

## City Research Online

### City, University of London Institutional Repository

**Citation:** Yang, W., Ye, Q., Ascigil, O., Sokoto, S., Balduf, L., Krol, M. & Tyson, G. (2025). Beyond Single-Tokenomics: How Farcaster's Pluralistic Incentives Reshape Social Networking. Paper presented at the ACM SIGMETRICS, 9-13 Jun 2025, New York, USA.

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: https://openaccess.city.ac.uk/id/eprint/36074/

Link to published version:

**Copyright:** City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

**Reuse:** Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online: <a href="mailto:http://openaccess.city.ac.uk/">http://openaccess.city.ac.uk/</a> <a href="mailto:publications@city.ac.uk/">publications@city.ac.uk/</a>

# Beyond Single-Tokenomics: How Farcaster's Pluralistic Incentives Reshape Social Networking

WEN YANG, Hong Kong University of Science and Technology (Guangzhou), China QIMING YE, Hong Kong University of Science and Technology (Guangzhou), China ONUR ASCIGIL, Lancaster University, UK SAIDU SOKOTO, City St George's, University of London, UK LEONHARD BALDUF, TU Darmstadt, Germany MICHAŁ KRÓL, City St George's, University of London, UK GARETH TYSON, Hong Kong University of Science and Technology (Guangzhou), China

This paper presents the first empirical analysis of how diverse token-based reward mechanisms impact platform dynamics and user behaviors. For this, we gather a unique, large-scale dataset from Farcaster. This blockchain-based, decentralized social network incorporates multiple incentive mechanisms spanning platform-native rewards, third-party token programs, and peer-to-peer tipping. Our dataset captures token transactions and social interactions from 574,829 wallet-linked users, representing 64.25% of the platform's user base. Our socioeconomic analysis reveals how implementation choices (e.g. eligibility criteria, redistribution mechanism) shape varying inclusion rates (7.6%-70% new participants) and wealth concentration patterns (Gini coefficients 0.72-0.94). While tipping exhibits echo chamber tendencies (67%), substantial cross-community transactions (48%) among non-following pairs suggest potential for broader value exchange. Our causal analysis further uncovers several critical trade-offs: (1) while most tokens boost content creation, they often fail to enhance—sometimes undermining—content quality; (2) token rewards increase follower acquisition but show neutral or negative effects on outbound following, suggesting potential asymmetric network growth; (3) repeated algorithmic rewards demonstrate strong cumulative effects that may encourage strategic optimization over authentic engagement. Our findings advance understanding of cryptocurrency integration in social platforms and highlight challenges in aligning economic incentives with authentic social value.

CCS Concepts: • Information systems  $\rightarrow$  Social networks; • Human-centered computing  $\rightarrow$  Empirical studies in collaborative and social computing; • Networks  $\rightarrow$  Network economics.

Additional Key Words and Phrases: Decentralized Social Networks, Blockchain Rewards, Token Incentives, User Behavior Analysis, Causal Inference

#### 1 Introduction

The emergence of Decentralized Online Social Networks (DOSNs) marks a shift in social networking, emphasizing user autonomy, data sovereignty, and censorship resistance [5, 60]. Despite this, most DOSNs have struggled to incentivize high-quality content, large-scale user uptake, and sustained engagement [93, 112]. Most notably, their emphasis on user sovereignty has limited the adoption of commonly used monetization models [14, 68, 94, 112], often resulting in insufficient funding being available to compete with larger players.

Consequently, some have attempted to integrate cryptocurrency-based token incentives to encourage participation by both content creators and infrastructure operators [23, 65, 101, 119, 120]. This, however, comes with key challenges, most notably the reliance on a single, platform-issued token incentive mechanism. For instance, Steemit [101], a token-based DOSN launched in 2016,

Authors' Contact Information: Wen Yang, Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China, wyang330@connect.hkust-gz.edu.cn; Qiming Ye, Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China, qiming@connect.hkust-gz.edu.cn; Onur Ascigil, Lancaster University, Lancaster, UK, o.ascigil@lancaster.ac.uk; Saidu Sokoto, City St George's, University of London, London, UK, saidu.sokoto@city.ac.uk; Leonhard Balduf, TU Darmstadt, Darmstadt, Germany, leonhard.balduf@tu-darmstadt.de; Michał Król, City St George's, University of London, London, UK, michal.krol@city.ac.uk; Gareth Tyson, Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China, gtyson@ust.hk.

utilizes its self-issued token for content interaction incentives. However, the failure of such a token renders the rewards worthless. Furthermore, research has revealed that Steemit's single designated token incentive mechanism is susceptible to token price fluctuations [8], and has suffered from gaming and farming (*i.e.* strategic interactions between colluding users designed to exploit reward systems), [66] and bot-driven adversarial manipulation [22]. This led to reward concentration among a small group of colluding users, increasing centralization and economic inequality while losing its effectiveness in promoting social engagement [52].

In response to this, a new DOSN called Farcaster was publicly launched in 2023 to support multiple incentive mechanisms [35]. Functionally similar to X (Twitter), Farcaster stands out from current DOSNs in two key aspects. First, Farcaster supports "modular" wallet binding — unlike platforms constrained by primary account-bound blockchain addresses [13, 65, 101, 119], Farcaster enables users to link any external Ethereum-compatible addresses [77], functioning as on-chain transaction wallets, alongside their user accounts (termed Farcaster Identifiers (FIDs)) [33], providing greater economic flexibility and autonomy. Second, Farcaster is the first to implement a "pluralistic" token incentive ecosystem. We refer to it as pluralistic because, unlike existing DOSNs, Farcaster does not have an officially issued token or a centrally designated incentive mechanism. Instead, Farcaster allows any token or incentive mechanism to coexist within the ecosystem, regardless of the token used (medium of reward) or the eligibility criteria designed. This opens up incentive design to users, third-party developers, or the platform's administrators themselves.

Thus, Farcaster enables users and developers to easily create and distribute their own tokens, creating an entirely decentralized reward ecosystem rather than a fixed incentive paradigm managed centrally. Such tokens can be used for any purpose deemed appropriate, including tipping content creators and operators who manage the infrastructure. Moreover, third-party developers can create custom applications (*mini-apps*) with algorithmic token reward distribution mechanisms [37], supporting a more community-driven incentive paradigm. We believe this presents a unique use case for studying the feasibility of a system where multiple tokens and diverse incentive mechanisms coexist to incentivize positive user behavior within social networks.

To understand its broader implications, this paper empirically examines how Farcaster's pluralistic incentive paradigm shapes platform dynamics and user behaviors. We gather both on-chain token transactions and off-chain social interactions relevant to Farcaster. As of April 27, 2025, our dataset covers 574,829 (64.25% of the user base) users who have at least one Ethereum-compatible wallet bound to their FIDs, with 5,878 unique tokens traded between users (far surpassing other DOSNs) [46, 61, 72]. Exploiting this data, we study the impact of multiple incentive mechanisms within the ecosystem. Specifically, we explore the following three research questions:

**RQ1:** How widespread and diverse is the token economy within Farcaster's ecosystem, specifically regarding: (1) the temporal dynamics of people binding their external cryptocurrency wallets to their Farcaster accounts, (2) how prevalent the various available tokens are, and (3) how these tokens serve different social functions through their incentive mechanisms?

**RQ2:** What socioeconomic risks are inherent in Farcaster's incentive system, specifically concerning: (1) disparities in new user participation rates across different token rewards, (2) inequity in reward distribution, alongside (3) echo chamber effects in tipping?

**RQ3:** What causal relationships exist between token incentives and subsequent social activities, and how do these dynamics vary across: (1) different token categories (volatile tokens vs. stablecoins), (2) distinct incentive mechanisms (user-to-user tipping vs. algorithmic rewards)?

To the best of our knowledge, we are the first to empirically study Farcaster's pluralistic incentive ecosystem. Our contributions are as follows:

- We reveal how specific eligibility criteria designs (*e.g.* nomination-based vs. behavioral scoring) and reward distribution structure (*e.g.* bot-driven tipping, redistribution mechanism) significantly impact both user inclusion (70% vs 7.6% new participants) and income equality (Gini coefficients 0.72-0.94) (see Sections 5.1 and 5.2).
- We demonstrate that, while user-to-user tipping represents the most flexible incentive mechanism, it is predominantly unidirectional (with less than 10% of users acting as both tip receivers and senders) (see Section 5.1). Additionally, 52–75% of tips occur across community boundaries, and 32.42% between non-following pairs. This suggests that token incentives can facilitate value exchange beyond established social community structures (see Section 5.3).
- We reveal trade-offs in incentivised social activities: while algorithmic rewards leveraging volatile tokens as the medium effectively increase content quantity, they show limited or negative effects on content quality (see Section 6.3).
- We uncover that repeated algorithmic rewards correlate with asymmetric social network growth (increased follower acquisition but decreased outbound following) and strategic engagement optimization (prioritizing immediate reactions over share-worthy content creation), highlighting risks in token-incentivized social platforms (see Section 6.3).

These findings advance both the theoretical understanding of token-based incentive design and provide practical guidance for implementing sustainable reward mechanisms in social platforms.

#### 2 A Primer on Farcaster

We begin by outlining the core design of *Farcaster*. Below, we provide brief descriptions of: (1) social interactions; and (2) token transactions. For full technical details, we refer readers to the official documentation [35].

Social Interactions. Upon registration, Farcaster users receive an on-chain identifier (an Ethereum custody address) anchored on the Optimism Layer-2 chain<sup>1</sup> [87] and managed through Farcaster's smart contracts [34]. Users must pay an annual storage fee [39] to rent network storage capacity during registration.<sup>2</sup> Users maintain exclusive control over their account's private key. To facilitate network interaction, each address is associated with both a unique numeric identifier (FID) and a human-readable username (Farcaster User Name (Fname), e.g., @vitalik).

The off-chain social interactions—referred to using Farcaster-specific terminology as "casts" (posts and replies), "reactions" (likes and re-posts), and "links" (follow actions)—are exchanged through a peer-to-peer (P2P) network of independently operated servers called *hubs* [36]. Each Hub maintains a complete copy of the interaction data and synchronizes with peers using the GossipSub [54] and Diff Sync protocol [31].<sup>3</sup> The system demonstrates robust fault tolerance: network functionality remains intact as long as a single Hub remains operational [29].

All social interactions (*e.g.* casts, links, and reactions) require a digital signature using the private key corresponding to the custody address. These signed actions are broadcast across the network, where participants (*i.e.* hubs, clients, and third-party applications) verify message authenticity by checking the digital signature against the on-chain registered public key for that FID. This hybrid (*i.e.* on-chain/off-chain) architecture preserves user ownership and interoperability while circumventing the scalability and cost constraints inherent in fully on-chain systems [23, 76, 101].

 $<sup>^{1}</sup>$ A Layer-2 (L2) is a scaling solution atop a Layer-1 (L1) blockchain (e.g. Ethereum), enabling faster and cheaper transactions while inheriting its security.

<sup>&</sup>lt;sup>2</sup>Farcaster's storage fee has been reduced three times since launching in October 2023, from \$7 to the current \$2 [27]

<sup>&</sup>lt;sup>3</sup>Note, Farcaster is transitioning to a new P2P coordination layer called *Snapchain*. Built upon GossipSub, Snapchain replaces full replication with a partitioned model, where each hub stores only a subset of data based on user FIDs [41].

Token Transactions. Custody addresses linked to FIDs are primarily intended for account management (e.g. signing social actions) rather than token transactions [35]. Farcaster enables users to bind external Ethereum-compatible addresses to their FID as transaction wallets [33], allowing for trading, rewarding, or payment activities. We refer to this flexibility as a modular wallet architecture. This architecture facilitates broader token interoperability and economic autonomy. By isolating user accounts from token transactions, it also enhances security and reduces risks associated with private key exposure (e.g. phishing/scam attacks [105]).

Note, since February 22, 2025, Farcaster has implemented a phased roll-out of official Ethereum-compatible wallets. This provides users with both optionally bound and officially issued Farcaster transaction wallets, along with the flexibility to designate any wallet as their primary wallet [33].

Note, while Farcaster allows users to bind both Ethereum [28] and Solana wallets [100] to their FIDs, Ethereum addresses significantly outnumber Solana addresses (794, 386 vs. 186, 434 as of May 2025). Moreover, Farcaster only introduced Solana Wallet Standard integration for Mini-apps on May 21, 2025 [44], beyond our study period. Farcaster users can exchange tokens across over 50 Ethereum-compatible L1 and L2 chains (e.g. Base, Optimism, Polygon, and BSC). However, we find that: (1) all top-ten tokens by daily transaction volume originate from Base chain deployments [72]; (2) Base chain transactions constitute nearly 90% of total activity among Farcaster users [91]; and (3) Farcaster's native reward mechanism exclusively employs Base chain USDC for weekly distributions to qualified users and builders [38]. We thus focus our analysis on Base chain alone. All subsequent references to "wallet addresses" in this paper denote Ethereum-compatible wallet addresses, distinct from both custody addresses and Solana-compatible wallets.

#### 3 Data Collection Methodology

Farcaster's hybrid data architecture necessitates both on-chain and off-chain data collection: (1) *Off-chain Data*: we gather a complete snapshot of Farcaster's Hub data as of April 27, 2025, including all user profiles (*i.e.* FIDs, user names, FID-bound wallet addresses) and social interactions (*i.e.* followings, posting, liking, replying, and re-posting.) with their creation timestamps. (2) *On-chain Data*: We use Alchemy APIs<sup>4</sup> to collect Farcaster's token transaction data from the Base chain (an Ethereum Layer-2 network [10]) and construct transaction graphs between users' wallet addresses.

#### 3.1 Off-chain Data (User Profiles and Social Interactions)

Following Farcaster's official documentation [40] and code-base [30], we deploy two Hub server instances (one in Asia and one in Europe) to synchronize the off-chain data.

FID Registration and Wallet Binding Records. As of April 27, 2025, our dataset encompasses 1,059,655 registered FIDs, with 894,678 valid FIDs. Note that when analyzing the timestamp data from the hub, we discovered that all FID registered before November 7, 2023 were aggregated to November 7, 2023. Therefore, in our analyses requiring FID registration timestamps, we set November 7, 2023 as the starting point.

While associations between Farcaster-issued wallets and FIDs are recorded both in the *KeyRegistry* smart contract's transaction logs [34] and Hub data, users' optionally bound external wallets are recorded solely in the Hubs and not on-chain [33]. However, Hubs periodically purge old data [30], resulting in the loss of information about wallets that were previously associated with an FID but were later unbound. To recover a complete list of external wallets bound to each FID, we query Neynar's API [84]. Since this API only provides mappings of historical bound wallets and FIDs,

<sup>4</sup>https://www.alchemy.com/

<sup>&</sup>lt;sup>5</sup>We exclude invalid users by identifying FIDs without historical storage units.

without any binding and unbinding timestamps, we must rely on the incomplete wallet records in Hubs with timestamps for data analyses where binding time is necessary.

For wallets recorded in hubs, we discover that 574,829 (64.25%) of FIDs have at least one transaction wallet, whether optionally bound or officially issued, totaling 794,386 Ethereum-compatible wallets. After retrieving the complete historical bound wallets, we identify a total of 1,282,783 external wallets bound to 606,827 (64.5%) FIDs. We find that 488,397 (38%) wallets were unbound before June 2025 after their initial binding.

Social Interactions. The social interaction data provided by our Hub contains 159,539,953 unique following relationships, 164,984,116 casts (comprising 36,646,412 posts (22.21%) and 128,337,704 replies (77.79%)), and 299,079,720 reactions (consisting of 252,771,162 likes (84.52%) and 46,308,558 re-posts (15.48%)). For clarity and consistency with conventional terminology in the literature [106], we use standard terms such as "follow", "post", "reply", "like", and "re-post" to denote these social interactions throughout the remainder of this paper.

#### 3.2 On-chain Data (Token Transactions)

Recall, we identified that Farcaster's token transactions predominantly happen on Base chain (see Section 2). Therefore, we extract all token transfer records on Base involving Farcaster users' wallets. To do so, we use the Alchemy API [2] to retrieve historical transfer data for all 1,282,783 wallet addresses, as of April 27, 2025. Additionally, to capture transactions involving smart contracts and other non-user wallet interactions, we include transactions where at least one party (either sender or recipient) is a user wallet. We collect a total of 87,687,791 transaction records, encompassing 5,878 distinct tokens (1.34% of all 440,274 tokens that have appeared in all user wallets but may not necessarily have been traded between users) transferred between users' Ethereum wallets.

#### 3.3 Dataset Availability and Reproducibility

All data used in this paper is derived from publicly accessible sources. The off-chain data (including user profiles, FID-linked wallet addresses, and social engagement) can be reproduced by participating in the Farcaster node network following the official documentation. Note that our study utilize the now-deprecated hub architecture. Following June 2025, Farcaster has fully transitioned to a new blockchain-style architecture called Snapchain (see more details in Section 2); subsequent researchers will need to deploy a Snapchain node to access the data that is previously available via hubs. The on-chain token transaction data is retrievable via Ethereum-compatible node API providers (e.g. Alchemy) using the wallet lists linked to FIDs. We release the code to perform this dataset reconstruction.

In light of ethical considerations (user profiles and content) and the substantial footprint of the raw datasets (Farcaster social engagement  $\approx 172$  GB; token transactions  $\approx 37$  GB), we do not mirror or redistribute the raw data. Instead, we publicly release weekly aggregated metrics (social engagement and token receipt) derived from the raw data that underpin our causal analyses in Section 6. These aggregated datasets provide sufficient granularity for independent replication of all results reported in Section 6, while substantially reducing the risk of reidentification.

In summary, we release: (i) the weekly aggregated metrics dataset, and (ii) the scripts and configuration templates to reproduce the Base-chain transfer retrieval pipeline, and to generate the weekly aggregates from the freshly syncronized raw data. These materials are accessible at: https://zenodo.org/records/17317121.

#### 3.4 Potential Biases from Hub Data Collection

Farcaster's hub architecture performs *full* replication of off-chain social interaction data across each hub. In our setup, two hubs (geographically distributed) are bootstrapped from an official snapshot and then continuously synchronized via GossipSub and Diff Sync (see Section 2). While transient discrepancies can occur across hubs, our empirical checks confirm these are trivial: over 99.9% of records converge across our two hubs within the observation window. For on-chain data, we retrieve all historical transfers for all FID-bound wallets on the Base chain; due to the immutability of blockchain ledgers, this component is complete by definition. Taken together, we do not find evidence of sampling biases that would alter our findings or their interpretation.

#### 3.5 Ethical Considerations

Our dataset includes publicly available casts and on-chain transactions. To address privacy concerns, we strictly follow established ethical standards [25], collect only public data, and operate Hubs non-intrusively at our own expense and following the guidelines of the Farcaster creators [36]. Notably, wallet addresses offer stronger pseudonymity than social identifiers like FIDs or Fnames, making it harder to link transaction histories to personal identities. This study was reviewed and received a waiver from the authors' institutional ethics committee.

#### 4 Token Economy Scale and Token Incentive Diversity.

We answer **RQ1** by exploring the scale and diversity of the Farcaster token economy. First, we assess the role tokenomics play in Farcaster's growth and user activity. We then identify the most popular and impactful tokens. Finally, we analyze the incentive mechanisms that use these tokens.

#### 4.1 Token-related Initiatives Driving Wallet Bindings.




Fig. 1. Daily engagement metrics and user growth on Farcaster.

Figure 1 presents daily user activity, platform growth, and user involvement in tokenomics. The platform experienced a steady activity growth, reaching a maximum of 73,180 daily active users (DAU) on July 2, 2024. Since then, the DAU stabilized at  $\approx 42k$ .

The new FID registrations and wallet binding show highly bursty behavior. Registration/binding spikes occur during token-related events or new platform feature introductions. This includes DEGEN airdrops<sup>6</sup> announcement [21, 110] (at ② and ⑤) or launch of new tokens that went viral (MOXIE [82] at ⑦ and DRB [88] at ①). At ④, Farcaster launched its token-focused mini-apps [37], while at ⑨, the platform introduced its official crypto wallet [33]. The only token-unrelated event with a significant impact occurred at ⑥, when Farcaster raised \$150M in funding [74], following an advertisement campaign (Farcaster Conference 2024) [32].

This suggests that tokenomics is an important factor driving the Farcaster userbase. The platform decreased its registration fees multiple times from \$7 to \$5 in December 2023 (1), to \$3 in January 2024 (3), and to \$2 in August 2024 (8) [27]. Surprisingly, those reductions did not significantly impact the new user registrations. The exact amount of the fees seems irrelevant for the new users who are mostly attracted by new features or the possibility of obtaining valuable tokens. This is further confirmed by the high rate of users who bound a token wallet to their account. The ratio is steadily increasing since late January 2024, reaching 64.25% in April 2025. We provide an expanded correlation analysis between platform growth and real-world events in Appendix A.

#### 4.2 Prevalent Token Detection.

While flexible wallet binding to user accounts enhances economic autonomy and interoperability, one trade-off is that it simultaneously floods user wallets with numerous tokens unrelated to the Farcaster ecosystem. This introduces significant noise into our Farcaster incentive analysis. Therefore, we next examine the tokens circulating within the Farcaster ecosystem to discover methods for filtering this noise and identifying prevalent tokens that are genuinely relevant and impactful to the Farcaster social network.

We identify 440,274 distinct tokens held in FID-linked wallets. Yet, most exhibit limited activity: 99% (435,871) tokens have fewer than 390 holders (by FIDs) and fewer than 1,065 transactions, while the remaining 1% (4,403) tokens account for 93.35% of all holders and 94.58% of all transactions (detailed in Appendix B). Furthermore, many tokens are widely distributed by just a small number of wallets, indicating a spam-like behavior without community adoption [105]. This is common when token creators airdrop tokens to expand their popularity [3, 73, 109]. 60% (258,138) of tokens were never sent by a single FID-bound wallet, and >99% of tokens (434,094) involve fewer than 191 unique FID senders.

These findings suggest that most tokens are passively received with limited social utility. We therefore strive to focus on sending activity to identify the platform's most socially engaged tokens. For brevity, we summarize the process of selecting these tokens below, and provide a detailed description and justification in Appendix C: (1) We filter the tokens with inter-FID transfers (*i.e.* transacted between at least one pair of Farcaster users); (2) we apply normalized Shannon entropy [69] to temporal transaction frequencies to filter out tokens with bursty, short-lived activity; (3) we retain tokens above the 99th percentile in unique FID senders (>254)<sup>7</sup>, filtering out those primarily distributed via airdrops rather than active social engagement. (4) based on the transaction graph, we calculate the clustering coefficients [98] and select 0.3-0.6 as criteria [108] to verify community-driven usage patterns.

Following this four-step process, we identify four prevalent tokens (DEGEN, MOXIE, HIGHER, and TN100X) issued by third-party developers as social rewards [21, 53, 55, 82]. We additionally include USDC, a stablecoin incorporated into Farcaster's official reward mechanisms [38, 109]. It is

 $<sup>^6</sup>$ Airdrop is the free distribution of cryptocurrency tokens to eligible wallets, often to promote token adoption or reward early users.

<sup>&</sup>lt;sup>7</sup>Note that this threshold of 254 unique FID senders is derived from the 99th percentile of inter-FID transfers and therefore differs slightly from the threshold of 191, which is the 99th percentile for overall transfers.

also used for user-to-user tipping as part of the platform's official design [11]. For our subsequent investigation, we use these five tokens as the primary subjects of study.<sup>8</sup>

Token	Holders Total Txns Inter-FID Txns FID Se		FID Sender	Clustering Coeff.	Token Age (wks)	Entropy (Norm)	
				≥ 254 (99 <sup>th</sup> m)	∈ [0.3, 0.6]	≥ 26	≥ 0.9
DEGEN	152,908	3,337,952	173,772	27,723	0.32	73	0.93
MOXIE	43,742	1,810,849	138,728	9,002	0.58	44	0.92
HIGHER	32,692	320,749	51,596	1,153	0.41	65	0.90
TN100X	16,409	193,678	9,996	1,838	0.36	69	0.92

29,464

0.23

0.93

473,801

Table 1. Prevalent tokens meeting the filtering criteria, sorted by overall transaction count (frequency).

Table 1 presents the primary transaction metrics and filtering criteria assessment for these five tokens. Notably, USDC only fails to meet the clustering coefficient criterion, with a value of 0.23 slightly below the lower threshold of 0.3, while satisfying all other three criteria. This indicates that USDC exhibits a relatively looser community structure compared to the four social reward tokens, which may be attributed to its additional use case as a stablecoin in payment scenarios rather than social interactions. Furthermore, it is worth noting that DEGEN's holder count ( $\approx$  153k) ranks second only to USDC ( $\approx$  216k), surpassing the other three social tokens by a considerable margin (by  $\approx$  3.5 to 9.3 times). Similarly, the number of FID senders for DEGEN approaches that of USDC (27,723 vs 29,464). These metrics demonstrate that DEGEN, being the earliest launched among the four social reward tokens, along with USDC (13 weeks older than DEGEN), has achieved the strongest network effects and highest community recognition among all tokens on Farcaster.

#### 4.3 Categorizing Incentive Mechanisms.

8,768,648

**USDC** 

216,050

Finally, we investigate the incentive mechanisms that use these five tokens. We analyze the official documentation [21, 82, 90], transaction history, and the smart contracts used for token distribution (detailed in Appendix D). We then classify the incentive mechanisms into two main categories—tipping and algorithmic rewards, with algorithmic rewards further subdivided into third-party and official-led initiatives.

Inter-FID Tipping. In this mechanism, users directly send each other tokens using direct transfers. (1) direct blockchain transfers to the wallet address displayed on a recipient's profile; or (2) intermediary mini-apps (e.g. @paybot [90]) that enable socially-driven interactions (similar to the donate function in YouTube). All 5 prevalent tokens are used in this mechanism.

*Third-party Algorithmic Rewards.* Farcaster enables any third party to launch tokens with bespoke distribution rules. These tokens are typically distributed via dedicated smart contracts designed to enhance user engagement. Such contracts often incorporate staking-based mechanisms<sup>10</sup> to

<sup>&</sup>lt;sup>8</sup>To avoid confusion caused by the numerous tokens sharing the same name and counterfeit tokens, the contract addresses for these five tokens are provided in Appendix D.


<sup>&</sup>lt;sup>9</sup>YouTube's fan funding feature: https://www.youtube.com/intl/en/creators/fanfunding/


<sup>&</sup>lt;sup>10</sup>Locking tokens in smart contracts for a set period to qualify for rewards or receive benefits like boosted scores.

mitigate undesired behaviors, including reward farming  $^{11}$  and sell-off pressure.  $^{12}$  We observe DEGEN $^{13}$  [21] and MOXIE $^{14}$  [82] being distributed through this mechanism.

Official Algorithmic Rewards. We distinguish the official algorithmic reward mechanism, implemented by Farcaster's administration through the USDC stablecoin [38].  $^{15}$  The mechanism provides weekly rewards to top-performing users based on engagement metrics.  $^{16}$ 

Mechanism Comparison. We first analyze user coverage and temporal transaction dynamics for the above three reward mechanisms. Collectively, these mechanisms reach a total of 103,666 unique recipients, accounting for 11.59% of all FIDs. More specifically, this figure corresponds to 17.56% of active users, defined as individuals who have posted at least once. This indicates a relatively high adoption rate given the diversity and scale of the user base, suggesting these incentive mechanisms play a substantial role in overall system usage. Interestingly, Inter-FID Tipping and Third-party Algorithmic Reward mechanisms reach 6.01% and 6.43% of all FIDs, respectively, surpassing the Official Algorithmic Reward mechanism (3.15%). This suggests that community-driven incentive mechanisms may be more effective in engaging users than centralized, protocol-driven rewards.





- (a) Tipping frequencies for five prevalent tokens.
- (b) Rewarding frequencies for algorithmic tokens.

Fig. 2. Stacked area charts of weekly aggregated transaction frequencies by token across mechanisms.

We next examine the individual tokens underpinning these mechanisms. Figure 2a depicts the transaction frequencies of Inter-FID Tipping across five major tokens (as referenced in Section 4.2), while Figure 2b illustrates the transaction frequencies for both third-party and official algorithmic reward mechanisms (note that only three out of five major tokens (DEGEN, MOXIE and USDC) are distributed within the algorithmic mechanism).

From the figures, we find that algorithmic rewards exhibit temporal patterns distinct from those of tippings. The frequencies of algorithmic rewards demonstrate pronounced episodic spikes, each corresponding to the initiation and duration of reward projects. By contrast, tipping frequencies display a more consistent and sustained temporal profile, closely tracking the fluctuations in daily

<sup>&</sup>lt;sup>11</sup>A small group of users engages in circular reward-giving amongst themselves to exploit token reserves.

 <sup>12</sup> Upon receiving token rewards, users immediately exchange them for more established cryptocurrencies (e.g. BTC, ETH).
 13 DEGEN uses a nomination-based system where users reply to posts with messages like "100 \$DEGEN" to nominate others.
 These are collected monthly to determine token rewards for post creators, resulting in the spike pattern shown in Figure 2b.
 14 MOXIE's algorithm linearly weights posting, replying, and token staking in its reward function, making it more prone to metric gaming [81, 82].

 $<sup>^{15}</sup> Far caster$  uses a black-box algorithm to mitigate farming and gaming behaviors, as noted by the co-founder: https://far caster.xyz/v/0x0e31071c

<sup>&</sup>lt;sup>16</sup>These rewards follow a tiered structure, ranging from \$1 to \$300, allocated to qualified users across different ranking tiers.

Metric	Pre-launch Mean	Pre-launch Median	Post-launch Mean	Post-launch Median
Unique Senders	1,823	1,779	7,960	8,014
Unique Receivers	1,220	1,117	6,849	7,198
New Senders	528	407	2,480	1,001
New Receivers	513	334	1,975	1,687

Table 2. Weekly user tipping statistics surrounding the Farcaster wallet launch.

active user (DAU) (as shown in Figure 1). This contrast underscores the project-driven nature of algorithmic rewards versus the organic, user-driven dynamics of tipping.

#### 5 Socioeconomic Risks in Farcaster's Incentives

Previous studies have shown that financial rewards, despite their potential to boost engagement, can inadvertently encourage negative behaviors (e.g. farming, whereby users collaborate to mass-produce content and artificially amplify engagement to accumulate rewards) [8, 52, 66]. Following **RQ2**, we investigate whether Farcaster's incentive mechanism exhibits similar socioeconomic risks. Particularly, we focus on three example behaviors: new user participation, potential reward concentration, and echo chamber formation across incentive mechanisms.




Fig. 3. Temporal dynamics of new user participation in (a) tipping and (b) algorithmic rewarding mechanisms.

#### 5.1 New User Participation Rates

To evaluate whether incentive mechanisms encourage broader participation (inclusivity) for new-comers or create barriers to entry (thereby offering more reward opportunities to incumbent recipients), we analyze the temporal patterns of user inclusion. This inclusion is measured by the rate of weekly new reward receivers to weekly total receivers.

Data and Methodology. We perform a temporal analysis by calculating the weekly counts of unique senders and receivers (by FIDs) for each type of reward, including eight token-mechanism pairs defined in Section 4.3. We also identify users who act as both senders and receivers within the same week. Additionally, we track the weekly influx of new receivers and senders, defined as those receiving or sending the specific reward for the first time that week. This longitudinal analysis reveals new user participation patterns across different incentive mechanisms.

*Results.* Figure 3 shows the weekly count of FIDs by their types. The stacked area charts display the number of unique senders (diagonal), receivers (solid), and users acting as both (horizontal), with overlaid lines representing weekly new receivers (solid red) and new senders (dashed red).

For *Inter-FID Tipping* (see Figure 3a and Table 2), the total sender and receiver counts fluctuate synchronously and the tipping user base is significantly enlarged due to the Farcaster wallet launch. An anomaly occurred during the significant tipping surge following the Farcaster wallet launch in late February 2025: the weekly new sender count spiked sharply (reaching 9,023 for the week of March 3 2025), far exceeding new receivers (which remained low at 270, similar to pre-launch levels). The following week, new senders dropped to 3,772 while new receivers rose to 3,611, and both metrics quickly resynchronized. Further breakdown (see Figure 9 in Appendix E) reveals this spike was mainly driven by USDC tipping. We conjecture this is due to official campaigns encouraging users to send USDC to activate wallet features or qualify for airdrops [109].

It is also worth noting that each week, only a small fraction (Mean = 9.76%, SD = 5.83%, median = 8.84%) of users engage in both sending and receiving. This indicates that most tipping flows are unidirectional instead of reciprocal.

For *Algorithmic Rewards* (see Figure 3b), recall that the three tokens (DEGEN, MOXIE, and USDC) were distributed in distinct time windows (as shown in Figure 2b), allowing the aggregated data to still reveal clear trends. Before August 2024, DEGEN's algorithmic rewards followed a monthly claim pattern (as mentioned in Section 4.3). For DEGEN, both weekly total receivers and weekly new receivers increased in the first four months (Jan-April 2024). The peak of new user participation rate (59.24%) was reached in late April (total 25,475; new 15,091), after which both weekly total and new receivers declined — with new receivers dropping more rapidly. By late May, new receivers accounted for only 30% of total receivers (23,005 vs. 6,908), and by June and July, this dropped to  $\approx$  16%, after which DEGEN algorithmic rewards ceased.

During the subsequent MOXIE reward period (about 30 weeks from Aug 2024 to Jan 2025), the weekly gap between total receivers (mean 4,293) and new receivers (mean 328) was much wider: new receivers account for only 7.6%, indicating far less inclusion in MOXIE's algorithmic reward project, with more reward opportunities offered to incumbent receivers. This stands in sharp contrast to DEGEN, suggesting the user-driven nomination-based reward design of DEGEN was more inclusive than MOXIE's behavioral scoring approach (detailed in Section 4.3). Finally, the official USDC algorithmic reward (from Feb 2025 to present) also exhibited acceptable inclusivity: new receivers accounted for 48.45% of total receivers on average (2,477 vs. 5,112), with both metrics (weekly total and new receivers) moving in parallel. This suggests that each token is used in quite distinct manners, with different degrees of inclusivity for attracting new users. This arguably highlights the benefits of the pluralistic approach taken by Farcaster.

#### 5.2 Income Inequality and Wealth Concentration

Previous research has shown that single-token incentive mechanisms can lead to income inequality and wealth concentration (Gini coefficient > 0.9) across decentralized networks [66, 79, 120]. This motivates us to assess whether Farcaster's pluralistic token incentive ecosystem (where user-to-user tipping and developer-led algorithmic rewards coexist) also faces the same challenges, resulting in wealth concentration.

Data and Methodology. To measure wealth concentration, for tokens other than USDC (a stablecoin), we first collect daily average price data for DEGEN, MOXIE, HIGHER and TN100X. We estimate each user's income by multiplying the received token amount by the average daily USD price on the day of receipt. This provides a practical approximation, as users may exchange their tokens at any time. Finally, we measure the concentration of each token's value (in USD).

Table 3	Income distribution and ine	equality metrics acro	es three incentive	mechanisms in Farcaster
Table 3.	miconie distribution and me	quality illetites acid	,55 tillee illetillive	illectiallisilis ill i alcastel.

		Inter-FID Tipping Reward				3rd-Party Algo	Off. Algo. Reward	
Metric	Degen	Higher	USDC	Moxie	Tn100x	Degen	Moxie	USDC
Gini Coeff.	0.8304	0.9382	0.8631	0.7246	0.8277	0.8433	0.9248	0.8598
Total	99,141.39	4,132.93	94,788.73	86,136.68	5,543.84	49,612,724.35	1,657,380.21	517,831.34
Max	2,061.85	616.44	999.99	2,040.63	457.78	492,133.29	15,542.65	1,772.83
Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
Mean	3.13	0.29	2.73	9.97	3.11	1,050.81	77.32	15.49
Median	0.19	0.02	0.11	3.89	0.07	111.03	1.95	1.00




Fig. 4. Income distribution (USD) across incentive mechanisms.

Results. Table 3 summarizes key statistics for all major incentive mechanisms and Figure 4 shows the Empirical Cumulative Distribution Function (ECDF) of user income (in USD) for each mechanism. It reveals that despite Farcaster's pluralistic approach (designed to potentially mitigate income inequality by offering more reward-receiving opportunities to a broader user base), significant wealth concentration persists across all token-mechanism pairs. The consistently high Gini coefficients (0.72-0.94 in Table 3) suggest that both tipping and algorithmic rewards may recreate the centralization issues observed in earlier token-based social platforms like Steemit (Gini coefficient ≈ 0.99) [52]. For *Inter-FID Tipping*, income from the HIGHER token exhibits a right-skewed distribution (more low-income users), with a Gini coefficient of 0.94. This high degree of inequality is further underscored by the fact that 98% of users received less than \$1, and 80% received less than \$0.05. This imbalance is consistent with HIGHER possessing the lowest number of unique senders and the most skewed sender-to-receiver ratio (1:12.5) among all tipping tokens examined (Table 4). Critically, we notice that 92.4% of HIGHER tipping transactions originate from only two bot accounts. These bots reward trivial amounts of HIGHER to users during specific interactions, such as content replies or lottery drawings. Consequently, the distribution of HIGHER is heavily concentrated among low-value recipients and is primarily driven by automated bot activity rather than organic peer-to-peer engagement.

For *Algorithmic Rewards*, the distribution patterns reflect a more structured approach than tipping. USDC, for example, uses a tiered reward scheme based on weekly behavioral rankings Section 4.3, distributing between \$1 to \$300. Notably, 75% of users receive the minimum reward of \$1. While both USDC and MOXIE rely on similar behavioral scoring algorithms (detailed in Section 4.3), their metric selection and openness differ significantly. MOXIE's algorithm explicitly weights posting,

	Inter-FID Tipping Reward				3rd-Party Alg	go. Reward	Off. Algo. Reward	
Metric	DEGEN	HIGHER	USDC	MOXIE	TN100X	DEGEN	MOXIE	USDC
Unique Sender Count	27,519	1,267	30,004	9,212	1,859	_	-	_
Unique Receiver Count	40,836	15,849	15,872	5,252	3,459	47,748	21,505	28,181
Total Transaction Count	196,555	58,449	596,213	139,516	11,026	101,232	354,952	81,812

Table 4. Token activity metrics across three incentive mechanisms on Farcaster.

replying, and token staking, making it more susceptible to metric gaming and the rich-get-richer phenomenon [81, 82]. In contrast, USDC's scoring algorithm remains opaque but predominantly includes social behavior signals, without any wealth status metrics, making it more resistant to gaming while ensuring more opportunity to baseline rewards for a broader user base [38] This aligns with research showing that modest, guaranteed incentives can outperform larger, uncertain rewards in driving participation [57].

Consequently, MOXIE's algorithmic rewards show more pronounced income inequality (Gini: 0.92) contrasted with USDC (Gini: 0.86). This is likely due to MOXIE's transparent scoring system that allows strategic users to optimize their behavior for maximum rewards, as well as its token-stake boosting scores, resulting in the rich-get-richer effect. This extreme concentration in MOXIE's algorithmic rewards echoes its poor inclusivity metrics observed in Section 5.1, as new users face barriers to participation while rich incumbents monopolize the reward opportunities.

That said, MOXIE also presents an interesting case where its redistribution mechanism effectively mitigates initial wealth concentration. While its algorithmic rewards show high inequality (Gini: 0.92) in initial distribution, its unique follower-followee redistribution mechanism [82]—where a portion of rewards (designated by the followee, *e.g.* 20%) received by followees automatically flows to followers in the form of inter-FID tipping—contributes to more balanced secondary distribution for MOXIE tipping (Gini: 0.72). This suggests that carefully designed redistribution rules *can* help address wealth concentration issues even when primary reward allocation is highly skewed.

#### 5.3 Echo Chamber Effect in Tipping

While inter-FID tipping mechanisms facilitate value exchange and content monetization, they may inadvertently amplify echo chamber effects within social networks. This might drive users to create smaller social communities, driven by trends in reward-giving. Thus, in the context of tipping behavior, we define an *echo chamber* as a closed loop of economic value circulation where tipping flows predominantly remain within tight-knit communities rather than across diverse user groups [24]. Such economic echo chambers could potentially lead to the concentration of tipping flows among a small subset of users, reducing exposure to diverse social content and limiting the platform's ability to sustain a broad and inclusive incentive model.

Data and Methodology. To investigate potential echo chamber effects, we begin by examining the temporal dynamics between following and tipping relationships (Table 5). Specifically, we compare the timestamp of the first tip between pairs of users with the timestamp of their follow relationship (if any). We classify tipping interactions into three categories based on the timing of follow relationships: (1) Followed before first tip, (2) Followed after first tip, and (3) Never followed, *i.e.* tipping between users who never established a follow relationship, accounting for 55.61%, 11.97%, and 32.42% of all tips, respectively. This distribution motivates a further analysis of whether tipping interactions, especially those without underlying social relationships (*i.e.* Never Followed), tend to occur within existing echo chambers or bridge across them.

Table 5. Tipping and following relationship.

# Louvain # % Inter-comm. Ratio Never Followed 55,847 32.42% [45.06%, 51.48%] Followed Before First Tip 95,790 55.61% [22.24%, 26.13%] Followed After First Tip 20,621 11.97% [22.47%, 25.26%] Infomap Never Followed 74.68% Followed Before First Tip 56.10% Followed After First Tip 59.66%

Table 6. Network overlap metrics.

Louvain			
Network Pair	NMI	AMI	Max Overlap
Follow vs. Combined	0.73	0.73	0.925
Follow vs. Tip	0.33	0.26	0.380
Tip vs. Combined	0.31	0.24	0.694
Infomap			
Follow vs. Combined	0.91	0.91	0.928
Follow vs. Tip	0.19	0.13	0.519
Tip vs. Combined	0.21	0.15	0.227

To explore this, we construct the Farcaster social graph based on follow relationships, resulting in a directed network with 883,712 nodes and 159 million edges—we refer to this as "Follow network" (see Table 10 in Appendix E for more details). We then incorporate the tipping relationships between pairs of users onto this network as additional edges to form the combined "Follow + Tip" network. The tipping relationships correspond to 55,847 edges. <sup>17</sup> To assess whether tipping rewards circulate within or across echo chambers, we identify communities within the follow network—*i.e.* groups of users with dense follow relationships each serving as a potential echo chamber.

We use two community detection approaches: *NetworKit*'s Louvain modularity optimization [83] and *Infomap*'s information flow-based partitioning [85]. Due to the inherent randomness in NetworKit's implementation, metrics subject to variation are reported as either means or ranges from three independent runs. Finally, we map tipping relationships (tip edges) onto the community structure to evaluate whether economic rewards tend to remain within follower communities or flow across them. Since tipping edges are overlaid on top of the follow network, we also assess whether they substantially alter the underlying community structure. To quantify the extent to which community structures persist across different network configurations (*i.e.* Follow vs. Follow+Tip), we use two standard metrics: Normalized Mutual Information (NMI) and Adjusted Mutual Information (AMI).<sup>18</sup>

*Results.* Using Louvain (Infomap) detection, we find that 52% (75%) of tips cross community boundaries while the remaining 48% (25%) stay within communities. We show full results for the relationship between tipping behavior and communities in Table 11 in Appendix E.

The stronger inter-community tipping observed under Infomap reflects its finer-grained community resolution. Infomap produces hierarchical clusters at multiple levels, *i.e.* Levels 0, 1, 2, and 3. We focus on Level 1 for our analysis, because it strikes a balance between overly coarse groupings (*e.g.* a single dominant community at Level 0) and overly fragmented structures at finer levels. At this level, the largest community detected by Infomap contains 330,867 users, compared to approximately 462,230 in Louvain.

We observe a clear difference in tipping behavior based on the underlying follow relationship between users. As shown in Table 5, tipping between users who never followed each other is substantially more likely to cross community boundaries: 45–51% under Louvain and 74.68% under Infomap. In contrast, tips between users with an existing follow relationship are more likely to

 $<sup>^{17}</sup>$ To ensure robust analysis, we exclude the lottery tipping bot (FID: 987581, Fname: Warpslot) to focus on organic content-driven tipping interactions.

<sup>&</sup>lt;sup>18</sup>NMI measures the similarity between two clusterings but may overstate agreement by not accounting for chance overlap. AMI corrects for this by adjusting for the expected similarity under random labelings to yield a more conservative measure of structural alignment.

remain within the same community—only 22–26% cross-community under Louvain and 56.1% under Infomap.

These differences in community-level tipping behavior are consistent with structural differences in how communities are formed under each network. As shown in Table 6, Louvain and Infomap produce highly similar communities when applied to the follow and combined graphs (NMI = 0.73 and 0.91, respectively), suggesting that tipping edges have limited impact on the overall community structure. This is also shown by the similar network metrics between Follow-only and Follow+Tip networks (Table 10 in the Appendix E.). However, both algorithms yield substantially lower overlap between follow and tip networks (*e.g.* AMI = 0.26 for Louvain, 0.13 for Infomap), indicating that tipping relationships form a distinct layer of interaction. Thus, while follow links defines stable community boundaries, tipping behaviors can cross these boundaries, particularly under finer-grained community partitions.

#### 6 Effectiveness of Token Incentives on Social Activities

Building upon our findings from Section 4 (RQ1) regarding the prevalence and diversity of token adoption, and Section 5 (RQ2) concerning socioeconomic risks, we finally investigate whether token incentives effectively encourage subsequent social engagement (RQ3), as this is the ultimate goal of the incentive design.

Given the criticism faced by previous platforms (e.g. Steemit) for coordinated low-quality content farming [8, 66], we specifically focus on whether Farcaster's pluralistic token incentive ecosystem fosters greater user engagement. To answer this, we investigate the causal impact of Farcaster's token incentives on social behavior through two complementary approaches [4, 9]: binary treatment analysis and continuous treatment analysis.

#### 6.1 Binary Treatment: Recipients vs. Non-Recipients.

Overview. We begin by using the binary treatment (e.g. receipt of a token reward) to compare reward recipients versus non-recipients, to measure the social impact of token rewards. Our analysis spans November 7, 2023, to April 27, 2025, examining five tokens (DEGEN, MOXIE, HIGHER, TN100X, USDC) across the three incentive mechanisms described in Section 4.2: user-to-user tipping, third-party rewards, and Farcaster's official algorithmic rewards. This generates eight token-mechanism pairs (as shown in Section 4.3), whose effects we analyze on nine social activities: posting, and bidirectional interactions in replying, liking, re-posting, and following. <sup>19</sup>

We implement Propensity Score Matching (PSM) and Difference-in-Differences (DID) analysis using a temporal alignment approach: each user's first reward reception or wallet binding timestamp is designated as T=0, with a four-week observation window before and after. We can then compare activity levels before vs. after. This window size aligns with established practices in previous causal inference studies and provides sufficient time to observe behavior changes while minimizing confounding temporal effects [9]. We next explain how we implement PSM and DID.

Propensity Score Matching (PSM). To compare the impact of receiving rewards, we employ PSM to construct comparable treatment and control groups by matching users with similar pre-treatment characteristics. We validate matching quality by examining standardized mean differences (SMD) of covariates (i.e. observed pre-treatment characteristics that may influence treatment or outcome) between matched groups. SMD is calculated as the difference in means between treatment and control groups divided by the pooled standard deviation, with values below 0.1 indicating successful

 $<sup>^{19}</sup>$ We include wallet binding as a baseline binary treatment to assess how participation in the token economy affects user behavior (see Table 7) .

matching in relevant studies [6].<sup>20</sup> Our matching incorporates comprehensive covariates, covering social activity metrics (account age when receiving the token reward, weekly aggregated posting frequency, bidirectional following, replying, liking, and re-posting frequencies) and token reward features (weekly aggregated reception frequencies across all token-mechanism pairs).

Our primary specification includes all available covariates in the PSM to ensure optimal matching between control and treatment groups. Diagnostic assessments demonstrate successful matching outcomes, with most covariates achieving SMD < 0.1 (see Figure 10) and matched pair sizes representing approximately 50% of their corresponding populations across different token-mechanism pairings (see Table 7). Our Average Treatment Effect on the Treated (ATT) and DID regression models below incorporate time fixed effects but exclude user fixed effects, as PSM already ensures group comparability.

Difference-in-Differences (DID). Beyond the PSM, to strengthen causal identification and account for time-varying confounders, we implement a DID analysis with parallel trend validation. The parallel trends assumption, which is fundamental to DID, requires that treatment and control groups exhibit similar outcome trajectories during the pre-treatment period [64].

Our validation approach divides the event timeline into pre-treatment (T-4 to T-1) and post-treatment (T+1 to T+4) windows, where the number following T denotes the number of weeks relative to the treatment day (T+0). For each pre-treatment window, we estimate differential coefficients between treatment and control groups. These coefficients measure the additional differences in outcome variables (such as posting frequencies) between treatment and control groups at each time window t.

In a valid parallel trend test, pre-treatment coefficients should be statistically insignificant (p-value > 0.05) [96]. We additionally adopt a 25% tolerance criterion: the parallel trends assumption is considered to hold if statistically significant pre-treatment differences appear in no more than one quarter of the pre-intervention windows. This allowance accounts for behavioral adjustments in anticipation of reward eligibility—such as increased engagement aimed at maximizing reward probability—while preserving the integrity of the identification strategy. This approach aligns with context-aware thresholds discussed in prior methodological work [96].

Covariate Adjustment. Due to high inter-correlations among social behaviors [80, 106], unadjusted analyses risk inflating treatment effects ATT by confounding concurrent activities (e.g. an increase in posting may naturally correlate with a rise in likes and replies). Our initial result exhibits this, showing broad positive impacts of token incentives on most social activities.

To mitigate this, we use a covariate-adjusted method that controls for both pre- and post-treatment social and token reward features, such as controlling for all other reward receptions and social activities when analyzing DEGEN tipping's effect on posting. Therefore, we further employ a covariate-adjusted method that accounts for both pre- and post-treatment social and token reward features as potential confounders. For example, when analyzing the impact of DEGEN tipping on posting behavior, we control for all other token rewards and social activities (both pre- and post-treatment) as confounders. This comprehensive approach reveals that the estimated effects of token rewards (*i.e.* net effects ATT) often become smaller—and sometimes reverse direction. These findings suggest a substantial correlation among social behaviors. Therefore, with the covariate-adjusted model accounting for additional social activity as confounding factors, net effects ATT more accurately reflect the independent impact of token rewards on specific behaviors, rather than capturing spillover effects through correlated activities. This net effect approach provides

 $<sup>^{20}</sup>$ In the literature, higher SMD thresholds (0.25) are also proposed [7].

deeper mechanistic insights, enabling us to identify which token rewards drive low-quality content farming versus high-quality engagement.

#### 6.2 Continuous Treatment: Reward Reception Frequency.

To quantify the intensity effect of each additional reward on social behaviors beyond binary treatment, we analyze how reward frequency affects behavioral changes through Ordinary Least Squares (OLS) regression:

$$\Delta Y_i = \alpha + \beta \cdot \log(RF_i) + \gamma \cdot C_i + \epsilon_i \tag{1}$$

where i indexes users,  $\Delta Y_i$  represents the change in social behavior metrics (calculated as post-treatment minus pre-treatment social activity frequencies),  $\log(RF_i)$  is the log-transformed frequency of a certain type of token reward received, and  $C_i$  includes all available pre-treatment covariates (e.g. posts and replies). For instance, when analyzing DEGEN tipping's impact,  $\Delta Y_i$  measures the change in weekly posting frequency, while  $RF_i$  counts the frequency of DEGEN tips received. Using this methodology, we analyze users who have received at least one instance of the relevant token reward, focusing on 4-week windows before and after alignment points.

#### 6.3 Results and Findings

Table 7 presents the results of our binary treatment causal analysis. We use colored symbols to denote significant effects that pass the parallel trends test (including tolerance cases): green + for positive effects and red – for negative effects, with the number of symbols indicating significance levels (e.g.+:p<0.05,++:p<0.01,+++:p<0.001). Non-significant effects are marked with "N". In Table 7, among the 81 treatment-outcome pairs (9 social activities × 9 treatments), we denote 6 cases (7.41%) passing with tolerance as "C", 12 failing cases (14.81%) as "F", and leave complete passes unmarked (63 cases, 77.78%).

Moreover, both effects must be in the same direction (either both positive or both negative). The results are summarized in Table 8, where the regression coefficients satisfying these criteria are highlighted with color (green for positive effects, red for negative effects). A more detailed result table, including  $\mathbb{R}^2$  and standard errors (SE), is provided in Table 12 in Appendix F.

Our causal analyses reveal several key patterns in how different token incentive mechanisms shape user behavior on Farcaster. These findings span three main dimensions: the quantity-quality trade-off in content engagement, the dynamics of social network growth, and the intensity effects of repeated token rewards. Through these analyses, we uncover both intended and inadvertent consequences of token-based incentive mechanisms.

Trade-off between Engagement Quantity and Quality. Our binary treatment analysis employing PSM and DID reveals that the initial reception of token incentives generally increases content engagement quantity (posts and replies) while showing insufficient effectiveness in improving quality (likes and re-posts).

To illustrate these findings, we present DID visualizations (Figure 5) for cases that both pass the parallel trends test and show significant positive ATT results (detailed in Section 6.1). We focus on three representative token-mechanism pairs as binary treatments: DEGEN tipping, DEGEN algorithmic reward, and USDC algorithmic reward. This selection enables us to compare both the effectiveness of different mechanisms (tipping vs. algorithmic) for the same token (DEGEN) and different token types (volatile DEGEN vs. stablecoin USDC) under the same mechanism (algorithmic). Complete results are available in Table 7.

 $<sup>^{21}\</sup>mathrm{Due}$  to the complexity of comparing the amount of USD value across different tokens, we only measure and compare the token reward frequencies.

Table 7. Causal effect summary of binary treatments on s
--

	Wallet		Inter-FID Tipping					Algo. Rewards			
Action	Binding	DEGEN	HIGHER	MOXIE	TN100X	USDC	DEGEN	MOXIE	USDC		
post	F,+++	+	C,+++	+	N	N	C,+++	F, + + +	N		
reply_out	N	+	_	+	N	N	N	N	N		
reply_in	N	N	N	N	+++	N	+++	+++	N		
like_out	F, N	N	N	_	N		+++		+++		
like_in	N	+++	N	N	C,	N	N	F, N	+		
repost_out	C, + + +	_	N	N	N		N	N			
repost_in	N	N	N	N	N	N					
follow_out	F, + + +	F,	F,	C, <b>–</b>	F,		F, + + +	F, + + +	F,		
follow_in	+++	C, + + +	N		N	+++	+++	F, N	+++		
populaion_si	ze 574829	40836	15849	5252	3459	15872	47748	21505	28181		
matched_pair	rs 48799	16817	7795	2643	2257	7119	15100	12260	7217		

#### Symbols:


- + & -: Positive & negative causal effects (measured by ATT significance levels (p < 0.05, p < 0.01, p < 0.001));
- + & -: Positive & negative causal effects with parallel trend pre-test passed (including deviation tolerance);
- N: Average Treatment Effect on the Treated (ATT) not significant in post-treatment period;
- C: Deviation exists in parallel trend pre-test (only 1 week deviation);
- F: Parallel pre-test fails (more than 1 week deviation).

Table 8. Regression summary of continuous treatment intensity with social activities.


		Inter-FID Tipping					Algo. Rewards			
Action	DEGEN	TN100X	HIGHER	MOXIE	USDC	DEGEN	MOXIE	USDC		
post	1.5703***	5.0001	1.9701***	-0.0517	-0.0805	12.7407***	15.2254***	8.2309***		
reply_out	-4.4867	8.6610	-3.8922	10.6198	-0.3703	9.4023	59.2029***	-19.1090***		
reply_in	1.7207	-35.1615	-2.4438	10.4600	-10.2936**	-4.8940	57.1112***	5.1841		
like_out	-1.4999	6.4651	1.4243	-4.1602	7.5337**	-9.7807	22.8327***	32.0449***		
like_in	-1.2148	25.2958	17.9290***	-27.5476***	0.5102	-77.0576***	7.7791**	73.9531***		
repost_out	0.4519	7.1819	2.4716*	-0.9295	-0.0732	-5.6276	-0.1039	-9.1389***		
repost_in	1.7003	-5.8862	-5.2017***	5.6276**	8.7852*	5.0918	-6.2899***	-30.2670***		
follow_out	0.2749	-0.9579	0.0211	14.5753***	0.6082	0.9927	20.3066***	-54.1936***		
follow_in	-12.4111	-57.8221	-19.9274***	9.9655***	15.6006	196.8301***	6.3755**	242.8866***		
population_size	40836	3459	15849	5252	15872	47748	21505	28181		
sample_size	38532	3438	13854	5041	13879	47748	21482	27484		

**Symbols:** \*\*\* p < 0.001, \*\* p < 0.05, \* p < 0.01. Significant coefficients with corresponding significant causal effects are bolded and colored


Figure 5a illustrates that DEGEN tipping positive, albeit delayed, effects (+, p < 0.05) on weekly post frequencies after the tipping reception (T=0), with the increase beginning around T+2. The parallel trend assumption is satisfied (see the shaded pre-treatment gap between the control group line and the treatment group line). Similar positive effects are observed for other tipping tokens, with HIGHER and MOXIE showing significant positive impacts (+, p < 0.05), while TN100x and USDC show no significant effects. Figure 5c demonstrates that DEGEN as an algorithmic reward shows even stronger positive effects on weekly post frequencies (+++, p < 0.001). The visualizations (Figures 5a and 5c) clearly show a more pronounced treatment effect for DEGEN algorithmic rewards




(a) DEGEN (tipping)'s positive impact (+) on weekly post frequencies.



(c) DEGEN (algorithmic reward)'s positive impact (+++) on weekly inbound post frequencies.



(b) DEGEN (tipping)'s positive impact (+++) on weekly inbound like frequencies.



(d) USDC (algorithmic reward)'s positive impact (+) on weekly inbound like frequencies.

Fig. 5. Difference-in-Differences (DID) visualizations.

compared to DEGEN tipping. While Moxie's algorithmic reward also shows positive ATT (+++, p < 0.001), it fails the parallel trends assumption test (see Table 7). In contrast, Farcaster's official USDC algorithmic rewards show no significant effects on posts and replies. According to these observations, third-party algorithmic rewards using volatile tokens as the medium (DEGEN and MOXIE) actually demonstrate stronger positive effects on content engagement quantity compared to both tipping mechanisms (all five tokens) and the official algorithmic rewards distributing stablecoins (USDC).

Compared to quantity metrics, token incentives show very limited improvement in content quality, measured by received likes and reposts. Only DEGEN tipping (Figure 5b (+++, p < 0.001)) and USDC algorithmic rewards (Figure 5d (+, p < 0.05)) significantly increase recipients' received likes. This may be attributed to DEGEN and USDC's network effects as the most-traded reward token on Farcaster (discussed in Section 4.2). However, these effects do not generalize to other tipping or third-party algorithmic tokens. Moreover, no tipping tokens show significant effects on re-post gains, while algorithmic reward tokens even show negative effects (see Table 7). This suggests token incentives not only fail to promote high-quality and share-worthy content, but may even have counter-productive effects, potentially echoing previous literature's findings on financial rewards' crowd-out effects on quality content due to strategic farming behaviors prioritizing quantity over quality [49, 71, 92, 116].  $^{22}$ 

 $<sup>^{22}</sup>$  These malicious behaviors have been reported by the co-founder of Farcaster (https://farcaster.xyz/v/0x0e31071c) and both DEGEN and MOXIE developers (https://x.com/degentokenbase/status/1802985205021466790, https://farcaster.xyz/dwr.eth/0x8bfde087)

Effects on Follower Growth. Beyond content interactions, we next examine follower growth. Wallet binding, the pre-requisite token economy participation behavior, serving as a baseline binary treatment, shows significant positive effects on follower growth (+++, p < 0.001), indicating users participating in Farcaster's token economy gain more followers than non-participants. This effect extends to both DEGEN and USDC across tipping and algorithmic mechanisms, reinforcing their unique network effect and social recognition status. However, all token rewards, including wallet binding, show neutral or negative effects on follow-out behavior, suggesting token economy participants are more likely to focus on self-promotion than expanding social connections, potentially exacerbating echo chamber effects.

Token Incentive Intensity Effects. Our continuous treatment analysis employing OLS regression provides insights into the cumulative effects of token rewards. Comparing causal analysis results (Table 7) with significant and directionally consistent OLS regression coefficients, we highlight significant intensity effects in Table 8 — examining whether higher reward frequency correlates with stronger social behavior impacts. The tipping mechanisms show minimal intensity effects, with only DEGEN and HIGHER showing slight positive effects on posting (coefficients: DEGEN 1.57, HIGHER 1.97, indicating less than 2 additional weekly posts per reward).

In contrast, algorithmic rewards demonstrate substantial intensity effects across most social behaviors. DEGEN and USDC algorithmic rewards show particularly strong follower growth effects ( $\approx$  197 and 243 additional weekly followers per reward respectively), while MOXIE shows no such effect. DEGEN and MOXIE algorithmic rewards show significant positive intensity effects on content quantity ( $\approx$  13 additional weekly posts per DEGEN reward,  $\approx$  57 additional replies per MOXIE reward) but no effects on quality metrics.

USDC algorithmic rewards demonstrate a more nuanced impact pattern: while showing strong positive intensity effects on like-based interactions ( $\approx +32$  likes given and  $\approx +74$  received per reward), they simultaneously exhibit significant negative effects on content sharing ( $\approx -9$  re-posts given and  $\approx -30$  received). Combined with the neutral effects on posting frequency, this pattern suggests that Farcaster's official USDC algorithmic rewards may shift user behavior toward producing content that attracts quick, surface-level engagement (likes) rather than content worthy of redistribution (re-posts). This behavioral shift aligns with previous research on monetary incentives in social platforms [56, 63, 92, 113, 116], where extrinsic rewards can potentially alter content creation motivations from intrinsic quality pursuit to reward optimization. The divergence between like-based and repost-based engagement particularly highlights how token incentives might inadvertently promote content optimized for immediate reaction rather than lasting value that users want to preserve and share with their networks.

#### 7 Related Work

The study of how incentives influence user behaviors and network dynamics on social platforms is a well-established area, situated at the intersection of behavioral economics and online social networks [70, 78, 99]. Prior research frameworks have examined how social and financial incentives shape user participation and network evolution [1, 59], utilizing controlled field experiments [102], laboratory simulations [50, 62], observational data analyses [8, 70], and quasi-experimental approaches [12, 115]. Measurement metrics include engagement indicators (likes, re-posts, replies) [26, 48, 106], content quality (accuracy, complexity, informativeness) [16, 89], and network-level effects such as clustering and propagation [15, 95, 118]. Within this context, our study leverages observational data from Farcaster, employing PSM and DID as quasi-experimental approaches to examine the influence of token incentives on social engagement indicators, using likes and re-posts as proxies for content quality.

Studies of traditional centralized social platforms demonstrate that monetary incentives reliably increase the quantity of social engagement behaviors, particularly under performance-contingent schemes [16, 50, 62, 111, 114, 115], though effects on content quality and novelty are mixed [56, 63, 92, 113, 116]. Moderating factors such as demographics, user characteristics, social status, and platform context critically shape incentive effectiveness [4, 51], while combined monetary and social incentives often yield superior outcomes [78, 89, 97]. Temporal analyses reveal strong short-term engagement boosts but potential long-term habituation effects (*i.e.* frequent users develop reduced sensitivity to social rewards over time, while occasional users remain highly responsive [4]), crowding-out effects (*i.e.* monetary incentives reduce intrinsic motivation, negatively impacting content quality) [49, 71, 92, 116], and inequality amplification [24]. These findings underscore the complexity of incentive design and user heterogeneity in digital environments.

In blockchain-based decentralized social platforms, Steemit [101] remains the most studied [8, 22, 52, 66, 79]. Steemit's proprietary blockchain and platform-mandated token mechanism enabled early advances in decentralized incentive design, eliminating transaction fees and facilitating high-throughput reward distribution. However, these design choices unintentionally introduced critical vulnerabilities, including susceptibility to farming and collusion [8], bot misuses [22], centralization of rewards, and exacerbation of economic stratification [52, 66, 79]. Research by Li et al. [66] and Ba et al. [8] indicates that successful users adapt their content strategies to maximize rewards, often focusing on content promotion rather than creation. This finding raises questions about whether financial incentives optimize for platform goals or user gaming. Ba et al. [8] further reveal strong correlations between cryptocurrency prices and user activity levels on Steemit: when token values increase, posting activity and user engagement spike correspondingly.

These prior works have focused on examining single mechanisms. However, studies suggest that the most effective incentive systems should combine multiple types of rewards rather than relying on single mechanisms [78, 89, 97]. Thus, our work differs from the above in that we move beyond single-incentive vulnerabilities. Instead, our research offers the first empirical analysis of Farcaster's *pluralistic* incentive ecosystem—integrating multiple tokens and diverse reward mechanisms through modular wallet binding and third-party reward projects [33, 35, 37]. Notably, we find that despite individual mechanisms retaining some prior identified shortcomings, their coexistence and complementarity show the potential to mitigate platform-wide risks.

#### 8 Discussion and Conclusion

#### 8.1 Summary of Main Findings

We have presented the first large-scale empirical analysis of Farcaster's pluralistic token incentive ecosystem, examining how diverse reward mechanisms shape user behavior and social network structure. Through the analysis of 574,829 wallet-linked users (64.25% of the user base), we have revealed several critical insights about token-based incentives in decentralized social networks. Our analysis demonstrates that while token incentives effectively drive platform growth and user participation, their differences in eligibility criteria, reward distribution structure and token types significantly impact socioeconomic outcomes.

While *user-to-user tipping* represents the most flexible and common incentive mechanism, it is predominantly unidirectional (with less than 10% of users acting as both tip receivers and senders) (see Section 5.1). Additionally, 52–75% of tips occur across community boundaries, and 32.42% between non-following pairs. This suggests that token incentives can facilitate value exchange beyond established social community structures (see Section 5.3).

Examining *algorithmic reward* mechanisms, we observe notable differences in inclusivity (Section 5.1). DEGEN, which relies on user-driven nominations, reaches up to 70% new participant rates.

In contrast, MOXIE relies on an open-source behavioral scoring algorithm and includes only 7.6% of new participants. This contrast suggests that transparent scoring systems are more susceptible to exploitation, reducing entry opportunities for new users.

However, wealth concentration persists across mechanisms (Gini coefficients: 0.72-0.94) (see Section 5.2). Compared to user-to-user tipping, algorithmic rewards demonstrate greater inequality, primarily due to: (1) the token staking model (e.g. MOXIE and DEGEN), which amplifies incumbent advantages; and (2) increased vulnerability to strategic farming and gaming (Section 4.3). Notably, MOXIE's innovative follower-followee redistribution mechanism alleviates initial concentration effects, suggesting that well-designed secondary distributions can help address wealth inequality.

Moreover, our causal analysis (see Section 6) uncovers fundamental trade-offs in promoting social activity via token incentives: while most rewards effectively boost content creation quantity (posts and replies), they often fail to enhance—and sometimes undermine—content quality measured by likes and re-posts. Furthermore, we identify potentially problematic patterns in repeated reward effects. Users receiving multiple algorithmic rewards show asymmetric network growth, seemingly prioritizing follower acquisition over outbound following. They also demonstrate strategic optimization behaviors, favoring immediate reactions over share-worthy content creation. These findings suggest that while token incentives can drive engagement, their current implementations inadvertently encourage superficial participation over meaningful social interaction.

#### 8.2 Design Implications and Mitigation Strategies

Our results yield three implications for sustainable incentive design.

First, coordination across pluralistic rewards is limited. Programs on Farcaster are launched independently, target overlapping social metrics (e.g. posts, replies, and likes), and operate without shared objectives or interoperable scoring. This leads to inefficient spending and conflicting reputation signals. This suggests that there could be value in introducing a composable reputation layer, which could provide a common substrate by aggregating on-chain behavior, off-chain interactions, and temporal consistency, ideally in a privacy-preserving fashion. Standardized interfaces could then let different programs align on definitions of value (e.g. novelty, informativeness, cross-community engagement), avoid redundant payouts, and specialize on complementary goals.

Second, our results indicate that Farcaster's existing incentives inadequately foster network exploration and content quality. This is evidenced by a tension in tipping behaviors (see Section 5.3): despite a substantial proportion (48%) of tips between non-following pairs occurring across community boundaries, tipping overall still exhibits echo-chamber tendencies, with 67% of all inter-FID transactions occurring within detected communities. Simultaneously, a quantity-quality trade-off exists (see Section 6.3): while most tokens boost content creation (measured by posts and replies), they often fail to enhance content quality (measured by likes and re-posts). We argue that dedicating budget to exploration-aware payouts (e.g. cross-community edges, first-time ties) and to quality-weighted rewards (e.g. delayed distributions conditioned on downstream popularity or more advanced quality metrics) is crucial to shift behavior toward new content discovery and quality social engagement.

Third, it is clear that robustness to gaming and manipulation is essential. The current static, transparent, count-based scoring invites farming, collusion and wealth concentration. Potential mitigations include periodically rotated and adversarially tested scoring functions; a mix of limited-opacity components with auditable elements; anti-collusion safeguards (staked attestations, adaptive rate limits); and redistribution mechanisms (e.g. MOXIE-style follower

sharing) to counter excessive wealth accumulation. Together, these measures could harden the system while preserving incentives for genuine, socially valuable activity.

#### 8.3 Limitations

This work carries several limitations that shape the scope of inference, and inform how our results may generalize beyond the studied setting.

First, our measurement of content quality relies on engagement-based proxies such as likes and re-posts. Whilst widely used in prior work [26, 48, 106], these signals imperfectly capture intrinsic quality. They are affected by users' network position, temporal visibility, and short-lived trends. As a result, our estimates may understate improvements in deeper dimensions—such as informativeness, novelty, or persuasiveness—that are less immediately reflected in reaction metadata (e.g. likes and re-posts).

Second, our DID design uses a four-week pre/post window around alignment events. Content with longer discovery cycles or diffusion paths may accrue recognition beyond this horizon. Consequently, our identification is more sensitive to short-term responses and may bias against detecting lagged, higher-fidelity effects. Methodological extensions with longer windows and complementary designs could better separate immediate reactivity from durable changes.

Third, Farcaster remains smaller and more crypto-centric than mainstream platforms (e.g. Twitter/X, Facebook) and some larger DOSNs (e.g. Mastodon, Bluesky). This composition likely raises activation thresholds and reduces sensitivity to marginal token incentives, which can attenuate absolute effect sizes. Nonetheless, the directional insights we uncover — wealth concentration risks, cross-community value flows, and the tension between content quantity and quality — are expected to carry over to other platforms that deploy analogous pluralistic incentive schemes.

Regarding external validity, our token selection follows a multi-criteria, robust filtering pipeline (see Section 4.2) that emphasizes sustained, community-relevant activity. This approach improves the representativeness of mechanisms most likely to shape ecosystem-level behavior, compared to analyzing arbitrary or transient tokens. Generalization to substantially less cryptonative populations or to ecosystems with different incentive mechanisms should be interpreted with the caveat that absolute magnitudes may differ, even if the underlying trade-offs persist.

#### 8.4 Future Directions

Building on our empirical findings, several avenues emerge for advancing research on token-based incentives in decentralized social networks.

First, future studies could develop and evaluate composable reputation systems that integrate multi-token signals, enabling coordinated reward allocation across mechanisms while preserving user privacy. Such systems would address the observed limitations in incentive alignment, potentially through machine learning models that optimize for long-term platform health metrics like sustained engagement and content diversity.

Second, research could investigate enhanced mechanism designs that explicitly incentivize network exploration and high-quality content creation. This might involve experimental deployments of rewards weighted by cross-community interactions or delayed propagation signals, testing their efficacy in mitigating asymmetric growth and superficial participation patterns identified in our analysis.

Third, advancing robustness against gaming requires systematic evaluation of adaptive scoring functions, anti-collusion protocols, and redistribution schemes. Longitudinal studies

could assess these interventions' impact on wealth concentration and inclusivity, particularly in platforms with varying scales and user demographics.

Finally, comparative analyses across diverse DOSNs would elucidate how contextual factors—such as user sovereignty levels and token volatility—influence incentive outcomes, informing generalizable design principles for sustainable token economies.

#### 8.5 Concluding Remarks

To conclude, our analysis reveals that, despite the persistent limitations of individual tokens or mechanisms, their combined presence and mutual reinforcement can effectively mitigate platform-wide vulnerabilities, such as high barriers to entry, excessive wealth concentration, the formation of echo chambers, and strategic farming behaviors that arise from a single tokenomic model. However, we also find that the mere coexistence of pluralistic incentives, without organic coordination and integration, fails to achieve the desired outcomes in terms of sustainable network growth and content quality.

Our findings advance understanding of token-based incentive design and provide future research directions and practical guidance for implementing reward mechanisms in social platforms. More broadly, our results suggest that while token incentives offer promising tools for decentralized platforms, their effectiveness depends on careful mechanism design that considers both immediate behavioral impacts and longer-term social dynamics.

#### Acknowledgments

This work was supported in part by the Guangzhou Science and Technology Bureau (2024A03J0684), Guangdong provincial project 2023QN10X048, the Guangzhou Municipal Key Laboratory on Future Networked Systems (2024A03J0623), the Guangdong Provincial Key Lab of Integrated Communication, Sensing and Computation for Ubiquitous Internet of Things (No.2023B1212010007), the Guangzhou Municipal Science and Technology Project (2023A03J0011), Guangdong provincial project (2023ZT10X009), and the 111 Center (No. D25008).

#### References

- Samuel F Acuff, James MacKillop, and James G Murphy. 2018. Applying behavioral economic theory to problematic Internet use: An initial investigation. *Psychology of Addictive Behaviors* 32, 7 (2018), 846.
- [2] Alchemy. 2025. Alchemy's API for Token Transfer History. https://www.alchemy.com/docs/reference/transfers-apiquickstart. Accessed: 2025-07-29.
- [3] Darcy WE Allen, Chris Berg, and Aaron M Lane. 2023. Why airdrop cryptocurrency tokens? *Journal of Business Research* 163 (2023), 113945.
- [4] Ian A Anderson and Wendy Wood. 2023. Social motivations' limited influence on habitual behavior: Tests from social media engagement. *Motivation Science* 9, 2 (2023), 107.
- [5] Nasreen Anjum, Dmytro Karamshuk, Mohammad Shikh-Bahaei, and Nishanth Sastry. 2017. Survey on peer-assisted content delivery networks. *Comput. Netw.* 116, C (April 2017), 79–95. doi:10.1016/j.comnet.2017.02.008
- [6] Peter C Austin. 2009. Balance diagnostics for comparing the distribution of baseline covariates between treatment groups in propensity-score matched samples. *Statistics in medicine* 28, 25 (2009), 3083–3107.
- [7] Peter C Austin. 2011. An introduction to propensity score methods for reducing the effects of confounding in observational studies. Multivariate behavioral research 46, 3 (2011), 399–424.
- [8] Cheick Tidiane Ba, Matteo Zignani, and Sabrina Gaito. 2022. The role of cryptocurrency in the dynamics of blockchain-based social networks: The case of steemit. PloS one 17, 6 (2022), e0267612.
- [9] Leonhard Balduf, Saidu Sokoto, Andrea Baronchelli, Ignacio Castro, Michał Król, Gareth Tyson, George Pavlou, Björn Scheuermann, and Onur Ascigil. 2025. Bootstrapping Social Networks: Lessons from Bluesky Starter Packs. In Proceedings of the International AAAI Conference on Web and Social Media, Vol. 19. 178–192.
- [10] Base. 2025. Base Chain Documentation. https://docs.base.org/. Accessed: 2025-07-29.
- [11] Bitget. 2025. Farcaster Launches "TIPS" Feature, Supporting One-Click Sending of USDC Rewards. https://www.bitget.com/news/detail/12560604678174.
- [12] Gordon Burtch, Yili Hong, Ravi Bapna, and Vladas Griskevicius. 2015. What are social incentives worth? A randomized field experiment in user content generation. (2015).
- [13] Jason C. 2025. Memo.Cash Official Documentation. https://jasonc.me/blog/what-is-memo. Accessed: 2025-07-29.
- [14] Shanglin Cai. 2024. Innovation of Media Business Models in the Digital Economy Era. SHS Web of Conferences (2024). https://api.semanticscholar.org/CorpusID:272071211
- [15] Damon Centola. 2010. The spread of behavior in an online social network experiment. science 329, 5996 (2010), 1194–1197.
- [16] Gizem Ceylan, Ian A Anderson, and Wendy Wood. 2023. Sharing of misinformation is habitual, not just lazy or biased. Proceedings of the National Academy of Sciences 120, 4 (2023), e2216614120.
- [17] Clanker. 2025. Clanker Official Document. https://clanker.gitbook.io/clanker-documentation/users/deploying-a-token. Accessed: 2025-07-29.
- [18] Clanker. 2025. Data Dashboard: Comparing Clanker with Other Token Launchpads. https://dune.com/clanker\_protection\_team/clanker-vs-others. Accessed: 2025-07-29.
- [19] Clanker. 2025. The Official Website of Clanker. https://www.clanker.world/. Accessed: 2025-07-29.
- [20] Coingecko. 2025. Farcaster Token List. https://www.coingecko.com/en/categories/farcaster-ecosystem. Accessed: 2025-07-29.
- [21] Degen. 2025. Degen Tokenomics Document. https://www.degen.tips/. Accessed: 2025-07-29.
- [22] Fatemeh Delkhosh, Ram D Gopal, Raymond A Patterson, and Niam Yaraghi. 2023. Impact of bot involvement in an incentivized blockchain-based online social media platform. *Journal of Management Information Systems* 40, 3 (2023), 778–806.
- [23] DeSo. 2025. DeSo Official Documentation. https://docs.deso.org/. Accessed: 2025-07-29.
- [24] Paul DiMaggio and Filiz Garip. 2012. Network effects and social inequality. *Annual review of sociology* 38, 1 (2012), 93–118.
- [25] D Dittrich and E Kenneally. 2012. The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. Technical Report. U.S. Department of Homeland Security. doi:paper/2012\_menlo\_report\_actual\_ formatted
- [26] Tim Donkers and Jürgen Ziegler. 2025. Understanding Online Polarization Through Human-Agent Interaction in a Synthetic LLM-Based Social Network. In *International Conference on Web and Social Media*. https://api.semanticscholar. org/CorpusID:279226933
- [27] Dune. 2025. Farcaster Storage Fee Changes. https://dune.com/data/farcaster\_optimism.storageregistry\_evt\_setprice. Accessed: 2025-07-29.
- [28] Ethereum. 2025. The Ethereum address format and why it matters when using MetaMask. https://ethereum.org/en/whitepaper/. Accessed: 2025-07-29.
- [29] Farcaster. 2022. Decentralization of Hubs. https://hackmd.io/@farcasterxyz/ry0QL4M4o#Hub-Costs.

[30] Farcaster. 2023. Farcaster Official Documentation - Hub Monorepo. https://github.com/farcasterxyz/hub-monorepo. Accessed: 2025-07-29.

- [31] Farcaster. 2024. Diff-Sync Description. https://github.com/farcasterxyz/protocol/discussions/163.
- [32] Farcaster. 2024. Farcaster Conference 2024. https://farcon.xyz/2024.html. Accessed: 2025-07-29.
- [33] Farcaster. 2025. External Wallet Connection. https://docs.farcaster.xyz/developers/guides/writing/verify-address. Accessed: 2025-07-29.
- [34] Farcaster. 2025. Farcaster Contract Addresses for Identity Registration and Management. https://github.com/farcasterxyz/contracts. Accessed: 2025-07-29.
- [35] Farcaster. 2025. Farcaster Docs. https://docs.farcaster.xyz/ Accessed: 2025-07-29.
- [36] Farcaster. 2025. Farcaster Hub Setup Tuturial. https://docs.farcaster.xyz/hubble/tutorials. Accessed: 2025-07-29.
- [37] Farcaster. 2025. Farcaster Mini-app Documentation. https://miniapps.farcaster.xyz/. Accessed: 2025-07-29.
- [38] Farcaster. 2025. Farcaster Offical USDC Reward. https://warpcast.notion.site/Warpcast-Rewards-15f6a6c0c10180339fcdee6c9a18c338. Accessed: 2025-07-29.
- [39] Farcaster. 2025. Farcaster Storage Registry Contract. https://optimistic.etherscan.io/address/0x00000000fcce7f938e7ae6d3c335bd6a1a7c593d. Accessed: 2025-07-29.
- [40] Farcaster. 2025. Hub Installation Configuration. https://docs.farcaster.xyz/hubble/install. Accessed: 2025-07-29.
- [41] Farcaster. 2025. Snapchain Official Document. https://docs.farcaster.xyz/hubble/migrating. Accessed: 2025-07-29.
- [42] Farcaster. 2025. Token Index. https://www.farcaster.in/. Accessed: 2025-07-29.
- [43] Farcaster. 2025. Token Links. https://warpcast.notion.site/Token-Links-17e6a6c0c10180a68c4dce055b64d12b. Accessed: 2025-07-29.
- [44] Farcaster. 2025. What's New: May 21, 2025 (0.0.49) Introduced Wallet Standard integration for Solana wallets. https://miniapps.farcaster.xyz/docs/sdk/changelog#may-21-2025-0049. Accessed: 2025-07-29.
- [45] Father. 2025. Farther Official Website. https://farther.social/. Accessed: 2025-07-29.
- [46] Filarm. 2025. Comparison of Lens and Farcaster. https://dune.com/filarm/lens-vs-farcaster. Accessed: 2025-07-29.
- [47] Folklore. 2024. The DEGEN Archives. https://folklore.mirror.xyz/q9fjoEXpgDH8YGm\_EaN3y8jhJfc\_kRz7kBW-4IEya7A.
- [48] Marco Furini. 2024. X as a Passive Sensor to Identify Opinion Leaders: A Novel Method for Balancing Visibility and Community Engagement. Sensors (Basel, Switzerland) 24 (2024). https://api.semanticscholar.org/CorpusID:267092962
- [49] Ina Garnefeld, Anja Iseke, and Alexander Krebs. 2012. Explicit incentives in online communities: boon or bane? *International Journal of Electronic Commerce* 17, 1 (2012), 11–38.
- [50] Laura K Globig, Nora Holtz, and Tali Sharot. 2023. Changing the incentive structure of social media platforms to halt the spread of misinformation. *Elife* 12 (2023), e85767.
- [51] Anatoliy Gruzd, Alyssa Saiphoo, and Philip Mai. 2024. The lure of decentralized social media: Extending the UTAUT model for understanding users' adoption of blockchain-based social media. Plos one 19, 8 (2024), e0308458.
- [52] Barbara Guidi, Andrea Michienzi, and Laura Ricci. 2024. Assessment of Wealth Distribution in Blockchain Online Social Media. IEEE Transactions on Computational Social Systems 11, 1 (2024), 671–682. doi:10.1109/TCSS.2022.3228925
- [53] Ham.fun. 2025. TN100X Tokenomics. https://docs.ham.fun/docs/tokenomics/. Accessed: 2025-07-29.
- [54] Andrew Heath et al. 2020. GossipSub: Attack-resilient message propagation in the Filecoin and Ethereum networks. arXiv preprint arXiv:2007.02754 (2020).
- [55] Higher. 2025. HIGHER Official Greenpaper. https://www.aimhigher.net/greenpaper. Accessed: 2025-07-29.
- [56] Li-Chun Huang. 2023. The Effect of Incentives on Facilitating User Engagement with Succulent Retailers' Social Media Pages. Horticulturae 9, 8 (2023), 849.
- [57] Megumi Ichimiya, Hope Muller-Tabanera, Jennifer Cantrell, Jeffrey B Bingenheimer, Raquel Gerard, Elizabeth C Hair, Dante Donati, Nandan Rao, and W Douglas Evans. 2023. Evaluation of response to incentive recruitment strategies in a social media-based survey. *Digital health* 9 (2023), 20552076231178430.
- [58] Itsmechaseb. 2024. The Dawn of AI Autonomy: How Two AI Agents Created a \$70M Cryptocurrency. https://x.com/itsmechaseb/status/1857511608810901941?s=46&t=ywALktyM5ZCyt-zigN5Big.
- [59] Matthew O Jackson. 2011. An overview of social networks and economic applications. Handbook of social economics 1 (2011), 511–585.
- [60] Ujun Jeong, Lynnette Hui Xian Ng, Kathleen M Carley, and Huan Liu. 2025. Navigating Decentralized Online Social Networks: An Overview of Technical and Societal Challenges in Architectural Choices. arXiv preprint arXiv:2504.00071 (2025).
- [61] Jhackworth. 2025. Zora Data Dashboard. https://dune.com/jhackworth/zora-network. Accessed: 2025-07-29.
- [62] Hansika Kapoor, Sarah Rezaei, Swanaya Gurjar, Anirudh Tagat, Denny George, Yash Budhwar, and Arathy Puthillam. 2023. Does incentivization promote sharing "true" content online? Harvard Kennedy School Misinformation Review (2023).

- [63] Warut Khern-am nuai, Karthik Kannan, and Hossein Ghasemkhani. 2018. Extrinsic versus intrinsic rewards for contributing reviews in an online platform. *Information Systems Research* 29, 4 (2018), 871–892.
- [64] Michael Lechner et al. 2011. The estimation of causal effects by difference-in-difference methods. Foundations and Trends® in Econometrics 4, 3 (2011), 165–224.
- [65] Lens. 2025. Lens Official Documentation. https://lens.xyz/docs/chain/overview. Accessed: 2025-07-29.
- [66] Chao Li and Balaji Palanisamy. 2019. Incentivized blockchain-based social media platforms: A case study of steemit. In Proceedings of the 10th ACM conference on web science. 145–154.
- [67] J. Liang, L. Li, and D. Zeng. 2018. Evolutionary dynamics of cryptocurrency transaction networks: an empirical study. *Plos One* 13 (2018), e0202202. Issue 8. doi:10.1371/journal.pone.0202202
- [68] Meng Liang. 2022. The end of social media? How data attraction model in the algorithmic media reshapes the attention economy. *Media, Culture & Society* 44, 6 (2022), 1110–1131.
- [69] Jianhua Lin. 2002. Divergence measures based on the Shannon entropy. *IEEE Transactions on Information theory* 37, 1 (2002), 145–151.
- [70] Björn Lindström, Martin Bellander, David T Schultner, Allen Chang, Philippe N Tobler, and David M Amodio. 2021.
  A computational reward learning account of social media engagement. Nature communications 12, 1 (2021), 1311.
- [71] Yuewen Liu and Juan Feng. 2015. Can Monetary Incentives Increase UGC Contribution? The Motivation and Competition Crowding Out. (2015).
- [72] Luc. 2025. Dune Dashboard of Farcaster Token. https://www.dune.com/0xluc/farconomy. Accessed: 2025-07-29.
- [73] Christos A Makridis, Michael Fröwis, Kiran Sridhar, and Rainer Böhme. 2023. The rise of decentralized cryptocurrency exchanges: Evaluating the role of airdrops and governance tokens. Journal of Corporate Finance 79 (2023), 102358.
- [74] MK Manoylov. 2024. Paradigm leads \$150 million raise for web3 social media platform Farcaster. https://www.theblock.co/post/295700/paradigm-150-million-raise-web3-social-platform-farcaster.
- [75] M. Mardan and I. Khosravipour. 2024. Dynamic evolution analysis of cryptocurrency market: a network science study. Journal of Interdisciplinary Economics (2024). doi:10.1177/02601079241265744
- [76] Memo.cash. 2018. memo.cash Blog. https://memo.cash/blog/introducing-memo. Accessed: 2025-07-29.
- [77] MetaMask. 2025. Ethereum Whitepaper. https://support.metamask.io/start/learn/the-ethereum-address-format-and-why-it-matters-when-using-metamask/. Accessed: 2025-07-29.
- [78] Leticia Micheli, Mirre Stallen, and Alan G Sanfey. 2021. The effect of centralized financial and social incentives on cooperative behavior and its underlying neural mechanisms. Brain sciences 11, 3 (2021), 317.
- [79] Andrea Michienzi, Barbara Guidi, and Laura Ricci. 2024. A wealth-driven analysis of user engagement in Blockchain Online Social Media. In 2024 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops). IEEE, 637–642.
- [80] Smitha Milli, Emma Pierson, and Nikhil Garg. 2023. Choosing the Right Weights: Balancing Value, Strategy, and Noise in Recommender Systems. ArXiv abs/2305.17428 (2023). https://api.semanticscholar.org/CorpusID:258960306
- [81] MOXIE. 2025. FarScores, FarBoost, and Cast Scores. https://github.com/Airstack-xyz/docs/blob/main/social-capital-value-and-social-capital-scores.md. Accessed: 2025-07-29.
- [82] MOXIE. 2025. MOXIE Official Website. https://www.moxie.xyz/. Accessed: 2025-07-29.
- [83] Networkit. 2025. Community Detection with NetworKit. https://networkit.github.io/dev-docs/notebooks/Community. html. Accessed: 2025-07-29.
- [84] Neynar. 2025. Farcaster Official Dashboard. https://dune.com/neynarxyz/farcaster. Accessed: 2025-07-29.
- [85] Infomap Online. 2025. Network community detection using the Map Equation framework. https://www.mapequation. org/infomap/. Accessed: 2025-07-29.
- [86] OpenRank. 2025. Farcaster Scoring System. https://docs.openrank.com/integrations/farcaster. Accessed: 2025-07-29.
- [87] Optimism. 2025. Optimism Official Documentation. https://docs.optimism.io/. Accessed: 2025-07-29.
- [88] Ournetwork. 2025. Coverage on Clanker, Farcaster, Lens, Kaito, World, Layer3 and Basenames. https://www.ournetwork.xyz/p/on-320-onchain-culture?utm\_campaign=ournetwork.
- [89] Mahsa Paridar, Mina Ameri, and Elisabeth Honka. 2023. More, faster, and better? effects of rewards on incentivizing the creation of user-generated content. Effects of Rewards on Incentivizing the Creation of User-Generated Content (December 31, 2023) (2023).
- [90] Paycaster. 2025. Paybot Documentation. https://paycaster.co/. Accessed: 2025-07-29.
- [91] Pixelhack. 2025. Dune Dashboard of Farcaster users transactions by chain. https://dune.com/queries/3027146/5036846.Accessed: 2025-07-29.
- [92] Dandan Qiao, Shun-Yang Lee, Andrew Whinston, and Qiang Wei. 2017. Incentive provision and pro-social behaviors. (2017).
- [93] Aravindh Raman, Sagar Joglekar, Emiliano De Cristofaro, Nishanth Sastry, and Gareth Tyson. 2019. Challenges in the decentralised web: The mastodon case. In *Proceedings of the internet measurement conference*. 217–229.

[94] Aravindh Raman, Sagar Joglekar, Emiliano De Cristofaro, Nishanth Sastry, and Gareth Tyson. 2019. Challenges in the Decentralised Web: The Mastodon Case. In *Proceedings of the Internet Measurement Conference* (Amsterdam, Netherlands) (IMC '19). Association for Computing Machinery, New York, NY, USA, 217–229. doi:10.1145/3355369. 3355572

- [95] David G Rand, Samuel Arbesman, and Nicholas A Christakis. 2011. Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences* 108, 48 (2011), 19193–19198.
- [96] Jonathan Roth, Pedro H. C. Sant'Anna, Alyssa M. Bilinski, and John Poe. 2022. What's trending in difference-in-differences? A synthesis of the recent econometrics literature. Journal of Econometrics (2022). https://api.semanticscholar.org/CorpusID:245668975
- [97] Ognjen Scekic, Hong-Linh Truong, and Schahram Dustdar. 2013. Incentives and rewarding in social computing. *Commun. ACM* 56, 6 (2013), 72–82.
- [98] ScienceDirect. 2025. Clustering Coefficient. https://www.sciencedirect.com/topics/computer-science/clustering-coefficient. Accessed: 2025-07-29.
- [99] Vivek K Singh, Ramesh Jain, and Mohan Kankanhalli. 2011. Mechanism design for incentivizing social media contributions. In Social media modeling and computing. Springer, 121–143.
- [100] Solana. 2025. EVM vs. SVM. https://solana.com/developers/evm-to-svm/smart-contracts. Accessed: 2025-07-29.
- [101] Steemit. 2025. Steemit Official Documentation. https://steemit.com/guide/@steemitblog/steemit-a-guide-for-newcomers. Accessed: 2025-07-29.
- [102] Tianshu Sun, Siva Viswanathan, Ni Huang, and Elena Zheleva. 2018. Designing promotional incentive to embrace social sharing: Evidence from field and online experiments. Forthcoming at MIS Quarterly (2018).
- [103] Bishenghui Tao, Hong-Ning Dai, Jiajing Wu, Ivan Wang-Hei Ho, Zibin Zheng, and Chak Fong Cheang. 2021. Complex network analysis of the bitcoin transaction network. IEEE Transactions on Circuits and Systems II: Express Briefs 69, 3 (2021), 1009–1013.
- [104] TechFlow. 2024. Vitalik and Jesse are buying one after another, what is the background of the ANON in the Base ecosystem? https://en.theblockbeats.news/news/55854.
- [105] Christof Ferreira Torres, Mathis Steichen, et al. 2019. The art of the scam: Demystifying honeypots in ethereum smart contracts. In 28th USENIX Security Symposium (USENIX Security 19). 1591–1607.
- [106] Mariapina Trunfio and Simona Rossi. 2021. Conceptualising and measuring social media engagement: A systematic literature review. *Italian Journal of Marketing* 2021, 3 (2021), 267–292.
- [107] Uniswap. 2025. The Decentralized Exchange Application. https://app.uniswap.org/. Accessed: 2025-07-29.
- [108] Vaiva Vasiliauskaite and Fernando E. Rosas. 2020. Understanding complexity via network theory: a gentle introduction. arXiv: Physics and Society (2020). https://api.semanticscholar.org/CorpusID:216867636
- [109] Victor. 2025. 3 Ways to Earn on Farcaster: USDC, Frames, and Airdrops. https://www.altcoinbuzz.io/bitcoin-and-crypto-guide/3-ways-to-earn-on-farcaster-usdc-frames-and-airdrops/.
- [110] Vladimir. 2024. What is DEGEN? The History of The Top Memecoin on Base. https://zerion.io/blog/the-true-history-of-degen/.
- [111] Chuansheng Wang, Xuecheng Tan, and Fulei Shi. 2023. Recognition of opinion leaders in blockchain-based social networks by structural information and content contribution. *PeerJ Computer Science* 9 (2023), e1549.
- [112] Yiluo Wei and Gareth Tyson. 2024. Exploring the Nostr Ecosystem: A Study of Decentralization and Resilience. *ArXiv* abs/2402.05709 (2024). https://api.semanticscholar.org/CorpusID:267548110
- [113] Zhijun Yan, Lini Kuang, He Huang, and Han Yang. 2017. Do financial incentives induce more online participatory behaviors? (2017).
- [114] Yinan Yu, Warut Khern-am nuai, and Alain Pinsonneault. 2018. The impact of performance-contingent monetary incentives on user-generated content contribution. (2018).
- [115] Yinan Yu, Warut Khern-am nuai, and Alain Pinsonneault. 2022. When Paying for Reviews Pays Off: The Case of Performance-Contingent Monetary Rewards. MIS Quarterly 46, 1 (2022).
- [116] Dongcheng Zhang, Hanchen Jiang, Maoshan Qiang, Kunpeng Zhang, and Liangfei Qiu. 2025. Time to Stop? An Empirical Investigation on the Consequences of Canceling Monetary Incentives on a Digital Platform. *Information Systems Research* 36, 2 (2025), 781–801.
- [117] J. Zhang and J. Shi. 2022. Asymptotic normality for plug-in estimators of generalized shannon's entropy. Entropy 24 (2022), 683. Issue 5. doi:10.3390/e24050683
- [118] Rui Zheng, Luis Ospina-Forero, and Yu-wang Chen. 2024. Implications of social network structures on socially influenced decision-making. *Decision* 51, 1 (2024), 85–103.
- [119] Zora. 2024. Zora Official Documentation. https://zerion.io/blog/guide-to-the-zora-ecosystem/.
- [120] Wenrui Zuo, Aravindh Raman, Raul J Mondragón, and Gareth Tyson. 2023. Set in Stone: Analysis of an Immutable Web3 Social Media Platform. In *Proceedings of the ACM Web Conference 2023* (Austin, TX, USA) (WWW '23). Association for Computing Machinery, New York, NY, USA, 1865–1874. doi:10.1145/3543507.3583510

#### A Token-related Events Driving User Wallet Binding.

To better understand how token activities drive token economy participation, this section plots the wallet binding dynamics along with influential token-related events in greater detail.




Fig. 6. Top 10 surges in 7-day moving average percentage change of daily wallet bindings (highlighted by color-coded dashed lines), annotated with platform milestones and token campaigns.

We observe a notable delay effect between the onset of growth and peak binding activities. After experimenting with various smoothing techniques including Exponential Moving Average (EMA) and different moving average windows (3-day, 5-day, 7-day, and 10-day), we find that the 7-day moving average most effectively captures well-distributed top 10 surges that align with visual inspection of the data.

Figure 6 presents the 7-day moving average percentage change for wallet bindings. We then annotate major spikes in activity with key events identified in the Farcaster ecosystem. Key event identification follows our mixed-methods approach: We do this by reviewing news from The Block Beats<sup>23</sup> and posts from Farcaster's hub dataset, combined with quantitative examination of token trading frequencies among newly bound wallets around surge dates.

Through this, we identify key wallet-binding-driving events covering two categories: platform-led events and social-token-driven events (further classified into reward-token and meme-coin<sup>24</sup> events). We next introduce each category of event and discuss the corresponding surges in wallet binding activity.

 $<sup>^{23} \</sup>mbox{The Block Beats:}$  www.the blockbeats.info.

<sup>&</sup>lt;sup>24</sup>Meme-coins are tokens typically created as jokes or for entertainment purposes, often inspired by internet memes or popular culture, with their value largely driven by community sentiment and social media trends.

#### A.1 Platform-led Events.

Four surges in wallet bindings closely align with core Farcaster milestones: the Mini-apps launch (106.74% surge, 2024-01-27), the \$150M funding announcement (47.06%, 2024-05-21), the Farcaster Conference (FarCon) Asia (53.39%, 2024-09-17), and the official Farcaster wallet launch (105.02%, 2025-02-22). Notably, these platform-driven events account for the 2nd to 5th largest surges among the top 10 observed, with the Mini-apps launch ranking 2nd, the official wallet launch 3rd, and the funding announcement and FarCon Asia ranking 4th and 5th, respectively. This pattern demonstrates that user growth on Farcaster is driven by feature releases and platform milestones, rather than by entry barriers (e.g. registration fee reductions).

#### A.2 Reward-Token Events.

In addition to platform-led milestones, token airdrop<sup>25</sup> announcements—such as DEGEN (26.62% surge, 2023-12-27; 37.34%, 2024-03-27), \$FARTHER (27.42%, 2024-04-30), and MOXIE (119.93%, 2024-08-21)—represent a distinct category of events that also drive wallet binding surges. These initiatives incentivize user engagement through mechanisms that allocate daily token allowances based on social interactions and third-party reputation scores (*e.g.* OpenRank scores [86]), without requiring direct financial expenditure from users. The distributed rewards are funded by project treasuries locked in smart contracts [21, 45, 82], encouraging both new and existing users to link wallets.

A particularly noteworthy development occurred on August 21, 2024, when MOXIE introduced a permissionless mechanism for users to issue and auction their own profile-tokenized Fan Tokens—a model closely aligned with Lens Protocol [65] and Zora [119]. These fan tokens can be freely traded, and holders are eligible to receive approximately 20% of the fan token issuer's daily MOXIE engagement rewards. Although this announcement led to the largest observed surge in wallet bindings ( $\approx$  120%), the underlying tokenomics and redistribution dynamics are beyond the scope of this work; in Section 4, we focus on incentive mechanisms for the initial allocation of tokens based on user engagement.

#### A.3 Meme-Coin Events.

While reward tokens inherently contain speculative elements [47, 110], their primary design is to foster social engagement. In contrast, the launch of @clanker—an AI-powered token issuance bot (FID = 874542)—on November 9, 2024, marked a significant shift by enabling an automated pipeline for meme-coin creation [17]. Users can deploy new meme-coins simply by posting with the desired token name and description while mentioning @clanker. The bot then deploys the meme-coin on the Base chain, establishes initial liquidity pools, <sup>26</sup> and facilitates instant trading via the Clanker platform [19] or decentralized exchanges such as Uniswap [107]. By June 5, 2025, Clanker had enabled the creation of 280,678 meme-coins, with 28,224 (about 10%) of these tokens observed in the on-chain transaction records of Farcaster users' wallets, attracting considerable attention for its rapid wealth effects and fee revenue model [18, 58].

Several key meme-coins issued via Clanker token produced pronounced spikes in wallet binding: (1) the launch of CLANKER (wallet binding increases 11.54%, 2024-11-09, the first eponymous meme-coin by @clanker followed by Ethereum founder Vitalik Buterin's [28] purchase of ANON token (a token representing anonymous internet culture) on November 19, 2024, jointly driving

<sup>&</sup>lt;sup>25</sup>Airdrop refers to the free distribution of cryptocurrency tokens or coins to eligible wallet addresses, often as a marketing strategy to increase protocol adoption and reward early users.

<sup>&</sup>lt;sup>26</sup>When Clanker launches a new meme-coin, it automatically creates a trading pair between ETH and the meme-coin as a liquidity pool, enabling users to trade the token.

a 28.9% wallet binding surge by December 19, 2024 [104]; (2) the launch of \$DRB by @grok (X's AI Agent account) on March 17, 2025, resulting in a 37.34%, which exemplifies AI-to-AI token interaction and triggered widespread discussion of the Farcaster ecosystem across X (Twitter) [88]. These events, amplified by social momentum and celebrity engagement, underscore the intricate relationship between platform growth and token-based speculation.

The above demonstrates that platform innovation and token-driven incentives are critical for driving deeper user engagement.

#### **B** Skewed Token Distribution

In Section 4.2, we aim to identify tokens that demonstrate sustained and widespread activity within the Farcaster ecosystem, rather than those irrelevant to Farcaster or exhibiting merely temporary bursts of activity (*e.g.* due to spam, speculation, or airdrops).

Among all FID-linked wallets in our dataset, we observe 440,274 distinct tokens. This substantial diversity stems from the interoperability between users' external wallets and the broader Ethereum ecosystem, resulting in the presence of many tokens that may have little to no direct connection to the Farcaster ecosystem. Therefore, in this section, we first examine the overall distribution of these tokens, which guides us in developing a systematic approach to identify prevalent tokens that maintain consistent usage patterns and meaningful relevance to social interactions on Farcaster.

#### **B.1** Trading Metric.

In traditional markets, trading activity is measured via trading volume (e.g. 1,000 Tesla shares) or dollar volume (total value at  $\approx$  \$300/share). Cryptocurrency markets similarly use token-denominated and fiat-equivalent volumes. However, both metrics pose challenges for cross-token analysis in blockchain systems. Token volumes cannot be meaningfully aggregated due to vast quantity differences across cryptocurrencies (e.g. 0.00001 BTC vs 10,000 DOGE, both  $\approx$  \$1 equivalent). Fiat-equivalent aggregation is complicated by high token price volatility and limited price data availability for illiquid tokens. We therefore focus on transaction frequency—the count of distinct transaction events per token.

We inspect three metrics: trading frequency, number of holders, and number of users as token sender (*i.e.* the wallet actively sending out the token is linked to an FID)—for all 440,274 tokens that have appeared in Farcaster users' Ethereum wallets. Figure 7 illustrates the ECDF for these three metrics. The overall distribution exhibits high skewness, with a small subset of tokens dominating these indicators.

#### **B.2** Transaction Frequency and Holder Count.

Based on a total token population of 440,274, our analysis reveals that 99% of tokens have no more than 390 holders ( $mean \approx 63.46$ , median = 1) and 1,065 transactions ( $mean \approx 207.86$ , median = 2)—notably low figures when compared to the potential market of 489,824 FID-linked wallets on Farcaster. The concentration of activity in a small number of tokens suggests that most tokens in users' wallets are not actively used for social interactions. Further, the low median values (1 holder, 2 transactions) suggest that many tokens never gained meaningful adoption.

#### B.3 FID as Token Sender.

Furthermore, this skewness is particularly pronounced for FID senders. Recall, an FID sender means the wallet is linked to an FID. Within the total population of 440,274 tokens, only 177,733 tokens (40.37%) exhibit at least one FID-sender interaction. This indicates that approximately 60% of all tokens have never been actively sent by any Farcaster user, suggesting they are only passively received by users and have never been employed in any use cases such as tipping other users




Fig. 7. ECDF of token metrics on Farcaster: transaction frequency, holder count (full population: 440,274 tokens), and FID-sender count (subset: 177,733 tokens).

and interacting with exchanges or smart contracts. Furthermore, among these 177,733 tokens with at least one FID-sender activity, 99% have no more than 191 unique senders ( $mean \approx 24.57$ , median = 1), revealing a highly concentrated distribution of active engagement.

This observation reveals critical insights into token circulation patterns: Due to the transparent and non-rejectable nature of blockchain transactions, users frequently become passive token recipients through promotional airdrops or potential phishing attempts. Consequently, active token sending behavior, particularly from FID-linked wallets, serves as a more reliable indicator of genuine user engagement and token utility. This becomes especially significant when considering the scale: among 489,824 FID-linked wallets, 99% of inter-user traded tokens engage fewer than 191 active senders (merely 0.039% of total FID-linked wallets), highlighting a striking disparity between trivial and influential token circulation in the ecosystem. This distribution pattern indicates that despite the presence of over 440K tokens in Farcaster users' Ethereum wallets, only a small subset demonstrates meaningful interaction initiated by Farcaster users, as evidenced by the highly skewed FID-sender distribution. This observation motivates us to introduce a systematic approach to automatically identify and analyze prevalent tokens in the next section.

#### C Prevalent Token Detection.

In the exploratory analysis of token metrics, we observe that commonly used indicators such as rankings of trading volumes, transaction frequencies, and holder counts may provide insufficient or potentially misleading signals regarding a token's genuine influence. Indeed, malicious actors could artificially manipulate these metrics through strategic token distributions targeting user wallet addresses, thereby fabricating an illusion of market popularity [105]. This phenomenon poses significant challenges for token valuation that rely on these surface-level metrics. Therefore, we propose a lightweight systematic method to differentiate between genuinely prevalent tokens and those potentially manipulated with artificially inflated indicators.

We formalize our detection approach as a four-step algorithm, structured in the following subsections.

#### C.1 Step 1: Inter-FID Transactions (5,878 tokens remained after screening).

We begin with user-to-user (inter-FID) transactions—by extracting all transactions where both sender and receiver wallets are explicitly linked to registered FID accounts on Farcaster. It is important to note that Farcaster allows users to link their existing Ethereum-compatible wallets to their accounts. Consequently, these wallets contain transaction records that extend beyond the Farcaster ecosystem, with inter-user transactions representing only a subset of total wallet activity. This initial filtering identifies tokens with at least one transaction between Farcaster users, excluding tokens solely traded with external smart contracts, exchanges, or wallets lacking Farcaster social context.

After restricting transactions where both recipient and sender wallets are linked to an FIDs, we identified a subset of 5,878 tokens, constituting 1.34% of the total 440,274 tokens, accounting for 3,354,378 transfers, representing 3.63% of the complete dataset containing 92,287,905 token transactions. This reveals a power-law distribution of token ecosystems, where a small fraction of tokens (1.34%) achieve meaningful social circulation, while the vast majority of tokens (98.66%) lack user-to-user activity and primarily operate in non-social contexts such as smart contract interactions and exchange swaps.

#### C.2 Step 2: Shannon Entropy (104 tokens remained after screening).

To address the ephemeral nature of most tokens, which typically show activity only in their initial weeks, we employ *Shannon entropy* to analyze weekly transaction distributions [69, 117].

*Shannon Entropy.* We compute Shannon entropy over the weekly transaction frequency distribution for each token. Specifically:

• **Input Data**: For each token, we construct a probability vector  $\mathbf{p} = (p_1, \dots, p_T)$  where:

$$p_t = \frac{n_t}{\sum_{i=1}^T n_i} \tag{2}$$

with  $n_t$  being the transaction count in week t, and T the token's lifespan in weeks.

• Entropy Calculation:

$$H(\mathbf{p}) = -\sum_{t=1}^{T} p_t \log_2 p_t \quad \text{(bits)}$$
 (3)

• Normalization:

$$H_{\text{norm}} = \frac{H(p)}{H_{\text{max}}}, \text{ where } H_{\text{max}} = \log_2 T$$
 (4)

Key properties:

- $H_{\text{norm}} \in [0, 1]$  with:
  - 1: Perfectly uniform distribution
  - 0: Single-week concentration
- Threshold  $H_{\text{norm}} \ge 0.9$  selects tokens with:

$$\frac{H(p)}{\log_2 T} \ge 0.9 \tag{5}$$

We compute normalized Shannon entropy over each token's weekly transaction frequency, retaining tokens with  $H_{\text{norm}} = H(\mathbf{p})/\log_2 T \ge 0.9$ . This threshold identifies tokens with temporal

uniformity (evenly distributed weekly transactions) and sustained vitality (activity sustained over a lifespan T long enough for stable entropy measurement), thus filtering for those that maintain consistent engagement beyond an initial launch phase. This yields 793 tokens exhibiting both temporal uniformity and sustained vitality.

Nevertheless, a substantial subset of 559 tokens (70.5% of the total 793 tokens), each with a lifespan not exceeding 5 weeks, exhibits high normalized entropy values (mean  $\approx$  0.967), despite their consistently low raw entropy (all values < 1.6, aligning with the mean raw entropy across the entire 793-token sample). To address this short-period bias, we add a minimum 26-week (half-year) lifespan requirement, yielding 104 tokens.

The choice of a 26-week minimum existence requirement is not arbitrary. We observe that among 969 tokens with raw entropy values above 3 (99.78th percentile, N=440,274), only 19 (1.96%, n=976) have existed for less than 26 weeks (compared to a mean existence of just  $\approx$  1.67 weeks across all 440K tokens). This threshold thus ensures both an adequate sample size ( $\approx$  1,000 tokens) and effectively excludes sampling insufficiency issues with newer tokens, giving appropriate weight to tokens with longer trading histories. In real-world applications, these calculations can be performed weekly to dynamically include tokens previously excluded by the 26-week requirement.

#### C.3 Step 3: FIDs as Token Senders (9 tokens remained after screening).

The 104 tokens identified in the previous steps exhibit a right-skewed distribution in their number of unique FID-linked senders. Here, an FID-linked sender is defined as a wallet address explicitly associated with a registered FID, ranging from 1 to 29,464 (mean  $\approx$  897). The number of FID-linked wallet senders serves as a crucial metric for evaluating token prevalence, as it more substantially reflects genuine social interactions rather than passive reception. This metric's significance stems from its ability to distinguish between tokens with meaningful user engagement and those with merely superficial circulation. Consequently, we employ the number of FID-linked wallet senders within each token's inter-FID-transactions as our final filtering criterion. Using the 99th percentile threshold (254 FID-senders), we finally identify nine tokens—four reward tokens and five blockchain network tokens—detailed in Table 1.

This process yields nine tokens. Notably, DEGEN and MOXIE correspond to significant user growth events (recall that we have discussed the top 10 events in Figure 6). TN100X and HIGHER, two other reward tokens launched in February and March 2024 respectively, did not trigger top 10 wallet binding surges. However, these tokens were identified through our screening process for long-term token popularity. Conversely, \$FARTHER, which appeared in the Top 10 events, was not selected by our screening criteria. Through analysis of community content and documentation [45], we discover that the \$FARTHER reward program was terminated by developers in August 2024 due to excessive user farming<sup>27</sup> and sell-offs. This finding suggests that while reward tokens may generate temporary enthusiasm and transaction bursts, only a select few achieve sustained user adoption and utilization. The remaining five are native tokens and stablecoins commonly used on blockchain main-net (Ethereum) and scaling layers (L2s and L3s): USDC, USDT, USDbC, WETH, and L3. These tokens were also identified by our screening methodology due to their broader market acceptance and high utilization rates.

 $<sup>^{27}</sup>$ User farming refers to the behavior where small groups of users or bots engage in circular reward-giving among themselves to exploit the project's token reserves

<sup>&</sup>lt;sup>28</sup>Sell-offs occur when users, upon receiving token rewards, immediately exchange them in the open market for more established cryptocurrencies (e.g. USDC, ETH, or BTC) instead of utilizing them within the ecosystem for services or tipping.

#### C.4 Step 4: Clustering Coefficient (4 tokens remained after screening).

We next calculate the average clustering coefficient for the transaction graph of each token. We do this for each transaction graph's largest connected component. *Average Clustering Coefficient (ACC)* effectively captures community network patterns and this metric has also been used in cryptocurrency analysis for identifying artificial transaction patterns [67, 75, 103]. We then select all tokens that have a clustering coefficient between 0.3 and 0.6. We choose this because previous studies present that real-world community has a clustering coefficient around 0.45 [108]. This leaves four remaining tokens that are considered prevalent.

Summary. Our methodology, based on social relationships (inter-FID transaction and FID sender), network spatial distribution (clustering), and transaction temporal distribution (entropy), successfully identifies 4 prevalent tokens within Farcaster, with strong social attributes (see Table 1). This approach is useful for identifying genuinely influential and commonly used tokens especially for permissionless ecosystems like Farcaster, where the ability to bind external wallets and support all Ethereum-compatible tokens necessitates robust filtering mechanisms to distinguish viable tokens from low-signal noise. This selection also aligns perfectly with our ground truth observations of the influential tokens that drive user growth in Section 4.1, validating our approach. Moreover, given the extreme sparsity of positive samples (*i.e.* few viable tokens among all tokens present in user wallets) and the fact that each filtering stage employs thresholds tailored to specific scenarios, our three-dimensional framework demonstrates superior operational practicality, performance, and contextual explainability compared to machine learning-based approaches. These advantages make it particularly suitable for real-world applications like platform built-in token linking algorithms [43] or index website ranking systems [20, 42].

#### **D** Token Incentive Distributions

#### D.1 Methodology for Tracing Reward Sources

To identify Inter-FID Tipping, we implement a triple-filtering process: (i) eliminating self-transfers within the same FID, (ii) transfers between different wallets bound to the same FID, and (iii) transfers between FIDs that have historically shared wallet bindings—indicating affiliated entities.

To identify third-party algorithmic rewarding, we filter transactions originating from token issuers' official wallets/dedicated smart contracts<sup>29</sup> to FID-linked wallets. Similarly, the official USDC algorithmic rewarding involves transfers from Farcaster's official wallet cluster/dedicated smart contracts to FID-linked wallets. The source wallet addresses and smart contracts for distributing algorithmic token rewards, along with the five identified prevalent token contract addresses, are listed in Table 9.

#### **D.2** Intersection of Reward Participants


Figure 8 illustrates the upset plots showing the intersections of reward participants across different incentive mechanisms and tokens. Figure 8a shows the overlap of receivers among three Mechanisms—Inter-FID Tipping, Third-party Algorithmic Rewards, and Farcaster's Official Algorithmic Rewards; Figure 8b depicts the overlap of senders across the five tokens within Inter-FID Tipping; Figure 8c displays the overlap of receivers across the five tokens within Inter-FID Tipping. These plots highlight the extent to which participants engage with multiple mechanisms or tokens.

The upset plots help analyze user coverage across three reward mechanisms involving five tokens on Farcaster. Inter-FID tipping reaches 53,766 unique receivers through five tokens (6.01% of


<sup>&</sup>lt;sup>29</sup>The identification of these addresses is based on a combination of the following methods: examining token project whitepapers, leveraging third-party data dashboards, and analyzing the transaction history of a wallet used for field experimentation in relevant token reward projects.

Table 9. Tokens's name and contract addresses, and algorithmic token reward distribution addresses

Token's Namer	Token's Contract Address	Algorithmic Token Reward Source Address
DEGEN	0x4ed4e862860bed51a9570b96d89af5e1b0efefed	0x9f07f8a82cb1af1466252e505b7b7ddee103bc91 0x7134ddcd6bc3d3f29643ec8882cd4d7d0b38379d x88d42b6dbc10d2494a0c6c189cefc7573a6dce62 0xa2a5c549a454a1631ff226e0cf8dc4af03a61a75 0x0bf676823f958d0ae6af9860880fc7a327a0c582 0xf4c9fc64ed902b08c3397c48dca0342bdfd49a32 0x81ac315d7baf7fae2e8278dc33ec91b752560166 0xdfeddda2bc2e75524f470991f46404ee90ae2eea
USDC	0x833589fcd6edb6e08f4c7c32d4f71b54bda02913	0x2d93c2f74b2c4697f9ea85d0450148aa45d4d5a2 0xd15fE25eD0Dba12fE05e7029C88b10C25e8880E3 0xf6ea479f30a71cc8cb28dc28f9a94246e1edc492 0xf25b82277d9dfdbb341fd1bd9bf197b60ae09e37 0xed64b7fd10f6f06def48af99e719c3cd8e9c78c6
MOXIE	0x8c9037d1ef5c6d1f6816278c7aaf5491d24cd527	0xfe569dc5a297de1dfbcce19273fc8267f2c52c45
HIGHER	0x0578d8a44db98b23bf096a382e016e29a5ce0ffe	None
TN100X	0x5b5dee44552546ecea05edea01dcd7be7aa6144a	None



(a) Receivers across Inter-FID Tipping, B and C.



(b) Senders across five tokens in Inter-FID Tipping. (c) Receivers across five tokens in Inter-FID Tipping.

Fig. 8. Upset plots.

Farcaster's 894,678 valid registered users). The upset plots in Figure 8 (Appendix D) reveal minimal multi-token overlap—the largest intersection comprises 8,662 users (16.11% of tipping receivers) who received both DEGEN and HIGHER tokens, while only 385 users (0.716%) received all five tokens. For algorithmic rewards, third-party algorithmic rewards (via DEGEN and MOXIE) reach 57,562 receivers (6.43% of FIDs), while the official USDC rewards engage 28,181 users (3.15%). Combined, both algorithmic rewards reach 80,762 users (9.03%), indicating moderate overlap (20.64%) between

Metric Follow-only Follow + Tip Nodes 883,712 883,906 Edges 159,539,953 159,595,800 # Communities [332, 359] [334, 372] Largest SCC [173, 174] 172 Total SCCs 173 [174, 175] Avg. Deg. 361.06 361.11 Max Deg. 564.120 564,989 Modularity [0.54389, 0.54580] [0.53196, 0.54625]

Table 10. Network statistics comparison.

Table 11. Community detection (follow-only network).

Metric	Louvain	Infomap					
		Level 0	Level 1	Level 2	Level 3		
Communities Largest Comm. Median Comm.	≈ 462,230			330,165 23,127 1			
muu commi	bution: ≈ 48% ≈ 52%	99.37% 0.63%	24.92% 75.08%	3.84% 96.16%	0.26% 99.74%		

third-party and official incentive mechanisms. Across all three mechanisms, the total unique recipient count is 103,666 (11.59% of FIDs)—notably lower than the sum of individual mechanism counts (178,702) but higher than any single mechanism's reach. The primary cross-mechanism overlap consists of 11,637 users (11.23% of all recipients) receiving both tipping and DEGEN algorithmic rewards, while only 2,511 users (2.42%) received rewards from all three mechanisms.

These patterns demonstrate strong mechanism specialization, with users typically engaging with specific token incentives rather than participating broadly across the reward ecosystem. Notably, algorithmic rewards collectively reach more unique users (80,762) than tipping (53,766), suggesting that systematic incentive distribution achieves wider user coverage than peer-to-peer tipping.

#### E Socioeconomic Risks of Token Incentives.

Figure 9 visualizes the participation for each token-mechanism pair in the form of stacked area charts with overlaid lines. The shaded areas represent the weekly number of unique participants—either senders (diagonal pattern), receivers (solid fill), or both (horizontal pattern). The solid line indicates the number of new receivers appearing each week, while the dashed line indicates the number of new senders appearing each week. This visualization captures both the cumulative engagement and the dynamics of new participant inflow over time.




Fig. 9. Stacked area charts with overlaid lines visualize participation for each token-mechanism pair.




Fig. 10. Balance diagnostics for propensity score matching (PSM).

Table 12. Detailed regression summary of continuous treatment intensity with social activities.

		Temporal A	Alignment:	First Reward	l Reception	Date as T=0		
		Int	er-FID Tipp	ing		Algo	orithmic Rev	ward
Action	DEGEN	TN100X	HIGHER	MOXIE	USDC	DEGEN	MOXIE	USDC
post	1.5703***	5.0001	1.9701***	-0.0517	-0.0805	12.7407***	15.2254***	8.2309***
	(0.60)[0.16]	(3.08)[0.32]	(0.65) [0.28]	(0.92)[0.04]	(0.43) [0.01]	(2.93) [0.22]	(0.74)[0.12]	(1.20)[0.18]
reply_out	-4.4867	8.6610	-3.8922	10.6198	-0.3703	9.4023	59.2029***	-19.1090***
	(3.50) [0.56]	(30.62) [0.64]	(3.07) [0.54]	(2.57)[0.42]	(2.75)[0.08]	(17.01) [0.69]	(4.39) [0.27]	(5.25) [0.52]
reply_in	1.7207	-35.1615	-2.4438	10.4600	-10.2936**	-4.8940	57.1112***	5.1841
	(3.31) [0.78]	(29.81) [0.87]	(2.75) [0.83]	(3.87) [0.59]	(3.57) [0.24]	(12.67) [0.80]	(3.24)[0.56]	(5.53) [0.74]
like_out	-1.4999	6.4651	1.4243	-4.1602	7.5337**	-9.7807	22.8327***	32.0449***
	(4.81) [0.50]	(32.23) [0.50]	(4.24) [0.66]	(3.73) [0.26]	(3.54) [0.09]	(33.93) [0.57]	(3.02)[0.42]	(6.20) [0.60]
like_in	-1.2148	25.2958	17.9290***	-27.5476***	0.5102	-77.0576***	7.7791**	73.9531***
	(5.34) [0.80]	(23.53) [0.96]	(3.89) [0.91]	(6.55) [0.74]	(3.98) [0.09]	(24.70) [0.75]	(3.50) [0.70]	(6.19) [0.90]
repost_out	0.4519	7.1819	2.4716*	-0.9295	-0.0732	-5.6276	-0.1039	-9.1389***
	(1.52) [0.44]	(7.93) [0.43]	(1.36) [0.56]	(1.28) [0.23]	(1.08) [0.05]	(10.63) [0.45]	(0.70) [0.34]	(2.27) [0.46]
repost_in	1.7003	-5.8862	-5.2017***	5.6276**	8.7852*	5.0918	-6.2899***	-30.2670***
	(2.00)[0.76]	(8.61) [0.94]	(1.47) [0.80]	(2.05)[0.25]	(2.18)[0.06]	(8.19) [0.64]	(1.16) [0.57]	(2.19) [0.81]
follow_out	0.2749	-0.9579	0.0211	14.5753***	0.6082	0.9927	20.3066***	-54.1936***
	(3.29)[0.14]	(12.09) [0.17]	(1.77) [0.10]	(2.98)[0.45]	(1.46)[0.04]	(15.90) [0.16]	(2.40)[0.45]	(3.17) [0.15]
follow_in	-12.4111	-57.8221	-19.9274***	9.9655***	15.6006	196.8301***	6.3755**	242.8866***
	(8.22) [0.21]	(45.63) [0.56]	(6.43) [0.46]	(2.68) [0.40]	(5.43) [0.06]	(34.16) [0.10]	(3.21) [0.47]	(5.15) [0.56]
popsize	40836	3459	15849	5252	15872	47748	21505	28181
sample_size	38532	3438	13854	5041	13879	47748	21482	27484

Symbols: \*\*\* p < 0.001, \*\* p < 0.05, \* p < 0.01. Coefficients deemed statistically significant are presented in bold; Coefficients associated with statistically significant causal effects are highlighted in both bold and color.

#### F Effectiveness of Token Incentives on Social Activities

Figure 10 demonstrates the balance tables for propensity score matching (PSM) analysis across token incentive mechanisms (see Section 6.1). While most covariates achieve balance (SMD < 0.1), notable imbalances are observed in outbound follows. SMD denotes standardized mean difference. Token-related covariates on the Y-axis are defined as follows: "A in/out": inbound/outbound Inter-FID Tipping; "B in": inbound Third-party Algorithmic Rewards (3rd Algo. Reward); "C in": inbound Official Algorithmic Rewards (Off. Algo. Reward).

Table 12 illustrates a more detailed result table for the regression analysis in Section 6.2, including standard errors (SE) and  $R^2$ . Standard errors (SE) are reported within parentheses (), and the coefficient of determination  $R^2$  is enclosed in square brackets [].

Received July 2025; revised September 2025; accepted October 2025