Eisermann et al. *EPJ Data Science* (2025) 14:85
https://doi.org/10.1140/epjds/s13688-025-00610-5

EPJ.org

EPJ Data Science

RESEARCH

Open Access

Concentration in governance control across decentralised finance protocols

Thomas Eisermann[1,6], Carlo Campajola[2,3,4*], Claudio J. Tessone[4,5] and Andreia Sofia Teixeira[6,1]

Handling Editor: Francesco Gullo

*Correspondence:
c.campajola@ucl.ac.uk
[2]Institute of Finance and Technology, University College London, London, UK
[3]DLT Science Foundation, London, UK
Full list of author information is available at the end of the article

Abstract

Blockchain-based systems are frequently governed through tokens that grant their holders voting rights over core protocol functions and funds. The centralisation occurring in Decentralised Finance (DeFi) protocols' token-based voting systems is typically analysed by examining token holdings' distribution across addresses. In this paper, we expand this perspective by exploring shared token holdings of addresses across multiple DeFi protocols. We construct a Statistically Validated Network (SVN) based on shared governance token holdings among addresses. Using the links within the SVN, we identify influential addresses that shape these connections and conduct a post-hoc analysis to examine their characteristics and behaviour. Our findings reveal persistent influential links over time, predominantly involving addresses associated with institutional actors or smart-contracts, which hold significant fractions of token supplies across the sampled protocols. Finally, we observe that token holding patterns and concentrations tend to shift in with protocol valuations and dollar denominated Total Value Locked.

Keywords: Decentralised finance; Governance; Cross-protocol influence; Complex networks

1 Introduction

Decentralised Finance (DeFi) encompasses a family of services and protocols that replicate traditional financial functions, such as collateralised lending and asset trading, on permissionless distributed ledger technology (DLT) networks. These functions are executed through smart contracts: automated software that any DLT user can operate, ensuring transactions remain transparent, irreversible, and free from reliance on third parties (Werner et al. [62]).

Although the execution of smart contracts is inherently permissionless and decentralised, the configuration and governance of these functions are determined by entities with governance rights over the protocols. Governance in DeFi typically begins under the centralised control of a "benevolent dictator" or a small council, transitioning over time to a more decentralised structure. This shift is facilitated through governance tokens: these are digital assets which confer membership in a Decentralised Autonomous Organisation (DAO), which in turn manages protocol operations, and token holders have the ability to propose or vote on protocol changes. Upgrades are presented as executable code,

Springer

with predefined thresholds of token ownership required to propose changes and a quorum needed for their approval. While the rights associated with governance tokens vary across projects, they generally provide significant decision-making authority over protocol governance (Werner et al. [62]).

Since a governance token represents decision-making power, the token holder distribution is a relevant factor, among others, to determine the extent of decentralisation of a protocol (Axelsen et al. [10]). While previous studies have analysed token distribution within individual protocols (Barbereau et al. [13], Fritsch et al. [31], Jensen et al. [40], Nadler and Schär [48]), few have explored the extent to which actors wield influence across multiple protocols simultaneously.

This is not necessarily a straightforward task, due to the pseudonymity of DLT systems. Users interact with DLT systems through unique but anonymous identifiers known as addresses. These addresses are derived from public-key cryptography, where the owner holds a private key, and the public key is shared across the network. Addresses allow users to validate transactions, interact with smart contracts, such as those in DeFi applications, and enable transparent tracking of token ownership within the system.

In this paper, we inspect the dynamics of cross-protocol governance by analysing the ownership structures of governance tokens within DeFi. Governance token holders may be individuals, companies, associations, or even smart contracts. Therefore, we investigate the presence and impact of user groups who may hold disproportionate influence across different DeFi protocols. In our analysis we examine the existence of such groups, who those users are, how they behave and the degree of control they may exert within the DeFi ecosystem.

2 Background

Our work builds on existing research addressing the emergence of centralisation in DLTs and the structure of DeFi governance, incorporating methodological approaches from complex financial network analysis.

Decentralisation is often seen as a key value in Web 3.0 projects aiming to reshape power dynamics (Bodó et al. [18]). However, its definition varies depending on whether decentralisation is assessed through technological, economic, geographical, or social and governance metrics. In this context, over the last years, researchers have been exploring the centralisation of wealth in cryptocurrency markets. For instance, Makarov and Schoar [44] identifies growing wealth inequality on the Bitcoin blockchain, driven primarily by the specialization of entities and the emergence of a financial intermediation industry. Similar trends have been observed in other cryptocurrency systems (Campajola et al. [20, 21]).

Centralisation has also been measured in technological terms by looking at the peer-to-peer infrastructure (Gao et al. [32], Grundmann et al. [35]), where emerging hubs constitute potential vulnerabilities, or at the consensus protocols of layer-1 blockchains (Brown [19], Li et al. [42]), where very few mining and staking pools have accrued the majority of the validation authority. Geographic centralisation has been another focus of study. In Sun et al. [53] the authors investigate the geographic distribution of Bitcoin miners, finding that while operations span as many as 139 countries, the majority of computing power is concentrated in just a few locations.

We find a broad consensus in the literature that DLTs exemplify the natural tendency of open and loosely regulated markets to concentrate power and evolve toward oligopolistic structures (Arthur [7]). This is further complicated by the limited accountability of central entities, often shielded by the anonymity granted by the design of these systems (Walch [61]), making regulation and anti-trust action very hard to implement.

The literature on DeFi governance has sought to evaluate decentralisation across its various dimensions, ranging from technical architecture to social layers and communities associated with these systems (Axelsen et al. [10], Cong et al. [24], Gogol et al. [34], Ovezik et al. [49], Rüetschi et al. [50], Sai et al. [51]). Since governance tokens typically confer voting rights to their holders, several studies have focused on the distribution of these tokens and the corresponding voting power within DeFi communities to determine who ultimately controls the protocols (Cong et al. [24], Feichtinger et al. [28], Fritsch et al. [31], Jensen et al. [40], Nadler and Schär [48]). The study by Nadler and Schär [48] examines various types of tokens, including governance tokens within DeFi, and shows a pronounced concentration in their ownership structure, with often 5 to 10 addresses holding more than 50% of the supply. They also find significant variation in multi-token holding patterns depending on the specific token. Similarly, in Fritsch et al. [31] the authors report high centralisation in governance across their sample. Their findings suggest that a few addresses wield substantial control, analogous to major shareholders in traditional corporate governance. Despite this centralisation, they observe that influential entities frequently align their votes with the broader community.

In Barbereau et al. [13], the authors find that despite the theoretical decentralisation in DeFi protocols, voting rights tokens tend to lead to highly centralised control. Their analysis across multiple DeFi projects highlights that even token distribution strategies designed to enable "fair launches" have not prevented centralisation over time. In Feichtinger et al. [28], the authors show that despite their intent for inclusive decision-making, many DAOs exhibit centralised control and inefficiencies, with high costs and low community participation. Jensen et al. [40] find that in practice, token-based voting power can undermine the democratic ethos by enabling major holders to push through unpopular decisions. This study suggests a misalignment between declared values and actual governance, proposing a framework to evaluate emerging voting systems and their potential to address these discrepancies. Finally, in Kitzler et al. [41], the authors show that DAO contributors, including project owners and developers, often hold significant governance influence, with majority control in 7.54% of DAOs and last-minute token acquisitions affecting 14.81% of proposals. Nonetheless, they also observe limited evidence of contributors exerting influence across competing DAOs, despite the potential incentives for governance token holders to oppose proposals that benefit rival protocols.

This growing research sheds light on the unequal token distribution of governance tokens, which results in a concentration of decision-making power and, therefore, control over the project. However, to the best of our knowledge, apart from the indirect findings of Nadler and Schär [48] and Kitzler et al. [41], no study has explicitly investigated the phenomenon of *cross-protocol control*: the influence exerted by one or more entities holding governance tokens across multiple protocols, thereby impacting several communities simultaneously.

Network analysis has been widely applied to model financial systems, as a powerful method to make interdependencies and contagion pathways apparent across assets and markets (Fonseca and de França Carvalho [29], Siudak [52]). Among these methods, Statistically Validated Networks (SVNs) have proven particularly effective for analysing complex financial networks, by identifying statistically significant links against a background of randomly occurring connections. This approach, introduced by Tumminello et al. [55], is useful for isolating key dependencies within bi-partite systems and has been applied in various financial contexts. In this context, Bardoscia et al. [14] offers a comprehensive summary of the theory behind financial networks and their main applications, illustrating how their structures influence systemic risk and the transmission of financial shocks. Financial systems are highly interconnected, and understanding these connections is crucial for identifying points of vulnerability, especially during periods of market instability.

Particularly close to our work is Gualdi et al. [36], where the authors used SVNs to analyse portfolio overlaps and assess systemic risk. By examining binary holding matrices and accounting for individual stock position sizes, they identified significant dependencies and found that the proportion of validated links increased steadily leading up to the 2007–2008 financial crisis, showing that market participants exhibited herding behaviour that amplified the crash.

Inspired by the methodologies of Tumminello et al. [55] and Gualdi et al. [36], we adapt them to examine dependencies across DeFi protocols, focusing on shared overlaps in addresses holding governance tokens.

3 Methodology

3.1 Data collection and pre-processing

Based on Nadler and Schär [48], we identify the following criteria that inform our selection of governance tokens:

1. *Governance Qualification*: the token must qualify as a governance token, granting its holders the ability to influence decisions shaping the ecosystem's rules (Freni et al. [30]). Tokens classified purely as stablecoins, utility tokens, token wrappers, or token baskets are excluded;
2. *ERC-20 Compliance*: The token must adhere to the ERC-20 standard (Vogelsteller and Buterin [60]). This requirement ensures consistency in token behaviour, facilitates reliable data collection, and enables comparability across governance tokens. Empirically, non-ERC-20 governance tokens are rare among major DeFi Protocols.
3. *Market and Protocol Significance*: at least one of the following conditions is satisfied at the time of data collection (December 8, 2022) (Freni et al. [30], Nadler and Schär [48]):
   (a) The token has a significant circulating supply with a market capitalisation (MCAP) of over 200 million USD according to CoinGecko,
   (b) The protocol's contracts have a total value locked (TVL), the estimated value of assets stored in the protocol), excluding vested tokens, of over 300 million USD according to DeFiLlama, with at least 50% of the TVL on Ethereum Mainnet.

A token is included in our sample if and only if it satisfies the following combination of the above criteria:

$$(1) \wedge (2) \wedge \big[(3a) \vee (3b)\big].$$

Table 1 List of selected governance tokens from DeFi protocols used in this paper, including token names and contract addresses

Protocol	Contract Address
Uniswap	0x1f9840a85d5af5bf1d1762f925bdaddc4201f984
Aave	0x7fc66500c84a76ad7e9c93437bfc5ac33e2ddae9
Lido	0x5a98fcbea516cf06857215779fd812ca3bef1b32
Maker	0x9f8f72aa9304c8b593d555f12ef6589cc3a579a2
Curve	0xd533a949740bb3306d119cc777fa900ba034cd52
1Inch	0x111111111117dc0aa78b770fa6a738034120c302
Bitdao	0x1a4b46696b2bb4794eb3d4c26f1c55f9170fa4c5
Convex	0x4e3fbd56cd56c3e72c1403e103b45db9da5b9d2b
Compound	0xc00e94cb662c3520282e6f5717214004a7f26888
dYdX	0x92d6c1e31e14520e676a687f0a93788b716beff5
Balancer	0xba100000625a3754423978a60c9317c58a424e3d
Sushi	0x6b3595068778dd592e39a122f4f5a5cf09c90fe2
Yearn Finance	0x0bc529c00c6401aef6d220be8c6ea1667f6ad93e
Instadapp	0x6f40d4a6237c257fff2db00fa0510deeecd303eb
Aura Finance	0xc0c293ce456ff0ed870add98a0828dd4d2903dbf

To verify the MCAP and TVL criteria, we relied on data from CoinGecko[1] and DeFiLlama,[2] respectively. We refer the reader to methodologies of the respective data providers for the exact methodology of calculation of these metrics.

The final token selection for this study is displayed in Table 1. It is important to note that Governance tokens have different rights and obligations associated with them. In Appendix G we provide a short description of their function, governance process, and governance scope.

From an Ethereum Erigon Node using ethereum-etl (Medvedev and the D5 team [47]) we retrieved the addresses holding the selected tokens between 2021-01-15 and 2022-06-15. These were obtained by aggregating the historical token transfer events at monthly intervals, resulting in monthly snapshots of token ownership. Snapshot dates and block heights are displayed in Table 4 and in Appendix A. The chosen blocks are those added to the longest blockchain closest to 12 AM UTC on the day of each snapshot.

Addresses are enriched with real-world *entity labels*. We label only those addresses that (i) hold at least 0.01% of the circulating supply of *any* token in our sample, and (ii) hold at least two or more tokens. To enrich the token holder address data with meaningful entity labels, we leverage two primary sources: Etherscan, accessed via a curated community endpoint (Art [6]), and the Arkham Intelligence API (Arkham [5]).

The labelling procedure consists of the following steps. First, we construct a deduplicated list of all addresses that meet the criteria mentioned above. Each address is queried for available labels using both the Etherscan and Arkham Intelligence endpoints. In the case of conflicting classifications, we perform manual inspection of the contract and apply a majority rule across Etherscan, Arkham, and manual curation. Addresses without available labels are categorised as *EOA* (Externally Owned Account) if no deployed bytecode is detected, and as *Unknown Smart Contract* if deployed bytecode is detected. Addresses known to be burner or black-hole accounts (e.g., *0x0000…dead*) are excluded

[1] https://www.coingecko.com/.

[2] TVL figures used are taken from DeFiLlama's chain-level endpoint (https://api.llama.fi/v2/historicalChainTvl/Ethereum). This source reports TVL net of the most common forms of double counting (e.g. assets rehypothecated across protocols or via liquid staking), so the market-wide TVL series we correlate with in Sect. 4 reflects an aggregate dollar-denominated activity indicator rather than the sum of protocol-reported TVLs.

Table 2 Entity-level label definitions

Label	Definition
`Blockchain Scaling`	Addresses related to blockchain scaling solutions or layer 2 technologies
`Bridges`	Addresses involved in transferring assets across blockchains
`Custodian`	Addresses associated with custodian entities that hold and secure cryptocurrencies on behalf of their clients to ensure safety and compliance
`Decentralised Exchange (DEX)`	Addresses involved in decentralized exchanges which allow users to trade cryptocurrencies without intermediaries
`Fund`	Addresses associated with institutional funds that invest in cryptocurrency tokens and companies
`Fund Decentralized`	Addresses associated with decentralized entities operating as investment funds
`Hacker`	Addresses associated with reported hacks and exploits
`Individual`	Addresses believed to be controlled by a single, identifiable human user
`Lending Centralized`	Addresses associated with centralized lending and borrowing platforms
`Lending Decentralized`	Addresses associated with decentralized lending and borrowing platforms
`Liquid Staking`	Addresses associated with liquid staking protocols that allow staked assets to remain fungible and tradable
`MEV Bot`	Contract addresses or entities that exploit Maximal Extractable Value (MEV) opportunities
`Real World Assets`	Addresses associated with tokens representing real-world assets (e.g., real estate, bonds)
`Smart Contract Platform`	L1/L2 foundation or treasury addresses
`Stablecoin`	Issuers of fiat-pegged tokens
`Yield`	Yield-farming or auto-compounder platforms
`Misc`	Addresses that do not clearly fit into any specific category
`Unknown Smart Contract`	Deployed bytecode without a clear label
`Externally Owned Account (EOA)`	Externally Owned Accounts (EOAs) controlled by private keys, not smart contracts and not identifiable

from all token supply and voting power calculations. We adopt Arkham's category system (accessed under free login) and extend it with two technical buckets (*Unknown Smart Contract* and *EOA*). Table 2 summarises the final label set used throughout the analysis.

3.2 Statistically validated network projections

Based on the approach by Tumminello et al. [55], we constructed a bipartite graph $G = \{\mathcal{N}_t, \mathcal{N}_a, \mathcal{E}_G\}$, by grouping addresses in a node set $\mathcal{N}_a$ and the token nodes into the other set $\mathcal{N}_t$. A link $(t_i, a) \in \mathcal{E}_G$ is established between a token node $t_i \in \mathcal{N}_t$ and an address node $a \in \mathcal{N}_a$ if the latter holds the token associated with the former.

As in Tumminello et al. [55], we assume a hypothesis of random connectivity between addresses and tokens, accounting for the degree heterogeneity in each set. Specifically, the probability that two tokens t_i and t_j share $N^a_{i,j}$ addresses by chance is given by the hypergeometric distribution:

$$P(X = N^a_{i,j}) = \frac{\binom{N^a_i}{N^a_{i,j}}\binom{N_a - N^a_i}{N^a_j - N^a_{i,j}}}{\binom{N_a}{N^a_j}} \tag{1}$$

where N_i^a and N_j^a denote the number of addresses holding tokens t_i and t_j respectively (or in other words t_i and t_j's degree in the bipartite network), and $N_a = |\mathcal{N}_a|$ is the total number of unique addresses in the bipartite network.

We then proceed to test the presence of the link between tokens t_i and t_j against the hypothesis of random connectivity. This is done by computing the p-value $p(N_{i,j}^a)$, which measures the probability of observing $N_{i,j}^a$ or more shared addresses under the null hypothesis:

$$p(N_{i,j}^a) = 1 - \sum_{x=0}^{N_{i,j}^a - 1} P(X = x) \tag{2}$$

To control for multiple hypothesis testing, we apply a Bonferroni correction, adjusting the threshold level to α/T, where $\alpha = 0.01$ is the family-wise error rate and T is the number of token pairs tested. This sets a hard cap on the occurrence of type I errors. The outcome is a SVN $G_V = (\mathcal{N}_t, \mathcal{E}_V)$, where $\mathcal{E}_V$ is the set of statistically validated edges between tokens. Each edge $(t_i, t_j) \in \mathcal{E}_V$ indicates a statistically significant relationship between tokens t_i and t_j, based on the number of shared addresses that hold both tokens, as determined by the hypergeometric test and the corrected significance threshold.

3.3 Statistical analysis and validation

The primary objective of our research is to understand the prevalence of control exerted by addresses across DeFi protocols. As highlighted by Gualdi et al. [36], binary token holding matrices do not account for position size, which is crucial for assessing the control exerted by groups of addresses across protocols or, in their case, the concentration in financial portfolios. Validating the original weighted matrix is challenging due to the lack of an analytical null model.

To overcome this limitation, we identify all links in the SVN projections based on binary token-holding matrices. These identified links between tokens t_i and t_j are used to filter and identify the set of relevant addresses responsible for the formation of a given link within a token projection. We define this set of relevant addresses as the *link-defining addresses*, denoted by $\mathcal{A}_{i,j}$. For every statistically validated token–token link $(t_i, t_j) \in \mathcal{E}_V$, we define the corresponding set of link-defining addresses as:

$$\mathcal{A}_{i,j} = \{a \in \mathcal{N}_a \mid (t_i, a) \in \mathcal{E}_G \wedge (t_j, a) \in \mathcal{E}_G\} \tag{3}$$

These link-defining addresses represent the addresses that simultaneously hold both tokens t_i and t_j, thereby contributing to the formation of the statistically validated link between these tokens. By analysing the properties of the set $\mathcal{A}_{i,j}$, we can further investigate the characteristics of the addresses that bridge multiple protocols through token holdings.

We start by analysing the structural evolution of SVNs over time and report a set of standard network-level metrics for each monthly snapshot. These metrics provide a coarse-grained view of the changing topological features of the SVNs: *Nodes and Edges*, the number of active governance tokens and statistically validated links based on shared token holdings; *Density*, the proportion of observed edges to all possible token–token connections, indicating overall network connectivity; *Giant Component (%)*, the proportion of nodes contained in the largest connected subgraph, capturing the extent of overall network integration; *Average Clustering*, the tendency of token triads to form, with higher values indicating more locally clustered structures; *Assortativity*, the degree correlation among connected nodes, where negative values indicate that highly connected tokens tend to connect to less-connected ones; *Number of Communities*, the number of modular partitions detected using the Louvain method for community detection (Blondel et al. [17]); *Largest Community (%)*, the share of nodes in the largest detected community, reflecting the extent of modular dominance; *Average Jaccard Similarity*, the mean pairwise Jaccard coefficient between connected nodes, quantifying the extent of overlap in their neighbour sets; *Betweenness Centrality*, the extent to which nodes lie on shortest paths between other nodes, indicating potential bridge nodes that control information flow; *Degree Centrality*, the number of connections per node, highlighting both general connectivity and the most connected hubs in the network.

For the analysis of link-defining addresses, we restrict our analysis to link-defining addresses $\mathcal{A}_{i,j}$ that hold at least 0.0005 % of the total supply of any token in our sample during the observation window, i.e., addresses satisfying $\frac{t_j}{s(t_j)} > 5 \times 10^{-6}$ for some t_j. Focusing on supply-adjusted stakes ensures we capture holders whose voting power is large enough to influence governance outcomes. In previous works, Cong et al. [23] showed that the aggregate holding ratio governs both platform productivity and wealth concentration, while Gualdi et al. [36] demonstrated that statistically validated portfolio overlaps become influential only when positions represent a non-trivial fraction of an asset's float. Guided by these insights, we introduce three supply-normalised metrics to quantify cross-protocol control.

Average Token Holding Share: the average proportion of total token supply held by the set of addresses $\mathcal{A}_{i,j}$. It provides a symmetric measure of shared token holdings across the token pair (i.e., validated link E_V) by averaging the supply shares for each token held by the overlapping address set.

$$\textit{Average Token Holding Share} = \frac{\sum_{a \in \mathcal{A}_{i,j}} \left(\frac{q_a(t_i)}{s(t_i)} + \frac{q_a(t_j)}{s(t_j)} \right)}{2} \tag{4}$$

where $q_a(t_i)$ and $q_a(t_j)$ are the amounts of tokens t_i and t_j held by address a, and $s(t_i)$ and $s(t_j)$ are the total supplies of tokens t_i and t_j, respectively.

Directional Token Holding Share: the proportion of a specific token's total supply, $s(t_i)$, held by the set of addresses $\mathcal{A}_{i,j}$, that simultaneously hold both tokens t_i and t_j. It is computed by summing the raw quantities of token t_i held by addresses, denoted $q_a(t_i)$, and normalising by the total supply of t_i:

$$\textit{Directional Token Holding Share}_{t_i} = \frac{\sum_{a \in \mathcal{A}_{i,j}} q_a(t_i)}{s(t_i)} \tag{5}$$

Unlike the symmetric Average Token Holding Share, this provides an asymmetric measure of shared holding, isolating the token holdings of overlapping addresses on a single token. The metric is computed independently for each token in the pair of a validated link.

Label-specific Token Holding Share: the share of the Average Holding Share attributable to addresses with a specific label L, where each address $a \in \mathcal{A}_{i,j}$ is assigned a label $L(a)$

based on the type of entity (e.g., lending pool, institution, or individual). It reflects the relative contribution of different categories of actors to the overall shared ownership between tokens t_i and t_j, and it is defined as:

$$\textit{Label-specific Token Holding Share}_L = \frac{\sum_{a\in\mathcal{A}_{i,j},L(a)=L}\left(\frac{q_a(t_i)}{s(t_i)}+\frac{q_a(t_j)}{s(t_j)}\right)}{\sum_{a\in\mathcal{A}_{i,j}}\left(\frac{q_a(t_i)}{s(t_i)}+\frac{q_a(t_j)}{s(t_j)}\right)} \tag{6}$$

We also report: *Link Size*, the number of addresses constituting a validated link $(i,j) \in \mathcal{E}_V$; *Median Token Holding Value*, USD-denominated value of governance tokens held by addresses in $\mathcal{A}_{i,j}$; and *Token Holding Share Inequality*, measured via the Gini coefficient (Gini [33]) of token holding share distribution across addresses in $\mathcal{A}_{i,j}$.

We evaluate the results for each metric m, where indicated, using a permutation test to assess differences against a control group. The control group $\mathcal{A}_{\text{control}}$ consists of a set of randomly selected addresses, matched in size to the link-defining address set $\mathcal{A}_{i,j}$ and sampled from the tokens constituting the link.

In each permutation, we randomly shuffle the addresses from the combined set $\mathcal{A}_{i,j} \cup \mathcal{A}_{\text{control}}$. Let $\hat{m}^{(k)}_{\mathcal{A}_{i,j}}$ and $\hat{m}^{(k)}_{A_{\text{control}}}$ represent the recalculated metric of interest for the shuffled groups in the k-th permutation. This process is repeated for 1000 iterations to generate a distribution of the metric under the null hypothesis of no effect.

The significance of the observed difference $\Delta m = m_{\mathcal{A}_{i,j}} - m_{A_{\text{control}}}$ is evaluated by calculating its percentile rank in the distribution of $\Delta\hat{m}^{(k)} = \hat{m}^{(k)}_{\mathcal{A}_{i,j}} - \hat{m}^{(k)}_{\mathcal{A}_{\text{control}}}$ across the 1000 iterations.

4 Results

4.1 Token projections

In Fig. 1, we present the validated SVN projections over time. Each network reflects the statistically validated links between tokens based on shared wallet holdings for each monthly snapshot from January 2021 to June 2022. Nodes represent tokens, while edges capture significant co-holding relationships.

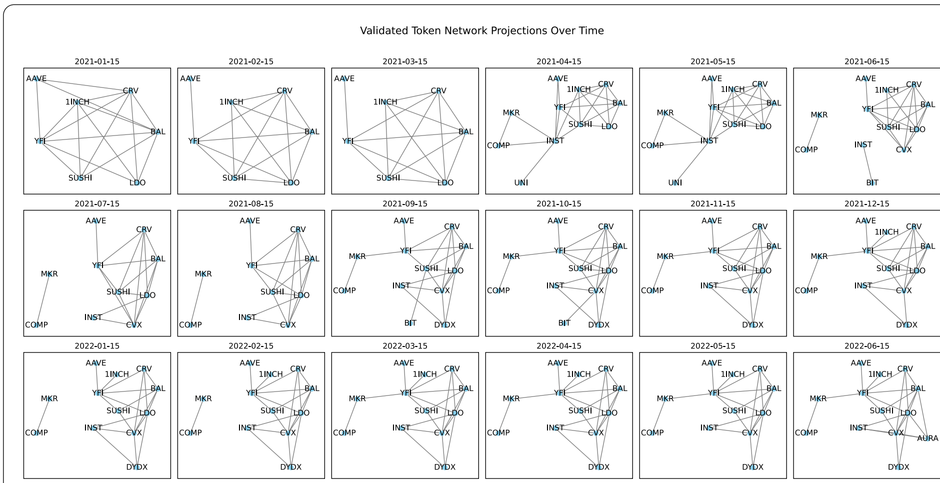

Figure 1 Validated token network projections across time

Table 3 Summary Network Metrics per Snapshot

Date	Nodes	Edges	Density	Giant Comp.%	Clust.	Assort.	Comms	Largest Comm.%
2021-01-15	7	17	0.810	100.000	0.871	−0.492	1	100.000
2021-02-15	7	15	0.714	100.000	0.743	−0.119	2	71.430
2021-03-15	7	15	0.714	100.000	0.743	−0.119	2	71.430
2021-04-15	11	22	0.400	100.000	0.696	−0.306	2	54.550
2021-05-15	11	23	0.418	100.000	0.714	−0.277	3	45.450
2021-06-15	12	19	0.288	66.670	0.444	0.443	3	66.670
2021-07-15	10	18	0.400	80.000	0.567	0.324	3	50.000
2021-08-15	10	18	0.400	80.000	0.567	0.324	3	50.000
2021-09-15	12	21	0.318	83.330	0.486	0.222	3	50.000
2021-10-15	12	23	0.348	100.000	0.508	−0.121	3	41.670
2021-11-15	11	21	0.382	81.820	0.588	0.186	3	45.450
2021-12-15	12	20	0.303	83.330	0.458	0.040	3	50.000
2022-01-15	12	20	0.303	83.330	0.456	−0.107	3	50.000
2022-02-15	12	19	0.288	83.330	0.428	−0.113	3	41.670
2022-03-15	11	18	0.327	81.820	0.479	0.080	3	54.550
2022-04-15	11	18	0.327	81.820	0.497	−0.116	3	45.450
2022-05-15	12	19	0.288	83.330	0.444	−0.231	3	41.670
2022-06-15	13	24	0.308	100.000	0.464	−0.352	3	46.150

In Table 3, we can observe that the early snapshots (e.g. January to May 2021) are small networks characterized by high density and full connectivity. However, from mid-2021 onward and with more tokens being added, the network becomes sparser and we see multiple communities start to appear. Clustering remains relatively high throughout, indicating persistent local cohesiveness. We find that assortativity is not uniquely positive or negative, alternating periods of disassortative and assortative network structure. Together with the relatively high clustering coefficients that we measure, disassortativity in this context would suggest that the validated links connect a high-degree node in the SVN (i.e. a token that is co-held with many other tokens significantly more than a random allocation would imply) to few other tokens which are not widely held but that form triangles with other low-degree nodes. Conversely, assortativity in the projection would suggest that most tokens are similarly widespread in the market. We report these projection-level statistics as we believe they are useful for summarising the topology of the SVNs; however, their interpretation is not straightforward in the economic context and is inherently dependent on the link-validation technique adopted. For this reason, in the next section we introduce a set of metrics designed to extract more economically relevant information from our analysis.

To inspect the persistence of network links over time, we constructed a Jaccard Similarity Matrix, shown in Fig. 2. The diagonal elements represent perfect similarity (1.0), while off-diagonal elements indicate the degree of similarity between SVN network snapshots at different points in time.

Figure 2 shows a relatively similar network structure of SVN snapshots before June 2021 and after June 2021. The break in similarity can be explained by the introduction of new governance tokens into our sample dataset, leading to the formation of additional validated links. In the appendix, Fig. 13, we show the absolute number of addresses and provide context for the introduction and increased presence of these tokens in our sample. Despite the apparent break in similarity, in Fig. 2 we can see that the overall links constituting the SVNs remain persistent and stable over time, resulting in a high similarity of the network typology. This persistence lends support to the robustness of our approach for

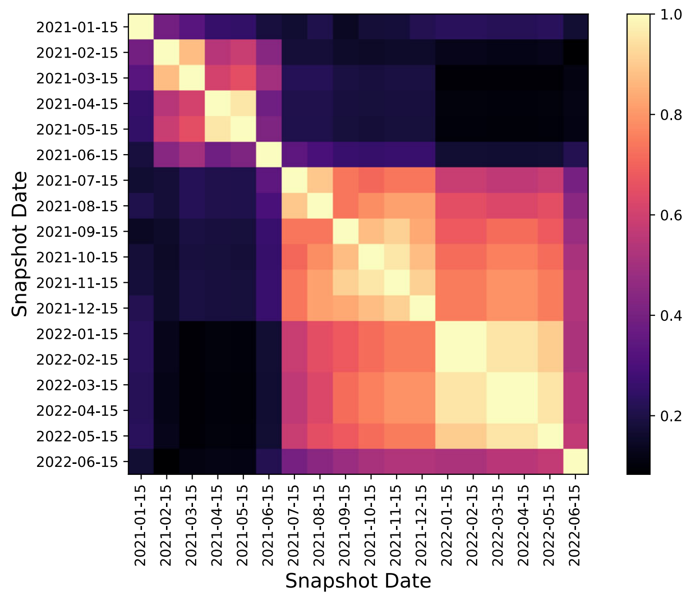

Figure 2 Jaccard similarity of the SVNs over time

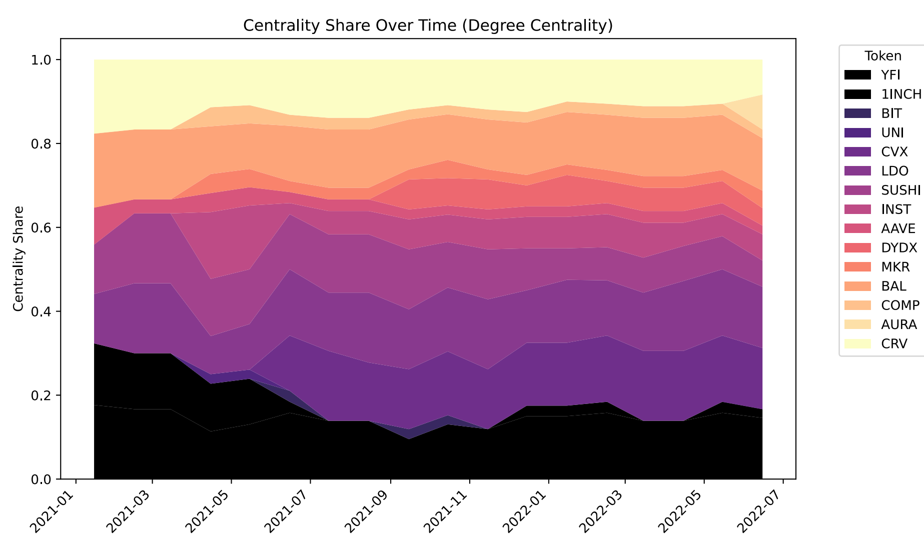

Figure 3 Degree Centrality Share of Tokens in SVNs over time

identifying and analysing these links and suggests that subsets of addresses contributing to the links are likely the same over time.

To better understand the evolving structure of cross-protocol governance networks, we analyse degree and betweenness centrality, which respectively quantify the amount of tokens co-held with a given token in users' wallets, and the extent to which a token acts as a bridge between different communities of users.

Figure 3 illustrates the relative degree centrality share of the sample tokens over time within the token-token governance network.

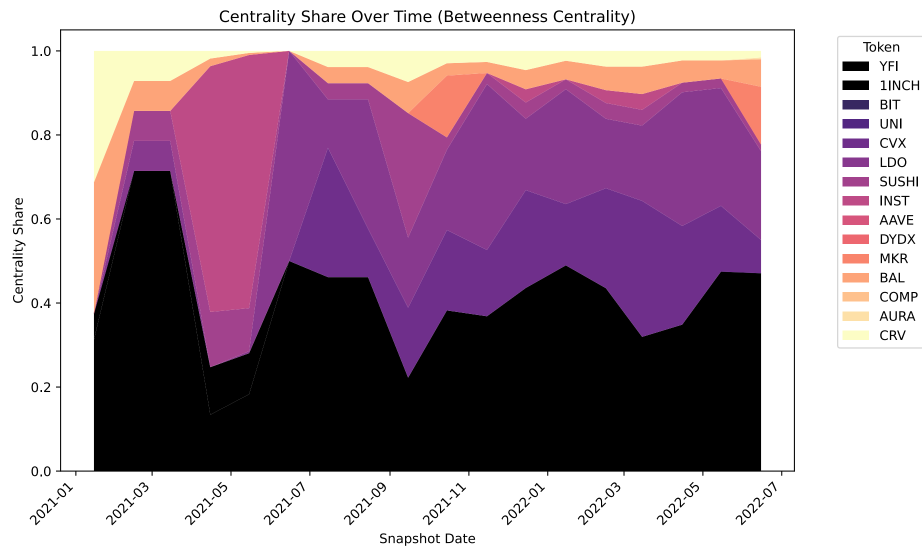

Figure 4 Betweenness Centrality Share of Tokens in SVNs over time

We observe that throughout the analysed period, tokens such as CRV, COMP, and BAL maintain consistently high degree centrality shares, indicating their broad presence in diverse portfolios. This suggests that these tokens are widely co-held with other governance assets, potentially reflecting their role as foundational DeFi primitives or popular collateral assets.

Tokens such as LDO, CVX, and UNI also show steady increases in centrality share over time, reflecting the growing prominence of the associated protocols in yield-bearing strategies.

Taken together, the degree centrality distribution suggests that while the network of governance tokens is increasingly integrated, the topology is shaped by a set of consistently central assets, whose influence is derived from widespread adoption rather than episodic bridging activity.

Transitioning to the relative betweenness centrality, in Fig. 4 we present the share of the sample tokens over time within the token-token governance network.

The distribution of betweenness centrality shares over time reveals shifting patterns in the structural influence of governance tokens within the cross-protocol holder network. In particular, YFI emerges as a consistently dominant bridge, occupying a central role throughout the observed period. This high centrality throughout the sample suggests that addresses holding YFI often connected otherwise disjoint groups of token holders, facilitating coordination or influence across protocol boundaries. Additionally, we can see that in mid-2021, CVX and LDO display a marked rise in centrality. Contrasting with its centrality in the DeFi infrastructure, CRV remains on the periphery in terms of betweenness centrality. This may reflect the fact that CRV governance participation is dominated by large, concentrated holders with limited overlap with other protocol token communities, reducing its role as a connector in the token-holder network.

At the same time, the overall distribution of centrality becomes more concentrated, pointing to a trend of increasing concentration of shared holding pairs. While a small number of tokens accumulate growing intermediary power, others exhibit short-lived surges in centrality, such as INST. These spikes, however, are not sustained, suggesting that the influence of these tokens was episodic rather than structural.

4.2 Link analysis

Building on the persistence of links observed in the SVNs, we now examine the nature of these validated links to understand the characteristics, behaviour and potential control exercised by the set of addresses constituting the links. In Fig. 5, we present the link size counting the addresses that make up the validated links, aggregated over the sampling period. Each row represents a different pair of tokens, with the x-axis indicating the address count.

The size of links across different token pairs shows high variability ranging from 1 to 320. Some token pairs show very compact distributions indicating the consistent size of $\mathcal{A}_{i,j}$ over the sampling period, such as ['SUSHI'-BAL'], ['YFI'-'BAL'], ['YFI'-'AAVE'], ['BAL'-'CRV'] and['CVX'-'SUSHI']. Other links have wide ranges, such as ['YFI'-'SUSHI'], ['YFI'-'CRV'], ['LDO'-'BAL'], ['YFI'-'SUSHI'] and ['CVX'-'CRV'], suggesting variability in how many addresses are part of the link. Regardless, most validated links are defined by a small number of addresses relative to the overall addresses in the sample (see Fig. 13, implying that a few entities repeatedly connect multiple governance communities, providing a first indication of concentrated cross-protocol exposure. In particular, the links that exhibit compact size distributions—such as YFI–BAL, YFI–AAVE, and BAL–CRV—are the same links that later show high average token holding share and strong null-model significance

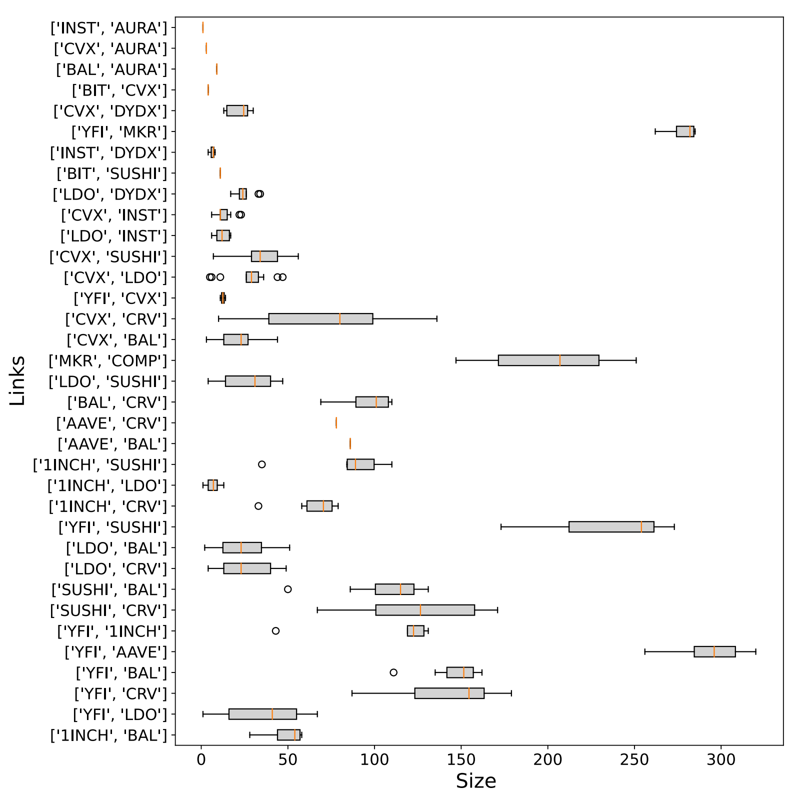

Figure 5 Distribution of the number of addresses across the sampling period that define validated links between governance tokens in validated token projections

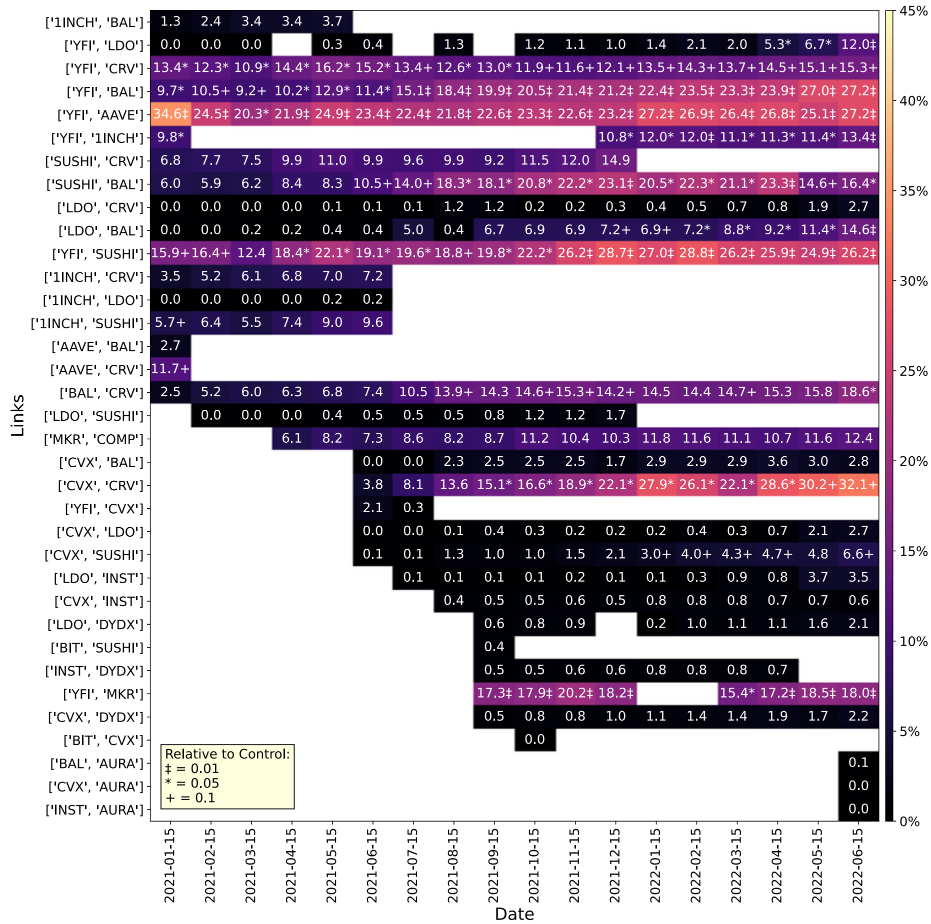

Figure 6 Influence exerted by addresses within token links

(see Fig. 6 and Appendix E). This consistency across metrics highlights a small set of token pairs that repeatedly define the network's structural core.

Building on the small link sizes observed above, we next ask how much influence these few addresses represent by measuring their average token holding share. The Average Token Holding Share (Equation (4)) provides a proxy of the potential governance control that exists between link-defining addresses within the token links they are part of.

In Fig. 6 we analyse the *average token holding share*. We show that across the identified links, the ranges from close to 0.004% to up to 34.4%.

Several links - ['YFI'-'CRV'], ['YFI'-'BAL'], ['YFI'-'AAVE'], ['SUSHI'-'BAL'], ['YFI'-'SUSHI'] and ['CVX'-'CRV'] - exhibit repeated significance across time under both validation procedures: a permutation test against size-matched random address sets and a degree-preserving null model that reshuffles token–holder associations while maintaining each token's holder count and total supply distribution (see Appendix E). These links, on average, also tend to have higher average token holding share, suggesting a potential relationship between persistent validation and concentrated governance control. This may suggest that in these links, the defining addresses may hold a non-negligible proportion of governance tokens in both tokens, implying potential governance control.

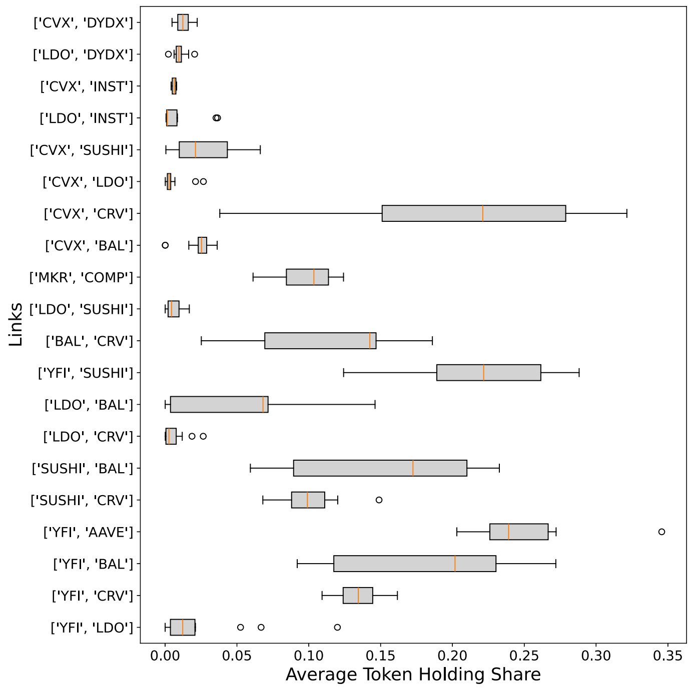

Figure 7 A collapsed time view of Average Token Holding Share metrics across token links, filtered for links that occurred at least nine times

In Fig. 7 we look at the distribution of *average token holding share* for token pairs that occur at minimum nine times across the validated projections representing half of the sampled time period. The goal of the filtering is to focus on the most persistent links in our sample and study the degree of holder overlap between them. High values, as in ['YFI'-'CRV'], ['YFI'-'BAL'], ['YFI'-'AAVE'], ['SUSHI'-'BAL'], ['LDO'-'BAL'], ['YFI'-'SUSHI'], ['CVX'-'CRV'] and ['BAL'-'CRV'], indicate substantial shared ownership, with mean overlaps exceeding 10% in some cases. We should note that this measure reflects shared token custody, not necessarily shared governance power, as many holdings may reside in smart contracts or wrappers without voting rights. The variation across pairs may highlight differing levels of interconnection between protocols, with more stable overlaps (e.g.,['YFI', 'AAVE']) suggesting more distinct holder and stable communities.

To assess the degree of inequality within validated links, we compute the Gini coefficient of the *average token holding share* of addresses. Following the same filtering procedure used in the previous analysis, in Fig. 8 we collapse the time dimension by retaining only those links that appear in at least nine snapshots with the rationale to highlight the most persistent links.

We find that the distribution of average token holding share within validated links tends to be unequal with most links achieving a Gini coefficient above 0.7. This highlights that even within links, the distribution of average token holding share is unequal, with a few meaningful addresses holding a disproportionately large share of link-defining tokens relative to the other addresses constituting the link. The null model (Appendix E), while randomizing token–holder associations, consistently produces similarly unequal distributions, suggesting that concentration is an inherent property of token holdings.

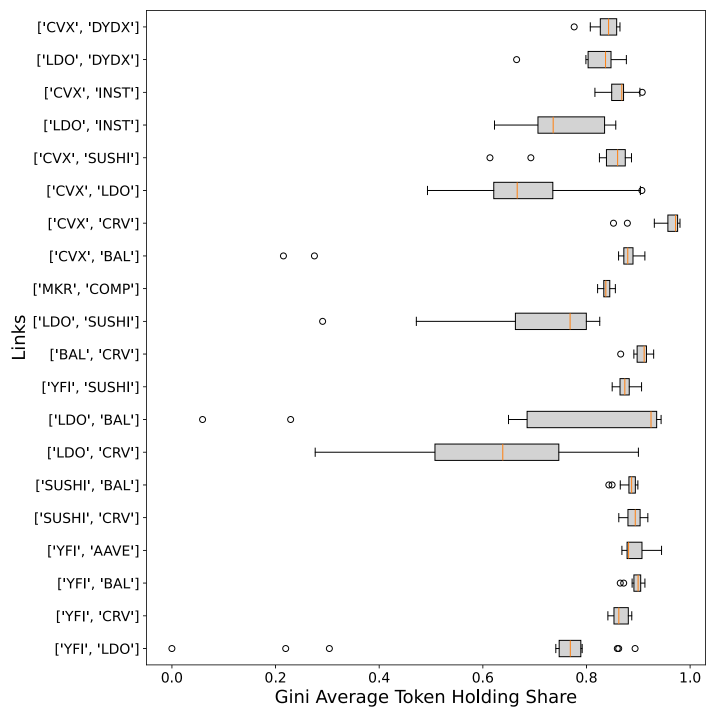

Figure 8 Gini coefficient showing the inequality of influence within token links occurring nine or more times

We now look into the median wealth of link-defining addresses for links, continuing to focus on the most persistent links. Figure 9 shows the distribution of *median wealth* (in millions of USD) for addresses holding both tokens of each pair. Wealth is estimated accounting for the value of all tokens in the sample held by these addresses. This serves as a proxy for the typical financial wealth in these wallets. Noteworthy, is that addresses have relative high median holding across addresses. Even links with smaller median wealth such as [‘CVX’-‘CRV’], [‘YFI’-‘AAVE’], [‘YFI’-‘BAL’], [‘YFI’-‘SUSHI’], [‘CVX’-‘SUSHI’] still suggests wallet holding in the six figure values.

We incorporate address labelling and analyse the relative average token holding share of link-defining addresses by entity type (as defined in Equation (6)). Figure 10 shows the average token share in tokens of a given pair held by entities with a given label. We report this statistic as it is a useful complement to the other concentration metrics to provide insight regarding the types of users that hold control across platforms.

Several patterns emerge. First, a large number of all token pairs is dominated by addresses linked to centralised exchanges. These are not directly governance relevant as they custody fund on behalf of users. This also suggests that these tokens imply no active governance participation from their holders (as delegation through exchanges is generally not available). Second, the overlap between [‘CVX’, ‘CRV’] is heavily dominated by smart contracts (87.8%), consistent with the fact that many Convex-controlled CRV positions are routed through staking and reward wrappers necessitating a deeper analysis of depositor to the underlying contracts to understand the effective governance control and entity distribution. [‘BAL’, ‘CRV’] and [‘LDO’, ‘BAL’] links indicate operational or liquidity integration, without a direct channel to governance as the majority of funds within DEX contract. Similarly, the [‘YFI’, ‘SUSHI’], [‘YFI’, ‘AAVE’], [‘YFI’, ‘BAL’], [‘YFI’, ‘CRV’] link provides

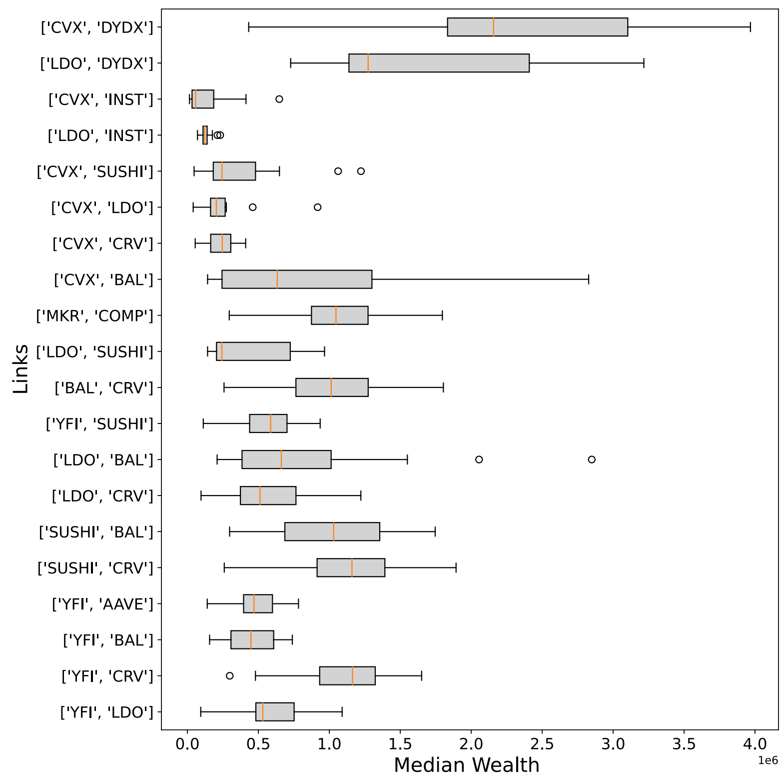

Figure 9 Median wealth held by link-defining addresses across all tokens in the sample

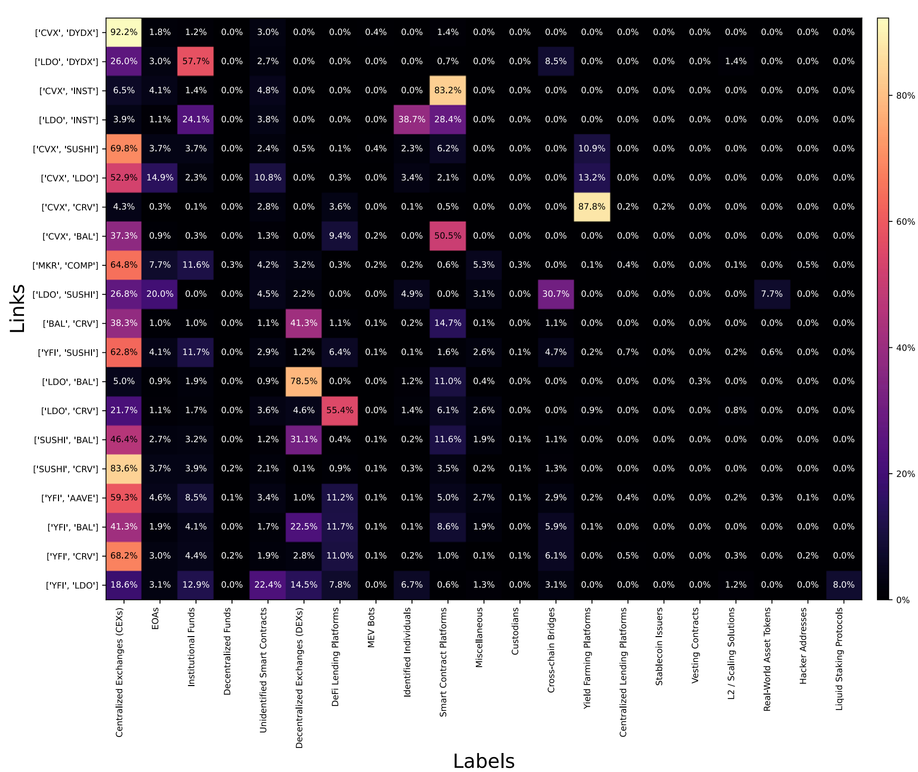

Figure 10 Relative Average Token Holding Share of link-defining addresses by entity type

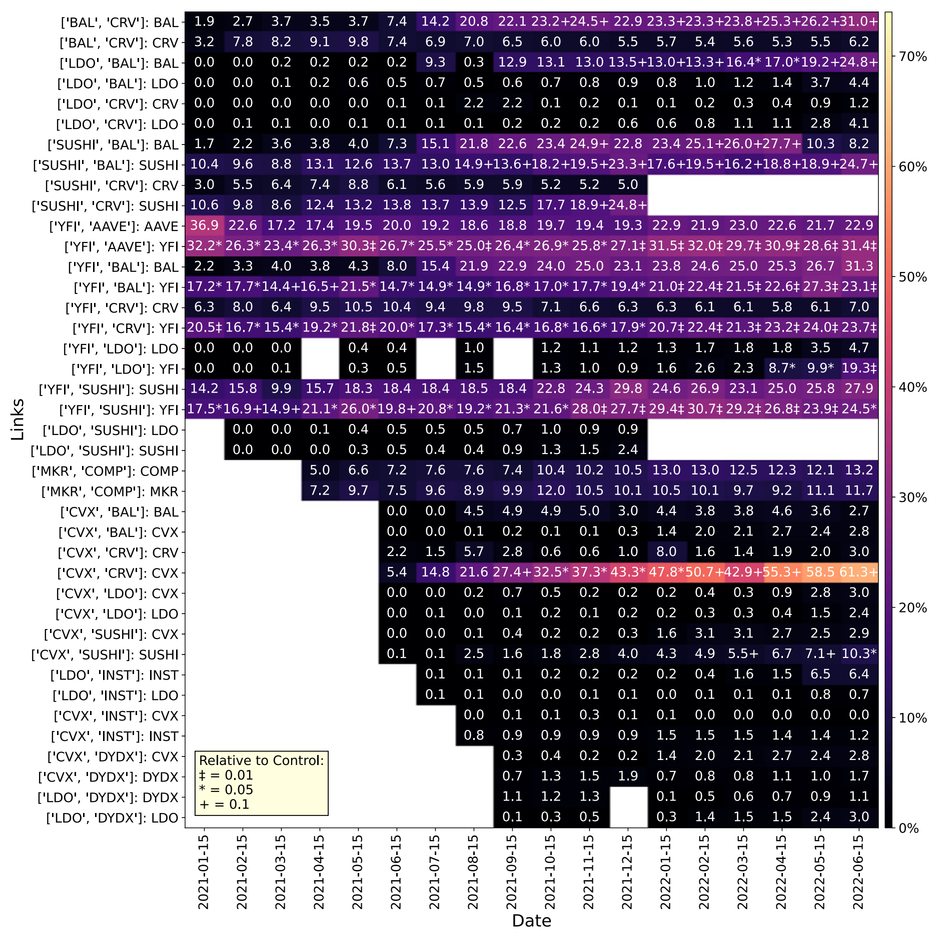

Figure 11 Directional Token Holding Share within token pairs in validated links on the respective token community

evidence of integration via lending infrastructure with around 10%. Pair such as ['YFI'-'AAVE'], ['YFI'-'BAL'] and ['YFI'-'SUSHI'], which displayed high Average Token Holding share, share shared of institutional funds with around 10%. Given that large shares are locked in governance irrelevant labels (such as CEX) it may indicate concentration into of substantial influence by a few powerful entities. These findings underscore that not all identified links are equal. Both the *type* of shared holder and their *economic profile* matter for assessing cross-protocol governance entanglement and resilience.

The asymmetry of the token holding within validated links is analysed with the Directional Token Holding Share (Equation (5)). We report these statistics in Fig. 11. The directional analysis refines these same persistent links by revealing which side of each pair holds greater weight. Thus, rather than introducing new relationships, this section explains the directionality of influence within the already-identified core pairs.

The most striking asymmetry occurs in when comparing ['CVX'-'CRV']: CRV to ['CVX'-'CRV']: CVX. From mid-2021 the Convex-Curve link-defining addresses set owns a much larger share of CVX tokens than of CRV tokens, reaching around 60% by June 2022. Because the vast majority of these link-defining addresses are smart contracts classified as

Yield Farming Platforms (87.8%) (see Fig. 10), their voting power depends on whether the contracts are whitelisted for governance within Convex.

As we saw before, YFI was identified as a central token in the SVN. When analysing directional holdings, we can notice that YFI pairs present significant imbalance (see for instance ['YFI'-'AAVE']: YFI vs. ['YFI'-'AAVE']: AAVE, ['YFI'-'SUSHI']: YFI vs. ['YFI'-'SUSHI']: SUSHI, ['YFI'-'CRV']: CRV vs. ['YFI'-'CRV']: YFI). We attribute this result to two factors. First, YFI was launched without pre-mining or allocations to insiders, so early liquidity providers across the DeFi landscape could acquire it on equal terms.[3] Secondly, Yearn's specialised vaults such as the "Backscratcher" vault incentivised CRV depositors – and, by extension, SUSHI liquidity farmers – to interact directly with Yearn. Those strategies may have pulled existing DeFi participants into the YFI holder set. The net result is a broadly asymmetric holding pattern of directional average holding share.

A similar, though smaller, imbalance appears in ['LDO'-'BAL']: BAL, ['BAL'-'CRV']: BAL and ['BAL'-'SUSHI']: BAL. Each link shows growth in directional influence around July 2021, which coincides shortly after with the launch of Balancer V2 in May 2021. As the majority of addresses are concentrated in DEX addresses – which include Balancer itself – a potential explanation lies in the Balancer 80:20 vault, leading to this asymmetric pairing. 80:20 are special AMMs which allow liquidity pools to be heavily weighted toward one asset (e.g. 80% BAL, 20% paired token).

Along with these asymmetries, we find that a small subset of token pairs display near-parity in their directional token-holding shares, as with the pairs ['MKR'-'COMP'], ['YFI'-'AAVE'] and ['YFI'-'SUSHI'], with differences below two percentage points in June 2022. The overlapping supply in these links is held by in two-diget by wallets labelled "Institutional Funds", with the remainder split largely with EOAs and other contracts. Such symmetry may imply a common investor cohort that allocates capital to both protocols in comparable size. This may create the pre-conditions for *coordinated* governance behaviour.

Finally, the sampling period coincides with two major sell-offs followed by renewed accumulation, namely around mid-May 2021 and mid-April 2022 (see Appendix C). With this information, we explore whether changes in link token holding share correlate with decline in market-wide dollar-denominated TVL movements to gain insights into potential holding behaviour.

In Fig. 12, we present the correlation between the percentage change in TVL and Average Token Holding Share across various link pairs over the sampling period, reflecting the strength of these correlations, with each horizontal line representing a specific link pair. The line's length reflects the correlation's magnitude, while the dot at the end of each line indicates the statistical significance of the coefficient.

We find that only a few correlations between Average Token Holding Share and changes in TVL achieve statistical significance with a negative coefficient (see Appendix, Table 5, adjusted for autocorrelation). This suggests that, in some cases, as speculative interest increases, proxied by rising TVL on Ethereum, the relative token holding share may decline. Conversely, during periods of declining TVL, the relative Token Holding Share of these addresses may increase, potentially re-concentrating governance power among long-term participants. Although limited in number, significant cases such as YFI-LDO highlight a

[3] https://docs.yearn.fi/resources/defi-glossary/, entry *Fair Launch*.

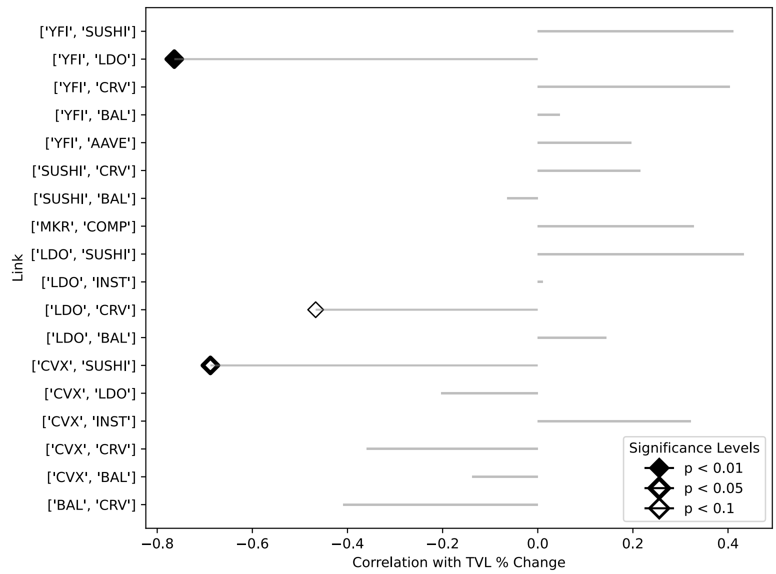

Figure 12 Correlation between Average Token Holding Share of Link Pairs and Percentage Change in Market-wide TVL (adjusted for auto-correlation by differencing)

potential vulnerability in open governance models like DAOs: during speculative inflows, governance may become more susceptible to short-term dynamics driven by newly arriving token holders.

5 Discussion

In this work, we explored how Statistically Validated Networks can offer valuable insight in the distribution of governance power in DeFi protocols. Our analysis reveals that governance token co-holdings across DeFi protocols are shaped by persistent, statistically validated relationships between token communities. These relationships, captured as links in SVNs, form a slowly evolving network topology, characterized by high clustering, shifting modular structure, and a consistently large connected component.

In several cases, such as with [‘CVX’, ‘CRV’] or [‘BAL’, ‘CRV’], a small set of addresses repeatedly co-hold substantial portions of both tokens, although our Gini coefficient analysis shows that this power is typically concentrated among a few entities. These patterns often align with known protocol integrations (e.g., Convex’s stake in Curve) or infrastructure dependencies (e.g., DEX vaults on Balancer), where the overlap reflects functional design rather than deliberate governance capture. However, other links, such as [‘YFI’, ‘SUSHI’], [‘YFI’-‘AAVE’], show high median wealth and a relatively large presence of institutional addresses (e.g., 11.7%, 8.5% respectively), raising the possibility that capitalised actors may possess substantial governance influence across both protocols.

While we do not observe widespread or universal governance centralisation across all links in our sample, the asymmetric distribution of token holding shares in several pairs indicates a directional dependence where addresses of one token are disproportionately embedded (structurally or through shared investor base) in another community. This may create incentives for coordinated behaviour on a technical or social level.

These findings align with previous work on governance centralisation within individual protocols (e.g. Fritsch et al. [31], Nadler and Schär [48]), which highlighted the concentration of token supply among a few influential entities. We expand this perspective by demonstrating how co-holding and potential influence extend across protocol boundaries.

We explored whether governance exposure fluctuates with market dynamics. Our results show only limited statistically significant correlations between market-wide TVL changes and average token holding share. However, in some links, we observed negative correlations, suggesting that governance exposure may dilute as speculative capital enter. This could reduce the relative influence of long-term more aligned holders during periods of appreciating prices.

Our findings suggest that entanglement in DeFi is not uniformly problematic but context-dependent, with structural, compositional, and directional factors all playing a role. A nuanced understanding of these dynamics will be critical as DeFi protocols increasingly rely on governance mechanisms to manage risk, incentives, and inter-protocol relationships as they decentralise.

We now point to some limitations that should be considered when interpreting these findings and that may be addressed by future research. Firstly, while our analysis measures concentration of token holdings across validated co-holding links, it does not differentiate eligibility to exercise votes of the tokens held in externally owned accounts (EOAs) and those held in smart contracts. In most DAOs, tokens held in smart contracts are unable to vote directly unless delegated, meaning that ownership concentration does not automatically imply governance power centralisation. Certain protocols in our sample such as YFI allowed multiple pre-approved contracts wherein users retained their voting power even after depositing tokens in them. A related limitation is that token holdings do not uniformly translate into governance power across protocols, as governance structures, voting mechanisms, and rights associated with tokens can differ substantially. While we provide a high-level overview of these differences in Appendix G, a more granular, protocol-specific analysis would be necessary to accurately assess the actual extent of control that token holders may exercise within each governance system. Another limitation, common to studies using raw blockchain data, is that blockchain addresses are inherently pseudonymous, and although we employed a structured approach to entity labelling, some misclassification may have occurred. Inaccuracies in entity categorisation (e.g., distinguishing EOAs from institutions or protocols) may distort the perceived distribution of Token Holding Share. Future work could benefit from the integration of semi-supervised address classification techniques (Béres et al. [15], Valadares et al. [58]) and clustering methods (Victor [59]) to improve the precision of the labels. Finally, our findings reveal that link-defining addresses span both entities and smart contracts, indicating two distinct types of interdependence: economic entanglement via smart contracts and institutional cross-holdings across protocols. Future research could explore whether these connections amplify systemic risk or foster coordination, drawing on insights from the literature on common ownership in traditional finance (Azar et al. [11], He and Huang [37]).

6 Conclusion

In this paper we explored the understudied risk vector of cross-protocol governance control in open token-based voting systems by analysing the characteristics of addresses that form statistically validated links between DeFi governance tokens. First, we find that influential links are often small in address count but may hold substantial token shares in both protocols. Second, these addresses are frequently institutional or contract-based, as reflected in label composition, median wealth, and token concentration. Third, many links exhibit asymmetry, where token holding is concentrated on one side of the token pair, potentially creating directional influence or structural dependency. Finally, we observe that core holder share may be diluted during periods of TVL expansion, suggesting that governance influence can shift with market cycles. These findings emphasise the importance of understanding not only protocol-internal token distribution, but also cross-protocol entanglements.

Appendix A: Block heights and snapshot dates

Table 4 Snapshot dates and corresponding Ethereum block heights for data collection, covering monthly intervals from January 15, 2021, to June 15, 2022

Block Height	Snapshot Date
11659570	2021-01-15
11861210	2021-02-15
12043054	2021-03-15
12244515	2021-04-15
12438842	2021-05-15
12638919	2021-06-15
12831436	2021-07-15
13029639	2021-08-15
13230157	2021-09-15
13422506	2021-10-15
13620205	2021-11-15
13809597	2021-12-15
14009885	2022-01-15
14210564	2022-02-15
14391029	2022-03-15
14589816	2022-04-15
14779829	2022-05-15
14967365	2022-06-15

Appendix B: Token ownership over time

Figure 13 shows the absolute number of addresses holding governance tokens over time. The introduction of new tokens after June 2021 contributes to the structural break observed in the Jaccard Similarity Matrix (Fig. 2).

Appendix C: TVL changes over the sampling period

Figure 14 shows the dollar-denominated Total Value Locked (TVL) in Ethereum DeFi protocols. Since the TVL is measured in dollars, a decline reflects the depreciation in the value of cryptoassets in Dollar more broadly, next to asset leaving the protocol. We retrieved the data for TVL from DefiLlama.[4]

[4] We use the following endpoint for this: https://api.llama.fi/v2/historicalChainTvl/Ethereum.

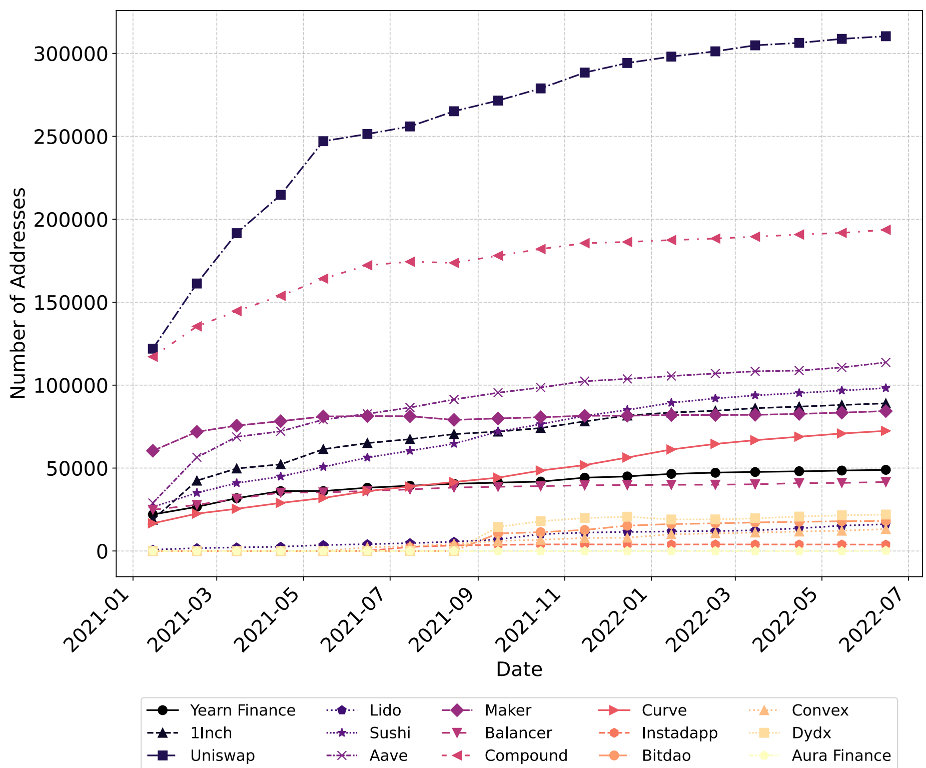

Figure 13 Absolute number of addresses holding each governance token over time

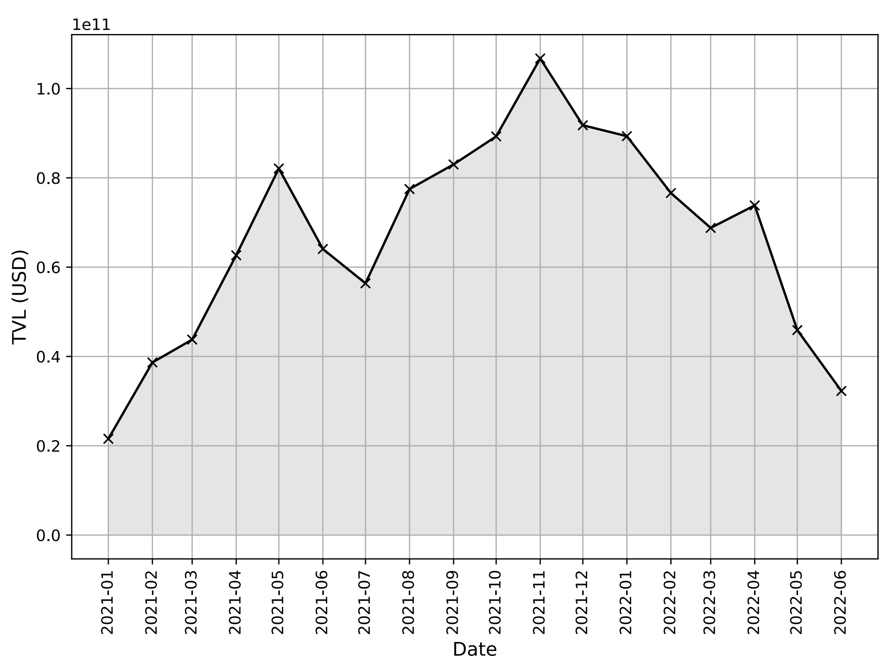

Figure 14 The figure shows the dollar-denominated total value locked (TVL) in Ethereum-based DeFi protocols over time. The chart illustrates short-term sell-offs and accumulations, with significant declines in May 2021 and April 2022, corresponding to broader market sell-offs

Appendix D: Summary statistics for correlation and autocorrelation analysis

Table 5 Summary Statistics for Correlation and Autocorrelation Analysis

Link Name	Correlation	P-value	Durbin-Watson	Ljung-Box P-value
BAL-CRV	−0.4095	0.1026	0.7388	0.8041
CVX-BAL	−0.1387	0.6674	2.1857	0.2479
CVX-CRV	−0.3596	0.2510	1.2972	0.5627
CVX-INST	0.3227	0.3631	2.3978	0.5616
CVX-LDO	−0.2033	0.5263	0.8803	0.6700
CVX-SUSHI	−0.6882	0.0134	1.0180	0.5627
LDO-BAL	0.1443	0.5806	2.6991	0.2802
LDO-CRV	−0.4669	0.0588	1.3538	0.5478
LDO-INST	0.0113	0.9737	2.0899	0.9814
LDO-SUSHI	0.4333	0.2110	1.0335	0.6265
MKR-COMP	0.3286	0.2514	2.1958	0.4494
SUSHI-BAL	−0.0649	0.8045	2.2561	0.2765
SUSHI-CRV	0.2158	0.5239	1.4347	0.7205
YFI-AAVE	0.1974	0.4475	0.9116	0.8980
YFI-BAL	0.0469	0.8581	1.5453	0.3596
YFI-CRV	0.4047	0.1071	1.5802	0.1811
YFI-LDO	−0.7641	0.0062	0.7247	0.3374
YFI-SUSHI	0.4116	0.1007	2.1141	0.8598

Appendix E: Null model significance persistence

To assess the robustness of observed cross-protocol concentration patterns, we apply a null-model test to each validated link across all snapshots. For each token pair, token ownership is randomly permuted 10,000 times while preserving each token's supply and address participation. The resulting empirical distributions define the expected metric values under random co-holding. We then compute one-sided p-values to identify links where observed values exceed the 95th percentile of this null distribution.

Tables 6–10 report, for each metric, the share of snapshots in which a link remains statistically significant (p smaller or equal to 0.05), alongside the median observed and null (p95) values. Persistently significant links (e.g., CVX–CRV, YFI–BAL, SUSHI–BAL) indicate enduring cross-protocol control structures unlikely to arise from random portfolio overlap.

Table 6 Appendix – Average Token Holding Share (significance persistence; min. sig. occurrences = 9). A link is counted as significant in a snapshot if the observed value exceeds the 95th percentile of the null model (empirical p(0.05))

Link	N	Sig.	% Sig.	Median Obs.	Median Null p95
YFI-AAVE	18	18	100.0	0.2499	0.1726
YFI-BAL	18	17	94.4	0.2194	0.1153
YFI-SUSHI	18	12	66.7	0.2317	0.1904
CVX-CRV	13	9	69.2	0.2371	0.2105
SUSHI-BAL	18	9	50.0	0.1992	0.1865

Table 7 Appendix – Directional Token Holding Share (token A) (significance persistence; min. sig. occurrences = 9). A link is counted as significant in a snapshot if the observed value exceeds the 95th percentile of the null model (empirical p(0.05))

Link	N	Sig.	% Sig.	Median Obs.	Median Null p95
YFI-AAVE	18	18	100.0	0.2848	0.2424
YFI-BAL	18	18	100.0	0.235	0.1428
YFI-CRV	18	18	100.0	0.2108	0.1687
YFI-SUSHI	18	15	83.3	0.253	0.2072
BAL-CRV	18	12	66.7	0.2334	0.1625
CVX-CRV	13	9	69.2	0.4442	0.4209
SUSHI-BAL	18	9	50.0	0.1611	0.1755

Table 8 Appendix – Directional Token Holding Share (token B) (significance persistence; min. sig. occurrences = 9). A link is counted as significant in a snapshot if the observed value exceeds the 95th percentile of the null model (empirical p(0.05))

Link	N	Sig.	% Sig.	Median Obs.	Median Null p95
YFI-BAL	18	12	66.7	0.2325	0.1561
LDO-BAL	18	12	66.7	0.155	0.0408
SUSHI-BAL	18	11	61.1	0.2305	0.2114
YFI-AAVE	18	11	61.1	0.207	0.2043
CVX-SUSHI	13	9	69.2	0.043	0.0295

Table 9 Appendix – Median Wealth of Link (significance persistence; min. sig. occurrences = 9). A link is counted as significant in a snapshot if the observed value exceeds the 95th percentile of the null model (empirical p(0.05))

Link	N	Sig.	% Sig.	Median Obs.	Median Null p95
YFI-AAVE	18	18	100.0	1198.848	965.2962
YFI-CRV	18	18	100.0	980.2129	640.5286
BAL-CRV	18	18	100.0	389.527	322.3645
LDO-BAL	18	9	50.0	766.3476	878.3195

Table 10 Appendix – Token Holding Share Inequality (Gini) (significance persistence; min. sig. occurrences = 9). A link is counted as significant in a snapshot if the observed value exceeds the 95th percentile of the null model (empirical p(0.05))

Link	N	Sig.	% Sig.	Median Obs.	Median Null p95

Appendix F: Sensitivity analysis

We conducted the sensitivity analysis on the supply threshold for the inclusion of link-defining wallets. In this analysis, we computed the minimum, maximum, and interquartile range. The threshold level throughout the analysis is 5e-06, meaning we consider all addresses that hold tokens with at least 0.0005% of the available supply of any token in our sample. The sensitivity analysis, depicted in Fig. 15, shows that results tend to be relatively stable for the metrics utilised in this research.

Appendix G: Overview governance token

Governance tokens and the rights associated with them are unique to each protocol, giving varying degrees of control and responsibility to token holders. To provide the reader with an idea of what governance rights entail we reviewed the protocol documentation.

Uniswap: Uniswap is a decentralised exchange protocol. It is governed by UNI token holders. The holder generally delegates to professional participants (either companies,

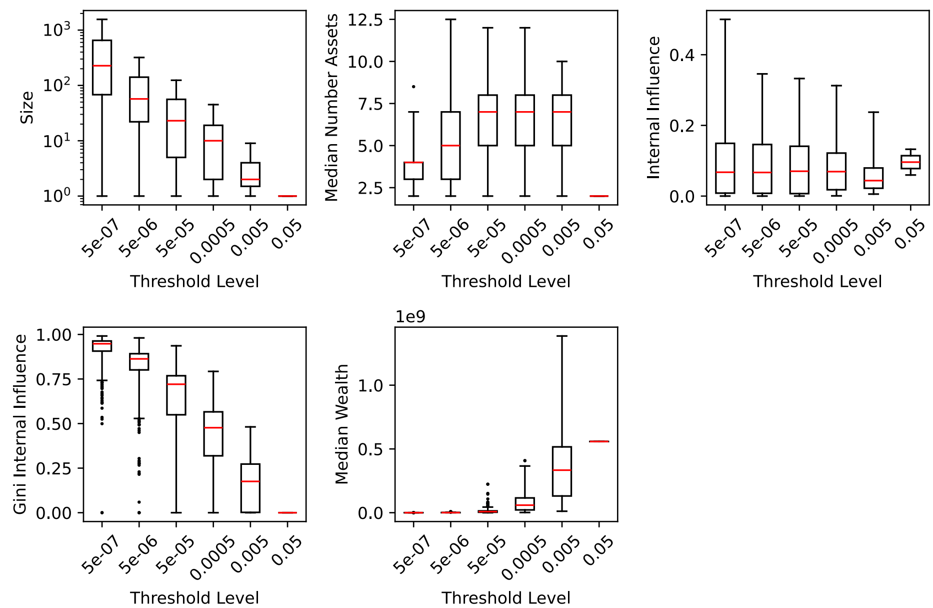

Figure 15 Sensitivity analysis of the inclusion threshold for link-defining wallets based on relative percentage supply held. The analysis evaluates the stability of the results across varying thresholds, highlighting that metrics for the chosen threshold of 0.0005% of the token supply are relatively insensitive to higher lower thresholds

university clubs or individuals) that are supposed to represent their vote. Decisions revolve around a broad range of topics such as fee parameters (Adams et al. [4]), delegation to different stakeholders, new protocol initiatives and features (Uniswap [57]). The Uniswap governance process begins with proposal creation, requiring 2.5 million UNI tokens. First, a "Temperature Check" gauges community interest via a forum poll with at least 25,000 UNI tokens needed to start. If successful, a "Consensus Check" follows, requiring 50,000 UNI in a formal Snapshot vote. The final stage is the on-chain "Governance Proposal," which undergoes a 7-day vote, needing 40 million UNI to pass and be executed. This structured process ensures thorough community involvement in decision-making (Uniswap [56]).

MakerDAO: MakerDAO is a decentralised credit platform that allows users to generate DAI, a stablecoin pegged to the US dollar, by locking up collateral. MakerDao Governance scope can be summarised with five areas: Stability, which centres on the financial stability and the Dai stablecoin; Accessibility, which targets frontends and distribution; Protocol, dedicated to technical development, maintenance, and security; Support, aimed at ecosystem support through tools and activities; and Governance, which deals with the interpretation of Alignment Artifacts and the balance of powers (MakerDAO [46]). Governance occurs through both off-chain and on-chain processes, with proposals discussed in the Maker Governance Forum before being voted on-chain. MKR holders have proportional voting power based on the amount of MKR they hold. It can be delegated to different users. The MKR token also acts as a backstop mechanism if credit becomes under-collateralised pushing MKR holders for poor governance (MakerDAO [45]).

Aave: Aave is a decentralised lending protocol where users can borrow and lend assets. Its governance is managed by AAVE token holders. The scope of governance can be summarised as enabling upgrades to governance itself, operational processes, voting on asset listings and risk parameters and treasury allocation (e.g. hiring, incentives) (Aave [2]) Proposals are discussed in public forums. Once a proposal garners enough support it is formalised into Aave Improvement Proposals (AIPs) and voted on. Successful proposals are implemented to enhance the protocol's functionalities and security. The protocol also includes a staking mechanism where AAVE holders earn rewards and help secure the protocol. The Aave governance structure allows risk admins to adjust risk parameters without requiring a vote for every change (Aave [3]).

Lido: Lido is a liquid staking protocol that enables users to stake assets and remain liquid by issuing a derivative token. Governed by LDO token holders, their responsibilities include approving smart contract upgrades, managing the treasury, setting staking parameters, and overseeing node operator management—this covers the onboarding, evaluation, and potential replacement of operators. They also allocate grants for community and development projects through the Lido Ecosystem Grants Organisation. The governance process engages the community in initial discussions and feedback on the research forum, followed by a Snapshot vote requiring participation from over 5% of LDO holders. If necessary, proposals move to an on-chain vote using the Aragon framework, needing a quorum of 5% and 50% approval to pass. In urgent situations, LDO holders can bypass standard procedures to make swift decisions (Lido [43]).

Yearn Finance: Yearn Finance is a yield optimisation protocol that aggregates yields from various DeFi platforms. It operates under a multi-DAO structure managed by constrained delegation. YFI holders are primarily responsible for proposing, discussing, and implementing changes through three main types of proposals: Yearn Improvement Proposals (YIPs), are formal proposals that can execute any power delegated to YFI holders or address issues outside the predetermined scope of powers; Yearn Delegation Proposals (YDPs), this type allows for the redistribution of decision-making powers among different operational teams within the ecosystem; Yearn Signalling Proposals (YSPs), are used to gauge community sentiment on various topics and are non-binding. YFI holders can: Manage and reallocate discrete powers among yTeams, specialised groups focusing on different aspects of the protocol, interact with the YFI token contract, including actions like minting new tokens or burning existing ones, setting and adjusting fee structures across the Yearn Protocol, select or change signatories of the multisig wallet, which holds significant operational powers including the execution or vetoing of on-chain decisions, ratify or deratify yTeams, thereby influencing which teams hold delegated powers, allocate and manage funds from Yearn's treasury (Yearn Finance [63]).

The governance flow typically involves yTeams proposing decisions to a transactional team (yTx), which then creates delegated transactions sent to the Multisig for execution or veto. This structured yet dynamic governance framework allows YFI holders to effectively oversee and direct the continuous development of the Yearn protocol (Yearn Finance [64]).

Aura Finance: Aura Finance is a protocol that aims to increase yields on Balancer liquidity pools. AURA token holders can lock their tokens to obtain veAURA, enhancing both their voting power and rewards, with incentives structured to favour long-term engagement up to 16 weeks. The voting power is proportionate to the amount of veAURA held (Aura Finance [9]). Governance participation is enabled by holding veAURA, which grants the right to propose and vote fee adjustments, tokenomics, and strategic initiatives (Aura Finance [8])

BitDAO: BitDAO is a decentralised autonomous organisation that supports DeFi projects through grants and partnerships, it now pivoted to become Mantle a product suite offering different decentralised technologies. BIT token holders engage in governance via delegated voting, allowing them to vote on proposals or delegate their voting power to others. Governance activities are primarily conducted through the Snapshot platform, ensuring transparency and efficiency by aggregating votes off-chain with potential future shifts to on-chain mechanisms. The proposal can cover a wide range of decisions from operational changes, and treasury management, to strategic initiatives, and requires sufficient community backing to be formalised into official votes. A multi-sig wallet executes these decisions (BitDAO [16]).

SushiSwap: SushiSwap is primarily a decentralised exchange protocol. It offers additional financial tools such as yield-generating instruments, bonds, and a platform for token streaming and vesting. SushiSwap's governance is powered by its community through a combination of forum discussions and Snapshot voting. Proposals can be submitted by any community member and, if they gather sufficient interest, are formalised through Snapshot where they are voted upon. The SUSHI token gives the holder voting power only if it is staked within a protocol-specific staking contract or deposited in the SUSHI-ETH pool. The process is underpinned by a multi-signature mechanism that involves prominent members from the DeFi community who execute or veto decisions based on the collective voting outcomes (SushiSwap [54]).

dYdX: The dYdX is a decentralised derivatives exchange, since the study period has concluded the protocol has moved from Ethereum. DYDX token holders were responsible for managing proposals affecting both strategic and operational aspects of the protocol, including key protocol amendments, liquidity, safety modules, and reward distributions. Proposals, categorised by their impact, pass through a lifecycle involving community discussion, off-chain drafting, on-chain voting, and execution via time-locked contracts. DYDX holders can delegate voting power on their behalf (dYdX [27]).

Instadapp: Instadapp is a middleware that aggregates multiple DeFi protocols into one upgradable smart contract layer. It uses the INST token for governance using a similar structure to Compound thus allowing delegation. Token holders can propose changes to the protocol, vote on upgrades, and influence the management of community treasury funds (Jain [39]). Governance starts with discussions in community forums, followed by an off-chain vote via snapshot followed by an on-chain vote via atlas (Instadapp [38]).

Curve: Curve is a decentralised exchange protocol, stablecoin provider, and lending platform on Ethereum and EVM-compatible chains. Curve's governance is managed by CRV token holders who can lock their tokens to receive veCRV, which grants voting power. Users need at least 2500 veCRV to create proposals, while anyone can vote with no minimum required. Proposals have a voting duration of seven days, and voting power decays linearly over time. Curve utilises three types of votes: ownership votes, parameter votes, and emergency votes, each with specific quorum requirements. The EmergencyDAO, a group of trusted agents, can intervene in critical situations, such as shutting down liquidity pools or gauges. The scope of governance ranges from incentive allocation on gauges to allocating future CRV emissions to liquidity pools, protocol upgrades and allocation of treasury and management (Curve Finance [26]).

Convex: Convex Finance is built on top of different DeFi protocols utilising a locked voting mechanism which enables the allocation of future token emissions (e.g. Curve, Frax

Finance). It allows liquidity providers and stakers to earn boosted rewards without needing to lock their tokens directly in the underlying protocols. By aggregating user stakes, Convex optimizes yield for users while offering its own native token (CVX) as an additional incentive. CVX, in addition, is used to vote on proposals. To participate in governance voting, users must lock their CVX tokens for a minimum of 16 weeks, creating vote-locked CVX with governance power, similar to Curve. This lock-in period aligns with the protocol's weekly epoch cycles. Locked CVX also accumulate rewards based on the protocol's earnings, providing financial incentives alongside governance influence. Convex's governance is further reinforced by a multi-sig that ensures that all CVX votes align with the interests of the protocols involved, blocking any harmful proposals (Convex Finance [25]).

Balancer: Balancer is a decentralised exchange protocol that allows for customizable liquidity pools and decentralised trading. Governance is handled through veBAL, where token holders lock BAL/WETH Balancer Pool Tokens for up to 1 year to participate in governance. Token holders can influence key decisions regarding protocol fees, and liquidity incentives by directing emissions of BAL, and other allocations of treasury resources to projects. The governance process begins with a decision on the forum, being formalised in a proposal, voted on and then implemented on-chain by a multisig arrangement of trusted members. The governance process is supplemented by an Emergency subDAO that can shut down or pause certain function-critical contracts (Balancer [12]).

Compound: Compound allows users to borrow and lend cryptocurrencies through a decentralised market. Users can earn interest on their deposits or take out loans against their crypto assets. The governance of Compound is managed by COMP token holders. COMP holders can delegate their vote or vote themselves. The governance process generally starts with a discussion on the forum, any address with more than 25,000 COMP delegates can vote on put a proposal forward. A proposal must achieve a majority of the votes with a minimum quorum of 400,000 COMP votes for approval. Approved proposals are queued in a Timelock contract for two days before implementation. The scope of governance includes adjustments to interest rates, adding or removing supported assets, protocol upgrades, and changes to the risk parameters. COMP holders have significant influence over the protocol's evolution and operational decisions, ensuring a decentralised and user-driven governance process (Compound [22]).

1Inch: 1inch operates as a decentralised exchange aggregator that optimises trades across multiple DEXes. The governance of 1inch is managed through a DAO, where 1INCH token holders can participate in governance decisions. To vote users need to either stake in the protocol's governance contract. User can increase their power by locking their token for a longer time. Votes can be delegated, both on-chain and off-chain. The scope of governance for token holders includes proposing and voting on protocol upgrades, fee changes, and feature additions (1Inch [1]).

Abbreviations

DeFi, Decentralized Finance; DLT, Distributed Ledger Technology; DAO, Decentralized Autonomous Organization; SVN, Statistically Validated Network; MCAP, Market Capitalization; TVL, Total Value Locked; DEX, Decentralized Exchange; EOA, Externally Owned Account; LP Tokens, Liquidity Provider Tokens; AIPs, Aave Improvement Proposals; YIPs, Yearn Improvement Proposals; YDPs, Yearn Delegation Proposals.

Acknowledgements

Not applicable.

Author contributions

T. E.: Conceptualization, Methodology, Software, Validation, Formal Analysis, Investigation, Data Curation, Writing - Original Draft, Visualization, Project Administration. C. C.: Conceptualization, Methodology, Resources, Writing - Reviewing and Editing, Supervision, Project Administration. C. J. T.: Conceptualization, Resources, Reviewing. A. S. T.: Conceptualization, Resources, Writing - Reviewing and Editing, Supervision, Project Administration, Funding Acquisition.

Authors' information
Not applicable.

Data availability
The dataset supporting the conclusions of this article is available in a GitHub repository at https://github.com/xm3van/research-project-erc20-governance.

Declarations

Ethics approval and consent to participate
This research did not involve human participants or animals, and no personal data were collected or processed. All data utilized were publicly available and do not contain identifying information. Consequently, no ethical approval was required for this study.

Consent for publication
We consent for this work to be published in EPJ Data Science.

Competing interests
The authors declare no competing interests.

Author details
[1]LASIGE, Faculdade de Ciências, Universidade de Lisboa, Lisboa, Portugal. [2]Institute of Finance and Technology, University College London, London, UK. [3]DLT Science Foundation, London, UK. [4]UZH Blockchain Center, University of Zurich, Zurich, Switzerland. [5]Blockchain & Distributed Ledger Technologies Group, University of Zurich, Zurich, Switzerland. [6]BRAN Lab, Network Science Institute, Northeastern University London, London, UK.

Received: 26 February 2025 Accepted: 9 December 2025 Published online: 29 December 2025

References
1. 1Inch (2024) Governance overview | 1inch documentation. https://gov.1inch.community/docs/governance/overview. Accessed 10 Oct 2024
2. Aave (2024) Governance. https://docs.aave.com/aavenomics/governance. Accessed 30 Sep 2024
3. Aave (2024) Governance process. https://docs.aave.com/governance/master/governance-process. Accessed 30 Sep 2024
4. Adams H, Zinsmeister N, Robinson M, et al (2021) Uniswap v3 core. https://uniswap.org/whitepaper-v3.pdf. Accessed 30 Sep 2024
5. Arkham (2025) Arkham intelligence tags (online resource). https://intel.arkm.com/tags. Accessed 8 Jul 2025
6. Art O (2025) Etherscan labels (online resource). https://octal.art/etherscan-labels/. Accessed 8 Jul 2025
7. Arthur WB (1994) Increasing returns and path dependence in the economy. University of Michigan Press, Ann Arbor
8. Aura Finance (2024) Governance. https://docs.aura.finance/aura/governance. Accessed 10 Oct 2024
9. Aura Finance (2024) Vote locking. https://docs.aura.finance/aura/usdaura/vote-locking. Accessed 10 Oct 2024
10. Axelsen H, Jensen JR, Ross O (2022) When is a dao decentralized? Complex Syst Inform Model Q 31:51–75. https://doi.org/10.7250/csimq.2022-31.04
11. Azar J, Schmalz MC, Tecu I (2018) Anticompetitive effects of common ownership. J Finance 73(4):1513–1565. https://doi.org/10.1111/jofi.12698
12. Balancer (2024) Governance - balancer docs. https://docs.balancer.fi/concepts/governance/. Accessed 10 Oct 2024
13. Barbereau TJ, Smethurst R, Papageorgiou O, et al (2022) Defi, not so decentralized: the measured distribution of voting rights. In: Proceedings of the 55th Hawaii international conference on system sciences (HICSS). https://doi.org/10.24251/HICSS.2022.734
14. Bardoscia M, Barucca P, Battiston S, et al (2021) The physics of financial networks. Nat Rev Phys 3(7):490–507. https://doi.org/10.1038/s42254-021-00322-5
15. Béres F, Seres IA, Benczúr AA, et al (2021) Blockchain is watching you: profiling and deanonymizing Ethereum users. In: 2021 IEEE international conference on decentralized applications and infrastructures (DAPPS). IEEE, pp 69–78. https://doi.org/10.1109/DAPPS52256.2021.00013
16. BitDAO (2024) Governance. https://docs.bitdao.io/litepaper-1/governance-overview. Accessed 10 Oct 2024
17. Blondel VD, Guillaume JL, Lambiotte R, et al (2008) Fast unfolding of communities in large networks. J Stat Mech Theory Exp 2008(10):P10008. https://doi.org/10.1088/1742-5468/2008/10/P10008
18. Bodó B, Brekke JK, Hoepman JH (2021) Decentralisation: a multidisciplinary perspective. Internet Policy Rev 10(2):1–21. https://doi.org/10.14763/2021.2.1563
19. Brown S (2023) Measuring the concentration of control in contemporary Ethereum. arXiv preprint. arXiv:2312.14562. https://doi.org/10.48550/arXiv.2312.14562
20. Campajola C, Cristodaro R, De Collibus FM, et al (2022) The evolution of centralisation on cryptocurrency platforms. arXiv preprint. arXiv:2206.05081. https://doi.org/10.48550/arXiv.2206.05081
21. Campajola C, D'Errico M, Tessone CJ (2022) Microvelocity: rethinking the velocity of money for digital currencies. arXiv preprint. arXiv:2201.13416. https://doi.org/10.48550/arXiv.2201.13416
22. Compound (2024) Compound governance. https://docs.compound.finance/v2/governance/. Accessed 10 Oct 2024
23. Cong LW, He Z, Tang K (2025) The tokenomics of staking. NBER Working Paper No 33640. Last revised Feb 2025

24. Cong LW, Rabetti D, Wang CC, et al (2025) Centralized governance in decentralized organizations. Available at SSRN 5168660. https://doi.org/10.2139/ssrn.5168660
25. Convex Finance (2024) Understanding cvx | convex finance. https://docs.convexfinance.com/convexfinance/general-information/understanding-cvx. Accessed 10 Oct 2024
26. Curve Finance (2024) Understanding governance - curve resources. https://resources.curve.fi/governance/understanding-governance/. Accessed 10 Oct 2024
27. dYdX (2024) dydx governance documentation. https://www.dydxdao.info/dydx-governance/governance-architecture. Accessed 10 Oct 2024
28. Feichtinger R, Fritsch R, Vonlanthen Y, et al (2024) The hidden shortcomings of (d)aos – an empirical study of on-chain governance. In: Essex A, et al (eds) Financial cryptography and data security. FC 2023 international workshops. FC 2023. Lecture notes in computer science, vol 13953. Springer, Cham. https://doi.org/10.1007/978-3-031-48806-1_11
29. Fonseca NC, de França Carvalho JV (2025) Analysis of financial contagion among economic sectors through dynamic Bayesian networks. Expert Syst Appl 260:125448. https://doi.org/10.1016/j.eswa.2024.125448
30. Freni P, Ferro E, Moncada R (2022) Tokenomics and blockchain tokens: a design-oriented morphological framework. Blockchain Res Appl 3(1):100069. https://doi.org/10.1016/j.bcra.2022.100069
31. Fritsch R, Müller M, Wattenhofer R (2024) Analyzing voting power in decentralized governance: who controls daos? Blockchain Res Appl 5(3):100208. https://doi.org/10.1016/j.bcra.2024.100208
32. Gao Y, Campajola C, Vallarano N, et al (2024) Heterogeneity-and homophily-induced vulnerability of a p2p network formation model: the iota auto-peering protocol. arXiv preprint. arXiv:2401.12633. https://doi.org/10.1109/ACCESS.2023.3336418
33. Gini C (1921) Measurement of inequality of incomes. Econ J 31(121):124–126. https://doi.org/10.2307/2223319
34. Gogol K, Killer C, Schlosser M, et al (2024) Sok: decentralized finance (defi)–fundamentals, taxonomy and risks. arXiv preprint. arXiv:2404.11281. https://doi.org/10.48550/arXiv.2404.11281
35. Grundmann M, Amberg H, Baumstark M, et al (2022) On the peer degree distribution of the bitcoin p2p network. In: 2022 IEEE international conference on blockchain and cryptocurrency (ICBC). https://doi.org/10.1109/ICBC54727.2022.9805511
36. Gualdi S, Cimini G, Primicerio K, et al (2016) Statistically validated network of portfolio overlaps and systemic risk. Sci Rep 6(1):39467. https://doi.org/10.1038/srep39467
37. He J, Huang J (2017) Product market competition in a world of cross-ownership: evidence from institutional blockholdings. Rev Financ Stud 30(8):2674–2718. https://doi.org/10.1093/rfs/hhx031
38. Instadapp (2021) Governance flow. https://gov.instadapp.io/t/governance-flow/31. Accessed 10 Oct 2024
39. Jain S (2021) Instadapp protocol & governance. https://blog.instadapp.io/protocol-and-governance/. Accessed 10 Oct 2024
40. Jensen J, Schirrmacher NB, Avital M, et al (2023) Tokens matter: how to win votes and influence daos. In: Rising like a Phoenix: emerging from the pandemic and reshaping human endeavors with digital technologies ICIS 2023. https://aisel.aisnet.org/icis2023/blockchain/blockchain/8
41. Kitzler S, Balietti S, Saggese P, et al (2024) The governance of decentralized autonomous organizations: a study of contributors' influence, networks, and shifts in voting power. In: Financial cryptography and data security 2024. Springer, Berlin. https://doi.org/10.48550/arXiv.2309.14232
42. Li SN, Spychiger F, Tessone CJ (2023) Reward distribution in proof-of-stake protocols: a trade-off between inclusion and fairness. IEEE Access 11:134136–134145. https://doi.org/10.1109/ACCESS.2023.3336418
43. Lido (2024) Lido liquid staking. https://lido.fi/governance. Accessed 10 Oct 2024
44. Makarov I, Schoar A (2021) Blockchain analysis of the bitcoin market. Available at SSRN 3942181. https://doi.org/10.3386/w29396
45. MakerDAO (2024) Maker operational manual. https://manual.makerdao.com/governance-processes/voting-in-makerdao. Accessed 30 Sep 2024
46. MakerDAO (2024) Mip101: define scope for makerdao subdaos. https://mips.makerdao.com/mips/details/MIP101#scope. Accessed 30 Sep 2024
47. Medvedev E, the D5 team (2018) Ethereum etl. https://github.com/blockchain-etl/ethereum-etl
48. Nadler M, Schär F (2020) Decentralized finance, centralized ownership? An iterative mapping process to measure protocol token distribution. arXiv preprint. arXiv:2012.09306. https://doi.org/10.48550/arXiv.2012.09306
49. Ovezik C, Karakostas D, Milad M, et al (2025) Sok: measuring blockchain decentralization. arXiv preprint. arXiv:2501.18279
50. Rüetschi M, Campajola C, Tessone CJ (2024) How do decentralized finance protocols compare to traditional financial products? A taxonomic approach. Ledger 9. https://doi.org/10.5195/ledger.2024.360
51. Sai AR, Buckley J, Fitzgerald B, et al (2021) Taxonomy of centralization in public blockchain systems: a systematic literature review. Inf Process Manag 58(4):102584. https://doi.org/10.1016/j.ipm.2021.102584
52. Siudak D (2022) A network analysis of the value migration process on the financial market. The effect of value migration network structure on stock returns. Expert Syst Appl 191:116129. https://doi.org/10.1016/j.eswa.2021.116129
53. Sun W, Jin H, Jin F, et al (2022) Spatial analysis of global bitcoin mining. Sci Rep 12(1):10694. https://doi.org/10.1038/s41598-022-14987-0
54. SushiSwap (2024) Sushiswap docs. https://docs.sushi.com/dao/governance. Accessed 10 Oct 2024
55. Tumminello M, Micciche S, Lillo F, et al (2011) Statistically validated networks in bipartite complex systems. PLoS ONE 6(3):e17994. https://doi.org/10.1371/journal.pone.0017994
56. Uniswap (2024) Governance reference | uniswap. https://docs.uniswap.org/concepts/governance/process. Accessed 30 Sep 2024
57. Uniswap (2024) Uniswap governance categories. https://gov.uniswap.org/categories. Accessed 30 Sept 2024
58. Valadares JA, Villela SM, Bernardino HS, et al (2023) Mapping user behaviors to identify professional accounts in Ethereum using semi-supervised learning. Expert Syst Appl 229:120438. https://doi.org/10.1016/j.eswa.2023.120438
59. Victor F (2020) Address clustering heuristics for Ethereum. In: Financial cryptography and data security: 24th international conference, FC 2020, Kota Kinabalu, Malaysia, February 10–14, 2020. Revised selected papers, vol 24. Springer, Berlin, pp 617–633. https://doi.org/10.1007/978-3-030-51280-4_33

60. Vogelsteller F, Buterin V (2015) Erc-20: token standard. Ethereum improvement proposals. https://eips.ethereum.org/EIPS/eip-20. Retrieved 31 Jan 2024
61. Walch A (2019) Deconstructing "decentralization": exploring the core claim of crypto systems. In: Brummer C (ed) Cryptoassets: legal, regulatory, and monetary perspectives. Oxford University Press, London, pp 39–68. https://doi.org/10.1093/oso/9780190077310.003.0003
62. Werner SM, Perez D, Gudgeon L, et al (2022) Sok: decentralized finance (defi). In: Proceedings of the 4th ACM conference on advances in financial technologies, ACM, pp 30–46. https://doi.org/10.1145/3558535.3559780
63. Yearn Finance (2024) Governance and operations. https://docs.yearn.finance/contributing/governance/governance-and-operations. Accessed 10 Oct 2024
64. Yearn Finance (2024) Proposal process. https://docs.yearn.finance/contributing/governance/proposal-process. Accessed 10 Oct 2024

Publisher's note