Decentralization in DeFi Lending: A Network Perspective

A Multi-Layered Analysis of Governance-Active Users

Sam Heslenfeld

Master Thesis

TU Delft

Decentralization in DeFi Lending: A Network Perspective

A Multi-Layered Analysis of Governance-Active Users

by

Sam Heslenfeld

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Friday July 11, 2025 at 09:30 AM.

Student number:	4697170	
Project duration:	September 23, 2024 – July 11, 2025	
Thesis committee:	Dr. H. Wang,	TU Delft, supervisor
	Dr. X. Zhang,	TU Delft

Cover: designed by Freepik

An electronic version of this thesis is available at `https://repository.tudelft.nl/`.

Preface

This thesis marks the final step in completing my Master's in Computer Science at TU Delft. Over the past year, I have researched lending platforms in Decentralized Finance. Starting with no knowledge of the topic or blockchain technology in general, just a strong desire to work on a network science project that matched my interests, and finishing with an entire thesis on this subject.

I would like to thank my supervisor, Dr. H. Wang, for her knowledge and guidance over the past year. It took time to find an interesting research direction and to get things moving, but her suggestions and original ideas helped shape this project. During her course Modeling and Data Analysis in Complex Networks, I first discovered network science, which immediately captured my interest and which ultimately led to this thesis.

I am also very grateful to everyone who supported me along the way, this journey was not always easy, but you all helped me stay motivated. Special thanks to my girlfriend for always providing love and support when I needed this, to my parents and my brother for constantly supporting me and encouraging me to pursue my ambitions ever since I was young, and to everyone I studied with at EWI during this past year. You all made working on my thesis much more enjoyable.

Sam Heslenfeld
Delft, July 2025

Abstract

Decentralized Finance (DeFi) lending platforms claim to use decentralized, community-driven governance. In practice, however, governance power remains concentrated among a limited number of users. This thesis investigates the behavior of users actively participating in governance in two major DeFi lending platforms, Aave and Compound, using a network science approach. Data covering governance actions and both governance and yield token transfers is combined in a multi-layered network model. By combining these layers, it is possible to investigate the characteristics of governance-active users. Specifically, how their governance activity relates to their behavior in the token transfer networks. The analysis studies correlations between token transfer patterns and governance participation, detects communities in governance token transfer networks, and uses a Susceptible-Infected (SI) process to estimate each user's structural influence in a spreading process to compare this to governance activity. The results show that governance token transfer volume strongly correlates with voting power, while other network features do not. Most governance-active users cluster together in a few communities in the governance token transfer network. Users with a greater simulated spreading potential only tend to partially hold higher actual voting power. Together, these findings suggest that DeFi governance is strongly shaped by token wealth and tight user clustering, highlighting that open design alone does not guarantee broad and equal participation.

Contents

Preface i

Abstract iii

1 Introduction 1
1.1 Research Questions . . . 2
1.2 Thesis Structure . . . 2

2 Background 3
2.1 Ethereum . . . 3
2.2 DeFi Mechanisms . . . 3
2.3 DeFi Incentives . . . 4

3 Related work 5
3.1 Governance in DeFi Platforms . . . 5
3.2 Decentralization Metrics and Dynamics in DeFi . . . 5
3.3 Temporal Network Analysis in Blockchain Networks . . . 6

4 Methodology 7
4.1 Data Collection . . . 7
4.2 Network Construction . . . 9
4.3 User Activity Features . . . 9
4.3.1 Features from the Temporal Network (G_T) . . . 9
4.3.2 Features from the Aggregated Network (G_C) . . . 10
4.3.3 Features from the Aggregated Network (G_V) . . . 10
4.3.4 Governance Activity Features . . . 11
4.4 Correlation Analysis . . . 11
4.5 Community Detection . . . 11
4.6 Simulation-Based Influence Scoring (SI Model) . . . 12

5 Results 15
5.1 Feature Correlations between Token Transfers and Governance Activity . . . 15
5.2 Community Structure and Governance Activity Distribution . . . 18
5.2.1 Aave . . . 19
5.2.2 Compound . . . 20
5.3 Simulation-Based Influence and Governance Power . . . 21
5.4 Cross-Platform Comparison . . . 23

6 Discussion 25

7 Conclusion and Future Work 27
7.1 Future Work . . . 27

A Supplementary Figures 29
A.1 Correlation Results . . . 29
A.2 Community Structure . . . 30
A.3 Influence Dynamics . . . 31

B AI Disclosure Statement 33

Bibliography 35

1 Introduction

Blockchain technology makes it possible to build applications that do not rely on a central authority, called decentralized applications (dApps). Instead of a bank or administrator keeping records, a blockchain stores transactions across a network of computers in a secure, transparent, and tamper-resistant way. Ethereum is the most widely used blockchain for this, because it supports smart contracts, which are small programs that automatically perform actions such as sending tokens, issuing loans, or recording votes when certain conditions are met. These smart contracts enable Decentralized Finance (DeFi), which offers financial services such as lending, borrowing, and managing assets without the need for banks or other intermediaries. Unlike traditional finance, DeFi is accessible to anyone with an internet connection and aims to be governed by its community rather than by a few centralized actors.

Within DeFi lending platforms, users supply liquidity by depositing their crypto assets into a lending pool. In return, they automatically receive platform-specific interest-bearing tokens, referred to here as yield tokens, which represent their share in the pool plus any earned interest. These tokens can be transferred, used as collateral themselves, or redeemed for the original assets at any time. Borrowers, in turn, must lock up their own crypto assets as collateral, usually worth more than the amount they borrow, to secure loans and protect lenders. In addition, they pay interest on the borrowed amount. This trustless, permissionless model enables anyone with crypto to lend or borrow without relying on a bank or intermediary.

DeFi lending platforms manage billions of dollars in assets through smart contracts and continue to grow, but the rules behind these contracts evolve through community governance. Proposals (suggested changes), voting (decisions on these changes), and delegation of voting power determine how key parameters or rules change, risks are managed, and new features are introduced. Governance takes place through its community of token holders. The platforms are governed through platform-specific governance tokens, which carry financial and governance value and can be freely traded. The voting power of users is generally proportional to their governance token holdings, which are issued in limited supply to maintain scarcity and ensure that decision-making remains meaningful.

Many prior studies have looked at how governance tokens, which determine voting power, are distributed among users. Despite the aim of decentralization, these studies consistently find that a limited number of users actively participates in DeFi governance, contrary to the platforms' stated goals. This raises concerns that a few users could steer important decisions, such as setting interest rates or collateral rules, for their own benefit. Other research analyzes token transfer networks to compare structural decentralization, but this mostly focuses on general network properties, rather than user-specific behavior. As a result, they do not examine who the governance-active users are, how they behave in detail, or how their voting actions relate to their token transfer behavior.

Obtaining a clearer view of how governance-active users interact within token transfer networks is important, since they are the ones directly influencing platform decisions. By measuring how governance actions align with token flows and network position, this research shows whether formal voting rights match practical behavior, or whether hidden structures allow a few users to concentrate control.

To address this gap, this thesis goes beyond simply counting participation or describing network structures in isolation. It focuses on the behavior of governance-active users, combining their actions;

proposing, voting, and delegating, with how they interact in governance and yield token transfer networks, and measuring how strongly these dimensions relate. A novel, multi-layered network science approach is applied to Aave and Compound, two leading DeFi lending platforms. Unlike prior work that studies governance or token flows separately, this study integrates three key dimensions into a unified analysis:

- **Governance Interaction Layer:** governance proposals, voting, and delegations.
- **Governance Token Transfer Layer**
- **Yield Token Transfer Layer**

This multi-layered view tests whether formal governance power matches actual network behavior. It also adds a dynamic perspective by detecting communities based on transfer activity and simulating each user's potential influence using a Susceptible-Infected (SI) spreading model on a temporal token transfer network. This helps reveal whether the same users who participate in governance also hold structural positions that amplify their influence. Together, these layers provide a realistic view on user interactions, network positions, and governance activity and help asses whether DeFi governance lives up to its open, token-based promise. The openly accessible DeFi data provides a unique opportunity to observe these interactions at scale.

1.1. Research Questions

This research uniquely combines multiple aspects of DeFi platforms, such as data on governance interactions (proposals, votes, delegations), governance token transfers, and yield token transfers. The study addresses the following research questions:

- To what extent does the governance activity of users (voting and delegating) relate to their properties in the governance and yield token transfer layers?
- Do nodes that cluster together in communities, based on transfer activity in the governance token transfer layer, tend to contribute similarly to governance activity?
- To what extent does the influence of users in a spreading process, as measured by the SI model, relate to their governance activity?
- How do Aave and Compound differ with respect to the questions above?

The core analysis uses three types of directed networks for each of the two token transfer layers (governance and yield tokens) on both platforms. In these networks, each node is an Ethereum address, which basically represents a unique account on the blockchain. This can belong to an individual user or a smart contract. The meaning of the edges differs per network. The first network, G_T, is a temporal network. Each single interaction (i.e., token transfer) is represented by an edge from address A to B at time t, with an associated weight indicating the amount of tokens transferred. From this, a weighted aggregated network can be constructed, in which a directed edge from A to B exists if A has transferred tokens to B at least once. Depending on the meaning of the edges, two static networks are created from this. In G_C, the weight of an edge from A to B denotes the number of separate transfers from address A to B. In G_V, the weight of an edge from A to B represents the total volume of tokens transferred from address A to B. All data used to build networks is collected over a period of a year.

All networks include all addresses and token transfer information, but the analysis is only performed on governance-active users; those who have proposed, voted, or delegated. This approach enables the application of network science techniques to study the behavior these users, by relating their governance activity to their properties in token transfer networks.

1.2. Thesis Structure

The remainder of this thesis is structured as follows. Chapter 2 introduces key concepts. Chapter 3 reviews prior work on DeFi governance, decentralization, and temporal network analysis. Chapter 4 details the research pipeline, including data collection, preprocessing, network construction, and analysis methods. Chapter 5 presents the findings, including structural analysis, community detection, and influence modeling. Chapter 6 interprets these results, and Chapter 7 concludes with reflections and directions for future work.

2
Background

This thesis uses network science to analyze DeFi lending platform data, with the aim of assessing whether their promised governance decentralization is truly achieved. To support this analysis, this chapter outlines essential concepts and technical foundations.

2.1. Ethereum

Ethereum is the most widely used blockchain for decentralized applications (dApps) [7, 33] because it supports smart contracts, which are small programs that run exactly as coded. These smart contracts make it possible to build DeFi platforms, where code replaces intermediaries such as banks. This is important because Aave and Compound, the platforms studied in this thesis, both run on Ethereum.

2.2. DeFi Mechanisms

DeFi lending platforms remove intermediaries by using smart contracts to automate financial services like lending and borrowing [35]. Key features include permissionless access, non-custodial control, transparency, and composability [29, 34]. By contrast, Traditional Finance (TradFi) refers to conventional banks and institutions that centrally manage funds, while Centralized Finance (CeFi) includes crypto-native platforms that still retain centralized custody and control of user assets. Unlike CeFi, DeFi lets users hold direct control without trusting an intermediary. In addition, DeFi lacks credit scoring and regulation, often requiring overcollateralization to manage risk [28]. This mechanism offers lenders protection against volatility, which is a common phenomenon in cryptocurrencies, ensuring the repayment of loans in cases of collateral value fluctuations. Benefits of DeFi include financial inclusion, censorship resistance, and low-cost global transactions [25]. Risks include smart contract vulnerabilities, regulatory uncertainty, and potential misuse for illicit activities.

DeFi platforms, such as Aave and Compound, facilitate the lending and borrowing of assets through smart contracts, eliminating the need for intermediaries. Typically, loans are overcollateralized and interest rates are algorithmically determined by supply and demand.

Yield tokens are created when users deposit assets, their increasing value reflects accrued interest. Conversely, borrowers are required to pay interest on the funds they borrow. Yield tokens are interoperable, meaning they can be transferred, used as collateral, or deployed in other DeFi protocols.

Governance tokens serve a dual purpose. On the one hand, they grant holders the right to participate in platform governance by proposing and casting votes on changes to rules, parameters, or upgrades. The proposing and voting power of a user is directly proportional to the number of governance tokens they hold, while the overall supply of these tokens is limited. Conversely, they are considered tradable assets and often carry monetary value. This duality can create a tradeoff between influence and profit, potentially affecting how decentralized and active governance actually is.

To address low participation, numerous platforms support delegation, where governance token holders assign their voting rights to more active or informed users, who can then vote on their behalf. The delegated voting power is proportional to the number of governance tokens possessed by the delegator. While this boosts engagement, it can concentrate power among a few delegates.

2.3. DeFi Incentives

DeFi attracts users with unique benefits compared to CeFi and TradFi. The most common motivations include the following, of which several are stated by [8]:

- Permissionless access: open to anyone with a wallet and internet access, without credit checks or Know Your Customer (KYC) regulations.
- Flash loans: arbitrage or strategy execution with no collateral.
- Transparency: smart contracts and transactions are fully auditable.
- Composability: assets can be used across platforms.
- Tax benefits: borrow without selling, deferring capital gains.
- Flexible terms: no fixed loan durations, collateralization governs lifespan.

DeFi lending is also used for traditional purposes, such as saving, liquidity access, and portfolio diversification. Its appeal lies in high yields, open access, and rapid innovation.

3

Related work

This chapter reviews prior research related to DeFi governance mechanisms, decentralization, and temporal network modeling techniques. It highlights gaps that this thesis aims to address.

3.1. Governance in DeFi Platforms

DeFi governance is often managed through Decentralized Autonomous Organizations (DAOs), where governance token holders can propose and vote on platform changes [13]. The process typically involves:

1. Proposal submission: a user submits a proposal to change the platform.
2. Community discussion: the community provides feedback and iterates on the proposal.
3. Token-weighted voting: token holders vote on the proposal, with voting power proportional to their token holdings.
4. Implementation via smart contracts: if approved, the change is automatically implemented via smart contracts.

Many existing studies focus on the distribution of governance tokens or voting power in DeFi lending platforms. They consistently find that governance often remains centralized in practice, with only a limited number of users actively proposing or voting, and a handful of large holders controlling most of the power. [12] report low voter turnout and a concentration of power among major voters. Other studies [4, 6, 17, 23] confirm token concentration among a few users and underused voting rights. [9] additionally show that many users who hold governance tokens rarely exercise their voting rights, highlighting a gap between formal rights and real participation. A high proposal acceptance rate is also common, often due to coordination and discussions that occur before voting. These discussions help quantify community sentiment and enable the selection of on-chain proposals that are more likely to secure majority support. While [24] argue that some concerns may be overstated, they still observe high concentration of governance power.

These studies show clear participation gaps, but do not examine how governance-active users behave in other parts of the platforms, such as token transfers. This thesis addresses this gap by combining the layers of governance actions and token transfers in one model, enabling a deeper understanding of the behavior of these users.

3.2. Decentralization Metrics and Dynamics in DeFi

Decentralization in blockchain means that control and decision-making are distributed among many users, but in practice, governance power can still concentrate in the hands of a few. Several studies have explored this by analyzing token transfer networks in DeFi. For example, [27] focused on Compound decentralization using inequality metrics such as the Gini and Nakamoto coefficients to show that voting power remains concentrated, but they did not analyze token flows. They suggest future research to examine what drives voting behavior and how voters interact with each other.

[36] argue that DeFi decentralization is multi-dimensional, identifying five key aspects of DeFi decentralization: consensus, network, wealth, governance, and transactions. They highlight that these aspects interact with each other and recommend future work to study how these different layers are connected in practice. This thesis follows their recommendation by looking at governance actions and token transfers together to assess governance decentralization in DeFi lending platforms.

Other studies [2, 37] use daily, undirected governance token transfer networks across DeFi platforms. They track global network features, such as the number of disconnected components, the relative size of the largest component, the modularity, and standard deviation of node degrees to construct timeseries that reflect overall network structure changes over time. Based on these timeseries, they assess and compare decentralization of these platforms. Their findings suggest that network structure can reveal concentration patterns. However, they do not analyze how these global network patterns relate to individual governance behavior.

This thesis builds on and extends these works by combining governance actions with governance and yield token transfers at the user level, using directed and temporal network models. By focusing on the characteristics of governance-active users and analyzing how their network positions and token activity relate to their governance participation, this study provides a more detailed, dynamic view of how governance-active users shape decision-making in practice and how token flows and network positions relate to their influence.

3.3. Temporal Network Analysis in Blockchain Networks

Most blockchain studies use static network analysis, where all interactions are aggregated into a single snapshot. While useful, this approach can miss important temporal patterns. To address this, recent research has focused on temporal network analysis, where interactions are modeled with timestamps, capturing how behavior, influence, or value flows evolve over time.

[19, 20] use temporal random walks, which only follow transfer paths in chronological order, to better model and predict user behavior on Ethereum. They introduce the Temporal Weighted MultiDiGraph (TWMDG), where each edge represents a single time-stamped transfer, allowing multiple directed, weighted edges between the same users. Together with [21], who apply temporal random walks to blockchain token transfer networks, these time-aware methods reveal richer interaction patterns than static graphs.

Another key development is the use of diffusion models, such as the Susceptible-Infected (SI) model, which simulate how information or influence might spread through a network. While common in epidemiology or social media analysis, SI models are only beginning to be applied in blockchain research. For instance, [30] model blockchain adoption in supply chain finance using a SEIR (Susceptible-Exposed-Infected-Recovered) model, and [1] simulate risk contagion across crypto exchanges using SI dynamics. However, these models have not yet been applied to blockchain token transfer networks, especially not in DeFi governance.

While temporal random walks are useful for predicting future interactions, they are less suited to modeling nodal influence in a network. In contrast, SI models simulate spreading processes to identify influential nodes in a network. Since this thesis aims to evaluate governance decentralization through influence dynamics, the SI model is a more appropriate fit.

This thesis contributes to this area by applying temporal network modeling to governance token transfers and simulating nodal influence using an SI model. This helps us understand the influence of nodes in the network, based on spreading processes, providing a dynamic view of governance decentralization that static models cannot offer.

4 Methodology

The preceding chapter, Chapter 3 showed that earlier studies consistently report centralized governance in DeFi lending, while others use network science techniques to measure platform-level decentralization. However, none of them analyze the platform-wide behavior and characteristics of governance-active users in detail. This is worth investigating, since these users directly influence platform decisions, while also participating in token transfers, jointly shaping how power is distributed. This thesis addresses this gap by combining governance activity and token transfer layers to reveal how governance-active user behavior relates across these dimensions.

This chapter outlines the methodology used to evaluate governance-active user behavior in Aave and Compound. The process involves collecting multi-layered data, constructing different types of token transfers, extracting user-level features, and applying network analysis techniques including correlation analysis, community detection, and nodal influence simulation. We define a feature as a single measurable behavior belonging to a user, such as total tokens transferred, number of transfers, or number of votes cast. Each data layer provides unique user features that can be used in the analysis, for example by computing correlations between them. Table 4.1 summarizes how each data layer contributes to the overall analysis by stating the source data for the layer, the networks constructed with the data and the analysis purpose for which the data is collected. In this study, users are considered governance-active if they created proposals, voted, or delegated their voting power during the study period.

Data Layer	Source Data	Networks	Purpose
Governance interaction	Proposals, votes, delegations	None	Governance features
Governance token transfer	AAVE & COMP transfers	G_T, G_C, G_V	Token transfer features, community detection, SI process
Yield token transfer	aWETH & cWETH transfers	G_T, G_C, G_V	Token transfer features

Table 4.1: Methodology per layer

4.1. Data Collection

The project relied on assembling a dataset from the Aave and Compound platforms, which forms the basis for building the networks and performing all analyses. This data is essential since it captures real user interactions. Compiling the dataset was challenging due to the required domain expertise before being able to identify the required data, locate and collect this data, interpret its meaning, handle the various formats from different sources, filter meaningful information from the large volume of records, and combine this into a usable and coherent piece. The data for this project was collected from the Ethereum blockchain through the following sources:

- Etherscan[1], Infura[2], The Graph[3], and Snapshot[4] for querying token transfer, governance proposal, voting, and delegation data.
- The Aave[5] and Compound[6] websites were consulted manually for supplementary information.

All data was collected over a fixed period of one year, ranging from October 2023 until September 2024. This approach ensured the collection of a sufficiently large, yet manageable dataset. All datasets, queries and Jupyter notebooks used and constructed for the project can be located in the project's repository, which is made publicly available on GitHub[7].

The dataset includes governance token transfers, AAVE for Aave and COMP for Compound, which grant holders voting power to participate in governance decisions. It also includes yield token transfers, specifically aWETH for Aave and cWETH for Compound, belonging to the Wrapped Ether (WETH) crypto asset. Yield tokens represent a user's share in the lending pool and earn interest automatically. DeFi lending platforms accept deposits in many different crypto assets, each having its own platform-specific yield token. The WETH crypto asset was chosen for its wide adoption and ease of comparison. Focusing on one representative yield token per platform keeps the analysis clear and comparable, without losing validity.

The following data was collected:

- Governance token transfers (AAVE, COMP): sender address, receiver address, amount, and timestamp.
- Yield token transfers (aWETH, cWETH): sender address, receiver address, amount, and timestamp.
- Governance proposals: proposer address, voting outcome, vote counts, and timestamp.
- Proposal votes: voter address, vote weight, and vote direction.
- Delegations: delegator address, delegate address, and timestamp.

More high-level information about the dataset used in this study is presented in Table 4.2, which highlights the limited number of users who participate in governance actions, such as proposing, voting, or delegating, compared to the total number of token holders and transfers in the networks. While our dataset includes addresses that delegated their voting rights, it does not contain information on the exact amount of voting power transferred in each delegation. Without the delegated token amounts, it is impossible to precisely calculate the delegated voting weight for each user. Furthermore, we observe that there are only very few users involved in creating proposals. Therefore, we do not consider proposal creation as a specific user governance feature in further analysis.

	Aave	Compound
Governance token transfer nodes	175,923	68,422
Governance token transfers	3,228,269	370,439
Yield token transfer nodes	95,745	11,433
Yield token transfers	1,143,544	25,752
Proposers	10	14
Proposals	260	153
Voters	2,136	450
Votes	9,642	3,201
Delegators	1,563	1,373
Delegates	670	1,182

Table 4.2: Dataset overview

[1] https://etherscan.io/
[2] https://www.infura.io/
[3] https://thegraph.com/
[4] https://snapshot.box/
[5] https://app.aave.com/
[6] https://compound.finance/
[7] https://github.com/SamHes/Master-thesis

4.2. Network Construction

For the analysis of both DeFi lending platforms, three distinct networks were constructed from token transfer data. Each network allows us to extract certain node-level features. Additionally, they facilitate community detection and employing an SI model. It is noteworthy that all three networks share the same set of nodes, denoted by N. Within this set, each node represents a unique Ethereum address, the semantics of the edges differ across the networks. To enhance interpretability, consistency, and quality and to eliminate disconnected nodes, only the largest Weakly Connected Component (WCC) is considered for each network. A WCC is a subgraph where all nodes are connected, regardless of edge directions, eliminating all disconnected components.

The first network is defined as a Temporal Weighted MultiDiGraph (TWMDG), which was first proposed by [19, 20] and discussed in Section 3.3. This network is denoted as $G_T = (N, E_T)$, where N is the set of nodes (with size $|N| = n$), $E_T = \{e(i,j,t), t \in [0, T-1], i, j \in N\}$ is the set of edges, each item $e(i,j,t)$ indicates an interaction between node i and j at time t. This network is also represented by a 3D adjacency matrix $A_{n\times n\times T}$, whose elements $A(i,j,t) = 1$ or $A(i,j,t) = 0$ represent, respectively, the presence or absence of a contact between node i and j at time t. Furthermore, a corresponding transfer volume matrix $W_{n\times n\times T}$ stores the amount of tokens transferred per interaction. This format captures detailed temporal dynamics of token transfers.

By aggregating the contacts between each node pair over the whole observation time $[0, T-1]$, a static aggregated network is obtained. In this network, two nodes i and j are connected, if there is at least one interaction between i and j over the observation time. There are two possibilities for the edge weights. To avoid misunderstandings, for both edge weights, we introduce a separate network. The first aggregated network, denoted by $G_C = (N, E_C)$, is weighted and directed, encoding transfer count or frequency. Each edge $l(i,j) \in E_C$ can be associated with a weight $C(i,j)$, which denotes the number of transfers from i to j. The corresponding weighted adjacency matrix $C_{n\times n}$ has elements $C(i,j) = \sum_{t=0}^{T-1} A(i,j,t)$. G_C therefore emphasizes how often users interact with each other.

$G_V = (N, E_V)$, is again a weighted, directed network, now containing transfer volume rather than interaction counts. Each edge $l(i,j) \in E_V$ can be associated with a weight $V(i,j)$, which represents the total amount of tokens sent from i to j. The corresponding weighted adjacency matrix $V_{n\times n}$ has elements $V(i,j) = \sum_{t=0}^{T-1} W(i,j,t)$. Thus G_V helps understand overall token flow volume.

For both Aave and Compound, the networks G_T, G_C, and G_V are constructed separately for governance and yield token transfers. These two types of tokens display very different structural patterns. Most yield token transfers, about 90% on average, go directly to or from a single central node that manages the lending pool, showing that value flows are highly concentrated around this central node. This centralization reflects the function of yield tokens, facilitating borrowing and lending, rather than peer-to-peer interaction. In contrast, governance token transfer networks are more distributed, as tokens are used for governance and value exchange, leading to more diverse transfer patterns. Due to this structural difference, yield token networks are excluded from community detection and SI modeling due to the limited insights these methods would provide on these tokens.

4.3. User Activity Features

To analyze governance-active user behavior, this study extracts a range of node-level features that capture different aspects of how users interact with the platform. These include activity levels, timing patterns, network position, and direct governance actions. Comparing and correlating these features helps reveal how governance-active users' behavior in governance activities relates to their role in token transfer networks. The following sections group the features by the network layer from which they were derived, each with an explanation and formal definition.

4.3.1. Features from the Temporal Network (G_T)

These features use temporal data or data per transfer:

- The burstiness $b(i)$ [3] of a node i is the degree to which the activity of this node is irregular and clustered in time, compared to a regular, evenly spaced pattern. Where the score ranges from -1 (perfectly regular) to 1 (highly bursty or irregular), defined as:

$$b(i) = \frac{\sigma_{\Delta t}(i)}{\mu_{\Delta t}(i)}$$

where $\sigma_{\Delta t}(i)$ is the standard deviation and $\mu_{\Delta t}(i)$ the mean of the inter-event times of node i, which is the time between consecutive transfers.

- The average per transfer $APT(i)$ of a node i is the average weight of all token transfers involving this node, defined as:

$$APT(i) = \frac{\sum_{j \in N} \sum_{t=0}^{T-1} (W(i,j,t) + W(j,i,t))}{\sum_{j \in N} \sum_{t=0}^{T-1} (A(i,j,t) + A(j,i,t))}$$

4.3.2. Features from the Aggregated Network (G_C)

This feature is based on aggregated transfer count data:

- The transfer count $TC(i)$ of a node i is the total number of combined incoming and outgoing transfers in which this node is involved, defined as:

$$TC(i) = \sum_{j \in N} (C(i,j) + C(j,i))$$

4.3.3. Features from the Aggregated Network (G_V)

These features are based on aggregated token volume data:

- The in-degree centrality $d^{in}(i)$ of a node i is the total number of incoming edges of this node, defined as:

$$d^{in}(i) = |\{j \in N \mid V(j,i) > 0\}|$$

- The out-degree centrality $d^{out}(i)$ of a node i is the total number of outgoing edges of this node, defined as:

$$d^{out}(i) = |\{j \in N \mid V(i,j) > 0\}|$$

- The eigenvector centrality [11, 26] $x(i)$ of a node i is the measure of the influence of this node based on connections to other highly connected nodes, defined as:

$$x(i) = \frac{1}{\lambda} \sum_{j \in N} V(j,i) x(j)$$

  where λ is the largest eigenvalue of the adjacency matrix.

- The clustering coefficient [10] $CC(i)$ of a node i is the sum of the weighted contributions of all triangles involving this node and two other nodes, divided by the number of all possible triangles that this node could form with other nodes, defined as:

$$CC(i) = \frac{1/2 \sum_j \sum_k (V(i,j)^{1/3} + V(j,i)^{1/3})(V(i,k)^{1/3} + V(k,i)^{1/3})(V(j,k)^{1/3} + V(k,j)^{1/3})}{[d^{tot}(i)(d^{tot}(i) - 1) - 2d_i^{\leftrightarrow}]}$$

  where $d^{tot}(i) = d^{in}(i) + d^{out}(i)$ is the total degree of node i, which is the sum of its in-degree and out-degree, and $d_i^{\leftrightarrow} = \sum_{i \neq j} [V(i,j) > 0 \wedge V(j,i) > 0]$ is the number of bilateral edges between i and its neighbors (i.e. the number of nodes j for which both $i \rightarrow j$ and $j \rightarrow i$ exist). The numerator sums the weighted contributions of all triangles involving i, while the denominator contains the number of all possible triangles that i could form.

- The 2-hop weight sum $W^{2hop}(i)$ of a node i is the total transfer volume in a 2-hop neighborhood of this node, defined as:

$$W^{2hop}(i) = \sum_{j \in N_2(i)} (V(i,j) + V(j,i))$$

  where $N_2(i) = \{j \in N | hops(i,j) \leq 2\}$, and $hops(i,j)$ counts the number of hops needed to get from node i to node j.

As explained in Section 4.2, the yield token transfer network is highly concentrated around a single central node, which is connected to many users. Because of this star-like structure, common structural metrics, like Eigenvector Centrality, Clustering Coefficient, and 2-hop Weight Sum, do not provide meaningful variation between users. These measures work best when there is rich peer-to-peer connectivity, which is not the case here. Therefore, for yield tokens, only activity-based features were calculated, as they still reveal useful patterns in how users interact with the platform.

4.3.4. Governance Activity Features

The following features capture user engagement in governance activities, which is of fundamental importance to be able to relate them to the previously defined features from the token transfer networks:

- Votes cast: the number of votes submitted by each user.
- Average vote weight: the average weight of all votes cast by each user.
- Delegations received: the number of times a user received delegated voting power from someone else.

In Table 4.2 we see that the number of users involved in creating governance proposals is minor, which means that this user-level information would likely not provide useful insights. Therefore, this governance feature has been excluded from the analysis.

4.4. Correlation Analysis

By examining the correlation between different user-level features, such as the volume or frequency of token transfers, their positions within the network, and governance actions, we can identify how the governance behavior of users relates to their activity in token transfer networks. This helps reveal which factors actually shape influence in DeFi governance and whether power depends more on wealth, network connections, or user engagement.

To explore potential relationships between node-level features, Kendall Tau-b rank correlation is used. The most recognized and used correlation coefficient is Pearson correlation, however, due to the presence of non-linear, heavy-tailed distributions, rank order information becomes more significant than the absolute values, making Kendall Tau-b a more adequate measure.

Kendall's Tau-b (τ_b) is derived from the comparison of pairs of observations [18] and measures the correlation between two ranked variables X and Y. It is especially robust to ties, which are common in sparse blockchain data. The formula is defined as:

$$\tau_b = \frac{n_c - n_d}{\sqrt{(n_c + n_d + t_x)(n_c + n_d + t_y)}}$$

where:

- n_c: number of concordant pairs: where the order of X and Y agrees,
- n_d: number of discordant pairs: where the order of X and Y disagrees,
- t_x, t_y: number of tied pairs in X and Y, respectively.

These correlations are computed between the extracted node-level features of governance token networks and yield token networks with the features from the governance activity to assess whether governance participation is linked to token transfer activity or network position. The identification of significant correlations helps understand which features may be indicative of influence or voting behavior.

4.5. Community Detection

Community detection is applied to the governance token transfer network, specifically to the aggregated network based on transfer count between users, G_C, to identify groups of addresses that interact frequently with each other. Here, a community means a subset of nodes with dense internal connections and relatively fewer connections to the rest of the network. Detecting such communities helps assess the positioning of governance-active users in the network: if governance-active users are spread across many communities, this indicates that these users do not interact more frequently with each other than with other nodes, however, if they are concentrated in a few communities, it may point to concentrated interactions.

We apply community detection only to the governance token transfer network, since the yield token transfer networks are structurally concentrated around a central node, the lending pool address. Their star-like structure means that almost all transfers are from or to this central node, with minimal peer-to-peer interaction. Since community detection algorithms look for groups with dense internal connections, they would not find meaningful partitions in such networks.

For this study, we use the Leiden algorithm [31], an improved version of the Louvain method [5] for modularity optimization in large, weighted, directed networks. The modularity score measures how strongly the network divides into distinct communities, with higher values (closer to 1) indicating denser connections within communities and sparser connections between them. Leiden avoids poorly connected or disconnected communities by refining and aggregating clusters through three phases:

1. **Local moving:** each node is moved to the community that maximizes modularity gain.

2. **Refinement:** disconnected parts within a community are split, a key improvement over Louvain.

3. **Aggregation:** each community is treated as a new node in the network.

These steps are repeated until the modularity no longer increases.

Since the Leiden algorithm is stochastic, it yields different community partitions across iterations. To ensure meaningful outcomes of community detection, its stability was assessed before analyzing the outcomes of the algorithm. By running it 100 times on each network, we measured the consistency of the resulting partitions using Variation of Information (VI) [22], Adjusted Mutual Information (AMI) [32], and Adjusted Rand Index (ARI) [14]. Lower VI and higher AMI and ARI scores indicate more stable partitions. Furthermore, we observe the average modularity score and average number of communities produced per network, which are important factors as well. The higher the modularity, the stronger the partitioning, while a lower number of communities makes the outcome easier to study. The results are reported in Table 4.3.

Platform	VI	AMI	ARI	Modularity	Number of communities
Aave	0.22	0.96	0.98	0.63	130
Compound	0.34	0.95	0.93	0.73	120

Table 4.3: Leiden stability analysis scores for Aave and Compound

The obtained scores highlight the high stability of Leiden on G_C, evidenced by its low VI and high AMI and ARI. Additionally, good modularity scores and manageable numbers of communities are achieved, facilitating further analysis of community sizes and internal features.

4.6. Simulation-Based Influence Scoring (SI Model)

To evaluate how governance-active users differ in their potential to reach other nodes in a simulated spreading process, a Susceptible-Infected (SI) model [15] was applied to the temporal governance token transfer network G_T. The resulting influence score, measured as average outbreak size, is treated as a nodal property and compared with governance participation features, which helps determine whether governance power of users is related to their position in the network. The selection of G_T is motivated by its time-sensitive structure, which enabled us to simulate how infections could theoretically propagate through governance token transfers over time, providing a dynamic view of users' structural positions in the network. By ranking users based on their ability to trigger wide-reaching cascades, we assess which nodes hold practical influence beyond our previously defined network features from Section 4.3.

As with community detection, the SI model is applied only to governance token transfers. The yield token transfer network is formed in a star-like structure centered around the lending pool node, meaning nearly all paths pass through this node. This produces unrealistic cascades that do not reflect user-to-user spreading potential relevant for governance.

In the model, each node can be either susceptible or infected. Initially at timestep t = 0, a single governance-active seed node was initialized as infected, leaving all others susceptible. At any timestep t, a susceptible node could get infected by an infected node if the two nodes have an interaction at that specific timestep. If a node becomes infected at time t, it could infect other susceptible nodes that it is connected to from time $t+1$ onward. The infection probability for an edge from node i to node j at timestep t is defined as:

$$p(i,j,t) = \beta \cdot \frac{\log(1 + W(i,j,t))}{\log(1 + W_{\text{max}})}$$

where:

- $\beta \in \{0.01, 0.1, 0.5, 1.0, 2.0\}$ is a tunable infection base rate,
- $W(i,j,t)$ is the token amount belonging to the edge from node i to node j at time t,
- $W_{\text{max}} = \max_{(i,j,t)} W(i,j,t)$ is the maximum token amount among all edges in the network.

Token transfer weights are heavily skewed. The log scaling compresses large values and raises small values, ensuring that even low-volume edges contribute meaningfully to the spread probability. The spreading influence of a node is defined as the average number of infected nodes (outbreak size) till time $t = T$ over all iterations. For each iteration of the experiment, the set of infected nodes and the timestamps of their infection are recorded. This methodology facilitates the quantification of the influence of each node, based on outbreak size, indicating how many nodes were infected in total through each seed node.

The correlations between the governance activity of users and the nodal influence, as modeled by the SI process, are used to measure how well this simulated influence mirrors governance power. This outcome is additionally assessed through the normalized Discounted Cumulative Gain (nDCG) [16]. Unlike Kendall's rank correlation, which evaluates the overall similarity between two ranked lists based on pairwise concordance, or recognition rate, which requires a binary ground truth, nDCG is top-heavy. This means that nDCG prioritizes correct ranking of highly relevant items (i.e., users with high average vote weight), which is particularly useful in governance contexts, where a few highly influential actors dominate decisions. Formally, Discounted Cumulative Gain (DCG) is defined as:

$$DCG_k = \sum_{i=1}^{k} \frac{rel_i}{log_2(i+1)}$$

where:

- rel_i: relevance (average vote weight) of the user at rank i,
- i: the rank based on SI influence in descending order.

nDCG normalizes this against the Ideal DCG (IDCG), which is the ideal ranking (sorted by average vote weight):

$$nDCG_k = \frac{DCG_k}{IDCG_k}$$

By including nDCG, we gain insight into how well the SI model highlights the most powerful governance actors, not just the overall alignment in rankings.

5
Results

This chapter presents the core results of this study, which examines the behavior of governance-active users across several layers of DeFi lending platforms, consisting of governance activity, governance token transfers and yield token transfers. To address the lack of integrated perspectives in prior research, the results combine correlation analyses, community detection, and simulation-based influence scores across multiple network layers. Together, these findings reveal the characteristics of governance-active users more closely, as well as how they relate to the behavior in token transfers. Specifically, the results show three key findings; first, governance power strongly correlates with governance token transfer volume, but not with the frequency of peer-to-peer activity or network positioning. Second, governance-active users are clustered within only a few communities, indicating they interact more with each other than with other nodes. Third, nodes with high simulated spreading potential partially overlap with those holding high voting power, although differences are observed between the platforms. Apart from the most interesting and informative results outlined in this chapter, supplementary figures and interpretations are provided in Appendix A. These materials offer additional examination of the collected data, the research methodology employed, and supplementary figures.

5.1. Feature Correlations between Token Transfers and Governance Activity

To explore how governance participation relates to user behavior, we first analyze correlations between governance activity (e.g., voting and delegating) and various user-level token transfer features. We use Kendall rank correlation, a metric that captures the strength of relationships between ranked variables. Throughout the results, we focus on correlations above 0.75 as strong. Figures 5.1 and 5.2 show how features from the AAVE and COMP governance token networks relate to governance behavior.

The standout finding is that users who participate in voting and governance token transfers show a very strong correlation between their average vote weight and their average amount of tokens transferred, with a Kendall Tau-b correlation of 0.88 and 0.89 Aave and Compound respectively. This suggests that higher governance power aligns with larger governance token flows, highlighting the central role of token wealth. In both Aave and Compound, correlations between other token transfer features do not exceed 0.60, implying that practical governance power is less dependent on users' network position or interaction frequency alone, and more on their token holdings.

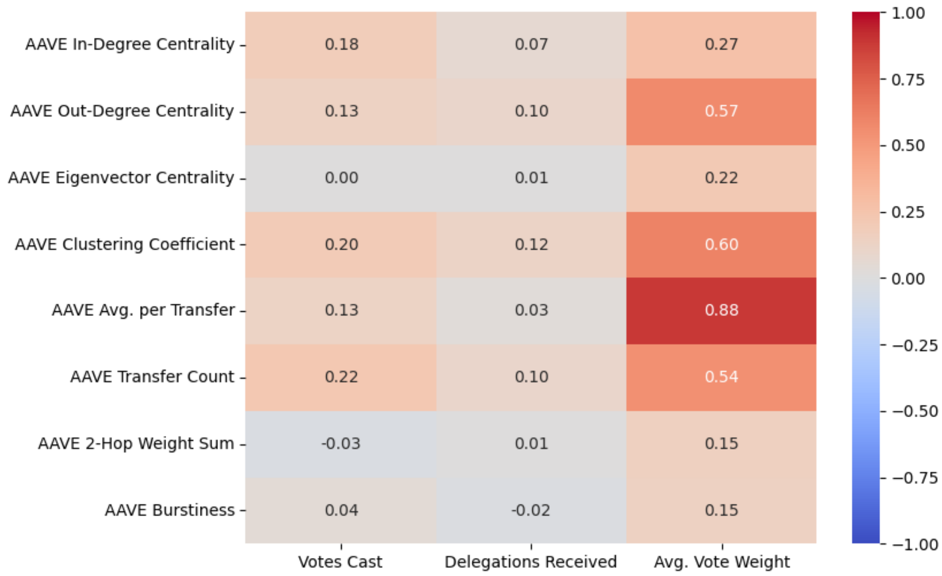

Figure 5.1: Kendall correlation between Aave governance token (AAVE) features & governance features

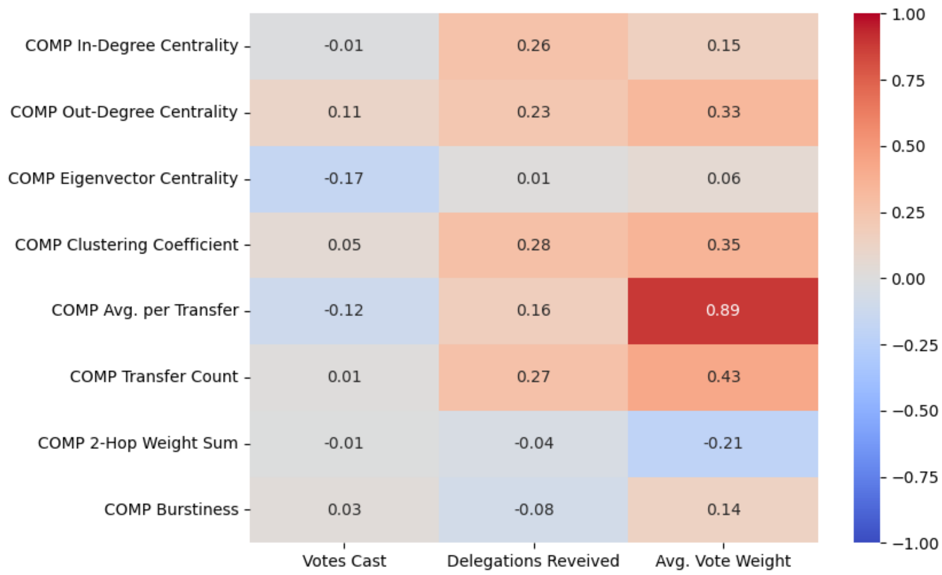

Figure 5.2: Kendall correlation between Compound governance token (COMP) features & governance features

In Figures 5.3 and 5.4 the strong correlation between the average amount of governance tokens per transfer and the average vote weight of users is illustrated through scatterplots, confirming the positive correlation between them. This reflects the token-based design of DeFi governance, where holding more tokens typically gives more influence.

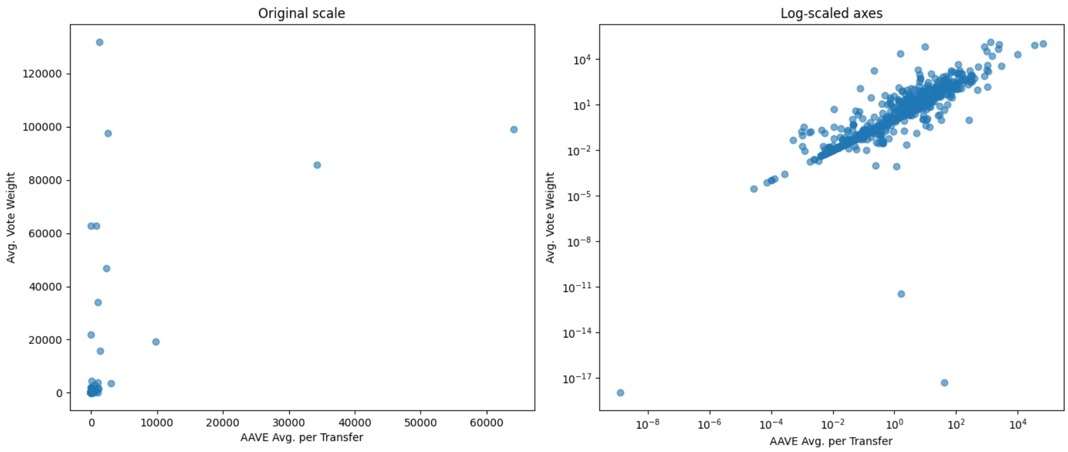

Figure 5.3: Scatterplot of ”AAVE Avg. per Transfer” with ”Avg. Vote Weight”

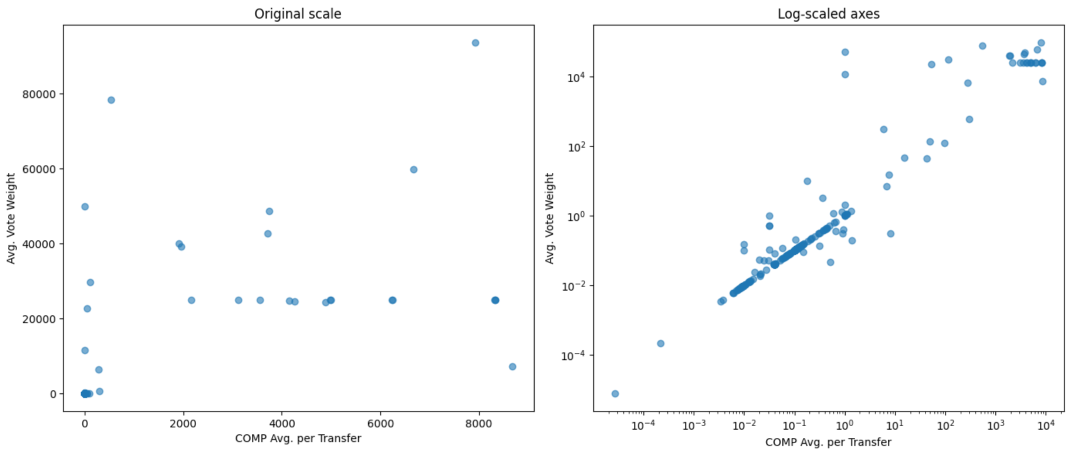

Figure 5.4: Scatterplot of ”COMP Avg. per Transfer” with ”Avg. Vote Weight”

In contrast, Figures 5.5 and 5.6 show that governance activity has no strong correlation with yield token activity (e.g., aWETH and cWETH transfers), with all correlation scores smaller than 0.50. This suggests that governance power is not clearly tied to how user interaction with the lending platform, but rather to the amount of governance tokens value they hold or move.

Figure 5.5: Kendall correlation between Aave yield token (aWETH) features & governance features

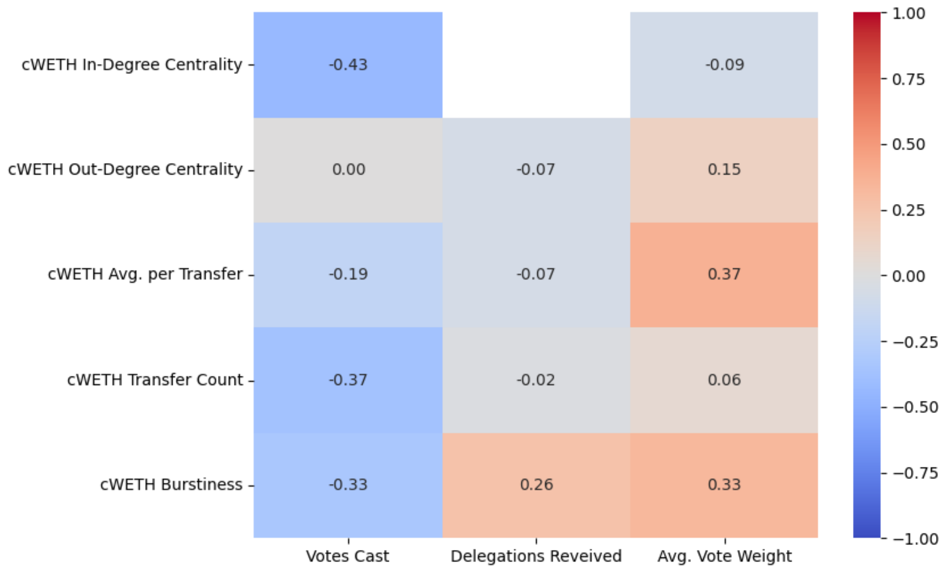

Figure 5.6: Kendall correlation between Compound yield token (cWETH) features & governance features

In both Aave and Compound, average vote weight per user, indicating governance power, correlates strongly with governance token amounts transferred. Other network features that are related to the activity level or position in the network of users, such as degree centralities, clustering coefficient, or burstiness, do not show high correlations. This reinforces the link between financial stake and governance power.

5.2. Community Structure and Governance Activity Distribution

To understand how governance activity is distributed across users, we used the Leiden algorithm, a method for detecting user communities. We applied this to the G_C network, which is based on the interaction count of governance token transfers between users, to reveal groups of users who often interact with each other. Within the communities, we then examined whether governance-active users, those who voted, delegated, or proposed, were spread out across the network or concentrated in specific communities.

5.2.1. Aave

Figure 5.7 shows the size of the 10 largest communities in terms of their number of nodes, highlighting governance-active nodes in orange. This partitioning consists of approximately 130 communities with a modularity score of 0.63, which indicates moderately strong community structure. At first glance, governance-active users seem highly concentrated in Communities 1 and 2.

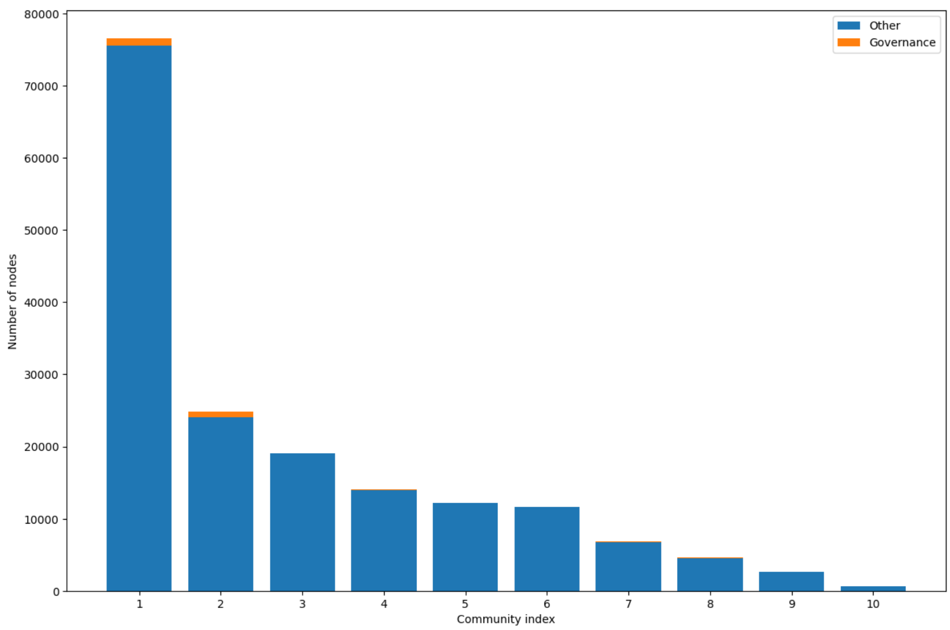

Figure 5.7: Sizes of the largest communities and governance-active nodes

Figures 5.8 and 5.9 confirm this pattern. These figures display the absolute number and relative percentage of governance-active nodes per community, with the red line denoting the overall percentage of governance-active users in the entire network. Community 2, and to a lesser extent Community 1, contain a disproportionately high number of governance-active users. This indicates that these users interact more often with each other than with the broader network.

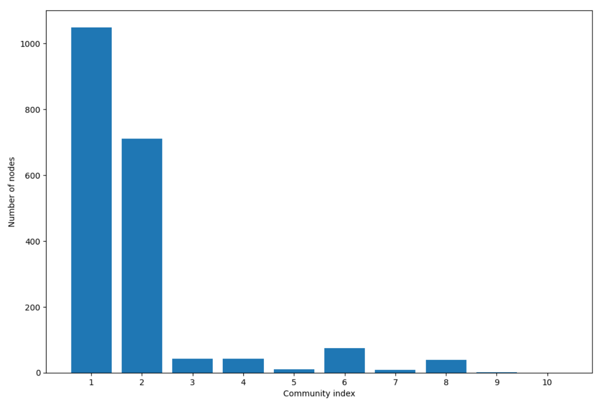

Figure 5.8: Amount of governance-active users per community

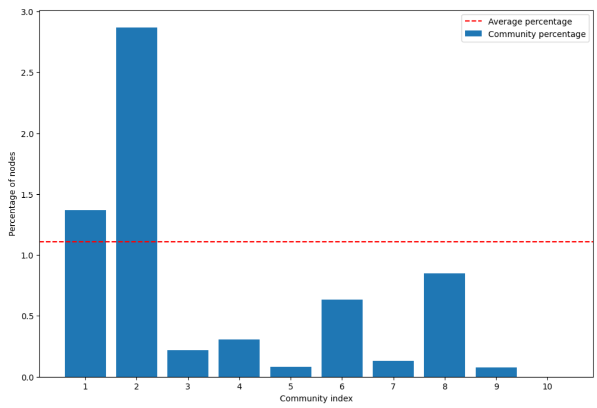

Figure 5.9: Percentage of governance-active users per community

For further understanding, we computed the share of governance-active users in Communities 1 and 2 as percentage of their overall amount in the entire network. Community 1 contains approximately 50% of all governance-active users, while this is 34% for Community 2, confirming the observed clustering.

5.2.2. Compound

In Compound, we observe an even stronger concentration of governance activity. Figure 5.10 illustrates the sizes of the 10 largest communities and the distribution of governance-active node over them. This figure resembles the characteristics previously examined in Aave. The partitioning contains approximately 120 communities, with a modularity score of 0.73, indicating even stronger community structuring than in Aave, which suggests that user interactions cluster tightly. Community 5 contains most of the active governance participants.

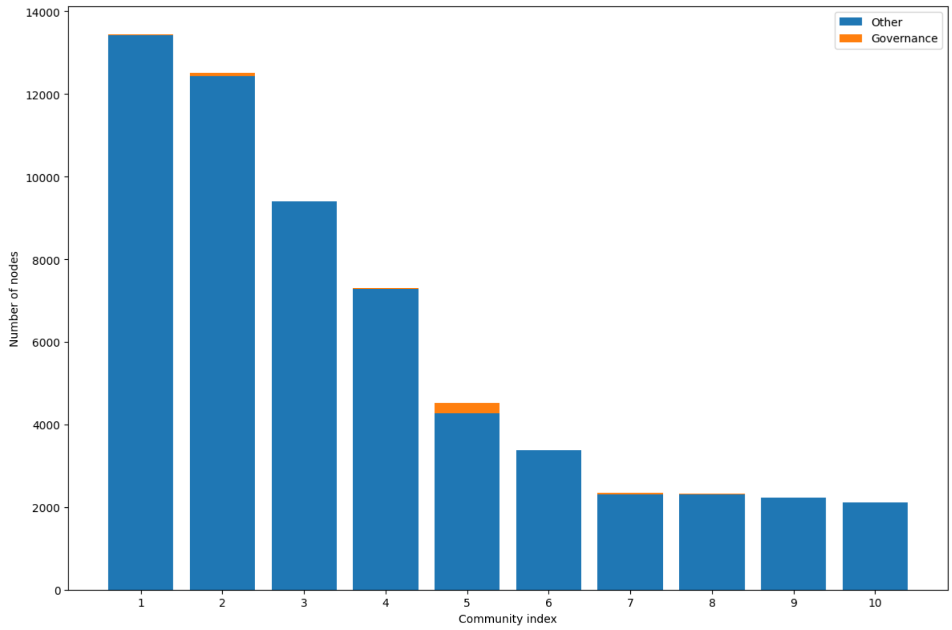

Figure 5.10: Sizes of the largest communities and governance-active nodes

Figures 5.11 and 5.12 confirm that governance-active users are highly clustered in Community 5. The figures again indicate that many governance-active users appear in the same community, therefore interacting more frequently with each other.

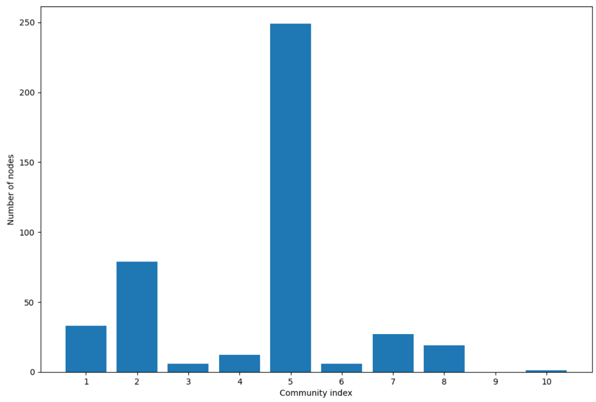

Figure 5.11: Amount of governance-active users per community

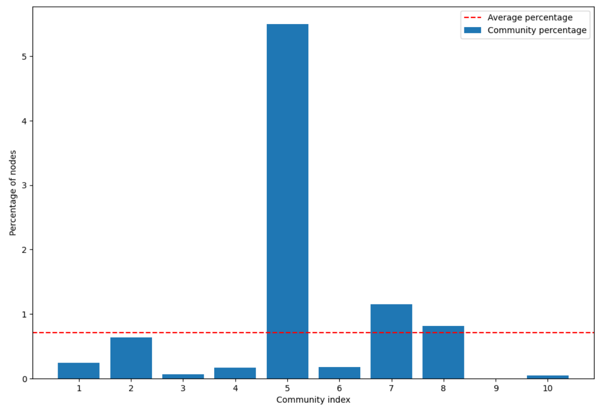

Figure 5.12: Percentage of governance-active users per community

We analyzed the percentage of all governance-active users clustered in Community 5. This percentage is approximately 52%, again indicating concentrated governance participation in a single community.

In both platforms, governance participation is not broadly distributed across communities in the network, but clustered within a small number of groups that interact more often with each other. This indicates that governance-active users tend to form connected subgroups that interact more among themselves than with other parts of the network, which helps explain how governance participation remains concentrated within relatively closed user groups that shape decision-making.

5.3. Simulation-Based Influence and Governance Power

To estimate each governance-active user's potential to reach others in the governance token transfer network, we calculate their nodal influence using the Susceptible-Infected (SI) model on the temporal network G_T. In this simulation, the influence is determined by the number of nodes that a single infected seed node is able to reach or infect, also known as the average outbreak size across different infection base rates β. We then compare these scores to actual governance power (primarily measured by average vote weight) through Kendall Tau-b rank correlation and nDCG ranking alignment.

Figures 5.13 and 5.14 contain Kendall rank correlations between the governance activity features and the obtained nodal-influence scores, corresponding to the average outbreak size per seed node. We observe only very limited correlations between the number of votes cast and the number of delegations received with the obtained influence scores. Stronger correlations are observed with respect to the average vote weight, which we therefore analyze further through computing the nDCG ranking alignment scores.

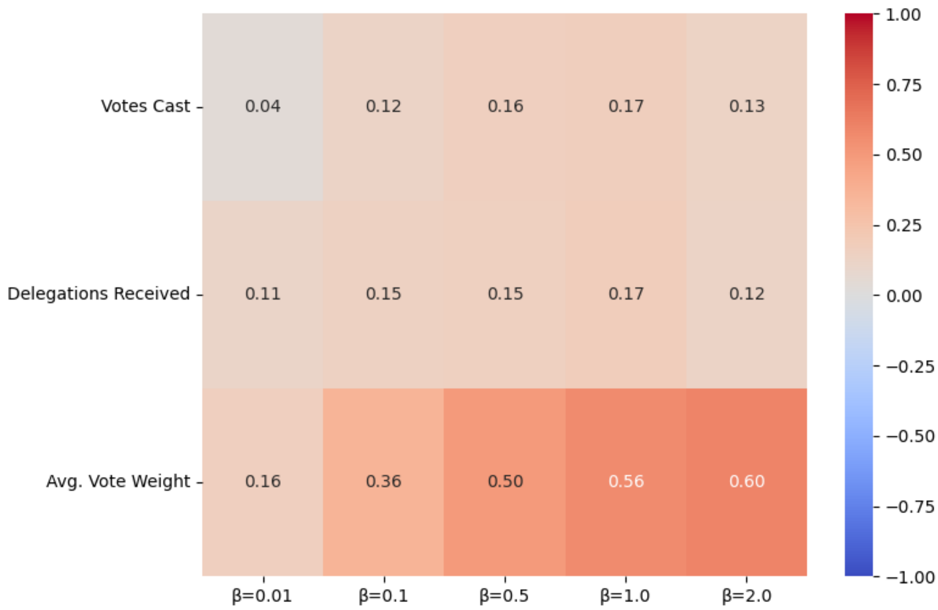

Figure 5.13: Kendall correlation between governance features and SI nodal influence per β in Aave

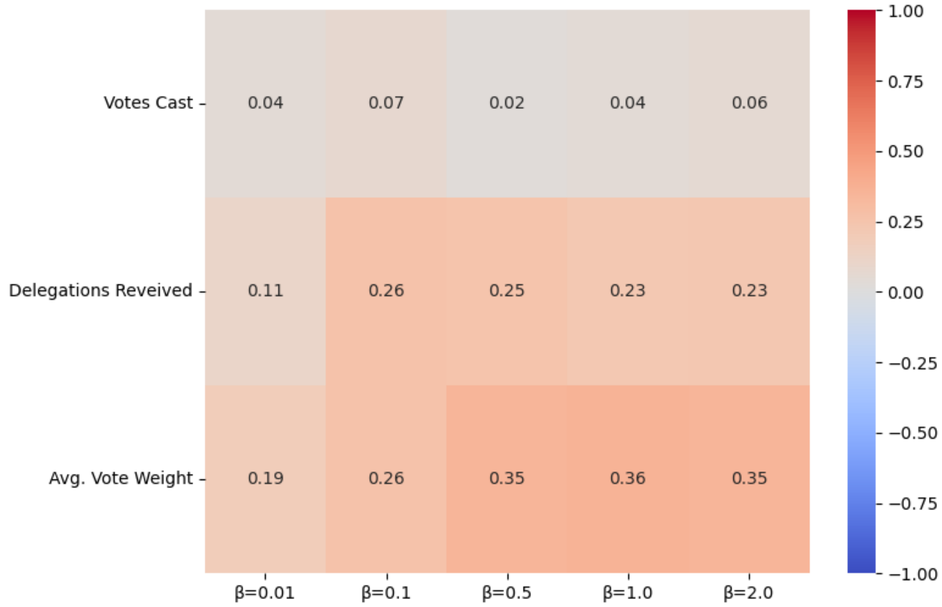

Figure 5.14: Kendall correlation between governance features and SI nodal influence per β in Compound

The Kendall Tau-b correlation and nDCG between the ranking of the simulated spreading influence scores and the average vote weight in Aave and Compound across different spreading rates (β) are displayed in Figure 5.15 and Figure 5.16 respectively.

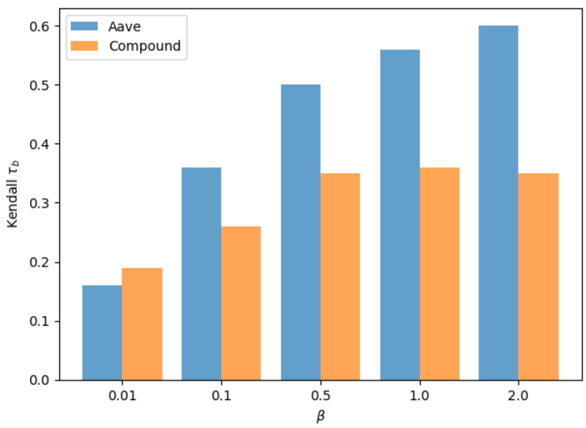

Figure 5.15: Kendall correlation between SI nodal influence and average vote weight

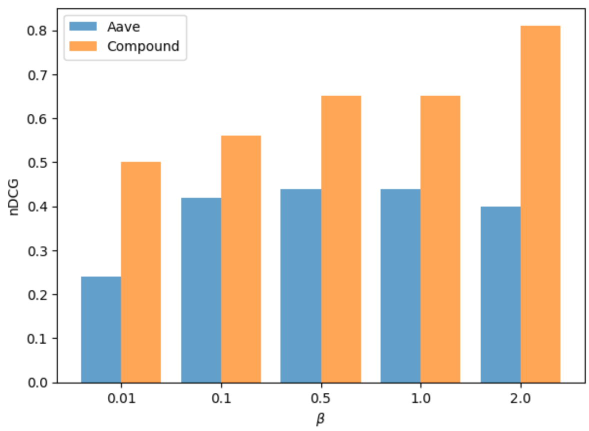

Figure 5.16: nDCG between SI nodal influence and average vote weight

These figures show that a user's potential to reach others through token transfers, as measured by SI influence scores, partly aligns with their practical governance impact. In Aave, the rank correlation reaches a peak of 0.60 at $\beta = 2.0$, while the nDCG score stays below 0.50. This means that, overall, nodes with higher simulated influence tend to have higher voting power, but the very top influential nodes do no perfectly match the top governance voters. In Compound, the overall correlation is more moderate, peaking at 0.36, but the higher nDCG scores, reaching 0.81 at $\beta = 2.0$, suggest that, although the general alignment is weaker, the top-ranked simulated influential users more closely match those with higher average vote weight. Together, these results suggest that governance-active users who are structurally well-positioned in the token transfer network can amplify their influence in practice, complementing their formal voting power, but the strength of this relationship varies by platform and infection base rates.

5.4. Cross-Platform Comparison

Although Aave and Compound differ in size, tokenomics (the economic aspects of a cryptocurrency or blockchain project), and user base, they show similar patterns in governance structure and nodal influence. Overall, the results show that governance-active users hold substantial voting power that aligns strongly with their governance token flows but less so with other network position metrics. They tend to cluster within a few communities and often hold network positions that help them reach other nodes, possibly increasing practical impact of their actions.

That said, there are some differences. Aave displays a more distributed community structure, with governance spread over two main groups, while this is more concentrated in Compound. Additionally, Aave shows a stronger alignment between actual vote weight and simulated influence, while Compound shows higher alignment in top-ranking (nDCG), suggesting different forms of power concentration.

6 Discussion

This study set out to better understand the behavior of governance-active users in DeFi lending platforms. While prior research showed that only a limited number of users actively contributes to DeFi governance, it did not examine how these users' governance actions relate to their behavior in token transfer networks. Although these users directly influence platform decisions, an analysis is missing of how the properties of governance-active users relate to their behavior in other layers. To address this, the thesis combined multiple data layers, consisting of governance actions, governance token flows, and yield token flows, across two major DeFi lending platforms, Aave and Compound. This study examined how the role of users that actively participate in governance relates to their positions in token networks using correlation analysis on nodal features, community detection, and a simulated spreading process.

A key finding is the strong correlation between the average vote weight and the average amount of governance tokens they transfer, while no significant correlations are observed between governance activities and other token transfer features. This shows that DeFi lending platforms are largely plutocratic, where governance token wealth is the main driver of governance power, while other network features, like connectivity, matter less. Community detection shows that governance-active users cluster into a few groups that interact more with each other than with the broader network. This suggests that even if many users hold tokens, only few connected subgroups in the governance token transfer network drive decisions, possibly indicating that token-based governance power remains concentrated in practice, despite being open in design. The Susceptible-Infected (SI) spreading scores, expressed as average outbreak size per seed node, partly align with the average vote weight of users. However, the correlations and nDCG values are moderate and vary per platform, meaning the network position of a node can help amplify governance impact, but does not guarantee governance power on its own. Most patterns are consistent across both platforms, with only minor differences, primarily in the outcomes of the SI process. This consistency suggests that other token-governed DeFi platforms may face similar issues.

That said, this study has several limitations. First, the analysis only includes users that actively participate in governance, which limits the scope. There are many addresses involved in DeFi lending platforms, but only a small portion of them participates in governance through proposing, voting, or delegating. This excludes insights about the full network context and could miss how passive users may shape the platforms indirectly. Second, token transfer volume may confound some results, since these volumes affect both the vote weight of users and the simulated influence measured by the SI process, which are compared in the analysis. Third, our dataset consists of only on-chain data, which can be extracted from the blockchain. However, there also exist off-chain interactions between users, mainly on forums where discussions and coordination on proposals takes place. Including such off-chain activity could offer a more complete picture of governance activity. Other limitations include the one-year observation window, the focus on just two platforms, and the lack of delegated voting weight data, which provides more information on the amount of power that is transferred through delegation.

Based on these findings, DeFi lending platforms should consider design changes to broaden participation and reduce power concentration. Simply having open access does not guarantee this. In practice, DeFi governance is vulnerable to concentration of power, despite its transparent and permissionless design. If decentralization is a core goal, more than open access and transparency are needed, designers must rethink the incentives, rules, and interfaces that shape governance behavior. Possible steps include lowering barriers for small token holders to vote, capping maximum vote weight per address, highlighting underrepresented proposals to encourage broader input, or experimenting with non-token-based governance models.

Overall, this study highlights that understanding who governance-active users are and how they behave across different layers in DeFi lending platforms provides important insights into whether DeFi's promise of community-led governance holds up in practice. It requires careful design and ongoing analysis to make sure influence is genuinely distributed.

7 Conclusion and Future Work

This thesis examines the behavior of governance-active users in DeFi lending platforms. By combining multiple data layers from Aave and Compound, consisting of governance actions and both governance and yield token transfers, the study provides a more complete picture of practical user influence.

The results show that formal governance power, measured by average vote weight, strongly aligns with the volume of governance tokens users transfer. Other network features, such as degree centrality, clustering coefficient, or burstiness, do not meaningfully explain governance participation. Community detection reveals governance-active users cluster into a few communities, indicating that decision-making is driven by connected subgroups rather than widely distributed across all users. Simulated spreading influence scores show that the position of a user in the network strengthens how effectively they reach others, but that this practical reach only partially aligns with governance power and differs per platform. These patterns are mostly consistent across Aave and Compound, suggesting that similar dynamics may exist in other token-governed DeFi platforms.

These findings directly address the research questions: governance actions, especially voting, strongly align with governance token flows, while other token network properties do not strongly relate to governance actions (RQ1); governance-active users cluster within only a few communities (RQ2); nodal influence from the SI model partially matches governance activity (RQ3); and both platforms yield similar findings (RQ4).

By combining layered network data and dynamic modeling, this work contributes a more realistic view of how governance activity of users in DeFi lending platforms relates to their behavior in other layers, specifically token transfers. It highlights that simply issuing governance tokens and providing open access is not enough to ensure broad participation. Designers of DeFi platforms should consider new incentives and design choices to broaden governance participation, such as lowering barriers for smaller holders to vote, capping maximum vote weight, making user influence more transparent, or experimenting with hybrid or non-token-based models.

7.1. Future Work

Future research could broaden the scope by not just analyzing the behavior of users actively participating in governance, but by also studying the behavior of other users that are involved in token transfers but not in governance. Longer observation periods could be used to show how the behavior dynamics evolve over time. Additional DeFi platforms and alternative blockchains could be included to test whether similar patterns hold across the ecosystem. Off-chain data could be integrated, this thesis completely relies on on-chain data, consisting of governance activity data and token transfer data. However, users also use forums and social media to discuss and interact with each other, which could provide more complete information about governance in DeFi platforms. The collection of such data should probably by done manually, which could be an issue. Methodologically, applying more advanced diffusion models could better capture nodal influence in a spreading process. Lastly, studying delegation dynamics in more depth could clarify how delegated voting power affects concentration and community structure.

A

Supplementary Figures

The following figures support the main analysis but were not included in the Results chapter due to their limited direct relevance or clarity of insight. While these elements may not be central to the primary conclusions, they could be useful for replication, extended analysis, or follow-up work.

A.1. Correlation Results

Figure A.1 and A.2 provide additional correlations between the features from governance tokens and yield tokens for Aave and Compound respectively. The most notable observation is the elevated correlation between the average governance tokens per transfer and the average yield tokens per transfer in Aave, indicating a positive relation between them.

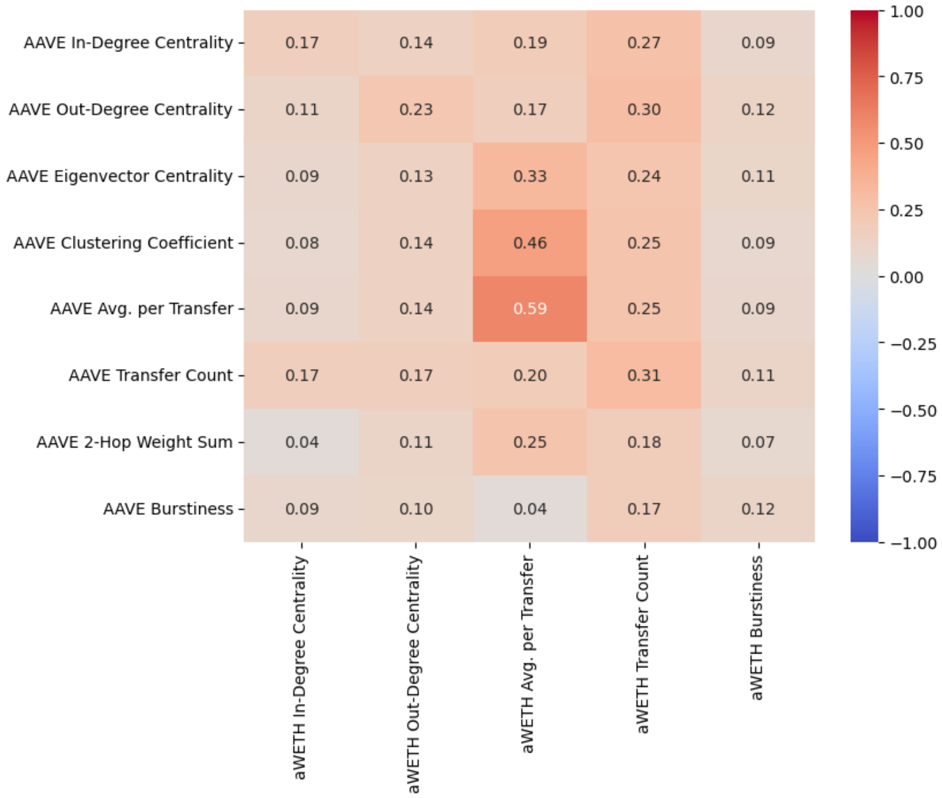

Figure A.1: Kendall correlation between Aave governance token (AAVE) features & yield token (aWETH) features

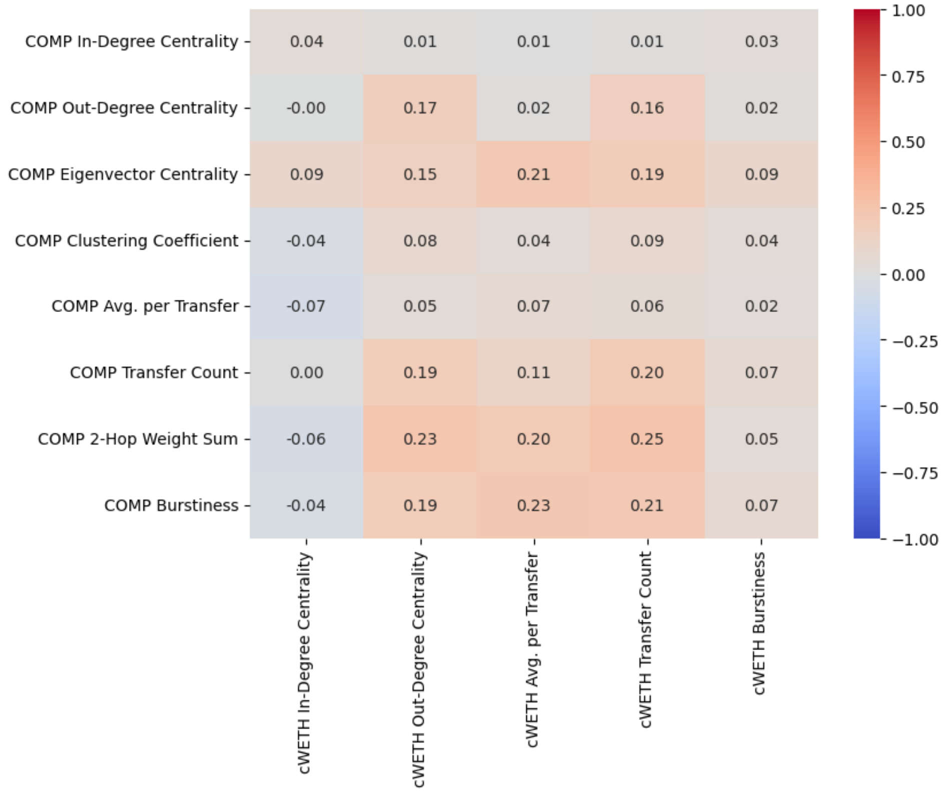

Figure A.2: Kendall correlation between Compound governance token (COMP) features & yield token (cWETH) features

A.2. Community Structure

As a result of the implementation of the Leiden algorithm for clustering the G_C governance token transfer networks into communities, it is possible to explore the averages of each node-level governance activity feature over these communities. The figures displaying these scores are visualized in Figures A.3 and A.4. The most notable observation are the high average feature scores for Community 6 in Compound, although this community only contains few governance-active users.

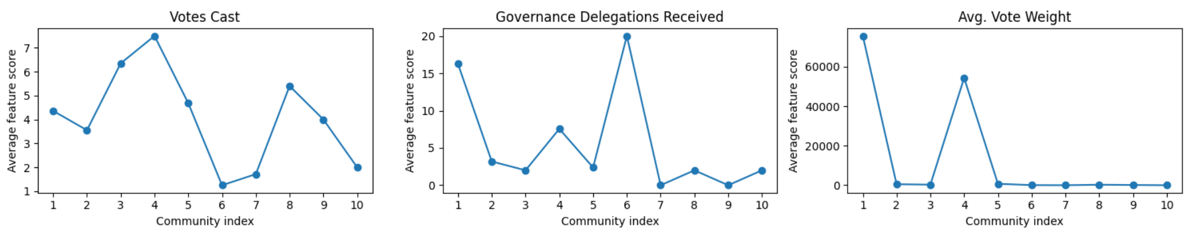

Figure A.3: Average feature scores per community in Aave

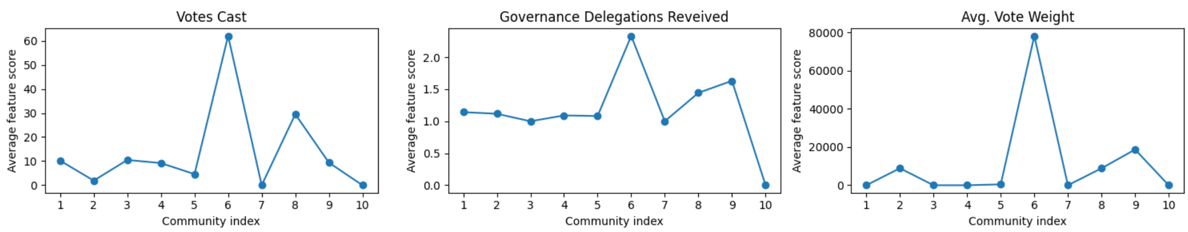

Figure A.4: Average feature scores per community in Compound

A.3. Influence Dynamics

We look at the outcomes of the SI model, employed on the G_T networks. We study the correlations between the outcomes of each β configuration, displayed in Figures A.5 and 5.13. Looking at these heatmaps, we can clearly see that larger values of β lead to higher correlations. This suggests that, as β increases, the outcomes become more similar to each other.

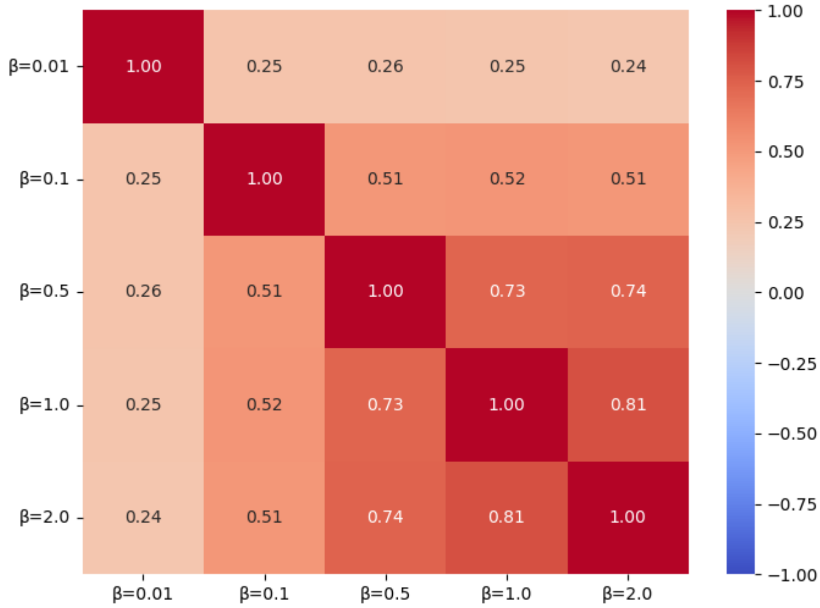

Figure A.5: Kendall correlation between SI nodal influence per β in Aave

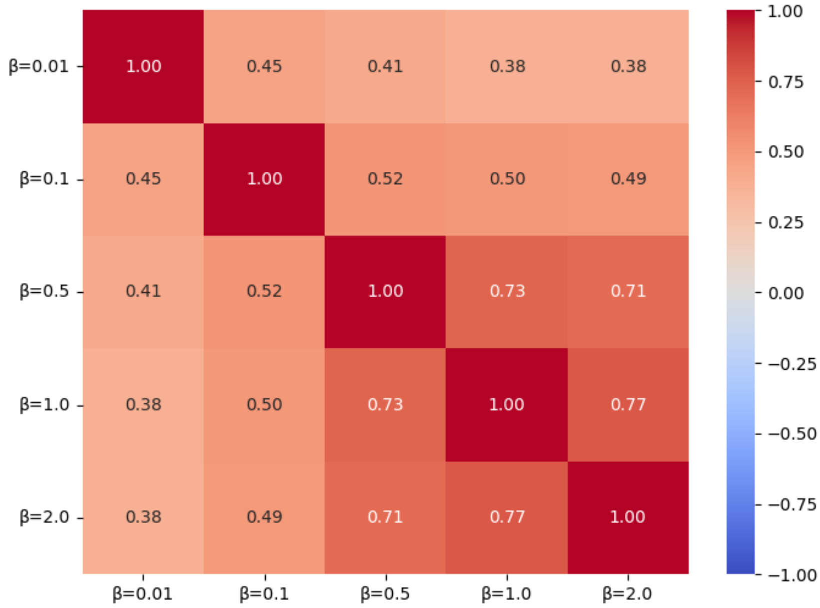

Figure A.6: Kendall correlation between SI nodal influence per β in Compound

Figures A.7 and A.8 show how influence scores correlate with other governance token transfer features. Users with high Out-Degree Centrality (many outgoing transfers), high Clustering Coefficient (interacting with dense groups), and high Transfer Count tend to have a higher spreading potential, meaning they can reach more nodes on average in a simulated spreading process, which makes sense, since these are all network features that facilitate the infection of neighboring nodes. The correlations become stronger with an increasing value of β. We do not observe strong correlations with other features, such as Eigenvector Centrality, 2-Hop Weight Sum, or Burstiness, indicating that spreading potential depends more on direct activity patterns and local connections than on broader, indirect network influence or timing irregularity of transfers.

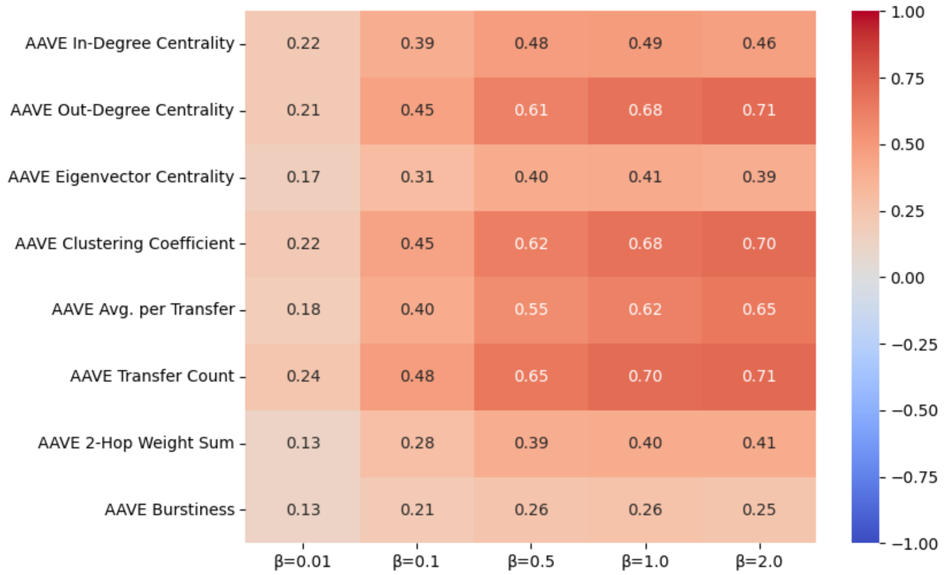

Figure A.7: Kendall correlation between Aave governance token (AAVE) features and SI nodal influence

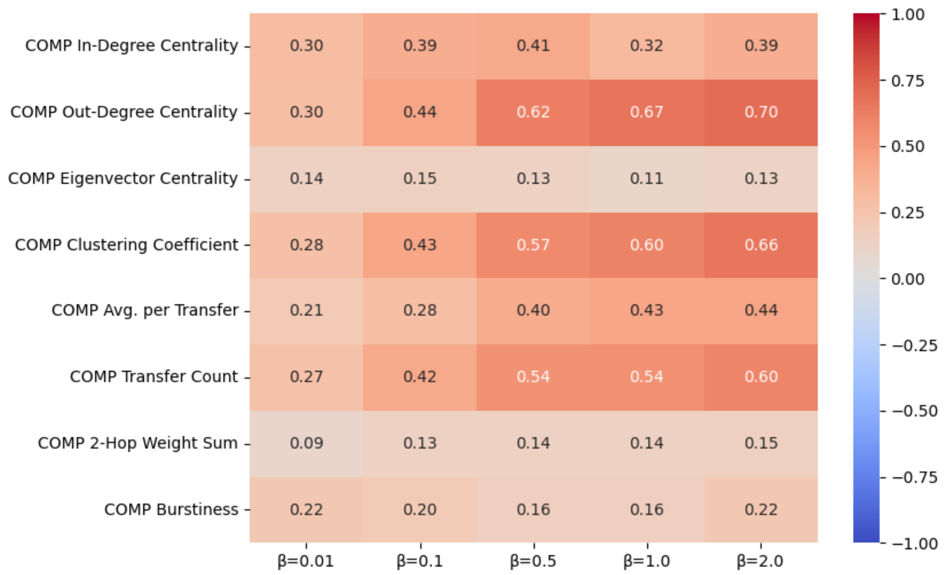

Figure A.8: Kendall correlation between Compound governance token (COMP) features and SI nodal influence

B
AI Disclosure Statement

This thesis includes content refined using artificial intelligence tools. Specifically, ChatGPT[1] and DeepL[2] were used to brainstorm ideas and to enhance clarity and readability. All analysis, interpretations, and conclusions are my own and I take full responsibility for the integrity and accuracy of the work.

[1] https://chatgpt.com/
[2] https://www.deepl.com/en/write

Bibliography

[1] Mauro Aliano and Stefania Ragni. “Game-based modeling of delayed risk contagion in cryptocurrency exchanges”. In: *Annals of Operations Research* (2025), pp. 1–23.

[2] Ziqiao Ao et al. “Is decentralized finance actually decentralized? A social network analysis of the Aave protocol on the Ethereum blockchain”. In: *arXiv preprint arXiv:2206.08401* (2022).

[3] Albert-Laszlo Barabasi. “The origin of bursts and heavy tails in human dynamics”. In: *Nature* 435.7039 (2005), pp. 207–211.

[4] Tom Barbereau et al. “Decentralised Finance’s timocratic governance: The distribution and exercise of tokenised voting rights”. In: *Technology in Society* 73 (2023), p. 102251.

[5] Vincent D Blondel et al. “Fast unfolding of communities in large networks”. In: *Journal of statistical mechanics: theory and experiment* 2008.10 (2008), P10008.

[6] Financial Stability Board. “The financial stability risks of decentralised finance”. In: *Basel. Financial Stability Board and International Monetary* (2023).

[7] Vitalik Buterin et al. “Ethereum white paper”. In: *GitHub repository* 1 (2013), pp. 22–23.

[8] Giulio Cornelli et al. *Why defi lending? evidence from aave v2*. Tech. rep. Bank for International Settlements, 2024.

[9] Maya Dotan et al. “The vulnerable nature of decentralized governance in defi”. In: *Proceedings of the 2023 Workshop on Decentralized Finance and Security*. 2023, pp. 25–31.

[10] Giorgio Fagiolo. “Clustering in complex directed networks”. In: *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics* 76.2 (2007), p. 026107.

[11] Linton C Freeman et al. “Centrality in social networks: Conceptual clarification”. In: *Social network: critical concepts in sociology. Londres: Routledge* 1.3 (2002), pp. 238–263.

[12] Robin Fritsch, Marino Müller, and Roger Wattenhofer. “Analyzing voting power in decentralized governance: Who controls DAOs?” In: *Blockchain: Research and Applications* (2024), p. 100208.

[13] Samer Hassan and Primavera De Filippi. “Decentralized autonomous organization”. In: *Internet Policy Review* 10.2 (2021).

[14] Lawrence Hubert and Phipps Arabie. “Comparing partitions”. In: *Journal of classification* 2 (1985), pp. 193–218.

[15] Maureen Hurley, Glen Jacobs, and Melinda Gilbert. “The basic SI model”. In: *New Directions for Teaching and Learning* 2006.106 (2006), pp. 11–22.

[16] Kalervo Järvelin and Jaana Kekäläinen. “Cumulated gain-based evaluation of IR techniques”. In: *ACM Transactions on Information Systems (TOIS)* 20.4 (2002), pp. 422–446.

[17] Johannes Rude Jensen, Victor von Wachter, and Omri Ross. “How decentralized is the governance of blockchain-based finance: Empirical evidence from four governance token distributions”. In: *arXiv preprint arXiv:2102.10096* (2021).

[18] Maurice G Kendall. “The treatment of ties in ranking problems”. In: *Biometrika* 33.3 (1945), pp. 239–251.

[19] Dan Lin et al. “Modeling and understanding ethereum transaction records via a complex network approach”. In: *IEEE Transactions on Circuits and Systems II: Express Briefs* 67.11 (2020), pp. 2737–2741.

[20] Dan Lin et al. “T-edge: Temporal weighted multidigraph embedding for ethereum transaction network analysis”. In: *Frontiers in Physics* 8 (2020), p. 204.

[21] Junliang Luo and Xue Liu. “Optimizing Blockchain Analysis: Tackling Temporality and Scalability with an Incremental Approach with Metropolis-Hastings Random Walks”. In: *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining*. 2025, pp. 410–418.

[22] Marina Meilă. “Comparing clusterings—an information based distance”. In: *Journal of multivariate analysis* 98.5 (2007), pp. 873–895.

[23] Johnnatan Messias et al. “Understanding blockchain governance: Analyzing decentralized voting to amend defi smart contracts”. In: *arXiv preprint arXiv:2305.17655* (2023).

[24] Matthias Nadler and Fabian Schär. “Decentralized finance, centralized ownership? an iterative mapping process to measure protocol token distribution”. In: *arXiv preprint arXiv:2012.09306* (2020).

[25] Peterson K Ozili. “Decentralized finance research and developments around the world”. In: *Journal of Banking and Financial Technology* 6.2 (2022), pp. 117–133.

[26] Lawrence Page et al. *The PageRank citation ranking: Bringing order to the web.* Tech. rep. Stanford infolab, 1999.

[27] Stamatis Papangelou, Klitos Christodoulou, and George Michoulis. “Exploring decentralized governance: a framework applied to compound finance”. In: *The International Conference on Mathematical Research for Blockchain Economy*. Springer. 2023, pp. 152–168.

[28] Kaihua Qin et al. “CeFi vs. DeFi–Comparing Centralized to Decentralized Finance”. In: *arXiv preprint arXiv:2106.08157* (2021).

[29] Kaushal Shah et al. “A systematic review of decentralized finance protocols”. In: *International Journal of Intelligent Networks* (2023).

[30] Ying Teng et al. “SEIR-diffusion modeling and stability analysis of supply chain finance based on blockchain technology”. In: *Heliyon* 10.3 (2024).

[31] Vincent A Traag, Ludo Waltman, and Nees Jan Van Eck. “From Louvain to Leiden: guaranteeing well-connected communities”. In: *Scientific reports* 9.1 (2019), pp. 1–12.

[32] Nguyen Xuan Vinh, Julien Epps, and James Bailey. “Information theoretic measures for clusterings comparison: is a correction for chance necessary?” In: *Proceedings of the 26th annual international conference on machine learning*. 2009, pp. 1073–1080.

[33] Shuai Wang et al. “Blockchain-enabled smart contracts: architecture, applications, and future trends”. In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 49.11 (2019), pp. 2266–2277.

[34] Sam Werner et al. “Sok: Decentralized finance (defi)”. In: *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*. 2022, pp. 30–46.

[35] Jiahua Xu and Nikhil Vadgama. “From banks to defi: the evolution of the lending market”. In: *Enabling the Internet of Value: How Blockchain Connects Global Businesses* (2022), pp. 53–66.

[36] Luyao Zhang, Xinshi Ma, and Yulin Liu. “Sok: blockchain decentralization”. In: *arXiv preprint arXiv:2205.04256* (2022).

[37] Yufan Zhang et al. “Blockchain network analysis: A comparative study of decentralized banks”. In: *Science and Information Conference*. Springer. 2023, pp. 1022–1042.