possess superior trading expertise or informational advantages, enabling them to achieve consistent profitability. Accordingly, we place particular emphasis on features capturing wallet experience (e.g., number of trades), as well as the magnitude (e.g., average return) and stability (e.g., return standard deviation, t -statistic) of historical profitability across long- (11th–15th previously traded coins), medium- (6th–10th), and short-term (1st–5th) horizons. A trader's market entry timing is another important factor, which we capture using features such as purchase price, traded quantity, and dollar amount. In addition, we include four dummy variables indicating the presence of manipulative bots, as identified by the detection methods described in §4.3, to explicitly account for the influence of bot activity on wallet profitability.

However, these features alone are insufficient to rule out the gradual bundle behavior described in §4.3.5 or to capture the sentiment conveyed by human-generated comments, given that meme coin valuations are largely driven by hype. Therefore, we additionally incorporate candlestick charts and comments observed at the time when the wallet trades meme coin into our MAS-based model. Tab. 1 explains the features included in the prediction model.

5.1.3 Model. We then describe the proposed **statistic-driven** and **MAS-based** models.

Statistic-driven Models. We select one representative model from each major class of machine learning methods: a linear model (LASSO), a tree-based model (XGBoost), and a neural network (NN). Hyperparameters are tuned using the training and validation sets, and final performance is evaluated on the test set. Details of pre-processing and hyperparameter tuning are provided in §A.3.

MAS-based Models. We introduce an LLM-powered MAS-based model for bot-resilient copy trading. The framework consists of three specialized LLM agents with CoT reasoning: a *wallet agent*, a *coin agent*, and a *timing agent*. We employ GPT-4o as the base model for all agents².

²The model's training data extends only up to October 2023, which predates all meme coin events evaluated in this study and thus precludes temporal **information leakage**.

Table 1: Feature descriptions, inclusion, and conditions, as described in §5.1.2.

Feature	Description	Feature Inclusion			MAS Condition
		GMGN	Statistic	MAS	
Return ($[x]$)	The average return across the previous $[x]$ participated meme coins. [†]	●	●	●	> 0
Number of Trades	The total number of trades executed across all participated meme coins.	●	●	●	> 25% Training Set
Return Standard Deviation	The standard deviation of returns across all participated meme coins.	○	●	●	< 1
t -stat	The t -statistic of the trader's mean return across all participated meme coins.	○	●	●	> 1.645
Time Since Last Trade	The seconds elapsed since the most recent trade before the current meme coin.	○	●	●	< 75% Training Set
Time Since First Trade	The seconds elapsed since the trader's very first trade.	○	●	●	> 25% Training Set
Time Since Launch	The seconds elapsed since the meme coin's launch.	○	●	●	< 75% Training Set
Trader Purchase Price	The execution dollar price at which the trader purchases the meme coin.	○	●	●	> 25% Training Set
Trader Purchase Amount	The total dollar value of the trader's meme coin purchase.	○	●	●	> 25% Training Set
Trader Purchase Quantity	The number of meme coin units purchased by the trader in the transaction.	○	●	●	> 25% Training Set
$[y]$ Bot	Dummy equal to 1 if $[y]$ bot is detected at the trader's first trade, and 0 otherwise. [‡]	○	●	●	$\mathbb{1}_{\text{Bump Bot}} = 0$
Candlestick	The candlestick chart of the meme coin when the trader makes the first trade.	○	○	●	No Mechanicality
Comments	The comments of the meme coin when the trader makes the first trade.	○	○	●	Organically Positive

[†] $x \in \{\text{all, 1st, 1st-5th, 6th-10th, 11th-15th}\}$.

[‡] $y \in \{\text{Bundle, Sniper, Bump, Comment}\}$.

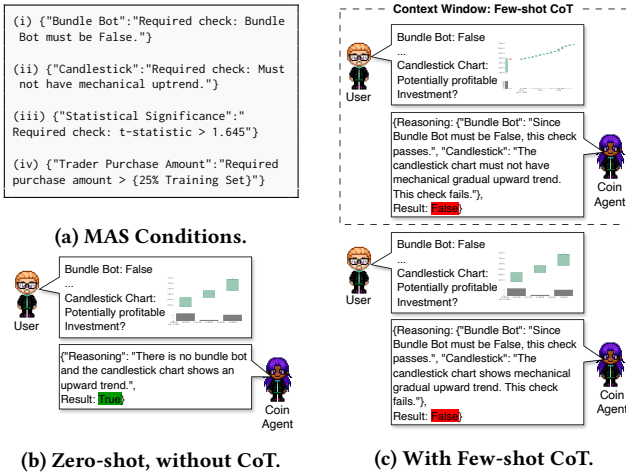


Figure 5: Comparison of model performance with and without CoT.

Fig. 5 illustrates how domain-specific knowledge is injected into each agent via few-shot CoT prompting. For each agent, we first decide the MAS conditions defined for each feature based on [6] and annotations from professional traders³. As shown in Fig. 5a and Tab. 1, we consider four types of conditions to make the prediction TRUE: (i) **Boolean** conditions enforcing binary constraints (e.g., the absence of bundle-bot behavior); (ii) **Visual or language-based** conditions based on candlestick pattern or textual sentiment cues; (iii) **Statistical** conditions imposing significance constraints (e.g., requiring t -statistics to exceed a threshold) or consistent directional effects; and (iv) **Continuous** conditions requiring feature values to fall within specific percentile ranges computed from the training set. We then select a small set of representative examples. Next, we verify whether these examples satisfy the conditions. From these, we retain a subset of profitable and unprofitable cases that satisfy all conditions and manually annotate them to construct agent-specific CoTs. These few-shot CoT exemplars are subsequently embedded into the prompts of the corresponding agents.

The MAS architecture is illustrated in Fig. 6. The step-by-step workflow of the MAS is as follows: ① We extract conditions for all features and incorporate these conditions into the construction of agent-specific CoTs. ② Each agent independently performs reasoning and produces a binary TRUE/FALSE decision together with the token logits. From each output, we compute the linear probability of the TRUE token and interpret it as a confidence score, following established confidence-based prompting techniques for LLMs [10, 11, 13, 22]. The three confidence scores form a probability vector. ③ We initialize a weight vector uniformly across agents. ④-⑤ The weight vector is optimized on the validation set to maximize the Area Under Curve (AUC) of the aggregated prediction. ⑥ The final prediction is obtained as a weighted combination of agent confidence scores and is evaluated on the test set. Detailed few-shot CoTs exemplars and prompts are provided in §A.5.

³The MAS condition design is also broadly aligned with the direction of feature contributions revealed by the SHAP analysis of the statistic-driven models in Fig. 9.

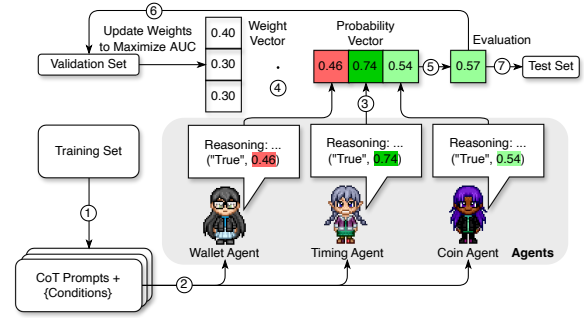


Figure 6: The architecture of our MAS-based copy-trading framework with CoT reasoning. Circled numbers indicate action order.

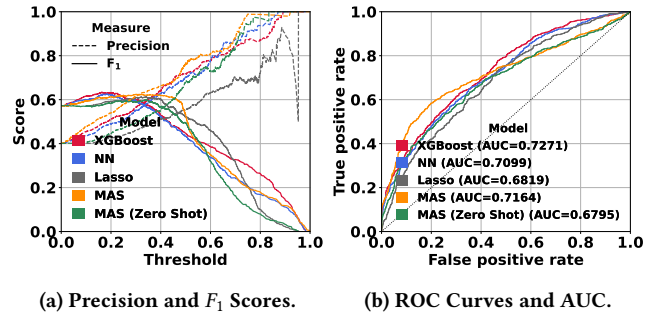


Figure 7: Test set performance comparison between statistic-driven and MAS-based models.

To assess the efficacy of CoT reasoning, we additionally train a zero-shot MAS baseline with all other components held constant, as illustrated in Fig. 5b.

5.2 Prediction Accuracy

Both statistic-driven and MAS-based models output a continuous likelihood score in $[0, 1]$ indicating whether a wallet is expected to be profitable. We obtain binary TRUE/FALSE predictions by applying a decision threshold. Fig. 7 reports out-of-sample performance on the test set as the threshold varies, using Precision, F_1 score, and AUC as evaluation metrics.

Fig. 7a shows Precision and F_1 scores across thresholds. Among statistic-driven models, XGBoost consistently dominates LASSO and performs comparably to the NN, particularly at moderate-to-high thresholds. LASSO exhibits unstable Precision and rapidly degrading F_1 scores as the threshold increases, reflecting its limited discriminative capacity in this setting. For MAS-based approaches, incorporating few-shot CoT reasoning leads to higher Precision and F_1 scores than the zero-shot MAS baseline across nearly all thresholds, indicating better calibration of confidence scores.

Fig. 7b presents receiver operating characteristic (ROC) curves and corresponding AUC values. XGBoost achieves the highest AUC overall, followed closely by the MAS with CoT reasoning. The NN and LASSO exhibit lower AUC values, while the zero-shot MAS performs worst, highlighting the importance of structured

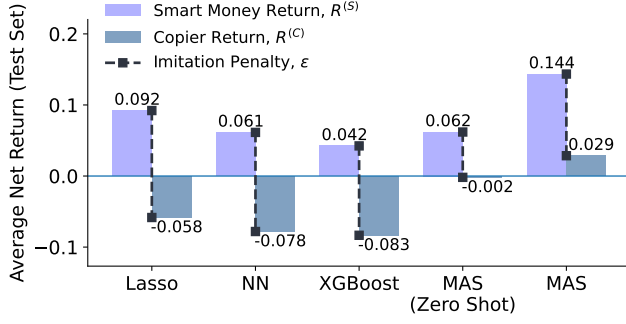


Figure 8: Test set average net return of smart money and copier with market frictions.

reasoning and agent specialization for robust ranking performance. In addition, we also present the feature importance in §A.4.

5.3 Economic Performance

In meme coin copy trading, identifying a profitable wallet is not sufficient to guarantee copier profitability, because execution occurs on a bonding curve where price impact depends on trade ordering. Even when a copier immediately replicates the same trades as a smart money wallet, the copier necessarily trades at a worse point on the curve. This creates a mechanical performance gap that we refer to as the imitation penalty.

To formalize this effect, we derive the returns of smart money and its one-to-one immediate copiers under the pump.fun bonding-curve execution model introduced in §A.1. The following lemma characterizes their respective returns.

LEMMA 1 (NET RETURNS OF SMART MONEY AND COPIER). Consider a smart-money wallet that executes a sequence of T trades indexed by $t = 1, \dots, T$, where trade sizes and timing are fixed and not affected by copier behavior. Each trade is characterized by a signed token quantity $q_t \in \mathbb{R}$, where $q_t > 0$ denotes a buy and $q_t < 0$ denotes a sell. Let Y_t^S denote the effective token reserve immediately before the smart money’s trade t . Feasibility requires $Y_t^S > |q_t|$. The net returns of smart money, $R^{(S)}$, and the corresponding one-to-one immediate copier, $R^{(C)}$, are

$$R^{(S)} = \frac{-\sum_{t:q_t < 0} \frac{kq_t}{Y_t^S(Y_t^S - q_t)}}{\sum_{t:q_t > 0} \frac{kq_t}{Y_t^S(Y_t^S - q_t)}} - 1, \quad R^{(C)} = \frac{-\sum_{t:q_t < 0} \frac{kq_t}{(Y_t^S - q_t)(Y_t^S - 2q_t)}}{\sum_{t:q_t > 0} \frac{kq_t}{(Y_t^S - q_t)(Y_t^S - 2q_t)}} - 1.$$

THEOREM 2 (IMITATION PENALTY). Under the above execution model, the copier strictly overpays on every replicated buy. Define the imitation penalty as $\epsilon = R^{(S)} - R^{(C)}$, then, $\epsilon > 0$.

Proofs of Lemma 1 and Theorem 2 are provided in §A.6.

Based on Lemma 1 and Theorem 2, Fig. 8 reports test-set average smart money wallet net returns, corresponding copier net return, and imitation penalty under market frictions. Although smart money returns are positive across all models, copier returns are consistently lower, indicating a persistent imitation penalty arising from price impact. For statistic-driven models, positive smart money returns are insufficient to yield positive copier profitability, with copier returns remaining negative despite accurate wallet

Table 2: Ablation study of the proposed MAS.

Ablation	AUC	Gross Return (Net Return+1)	
		Smart Money	Copier
w/o Wallet Agent	0.6382▼11.3%	1.0753▼6.0%	0.9327▼9.4%
w/o Coin Agent	0.6489▼9.8%	1.0837▼5.2%	0.9415▼8.5%
w/o Timing Agent	0.6943▼3.5%	1.1048▼3.4%	1.0124▼1.6%

identification. In contrast, the MAS-based model achieves the highest smart money returns and is the only approach that delivers a positive average copier return, demonstrating improved economic effectiveness under realistic market conditions.

5.4 Ablation Study

We conduct an ablation study to assess the contribution of each agent in the proposed MAS-based framework. Tab. 2 reports AUC and gross return (net return +1) for both smart money wallets and copiers when removing each agent. Removing the *wallet agent* leads to the largest degradation across all metrics, with an AUC drop of 11.3% and pronounced declines in both smart money (6.0%) and copier returns (9.4%), underscoring the central role of wallet-level signals in identifying profitable trading behavior. Excluding the *coin agent* also results in substantial performance losses (AUC down 9.8%), confirming that token-level characteristics and embedded manipulative behaviors materially affect copy-trading outcomes. In contrast, removing the *timing agent* yields a smaller but non-negligible decline, suggesting that entry timing provides complementary information but is less dominant than wallet- and coin-level assessments.

6 Conclusion

In this paper, we propose a LLM-powered multi-agent framework that explicitly reasons over multi-modal on-chain and off-chain signals. Empirical evaluation on over 6,000 meme coin projects shows that our approach outperforms zero-shot and most statistic-driven baselines in prediction accuracy and all baselines in profitability. Overall, our results demonstrate the effectiveness of agent-based defenses in adversarial meme coin markets and establish that trader profitability is predictable, providing a practical foundation for robust copy trading.

Acknowledgments

The authors gratefully acknowledge Ievgen Gerasymchuk, Wenzhi Ding, and Stefanos Chaliasos for their insightful comments.

This research is funded by the National Research Foundation, Singapore, and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-GC-2023-008), the National Research Foundation, Singapore, and the Cyber Security Agency under its National Cybersecurity R&D Programme (NCRP25-P04-TAICeN), the National Research Foundation, Prime Minister’s Office, Singapore under Campus for Research Excellence and Technological Enterprise (CREATE) programme, Engineering and Physical Sciences Research Council (EPSRC) under UKFin+ Programme, and Ripple under University Blockchain Research Initiative (UBRI) [8].

References

- [1] Arash Aloosh and Jiasun Li. 2024. Direct Evidence of Bitcoin Wash Trading. *Management Science* 70, 12 (3 2024), 8875–8921.
- [2] Jose Apesteguia, Jörg Oechssler, and Simon Weidenholzer. 2020. Copy Trading. *Management Science* 66, 12 (7 2020), 5608–5622.
- [3] Lin William Cong, Xi Li, Ke Tang, and Yang Yang. 2023. Crypto Wash Trading. *Management Science* 69, 11 (9 2023), 6427–6454.
- [4] Walter Hernandez Cruz, Firas Dahi, Yebo Feng, Jiahua Xu, Aanchal Malhotra, and Paolo Tasca. 2025. AMM-based DEX on the XRP Ledger. In *IEEE International Conference on Blockchain and Cryptocurrency*.
- [5] Anirudh Dhawan and Tălis J. Putniņš. 2023. A New Wolf in Town? Pump-and-Dump Manipulation in Cryptocurrency Markets. *Review of Finance* 27, 3 (5 2023), 935–975.
- [6] Wenzhi Ding, Chen Lin, Yichen Luo, and Jiahua Xu. 2025. Decompose Market Manipulation Strategies: Evidence from On-chain Meme Coin Market. *SSRN Electronic Journal*:5953738 (9 2025).
- [7] Yimika Erinle, Yathin Kethepalli, Yebo Feng, and Jiahua Xu. 2025. SoK: Design, vulnerabilities, and security measures of cryptocurrency wallets. *Computer Networks* 123 (12 2025), 111691.
- [8] Yebo Feng, Jiahua Xu, and Lauren Weymouth. 2022. University Blockchain Research Initiative (UBRI): Boosting blockchain education and research. *IEEE Potentials* 41, 6 (2022), 19–25.
- [9] Bo Gao, Yuan Wang, Qingsong Wei, Yong Liu, Rick Siow, Mong Goh, David Lo, and Bo Gao. 2025. AiRacleX: Automated Detection of Price Oracle Manipulations via LLM-Driven Knowledge Mining and Prompt Generation. *arXiv:2502.06348* (2 2025).
- [10] Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koeppl, Preslav Nakov, and Iryna Gurevych. 2024. A Survey of Confidence Estimation and Calibration in Large Language Models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1. 6577–6595.
- [11] Sanyam Kapoor, Nate Gruver, Manley Roberts, Katherine Collins, Arka Pal, Umang Bhatt, Adrian Weller, Samuel Dooley, Micah Goldblum, and Andrew Gordon Wilson. 2024. Large Language Models Must Be Taught to Know What They Don't Know. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, Vol. 37. 85932–85972.
- [12] Daisuke Kawai, Kyle Soska, Bryan Routledge, Ariel Zetlin-Jones, and Nicolas Christin. 2024. Stranger Danger? Investor Behavior and Incentives on Cryptocurrency Copy-Trading Platforms. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. New York, NY, USA, 1–20.
- [13] Reza Khanmohammadi, Erfan Miahi, Mehra Mardikoraem, Simerjot Kaur, Ivan Brugere, Charese H Smiley, Kundan Thind, Mohammad M Ghassemi, Jpmorgan Ai Research, and Henry Ford Health. 2025. Calibrating LLM Confidence by Probing Perturbed Representation Stability. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*. 10448–10514.
- [14] Yuan Li, Bingqiao Luo, Qian Wang, Nuo Chen, Xu Liu, and Bingsheng He. 2024. CryptoTrade: A Reflective LLM-based Agent to Guide Zero-shot Cryptocurrency Trading. In *Empirical Methods in Natural Language Processing*. 1094–1106.
- [15] Yinheng Li, Shaofei Wang, Han Ding, and Hang Chen. 2023. Large Language Models in Finance: A Survey. *arXiv:2311.10723* (2023).
- [16] Yueyao Li, Nanjun Yao, Yuhui Huo, and Wei Cai. 2025. Trust Dynamics and Bot-Driven Responses: An Approach to Rug Pulls in Solana Meme Coin Markets. In *Proceedings of the 17th ACM Web Science Conference*, Vol. 25. 106–116.
- [17] Hao Liu, Carmelo Sfrazza, and Pieter Abbeel. 2023. Chain of Hindsight Aligns Language Models with Feedback. In *12th International Conference on Learning Representations*.
- [18] Houwan Long, Hongyang Li, and Wei Cai. 2025. CoinCLIP: A Multimodal Framework for Assessing Viability in Web3 Memecoins. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*. 4986–4990.
- [19] Hou Wan Long, Nga Man Wong, and Wei Cai. 2025. Bridging Culture and Finance: A Multimodal Analysis of Memecoins in the Web3 Ecosystem. In *Companion Proceedings of the ACM Web Conference 2025*. 1158–1161.
- [20] Scott Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Vol. 2017-December. 4768–4777.
- [21] Yichen Luo, Yebo Feng, Jiahua Xu, and Paolo Tasca. 2026. Piercing the Veil of TVL: DeFi Reappraised. In *Financial Cryptography and Data Security*. 3–19.
- [22] Yichen Luo, Yebo Feng, Jiahua Xu, Paolo Tasca, and Yang Liu. 2025. LLM-Powered Multi-Agent System for Automated Crypto Portfolio Management. *arXiv:2501.00826* (1 2025).
- [23] Qian'ang Mao, Yuxuan Zhang, Jiaman Chen, Wenjun Zhou, and Jiaqi Yan. 2025. Know Your Intent: An Autonomous Multi-Perspective LLM Agent Framework for DeFi User Transaction Intent Mining. *arXiv:2511.15456* 1 (2025).
- [24] Bruno Mazorra, Victor Adan, and Vanesa Daza. 2022. Do Not Rug on Me: Leveraging Machine Learning Techniques for Automated Scam Detection. *Mathematics* 10, 6 (3 2022), 949.
- [25] Alberto Maria Mongardini and Alessandro Mei. 2026. A Midsummer Meme's Dream: Investigating Market Manipulations in the Meme Coin Ecosystem. *arXiv:2507.01963* (1 2026).
- [26] Minh Hoang Nguyen, Phuong Duy Huynh, Son Hoang Dau, and Xiaodong Li. 2023. Rug-pull malicious token detection on blockchain using supervised learning with feature engineering. In *Australasian Computer Science Week*. 72–81.
- [27] Natapong Nitarch, Warit Sirichotedumrong, Panop Pitchayarthorn, Pittawat Tavekitworachai, Potsawee Manakul, and Kunat Pipatanakul. 2025. FinCoT: Grounding Chain-of-Thought in Expert Financial Reasoning. *arXiv:2506.16123* (6 2025).
- [28] Yuanzheng Niu, Xiaoqi Li, Hongli Peng, and Wenkai Li. 2024. Unveiling Wash Trading in Popular NFT Markets. In *Companion Proceedings of the ACM Web Conference*. 730–733.
- [29] Daniel Perez, Jiahua Xu, and Benjamin Livshits. 2020. Revisiting Transactional Statistics of High-scalability Blockchains. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference*. 535–550.
- [30] Zhe Shen, Wei Jiang, and Zhiqiang (Eric) Zheng. 2025. Irrationality-Aware Human Machine Collaboration: Mitigating Afterfactual Irrationality in Copy Trading. *Information System Research* (4 2025).
- [31] Dianxiang Sun, Wei Ma, Liming Nie, and Yang Liu. 2025. SoK: A Taxonomic Analysis of DeFi Rug Pulls: Types, Dataset, and Tool Assessment. *Proceedings of the ACM on Software Engineering* 2, ISSTA (6 2025), 550–572.
- [32] Qian Wang, Yuchen Gao, Zhenheng Tang, Bingqiao Luo, Nuo Chen, and Bingsheng He. 2024. Exploring LLM Cryptocurrency Trading Through Fact-Subjectivity Aware Reasoning. In *ICLR 2025 Advances in Financial AI Workshop*.
- [33] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, Vol. 35. 24824–24837.
- [34] Zhiyuan Wei, Jing Sun, Yuqiang Sun, Ye Liu, Daoyuan Wu, Zijian Zhang, Xianhao Zhang, Meng Li, Yang Liu, Chunmiao Li, Mingchao Wan, Jin Dong, and Liehuang Zhu. 2025. Advanced Smart Contract Vulnerability Detection via LLM-Powered Multi-Agent Systems. *IEEE Transactions on Software Engineering* 51, 10 (10 2025), 2830–2846.
- [35] Cong Wu, Jing Chen, Jiahong Li, Jiahua Xu, Ju Jia, Yutao Hu, Yebo Feng, Yang Liu, and Yang Xiang. 2025. Profit or Deceit? Mitigating Pump and Dump in DeFi via Graph and Contrastive Learning. *IEEE Transactions on Information Forensics and Security* 20 (2025), 8994–9008.
- [36] Yuexin Xiang, Qishuang Fu, Yuquan Li, Qin Wang, Tsz Hon Yuen, and Jiangshan Yu. 2025. Measuring Memecoin Fragility. *arXiv:2512.00377* (11 2025).
- [37] Zhuohan Xie, Daniil Orel, Rushil Thareja, Dhruv Sahnani, Hachem Madmoun, Fan Zhang, Debopriyo Banerjee, Georgi Georgiev, Xueqing Peng, Lingfei Qian, Jimin Huang, Jinyan Su, Aaryamonvikram Singh, Rui Xing, Rania Elbadry, Chen Xu, Haonan Li, Fajri Koto, Ivan Kozychev, Tammy Chakraborty, Yuxia Wang, Salem Lahlou, Veselin Stoyanov, Sophia Ananiadou, Preslav Nakov, and St Kliment Ohridski. 2025. FinChain: A Symbolic Benchmark for Verifiable Chain-of-Thought Financial Reasoning. *arXiv:2506.02515* (6 2025).
- [38] Siheng Xiong, Ali Payani, and Faramarz Fekri. 2025. Enhancing Long Chain-of-Thought Reasoning through Multi-Path Plan Aggregation. *arXiv:2510.11620* (10 2025).
- [39] Siheng Xiong, Ali Payani, Yuan Yang, and Faramarz Fekri. 2025. Deliberate Reasoning in Language Models as Structure-Aware Planning with an Accurate World Model. *Proceedings of the Annual Meeting of the Association for Computational Linguistics* 1 (2025), 31900–31931.
- [40] Jiahua Xu, Yebo Feng, Daniel Perez, and Benjamin Livshits. 2025. Auto.gov: Learning-Based Governance for Decentralized Finance (DeFi). *IEEE Transactions on Services Computing* 18, 3 (2025), 1278–1292.
- [41] Jiahua Xu and Benjamin Livshits. 2019. The Anatomy of a Cryptocurrency Pump-and-Dump Scheme. In *Proceedings of the 28th USENIX Conference on Security Symposium*. 1609–1625.
- [42] Jiahua Xu, Krzysztof Paruch, Simon Cousaert, and Yebo Feng. 2023. SoK: Decentralized Exchanges (DEX) with Automated Market Maker (AMM) Protocols. *Comput. Surveys* 55, 11 (11 2023).
- [43] Jingfeng Yang, Hongye Jin, Ruixiang Tang, Xiaotian Han, Qizhang Feng, Haoming Jiang, Shaochen Zhong, Bing Yin, and Xia Hu. 2024. Harnessing the Power of LLMs in Practice: A Survey on ChatGPT and Beyond. *ACM Transactions on Knowledge Discovery from Data* 18, 6 (7 2024), 1–32.
- [44] Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. 2024. A Survey on Multimodal Large Language Models. *National Science Review* (11 2024).
- [45] Jingyi Zhang, Jiaying Huang, Sheng Jin, and Shijian Lu. 2024. Vision-Language Models for Vision Tasks: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 46, 8 (2024), 5625–5644.
- [46] Juantao Zhong, Daoyuan Wu, Ye Liu, Maoyi Xie, Yang Liu, Yi Li, and Ning Liu. 2025. Detecting Various DeFi Price Manipulations with LLM Reasoning. In *Proceedings of IEEE/ACM Automated Software Engineering*.

A Appendix

A.1 Bonding Curve Mechanism

Pump.fun adopts a constant-product bonding curve to govern the relationship between deposited SOL and received tokens [4, 42]:

$$y = y' - \frac{k}{x + x'}, \tag{2}$$

where x is the amount of SOL deposited, and y is the corresponding number of meme coins issued. The parameters $x', y' > 0$ denote protocol-fixed virtual reserves.

Define effective reserves $X \equiv x' + x$ and $Y \equiv y' - y$, so that the curve satisfies the invariant $XY = k = x'y'$.

Monotonicity and concavity. Taking derivatives with respect to cumulative demand x , we obtain

$$\frac{dy}{dx} = \frac{k}{(x' + x)^2} > 0, \quad \frac{d^2y}{dx^2} = -\frac{2k}{(x' + x)^3} < 0,$$

for all $x \geq 0$, since $k > 0$ and $x' > 0$. Hence, the issuance function $y(x)$ is strictly increasing and strictly concave: additional SOL deposits always yield more tokens, but at a diminishing rate.

The marginal price of the meme coin, denominated in SOL, is the reciprocal slope of the issuance function:

$$p = \frac{dx}{dy} = \frac{(x' + x)^2}{k} = \frac{k}{(y' - y)^2}. \tag{3}$$

Differentiating yields

$$\frac{dp}{dx} = \frac{2(x' + x)}{k} > 0, \quad \frac{d^2p}{dx^2} = \frac{2}{k} > 0,$$

implying that the price is strictly increasing and strictly convex in cumulative demand.

A.2 Comment Bot Detection Prompt

To complete the comment bot detection task in §4.3.4, we design a prompt to classify whether a given comment is human-generated or bot-generated based on its semantic content. The prompt provides the LLM with a small set of manually curated few-shot examples, including both human-generated comments and bot-generated comments exhibiting contextless, hype-driven, or slogan-like language. Each target comment is then presented independently, and the LLM outputs a binary classification indicating whether the comment is bot-generated. The prompt template is provided in Prompt 1.

Prompt 1: Comment Bot Detection.

You are a meme coin comment analyzer. Your task is to classify a given comment as bot-generated or human-generated. Bot-generated comments are often short, context-less, and mass-producible slogans that express hype or hostility. In contrast, human-generated comments tend to be more personalized and nuanced, containing context or opinion with reasoning. Comments that reference other users (e.g., @89009679) are typically human-generated, although not all human comments contain such references. Respond with id and true (if the comment is bot-generated) or false (if it is human-generated). Your response should follow this format: {"result": <true/false>}

TO THE MOON!!! READYD {"result": true}
 we'll get there! LFG {"result": true}
 @88857219 show screenshot as proof pls? {"result": false}
 Fake web bros, not same ca {"result": false}

{comment}

Table 3: Hyperparameters of statistic-driven models.

Model	Hyperparams	Search Space
LASSO	C	$\{10^1, \dots, 10^4\}$
	Solver	liblinear
	Class weight	balanced
NN	Hidden layers	$\{(32), (32, 16), (32, 16, 8), (32, 16, 8, 4), (32, 16, 8, 4, 2)\}$
	α	$\{10^{-5}, 10^{-4}, 10^{-3}\}$
	Learning rate η	$\{10^{-4}, 5 \times 10^{-4}, 10^{-3}, 10^{-2}\}$
	Optimizer	Adam
XGBoost	$n_{\text{estimators}}$	$\{300, 600\}$
	Max depth	$\{2, 3, 4, 6\}$
	Learning rate	$\{0.01, 0.05, 0.1\}$
	Subsample	$\{0.7, 0.9, 1.0\}$
	Column sample	$\{0.7, 0.9, 1.0\}$
	Min child weight	$\{1, 5, 10\}$
	λ (L2)	$\{1, 5, 10\}$

A.3 Preprocessing and Hyperparameter Tuning for Statistic-driven Models

This section documents the preprocessing procedures and hyperparameter tuning used for the statistic-driven models. Continuous features are first winsorized at the 2.5% and 97.5% percentiles to mitigate the influence of outliers, and subsequently standardized to zero mean and unit variance using statistics computed from the training set. The same preprocessing transformations are then applied to the validation and test sets. All model hyperparameters are tuned using the training and validation sets, and the corresponding hyperparameter search spaces are summarized in Tab. 3.

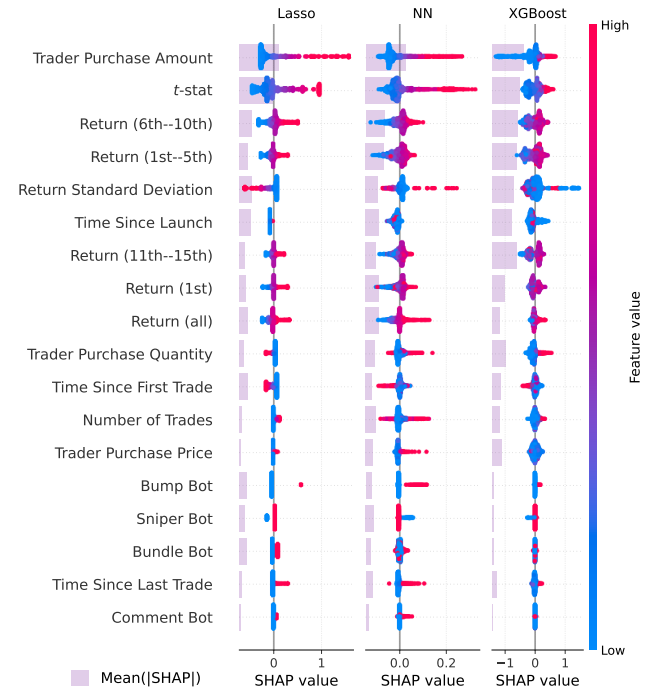


Figure 9: SHAP summary computed on the validation set.

A.4 Feature Importance

In this section, we measure feature importance using SHAP on the validation set [20], as summarized in Fig. 9. For LASSO, importance is computed using SHAP LinearExplainer values. For XGBoost, we use SHAP TreeExplainer to derive feature attributions. For the NN, we compute SHAP values using PermutationExplainer.

A.5 Few-shot CoTs and Prompts for the MAS

This section describes how we employ few-shot CoT prompting to guide the wallet, coin, and timing agents within the proposed MAS.

A.5.1 CoT. In this section, we present the few-shot CoT exemplars used to guide agent reasoning, as provided in CoT 1, 2, and 3. Each exemplar demonstrates how domain-specific signals should be interpreted, how intermediate reasoning steps should be structured, and how final decisions should be derived. By providing explicit reasoning patterns, the CoT examples improve decision consistency, enhance interpretability, and reduce spurious or unstable reasoning behaviors across agents.

CoT 1: Wallet Agent.

T-statistic of Returns: 24.39	Average Return: 1.25
Return Standard Deviation: 0.84	Number of Trades: 4114
Last Return: 0.12	Five-to-One Return: 0.14
Ten-to-Six Return: 0.79	Fifteen-to-Eleven Return: 0.48
Time Since Last Trade: 371	Time Since First Trade: 19118974

Listing 1: Wallet Agent Reasoning (Positive Case)

```
{
  "reasoning": {
    "Statistical Significance": "Gate check: require t-statistic > 1.645 (one-tailed 5%). Observed = 24.39. Since 24.39 > 1.645, the gate passes.",
    "Profitability": "Check average return > 1.00. Observed = 1.25. Since 1.25 > 1.00, this check passes.",
    "Risk Profile": "Check standard deviation < 1.00. Observed = 0.84. Since 0.84 < 1.00, this check passes.",
    "Temporal Robustness": "Require 5-to-1, 10-to-6, and 15-to-11 returns all > 1.00. Observed = (0.12, 0.14, 0.79, 0.48). All exceed 1.00, so this check passes.",
    "Statistical Robustness": "Require number of trades > {25% Training Set}. Observed = 4114. Since 4114 > {25% Training Set}, this check passes.",
    "Activity and Longevity": "Recency requires time since last trade < {75% Training Set}: 371 < {75% Training Set} passes. Longevity prefers time since first trade > {25% Training Set}: 19118974 > {25% Training Set}, so this check passes.",
    "Summary": "All checks pass (significance, profitability, risk control, temporal robustness, longevity, and sample size). The wallet is attractive for next-trade copy trading.",
    "result": true
  }
}
```

T-statistic of Returns: 0.45	Average Return: 0.00
Return Standard Deviation: 0.31	Number of Trades: 461
Last Return: 0.02	Five-to-One Return: -0.02
Ten-to-Six Return: 0.11	Fifteen-to-Eleven Return: -0.12
Time Since Last Trade: 10	Time Since First Trade: 15449

Listing 2: Wallet Agent Reasoning (Negative Case)

```
{
  "reasoning": {
    "Statistical Significance": "Gate check: require t-statistic > 1.645 (one-tailed 5%). Observed = 0.45. Since 0.45 < 1.645, the gate fails.",
    "Profitability": "Require average return > 0.00. Observed = 0.00. Since 0.00 is not greater than 0.00, this check fails.",
    "Risk Profile": "Check standard deviation < 1.00. Observed = 0.31. This passes, but cannot compensate for failed core profitability and significance gates.",
    "Temporal Robustness": "Require all horizons > 1.00. Observed = (0.02, -0.02, 0.11, -0.12). Multiple horizons are below 0.00, so this check fails.",
    "Statistical Robustness": "Require number of trades > {25% Training Set}. Observed = 461. Since 461 < {25% Training Set}, this check fails.",
    "Activity and Longevity": "Recency 10 < {75% Training Set} passes; longevity 15449 < {25% Training Set} not met. Auxiliary signals cannot offset failed core checks.",
    "Summary": "Core gates fail (statistical significance and profitability) and robustness is insufficient. The wallet is unsuitable for next-trade copy trading.",
    "result": false
  }
}
```

CoT 2: Coin Agent.

Transaction Features:

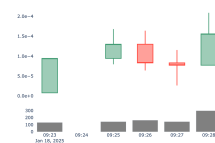
Bundle Bot: False
Sniper Bot: False
Bump Bot: True
Comment Bot: True

Comment History:

2025-01-17 15:06:36 - 3yxCdw: Makers missing? Upgrade your strategy now!

...

Candlestick Chart:



Listing 3: Coin Agent Reasoning (Positive Case)

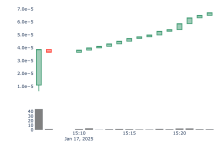
```
{
  "reasoning": {
    "Bundle Bot": "Required check: Bundle Bot must be False. Observed = False. Since False == False, this check passes.",
    "Candlestick Pattern": "Required check: chart should show gradual, sustained price discovery (no single-candle spike-and-dump). Observed pattern indicates a gradual increase, so this check passes.",
    "Sniper Bot": "Auxiliary check: Sniper Bot ideally False. Observed = False. This reduces early predatory trading risk.",
    "Bump Bot": "Auxiliary check: Bump Bot may be True as a weak visibility/support signal but is not required. Observed = True; treat as weakly supportive.",
    "Comments": "Auxiliary check: discount repetitive spam; prefer evidence of non-boilerplate engagement. Observed Comment Bot = True; treat as caution, but not decisive if other required signals are clean.",
    "Summary": "Required checks pass (no bundle and no pump-like candlestick signature). Auxiliary signals are not contradictory. Classify as a good investment opportunity.",
    "result": true
  }
}
```

Transaction Features:

Bundle Bot: True
Sniper Bot: True
Bump Bot: False
Comment Bot: False

Comment History:

Candlestick Chart:



Listing 4: Coin Agent Reasoning (Negative Case)

```
{
  "reasoning": {
    "Bundle Bot": "Required check: Bundle Bot must be False. Observed = True. Since True != False, this check fails.",
    "Candlestick Pattern": "Required check: chart should show gradual, sustained price discovery. Observed pattern resembles a spike/pump-like move, so this check fails.",
    "Sniper Bot": "Auxiliary check: Sniper Bot ideally False. Observed = True; negative signal indicating elevated predatory early trading risk.",
    "Bump Bot": "Auxiliary check: Bump Bot may be supportive if True. Observed = False; not supportive.",
    "Comments": "Auxiliary check: prefer sustained, organic engagement. Observed Comment Bot = False; absent bot flags alone is not supportive without additional evidence, treat as non-decisive.",
    "Summary": "One or more required checks fail (bundle present and pump-like price action). Auxiliary signals do not offset these failures. Classify as a poor investment opportunity.",
    "result": false
  }
}
```

A.5.2 Prompt. We then present the prompts used by each agent in the proposed MAS, as provided in Prompt 2, 3, and 4. The prompts specify the agent's role, available inputs, and decision objectives, and instruct the LLM to reason before producing a final decision. This design ensures consistent task alignment across agents while allowing each agent to specialize in its respective subtask.

A.6 Smart Money Return, Copier Return, and Imitation Penalty

In this section, we provide proofs of Lemma 1 and Theorem 2.

PROOF OF LEMMA 1. Under the constant-product invariant $XY = k$ derived in §A.1, executing a trade of size q at effective reserve Y induces the following SOL flow:

$$\Delta X(Y, q) \equiv \frac{k}{Y - q} - \frac{k}{Y} = \frac{kq}{Y(Y - q)}.$$

CoT 3: Timing Agent.

Time Since Launch: 3723
 Trader Purchase Price: 1.19e-05
 Trader Purchase Amount: 661.21
 Trader Purchase Quantity: 55153573.25

Listing 5: Timing Agent Reasoning (Positive Case)

```
{"reasoning":{"Time Since Launch":"Require time since launch < {75% Training Set}. Observed = 3723. Since 3723 < {75% Training Set}, this check passes.", "Trader Purchase Price":"Require purchase price > {25% Training Set}. Observed = 1.19e-05. Since 1.19e-05 > {25% Training Set}, this check passes.", "Trader Purchase Amount":"Require purchase amount > {25% Training Set}. Observed = 661.21. Since 661.21 > {25% Training Set}, this check passes.", "Trader Purchase Quantity":"Require purchase quantity > {25% Training Set}. Observed = 55153573.25. Since 55153573.25 > {25% Training Set}, this check passes.", "Summary":"All timing checks pass, indicating suitable entry timing for copy trading this wallet in the meme coin."},"result":true}
```

Time Since Launch: 85374
 Trader Purchase Price: 4.95e-06
 Trader Purchase Amount: 99.65
 Trader Purchase Quantity: 6026170.61

Listing 6: Timing Agent Reasoning (Negative Case)

```
{"reasoning":{"Time Since Launch":"Require time since launch < {75% Training Set}. Observed = 85374. Since 85374 > {75% Training Set}, this check fails.", "Trader Purchase Price":"Require purchase price > {25% Training Set}. Observed = 4.95e-06. Since 4.95e-06 < {25% Training Set}, this check fails.", "Trader Purchase Amount":"Require purchase amount > {25% Training Set}. Observed = 99.65. Since 99.65 < {25% Training Set}, this check fails.", "Trader Purchase Quantity":"Require purchase quantity > {25% Training Set}. Observed = 6026170.61. Since 6026170.61 < {25% Training Set}, this check fails.", "Summary":"One or more timing checks fail, indicating unsuitable entry timing for copy trading this wallet in the meme coin."},"result":false}
```

Prompt 2: Wallet Agent.

You are a professional on-chain analyst specializing in meme coin wallet trading behavior. You will be given features of a wallet's historical meme coin trading activity. Your task is to assess whether the wallet's next trade will be profitable so that a user can copy trade from it.

{CoT 1}

T-statistic of Returns: {t-stat}
 Average Return: {Return (all)}
 Return Standard Deviation: {Return Standard Deviation}
 Number of Trades: {Number of Trades}
 Last Return {Return (1st)}
 Five-to-One Return: {Return (1st-5th)}
 Ten-to-Six Return: {Return (6th-10th)}
 Fifteen-to-Eleven Return: {Return (11th-15th)}
 Time Since Last Trade: {Time Since Last Trade}
 Time Since First Trade: {Time Since First Trade}

Prompt 3: Coin Agent.

You are a professional on-chain analyst specializing in meme coin investment potential. You will be given various transaction features, candlestick chart, and comment history related to a meme coin. Your task is to assess whether the meme coin is a good investment opportunity.

{CoT 2}

Transaction Features:

Bundle Bot: {Bundle Bot} Sniper Bot: {Sniper Bot}
 Bump Bot: {Bump Bot} Comment Bot: {Comment Bot}
 Comment History: {Comments} Candlestick Chart: {Candlestick}

Prompt 4: Timing Agent.

You are a professional on-chain analyst specializing in meme coin wallet trading behavior. You will be given features of a wallet's historical meme coin trading activity. Your task is to assess whether the wallet's next trade will be profitable so that a user can copy trade from it.

{CoT 3}

Time Since Launch: {Time Since Launch}
 Trader Purchase Price: {Trader Purchase Price}
 Trader Purchase Amount: {Trader Purchase Amount}
 Trader Purchase Quantity: {Trader Purchase Quantity}

expenditure and proceeds across all trades are

$$X_{in}^{(S)} = \sum_{t:q_t>0} \Delta X(Y_t^S, q_t), \quad X_{out}^{(S)} = - \sum_{t:q_t<0} \Delta X(Y_t^S, q_t),$$

yielding the smart money's net return

$$R^{(S)} \equiv \frac{X_{out}^{(S)}}{X_{in}^{(S)}} - 1 = \frac{- \sum_{t:q_t<0} \frac{kq_t}{Y_t^S(Y_t^S - q_t)}}{\sum_{t:q_t>0} \frac{kq_t}{Y_t^S(Y_t^S - q_t)}} - 1.$$

A copier wallet replicates each trade immediately after the smart money. Because the smart money's trade updates the bonding curve state, the copier faces a strictly worse reserve:

$$Y_t^C = Y_t^S - q_t.$$

The copier's total SOL expenditure and proceeds are therefore

$$X_{in}^{(C)} = \sum_{t:q_t>0} \Delta X(Y_t^S - q_t, q_t), \quad X_{out}^{(C)} = - \sum_{t:q_t<0} \Delta X(Y_t^S - q_t, q_t),$$

with net return

$$R^{(C)} \equiv \frac{X_{out}^{(C)}}{X_{in}^{(C)}} - 1 = \frac{- \sum_{t:q_t<0} \frac{kq_t}{(Y_t^S - q_t)(Y_t^S - 2q_t)}}{\sum_{t:q_t>0} \frac{kq_t}{(Y_t^S - q_t)(Y_t^S - 2q_t)}} - 1.$$

□

PROOF OF THEOREM 2. Consider any buy trade of size $d > 0$ executed at effective reserve Y . Under the bonding curve, the required SOL outlay is

$$\Delta X(Y, d) = \frac{kd}{Y(Y - d)}.$$

Since the copier executes strictly after the smart money, it faces the updated reserve $Y - d$ and therefore pays

$$\Delta X(Y - d, d) = \frac{kd}{(Y - d)(Y - 2d)},$$

which is well-defined for $Y > 2d$. The relative execution cost is

$$\frac{\Delta X(Y - d, d)}{\Delta X(Y, d)} = \frac{Y}{Y - 2d} > 1,$$

implying that the copier strictly overpays on every replicated buy. By symmetry, for any sell trade ($d < 0$), the copier receives strictly less SOL than the smart money. Aggregating across all trades yields

$$X_{in}^{(C)} > X_{in}^{(S)}, \quad X_{out}^{(C)} < X_{out}^{(S)}.$$

Since the net return is defined as $R \equiv X_{out}/X_{in} - 1$ and is strictly increasing in the ratio X_{out}/X_{in} , it follows immediately that

$$R^{(C)} < R^{(S)}.$$

Therefore, the imitation penalty can be defined as:

$$\epsilon = R^{(S)} - R^{(C)} > 0.$$

Economically, ϵ captures the structural execution disadvantage of copy trading on bonding curves. □

For buys ($q > 0$), $\Delta X(Y, q) > 0$ represents SOL spent; for sells ($q < 0$), $\Delta X(Y, q) < 0$ represents SOL received. The smart money's total SOL