

The Governance of Decentralized Autonomous Organizations: A Study of Contributors’ Influence, Networks, and Shifts in Voting Power

Stefan Kitzler^{1,2}, Stefano Ballelli³, Pietro Saggese^{4,2,1}, Bernhard Haslhofer¹
and Markus Strohmaier^{5,3,1}

¹ Complexity Science Hub Vienna

² AIT Austrian Institute of Technology

³ University of Mannheim

⁴ IMT School for Advanced Studies Lucca

⁵ GESIS - Leibniz Institute for the Social Sciences

Abstract. We present a study analyzing the voting behavior of contributors, or vested users, in Decentralized Autonomous Organizations (DAOs). We evaluate their involvement in decision-making processes, discovering that in at least 7.54% of all DAOs, contributors, on average, held the necessary majority to control governance decisions. Furthermore, contributors have singularly decided at least one proposal in 20.41% of DAOs. Notably, contributors tend to be centrally positioned within the DAO governance ecosystem, suggesting the presence of inner power circles. Additionally, we observed a tendency for shifts in governance token ownership shortly before governance polls take place in 1202 (14.81%) of 8116 evaluated proposals. Our findings highlight the central role of contributors across a spectrum of DAOs, including Decentralized Finance protocols. Our research also offers important empirical insights pertinent to ongoing regulatory activities aimed at increasing transparency to DAO governance frameworks.

Keywords: DAO · Governance · Ethereum · Networks · Blockchain · Voting

1 Introduction

DAOs represent organizational structures designed to offer an alternative, decentralized form of governance for decentralized applications (dApps) operating on Distributed Ledger Technologies (DLTs). The intention of DAOs is to circumvent central authorities and hierarchical structures that are prevalent in traditional organizations, and democratize the decision-making process by distributing voting rights through so-called governance tokens to community members [9].

Anecdotal evidence suggests that these intentions are not always met in practice. For instance, there are signs of a centralized power circle that has emerged within the Decentralized Exchange (DEX) service Sushiswap [29]. Similarly, the governance of Arbitrum DAO proposed channeling tokens valued at 1 billion US

dollars into their own treasury [38]. The lending protocol Solend confiscated the funds of a prominent user who posed a risk to its financial stability [10]. In another instance, major cryptoasset exchanges, significant entities in this context, reportedly colluded and leveraged investors’ tokens to vote on the Steem platform [15,27]. Attempts at bribery have been noted among community members in governance forums [40]. Lastly, developers from the mixing service Tornado Cash are reportedly under investigation for financial crimes; it’s alleged they manipulated its governance to circumvent the introduction of rigorous anti-money laundering controls [18,13].

It is well known that governance tokens are distributed primarily to team members, early investors, or protocol treasuries [6], and decision-making power can be concentrated in the hands of a few [5]. Earlier research has provided preliminary evidence on the involvement of DAO team members and developers in DAOs decision-making processes [34,23,20]. However, there is a surprising gap in studies that systematically investigate the role of vested users in the governance of DAOs and how they determine their trajectories.

In this study, we focus on DAO *contributors*, encompassing project owners, administrators, and developers. These contributors are involved in the technical realization of the dApp overseen by a DAO and thus can be viewed as vested users. Our aim is to empirically examine their influence in decision-making processes, the structure of their co-voting network, and any sudden shifts in majorities just before voting takes place. Our contributions and findings can be summarized as follows:

1. We compiled a dataset comprising 986 557 voters across 872 DAOs with 7478 recognized contributions. Additionally, we cross-verified a subset of 438 668 votes from 8116 proposals against their on-chain records, determining that 97.48% of these were consistent.
2. We introduce a metric to measure the *involvement of contributors* in DAO voting: in 66 (7.54%) DAOs contributors held, on average, the necessary majority to steer governance. We also measured *contributor self-decisions*, discovering that their votes were decisive in 178 (20.41%) DAOs.
3. We analyze the co-voting structures of users through a network approach. Our findings indicate that contributors are more likely to be found towards the center of the DAO governance ecosystem. Furthermore, contributors are highly concentrated in a few communities formed by co-voting patterns.
4. We observed *majority shifts* in governance token ownership in 1202 (14.81%) out of 8116 proposals in the days preceding the votes. The number of majority shifts increases sharply prior to governance polls, indicating last-minute token acquisitions.

To the best of our knowledge, our study is the first to systematically investigate the role of contributors in the governance of DAOs. It underscores their pivotal role across various DAOs, including leading Decentralized Finance (DeFi) protocols. Beyond shedding light on centralization tendencies within DAO governance structures, our findings demonstrate that contributors possess the capability to effectively steer the direction of DAOs. These insights have significant

implications regarding accountability. Moreover, they are relevant in the context of current regulatory initiatives aimed at pinpointing the individuals who either control or exert notable influence over DeFi operations or structures [31].

We will release our dataset and the implementation of methodologies to ensure the reproducibility of our findings.

2 Background, Definitions and Related Work

2.1 Voting in Decentralized Autonomous Organizations

Decentralized Autonomous Organizations (DAOs) are a novel form of governance model that has become popular in the crypto ecosystem since 2020. They can govern decentralized applications (dApps) and their associated smart contracts [45]. Several Decentralized Finance (DeFi) protocols implement DAO governance models [2], e.g., MakerDAO [26,36], Uniswap [1], Sushiswap [37], and Compound [25]. DAOs can also operate without an underlying dApp [12].

DAO voting mechanisms can be divided into two primary categories: *on-chain* and *off-chain* voting. The former occurs directly on a DLT, through smart contracts implementing the voting logic. To vote, token holders delegate an address that can be controlled by another entity. This approach offers security and transparency, but transaction costs make it economically inefficient [17,16]. The latter takes place on centralized platforms like Snapshot [33], and only the voting outcome is stored on the DLT. This method is more scalable, accessible and efficient, at the cost of higher centralization (e.g., concerns that the DAOs might not enforce the decisions, or concurrent voting on different platforms, or non-tamper-proof databases). Our study focuses on the Snapshot platform, the largest off-chain governance platform covering more than 11,000 DAOs [43].

Decision-making in DAOs is executed through voting on so-called *improvement proposals* that can determine the evolution of the technical infrastructure [41], modify parameters affecting the economic incentives and design [14,11], or reallocate funds managed by a DAO [39,42]. Governance users can participate in the voting by possessing specific tokens, known as *governance tokens*, which represent their DAO membership and their proportional decision-making power.

2.2 Definitions

We now present a conceptual model of DAO voting and introduce the key terminology and notation used throughout this paper, referring to DAOs as *spaces* [33]. Figure 1 illustrates the entities: spaces, proposals and users; and it describes their relations of contribution and vote.

- Let \mathcal{U} be the set of all **users** exercising voting rights and \mathcal{S} be the set of all **spaces**. Users can also be denoted as **voters** in this context.
- A **contribution** is a relation $\mathcal{C} \subseteq \mathcal{U} \times \mathcal{S} \times \mathbb{P}(\mathcal{T})$, where $(u, s, T) \in \mathcal{C}$, if a user u contributes to a space s in one or more role types $T \subseteq \mathcal{T} = \{\text{owner, administrator, developer}\}$. Users can take multiple roles. A **contributor** is a user that has at least one contribution association to one space.




Fig. 1. Conceptualization of DAO voting. A proposal p introduces potential changes to a DAO space s , and users u can exert their decision-making power on them with their vote v (\rightarrow) and voting power w indicated by the arrow thickness. Governance users can be vested by their contribution c as *owner*, *administrator* or *developer* to a space ($- \rightarrow$). We denote their vested vote as V_{SS}^P when they are contributors of the *same-space* (\rightarrow) they are voting on, and V_{OS}^P when they are contributors of an *other-space* (\rightarrow).

- A **proposal** is a relation $\mathcal{P} \subseteq \mathcal{S} \times \mathbb{P}(\mathcal{O}) \times \mathbb{P}(\mathcal{F}) \times \mathbb{N}^+$, where $(s, O, F, h) \in \mathcal{P}$, if there is a proposed change to a space s providing a set of choices or options O to vote on, and a set of strategies F to be applied for determining the outcome of a vote at a given block height h . The sets of options and strategies are defined as follows:
 - $O^p \subseteq \mathbb{P}(\mathcal{O})$ denotes the set of possible options (choices) that can be selected during the voting phase on the improvement proposal p . In most cases, the alternatives are simply a yes/no answer (i.e., $O^p = \{Yes, No\}$).
 - $F^p \subseteq \mathbb{P}(\mathcal{F})$ denotes the set of strategy functions that are applied to compute the voting power for the governance user issuing a vote.
- A **vote** is a relation $\mathcal{V} \subseteq \mathcal{U} \times \mathcal{P} \times \mathcal{O} \times \mathbb{R}^+$, where $(u, p, o, m) \in \mathcal{V}$, if a user u votes on a proposal p by selecting an option $o \in O^p$, where O^p denotes the set of options published as part of a specific proposal p . In rare cases, o can become a vector, e.g., the associated voting strategies allow one to express multiple choices. Then, the magnitude vector m characterizes the weighted preference of each option, and $\sum_i m_i = 1$; if the vote expresses one single choice, m is a scalar equal to 1. We further denote as $V^p \subseteq \mathcal{V}$ the set of all votes related to proposal p . We distinguish between two types of votes:
 1. A user can contribute and vote on an improvement proposal of the same space. We, therefore, denote $V_{SS}^P \subseteq V^p$ as the set of **same-space votes**, where, for all tuples $(u_i, p_i, o_i, m_i) \in V_{SS}^P$, a tuple $(u_j, s_j, T_j) \in \mathcal{C}$ such that $u_i = u_j$ and $s_n = s_j$ for $p_i^{s_n}$ exists, i.e., the users u_i equals u_j and s_n of proposal p_i equals the space s_j .
 2. A user can also contribute to one space and vote on an improvement proposal for another space. We denote $V_{OS}^P \subseteq V^p$ as the set of **other-space votes**, where for all tuples $(u_i, p_i, o_i, m_i) \in V_{OS}^P$, a tuple $(u_j, s_j, T_j) \in \mathcal{C}$ such that $u_i = u_j$ exists, but there does not exist one where additionally $s_n = s_j$ is fulfilled for $p_i^{s_n}$. Note that: $V_{SS}^P \cap V_{OS}^P = \emptyset$.

3. Finally, we denote $V_C^P = V_{SS}^P \cup V_{OS}^P$ as the set of contributor votes.
- The **voting power** is the weight w assigned to an option o and characterizes the influence of a vote v . It is determined by the strategy function $f : \mathcal{V} \times \mathbb{N}^+ \rightarrow \mathbb{R}^+$ of the vote v at block height h . For proposals with multiple functions, the weight is defined by their sum $F^p(v, h) := \sum_{f \in F^p} f(v, h)$.
 - Finally, the options O^p can be ranked by aggregated voting power w . We denote as the *outcome* the options $\hat{O}^p = [\hat{o}_1^p, \hat{o}_2^p, \dots]$ ranked in descending order by voting power, and denote \hat{o}_1^p as the *decision*, i.e., the option having the highest accumulated voting power for the proposal p .

2.3 Related work

Prior research has extensively documented that the ownership of governance tokens is highly concentrated [34,23,28,6,16], as a result of intentional design decisions and market dynamics (governance tokens carry a market price and can be traded). Furthermore, their total supply, the monetary policy, and the initial token allocation affects their distribution; finally, mechanisms such as airdrops [21] further favor early participants and DAO members.

Studies focusing specifically on on-chain DAO voting confirm that governance tokens are highly concentrated. Furthermore, they show that users rarely exercise voting rights [5], and that individuals who possess the potential power to alter outcomes rarely exercise it [20,17]. Two related works identify the existence of voters’ coalitions in MakerDAO [36,35]. A preliminary study reports examples of voters who held governance tokens for the duration of a single proposal life-cycle [16].

Our work is closely related to studies on off-chain voting. Wang et al.[43] delivered a comprehensive overview of the voting platform Snapshot. Laturnus [24] utilized data from Snapshot and DeepDAO to investigate the economic performance of DAOs in relation to ownership concentration and voting participation.

While earlier research has provided preliminary evidence on the involvement of DAO team members and developers in DAOs decision-making processes, none of these studies has systematically investigated the role of vested users in DAOs and their influence in determining their trajectories. This knowledge gap serves as the motivation for our study.

3 Data

To analyze the involvement of contributors in DAO decision-making through voting, we gather data from the following sources: Snapshot, Ethereum blockchain, Ethereum Name Service (ENS) and The Graph. We combine them to identify contributions, as defined in Section 2.2. Then, we clean, verify, and validate our dataset, as summarized in Table 1. Additional details on the entire data preparation process and contribution identification are reported in Annex A.

	Raw	Cleaned (Sections 4 & 5)	Validated (Section 6)
Spaces \mathcal{S}	12 294	872	357
Voters \mathcal{U}	1 603 994	986 557	119 413
Contributions \mathcal{C}	11 949	7478	3927
Proposals \mathcal{P}	76 851	35 124	8116
Votes \mathcal{V}	8 365 707	5 240 622	438 668
Contributor votes \mathcal{V}_c	316 900	191 507	22 878

Table 1. Dataset summary. The *raw* dataset combines Snapshot data on voters \mathcal{U} with additional sources to identify the contributions \mathcal{C} and quantify their voting activity \mathcal{V}_c . Users vote (\mathcal{V}) on improvement proposals \mathcal{P} to DAO spaces \mathcal{S} . To focus on DAOs with mature governance structures, we *cleaned* and *validated* the data set using selected proposals with Ethereum on-chain data.

Raw dataset. We obtained 1 603 994 *DAO voters* and their wallet addresses from the Snapshot dataset. These have cast 8 365 707 votes from Nov-2020 to Dec-2022 on 76 851 *proposals*, using 208 distinct voting strategies in 12 294 *DAO spaces*. Next, we identify voters’ contributions to DAOs by joining their addresses with additional data. We extract their respective roles \mathcal{T} by retrieving the domain *owner* address from ENS references, the *administrators’* addresses from Snapshot, and the creators, or *developers*, of code accounts (CA) from the blockchain transaction for all space-related CA from Snapshot.

Cleaned dataset. We found that 42.67% of DAO spaces have one proposal only, 54.05% have less than five followers and 50.85% have at most two voters. We consider these as indicators of immature governance structures and, therefore, exclude them. We then remove non-final proposals and restrict to proposals using the *single-choice* voting, but explicitly avoid selecting proposals based on their relevance. Thus, we apply a cleaning procedure by incorporating related benchmarks that assess minimum requirements to include mature DAOs in the data set.

Validated dataset. Previous studies have shown inconsistencies between the reported and actual on-chain data, including instances of flawed data records within a Blockchain explorer [22], emphasizing the need for a validation framework. Therefore, we validate the consistency between the voting power values computed by Snapshot and the ground truth reflected in on-chain data.

We focus on the Ethereum Blockchain, the most relevant one for Snapshot [43] in terms of expressed voting power, and only consider proposals that are almost entirely⁶ covered by strategies bound to $F' \subseteq \mathbb{P}(\{f^{erc20}, f^{erc721}, f^{eth}\})$. With this approach, we could verify that 461 402 (97.48%) of 473 306 Ethereum Snapshot weights are correct.

⁶ The selected strategies cover more than 99% of the voting power in the proposals.




Fig. 2. Contributor involvement across DAO spaces. The DAOs are ranked by contributor involvement \bar{w}_C^s (●) from highest (left) to lowest (right). Some high-TVL dApps (●) are annotated for illustrative purposes and contributor involvement of more than 50% is colored (●).

4 Influence of contributors on DAO governance

4.1 Contributor involvement

Contributors are vested users, having intuitively higher incentives to be involved in the decision-making in DAOs. We analyze their involvement by measuring their voting power exercised in proposals. As discussed in Section 2.2, the voting power w_i is the weight assigned to an option o_i and characterizes the influence of a vote v_i . In most cases, it is equal to the amount of governance tokens held by the user. Recall that the weight w_i is the result of a proposal strategy function $F^p(v_i, h)$ applied on a vote v_i at a specific block height h .

We compute the involvement of contributors in a given space by averaging the share of voting power they have in proposals associated with that space. Since minted amounts of governance tokens, their prices, and their distributions across user vary, we normalize voting power by total voting power across proposals as follows: $\tilde{w}_i = w_i \times (\sum_{v_i \in V^p} w_i)^{-1}$. Next, we consider the set of contributors V_C^p , that includes *same-spaces* voters V_{SS}^p as well as *other-space* V_{OS}^p . Then, we compute the fraction of weights controlled by contributors \tilde{w}_C^p (1) and finally obtain the **contributor involvement** \bar{w}_C^s (2) as the average weights of contributors' votes for all proposals P in a DAO space s .

$$\tilde{w}_C^p = \sum_{v_i \in V_C^p} \tilde{w}_i, \quad (1) \quad \bar{w}_C^s = |P|^{-1} \sum_{p \in P} \tilde{w}_C^p. \quad (2)$$

To give an example, let's assume proposal p has four votes $\{v_1, v_2, v_3, v_4\}$ with normalized voting powers $\{.1, .4, .3, .2\}$, where $\sum \tilde{w}_i = 1$. Supposing that the first two voters are contributors, ie $V_C^p = \{v_1, v_2\}$, then $\tilde{w}_C^p = .1 + .4 = .5$.

We determine \bar{w}_C^s for all spaces on the cleaned data set, and show the results in Figure 2. The DAO spaces are ranked by contributor involvement in descending order. Thus, the involvement of contributors is high for the DAOs on the left-hand side and low for the DAOs on the right-hand side. For illustration purposes, we highlight the data points representing top DeFi protocols in terms of Total Value Locked (TVL), such as Aave, Uniswap or Instadapp.

Our results show that the involvement of contributors in terms of average voting power is relatively low for most DAOs. The median value is 4.26% and the standard deviation is 21.22. However, for 297 spaces, the relative voting power of contributors is higher than 10% and for 66 DAOs it is higher than 50%. In these spaces, the contributors have, on average, a majority of voting power and can determine single-handedly the outcome of proposals. In 9 spaces, the contributors were the only voters with 100% voting power.

4.2 Contributor self-decisions

Knowing that DAO contributors are involved in decision-making, we now investigate to what extent they decide on proposals related to their own spaces. Thus, we concentrate our analysis on the votes cast by users that contributed on improvement proposals of the same spaces (V_{SS}^p , see Section 2.2), which we herein denote as “self-votes”. Furthermore, we also consider the choices they made with their votes and their influence on the outcome of a proposal.

Recall that the *decision* of a proposal poll is determined by the option with the highest voting power \hat{o}_1^p within the ranked outcome $\hat{O}^p = [\hat{o}_1^p, \hat{o}_2^p, \dots]$. We denote the set of *decisive self-votes* as V_D^p , where the option o of V_D^p is the winning choice \hat{o}_1^p .

We are specifically interested in the self-votes where the decision-making was dominated by contributors, which we denote as *contributor self-decisions*. Intuitively, for a given space s , we determine the share of selected proposals based on two joint conditions. First, we consider the weight of contributor votes within a decision and select those proposals where contributors have a relative majority ($\geq 50\%$). Second, we consider also the second-ranked option and select those proposals where the weight of contributors in the decision is higher than the weight of the second ranked option. The underlying intuition is that in a head-to-head race between options, contributors might want to outweigh and overrule a leading option.

More formally, we define the set of *decisive self-votes* $V_D^p = V_{\hat{o}_1^p}^p \cap V_{SS}^p$ and also the complement set $V_{CV}^p = V_{\hat{o}_1^p}^p \setminus V_D^p$; we identify the fractions of relative voting power for decisive self-votes (3), for the complement set (4) and for the second choice \hat{o}_2^p (5) as

$$\tilde{w}_D^p = \sum_{v_i \in V_D^p} \tilde{w}_i, \quad (3) \quad \tilde{w}_{CV}^p = \sum_{v_i \in V_{CV}^p} \tilde{w}_i, \quad (4) \quad \tilde{w}_{\hat{o}_2^p}^p = \sum_{v_i \in V_{\hat{o}_2^p}^p} \tilde{w}_i. \quad (5)$$

For a given space s , we can define the *contributor self-decisions* δ^s as follows:

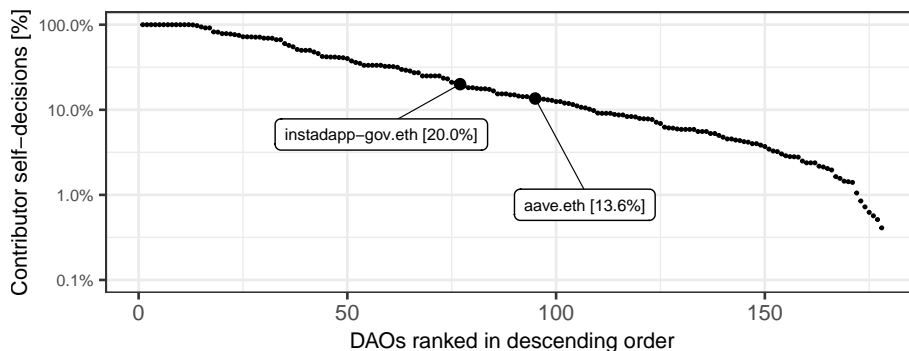


Fig. 3. Contributor self-decisions across DAO spaces. The 178 DAOs are ranked by contributor self-decisions δ^s (●) in descending order, with a threshold of 0.1%. The y-axis represents the fraction of proposals in which DAO contributors voted and decided their outcome with dominant voting power. Some high-TVL dApps (●) are annotated for illustrative purposes.

$$\delta^s := |P|^{-1} \sum_{p \in P} [(\tilde{w}_D^p > \tilde{w}_{CV}^p) \wedge (\tilde{w}_D^p > \tilde{w}_{\delta 2}^p)] \quad . \quad (6)$$

Figure 3 shows the results with DAOs ranked by self-decisions in descending order. Note that we introduced thresholds and only show spaces with self-decisions above 0.1%. This gives us 178 (20.41%) different spaces where contributors of the same DAO decided on at least one proposal on their own. In total 2100 out of 35 124 proposals were decided by governance users who contributed and voted on the same DAO. Annex B provides more details and analyses on involvement and self-decisions.

5 Co-voting networks

In Section 4, we measured the involvement of contributors in decision-making for each space separately; however, users can contribute to multiple proposals and DAO spaces. Herein, we conduct a network-based analysis of users’ co-voting patterns across DAOs. We analyze topological features such as centrality measures and community structures that may indicate whether contributors occupy a central role in the DAO voting ecosystem.

Networks construction. The basis for the investigation is the bipartite network G_{PU} that links users \mathcal{U} to the proposals \mathcal{P} they voted on, having options \mathcal{O} as edge features. We derive *co-voting networks* as a monopartite projection of the G_{PU} on voters, by creating a network of users with weighted links that represent

Network	G_{AA}	G_{AW}	G_{TA}	G_{TW}
Daos	All	All	Top-100	Top-100
Votes	All	Winning	All	Winning
Num Nodes	104 863	75 879	20 401	14 494
Num Edges	739 813 062	107 374 710	19 917 792	6 045 065
Avg. Degree	14 110.09	2830.16	1952.63	834.15
Contr. Nodes	1.29%	1.45%	3.25%	4.5%
Contr. Edges	1.61%	1.76%	3.4%	8.0%

Table 2. Network statistics of four co-voting networks. The top of the table defines the four networks as a unique combination of two features: DAOs and Votes. *Top-100* DAOs are ranked by total value locked (TVL); *Winning* votes are votes for the choice that ultimately won the majority of voting power.

the number of proposals they voted together. We introduce a global threshold T on links to focus on users that systematically voted together on the same proposals and for computational reasons.

Ultimately, we build four co-voting networks crossing DAOs and votes as shown at the top of Table 2, namely:

- G_{AA} is the entire co-voting network, containing all votes, regardless of users’ choices, and all spaces;
- G_{AW} is the co-voting network of decision-makers. It only takes into account votes v_i for the *winning* decision \hat{o}_1^p (we hypothesize that co-voting patterns may be especially relevant among users who voted for the winning outcome);
- G_{TA} is the co-voting network of all votes of the top-100 DAOs by TVL.
- G_{TW} is the co-voting network of decision-makers in the top-100 DAOs by TVL, was constructed in the same fashion of G_{AW} .

Network descriptive statistics. We utilize the cleaned data set of 5 240 622 voting relations on 872 DAOs and 35 124 proposals. We create the networks using the threshold $T = 10$ on the links among voters; despite this, given the inherent computational challenge on computing metrics on the large network G_{AA} , we focus mainly on the remaining three networks.

In all networks, we identify small-world features, i.e., they are characterized by the presence of several hubs conveying information rapidly across connected communities [44]. Interestingly, the share of contributor nodes and edges increases in the Top-100 networks (see bottom rows in Table 2), suggesting both that they are more active and that they tend to have a larger weight in shaping the outcomes of proposals of the most important DAOs, rather than the peripheral ones. The rest of the section tries to confirm whether this intuition is true. Appendix C contains additional network statistics and analyses.




Fig. 4. Pagerank and k-core statistics in the four co-voting networks. Contributors tend to have higher pagerank and k-core across networks. All centrality measures make use of edge weights and are applied to the giant component; k-core statistics use geometric mean to limit the effect of outliers. Error bars are 95% confidence intervals of the means.

5.1 Centrality of contributors

To understand the influence of contributors on governance voting, we computed several network centrality measures for contributor and non-contributor nodes, namely pagerank, closeness, eigenvector, and betweenness centrality, and a k-core analysis. We present pagerank in Figure 4 as well as the k-core. Across all four networks, contributor nodes score higher in centrality in all measures but eigenvector for G_{AW} . These differences are generally highly significant (t-test $p < 0.001$), only the betweenness centrality shows more variability ($p < 0.05$ for G_{TA} , and $p < 0.1$ for G_{TW}).

We also computed the k-core-ness of contributors. A high k-core indicates direct connections with other nodes with high k-core nodes, that is, nodes with at least degree k . Across all networks, contributors have, on average, significantly higher k-core (t-test $p < 0.001$). For these statistics we chose to use the geometric means because they are less sensitive to outliers. In fact, contributors are generally less frequent in the portion of the distribution with the lowest k-core, however, there exists a few clusters of mainly non-contributors with very high k-core, which would skew the results.

5.2 Communities of contributors

To understand the presence of hidden co-voting formations, we performed the Louvain community detection method on three co-voting networks. This method optimizes the modularity of the graph so that the connections within each community are dense, while the connections across communities are sparse. As a result, each node in the graph is uniquely assigned to a community. We then tested if contributors can be found with equal probability in all communities or whether they are more likely to cluster together in a few of them. To answer




Fig. 5. Concentration of contributors across network communities. The bar plots show the Herfindahl-Hirschman concentration index for the distribution of contributors (■) and non-contributors (■) to communities assigned by the Louvain community detection algorithm. The inset donut plots show the share of communities with at least one contributor; in all networks, contributors are concentrated in a few of them.

this question, we computed the Herfindahl-Hirschman index of market concentration on the distribution of contributors to communities. A higher value of this index indicates a more concentrated market, that is a distribution with fewer groups or communities dominating in size. Figure 5 indeed shows a very high concentration level for contributors in all networks with a peak above 7000 (below 1500 is considered well-mixed, between 1500 to 2500 is moderately concentrated, and above 2500 is highly concentrated). Counting the communities with at least one contributor (the donut plot inside each panel of Figure. 5), contributors are to be found only in about 21-50% of all detected communities. A Pearson’s Chi-squared test indicates a significant deviation from chance in all networks ($p < 0.001$, with 100 000 bootstrapped replicates). Figure 6 visually confirms this result for the G_{TA} network: contributors (dark red) tend to cluster in a few central communities.

6 Pre-voting power shifts

Governance tokens are cryptoassets and, consequently, can be purchased and sold. Furthermore, previous works provide preliminary evidence that users may hold their voting rights only for the duration of single proposals [16]. Therefore, we hypothesize that changes in the ownership distribution shortly before the voting power is determined could indicate attempts to acquire additional power to influence a proposal’s decision. We investigate to what extent users, and especially contributors, acquire voting rights shortly before the poll execution.

For proposals that rely on on-chain data, it is possible to access the current and historical token balances of voters. Thus, we determine their voting power at earlier points in time by re-implementing the proposal strategies on historical data and comparing it to their actual voting power. Note that we recompute it




Fig. 6. The co-voting network of the Top-100 DAOs by TVL (winning votes only). The colors identify the largest communities obtained by optimizing modularity, while smaller communities are in light gray; contributor nodes are colored in dark red (●). Contributors are not uniformly distributed across communities, but they tend to cluster in a few of them in the middle of the graph, suggesting a higher network centrality. Network plotted using the OpenOrd layout algorithm in Gephi [7], after removing redundant edges following [30].

considering that the voter v_i still selects the same option o_i . Assuming that users' holdings do not fluctuate with high frequency, we sample their token balances on a daily basis⁷. More formally, given a proposal p , for each voter u_i we denote the actual voting power $w_i(h_\tau)$ as the voting power at the block h_τ of the vote execution, and $\hat{O}^p(h_\tau) = [\hat{o}_1^p(h_\tau), \dots]$ as the actual ranked outcome. Next, for each of the 100 days preceding the vote on the proposal p , we recompute the users' historical voting power $w_i(h_{\tau-t})$ and the resulting hypothetical ranked outcome $\hat{O}^p(h_{\tau-t})$, where $h_{\tau-t}$ is a block representative of the t^{th} day before the poll. We thus compare $\hat{O}^p(h_{\tau-t})$ to $\hat{O}^p(h_{\tau-t-1})$ and determine whether there was a *majority shift* if $\hat{o}_1^p(h_{\tau-t}) \neq \hat{o}_1^p(h_{\tau-t-1})$. Finally, we measure the number of majority shifts across proposals. Since we are correlating against on-chain data, we utilize the validated dataset described in Section 3, covering 8116 (23.11%)

⁷ Days are approximated assuming that blocks are mined on average every 15 seconds, i.e., $1d \doteq (86400/15)$ blocks.




Fig. 7. Majority shifts occur in temporal proximity of polls. Majority shifts occur when the voters, shortly before a proposal, trade enough governance tokens to swing the final outcome of the poll. We focus on the *validated dataset* and identify majority shifts up to 100 days before the votes. In temporal proximity to the proposals, the number of shifts increases, indicating *last-minute* voting power acquisition.

proposals. We therefore emphasize that the findings reported in this Section are a lower boundary estimation.

In total, we found majority shifts for 1202 (14.81%) proposals in 229 DAOs in the 100 days before the poll. The median number of shifts per proposal is 1, with a standard deviation of 2.64, and the maximum number of shifts for a single proposal is 30. To investigate whether the majority shifts are more frequent in the proximity of vote executions, Figure 7 reports the aggregated count of majority shifts across proposals as a function of the time distance from the vote execution. We observe a constant or slightly increasing trend in farther dates from $-100d$ to $-50d$, and a clearly increasing trend the closer time gets to the vote date $0d$. This indicates that the trading of governance tokens increases shortly before polls and that users might trade voting power to decide the outcome of the proposal in their preferred way. We acknowledge, however that we only identify a pattern and further research is required to better investigate this phenomenon.

Finally, we examine the participation of contributors in the proposals with majority shifts. Out of 1202 proposals with majority shifts, 1362 contributors associated with 1457 different DAO spaces voted in 728 (60.57%) proposals.

7 Discussion and Conclusions

Our study augments the existing body of knowledge on decision-making in DAOs. It substantiates the results of previous studies, highlighting that the distribution of governance tokens is highly concentrated [6,20], and the exercise of voting rights is very low [5]. Going beyond these findings, we discovered evidence that contributors, who are essentially users vested in DAOs, are involved in decision-making and, in some cases, have the power to effectively influence

the trajectories of DAOs. Furthermore, we provide evidence that contributors are more likely found at the center of the DAO governance ecosystem and that majority shifts happen, especially shortly before the votes.

These findings have several implications. First, they suggest that contributors are overrepresented in the decision-making process of certain DAOs compared to other governance users. This is in line with known concerns that contributors may have differing interests and that users with smaller stakes might be discouraged from voting [8]. Second, we found only limited evidence for DAO contributors influencing other, possibly competing, spaces. This is relevant because a rational governance token holder vested in a space might be interested in voting against proposals benefiting the evolution or adoption of other, competing DAOs (see [19]). Third, we found evidence of co-voting patterns among contributors, which is an indicator of the existence of inner circles of power in DAOs. These findings refute the conventional wisdom that DAOs are decentralized and run autonomously without being under anyone’s control. This is relevant for resolving questions of accountability, as vested users and large governance token holders may be considered members of a legally recognized entity and therefore responsible for the underlying dApp. The ongoing Tornado Cash investigation claiming that the developers influenced its governance is a prime example of this line of argumentation.

Our work clearly faces some limitations and opens directions for further research. Currently, it focuses on off-chain voting and on one platform alone (Snapshot), whilst voting is executed also on-chain on multiple DLTs, as well as on other off-chain platforms. Extending the study to other governance platforms and to on-chain DLT voting would be a straightforward improvement. Furthermore, our results provide preliminary evidence that majority shifts take place before voting. It would be important to investigate and explain more formally the factors influencing governance participation and voting, also combining on-chain data with traditional methods surveying crypto users [4,3]. In addition, further research could distinguish contributors by their role. Lastly, we view the contributors in our dataset as a lower boundary. Incorporating more data sources, e.g. Github, would likely elevate this baseline.

Who governs DAOs? This question has been the driving force behind our research and is also a significant concern for regulators currently formulating policy recommendations for Decentralized Finance (DeFi) systems [31]. Although our study does not aim to unveil the identities of the responsible individuals, as regulatory efforts suggest, it offers a systematic investigation into the role of vested users in the governance of DAOs. This, in turn, can provide valuable insights and inform ongoing regulatory debates on that topic.

Acknowledgments

This work is partially funded by the Austrian security research program KIRAS of the Federal Ministry of Finance (BMF) under the project DeFiTrace (grant agreement 905300) and the COMET Centre ABC — Austrian Blockchain Center — managed by the FFG (grant agreement 909237).

References

1. Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. Uniswap v3 core. Technical report, Uniswap, 2021. Available at <https://uniswap.org/whitepaper-v3.pdf>.
2. Raphael Auer, Bernhard Haslhofer, Stefan Kitzler, Pietro Saggese, and Friedhelm Victor. The technology of decentralized finance (DeFi). *Digital Finance*, August 2023. doi:10.1007/s42521-023-00088-8.
3. Raphael Auer and David Tercero-Lucas. Distrust or speculation? the socio-economic drivers of us cryptocurrency investments. *Journal of Financial Stability*, 62:101066, 2022.
4. Stefano Ballelli, Can Celebi, Myriam Kassoul, and David Tercero-Lucas. State of crypto: The socio-political landscape of the community, 2023. Available at <https://stateofcrypto.net/results/>.
5. Tom Barbereau, Reilly Smethurst, Orestis Papageorgiou, Johannes Sedlmeir, and Gilbert Fridgen. Decentralised finance’s timocratic governance: The distribution and exercise of tokenised voting rights. *Technology in Society*, page 102251, April 2023. URL: <https://www.sciencedirect.com/science/article/pii/S0160791X23000568>, doi:10.1016/j.techsoc.2023.102251.
6. Tom Josua Barbereau, Reilly Smethurst, Orestis Papageorgiou, Alexander Rieger, and Gilbert Fridgen. Defi, not so decentralized: The measured distribution of voting rights. In *Proceedings of the Hawaii International Conference on System Sciences 2022*, page 10, 2022.
7. Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. Gephi: an open source software for exploring and manipulating networks. In *Proceedings of the international AAAI conference on web and social media*, volume 3, pages 361–362, 2009.
8. Vitalik Buterin. Moving beyond coin voting governance, 2021. Available at <https://vitalik.ca/general/2021/08/16/voting3.html>.
9. Usman W. Chohan. The Decentralized Autonomous Organization and Governance Issues. Technical Report 3082055, Centre for Aerospace & Security Studies (CASS); Critical Blockchain Research Initiative (CBRI); International Association of Hyperpolyglots (HYPIA); University of New South Wales (UNSW), Rochester, NY, December 2017. URL: <https://papers.ssrn.com/abstract=3082055>, doi:10.2139/ssrn.3082055.
10. Cointelegraph. What decentralization? solana lender solend approves whale wallet takeover to avoid defi implosion, 2022. Available at: <https://cointelegraph.com/news/what-decentralization-solend-approves-whale-wallet-takeover-to-avoid-defi-implosion>.
11. Compound. Improvement proposal: Compound v2 to v3 migration phase 1, 2023. Available at: <https://compound.finance/governance/proposals/152>.
12. ConstitutionDAO. Documentation. Technical report, ConstitutionDAO, 2021. Available at <https://www.constitutiondao.com/>.
13. United States District Court. plaintiffs’ reply in support of motion for partial summary judgment and response to defendants’ cross-motion for summary judgment, 2022. Available at: https://assets.ctfassets.net/c5bd0wqjc7v0/6UF4IrKU07F1tWhnBT678a/fea1f8fecc839b9f625519247e8e938e/Plaintiffs_-_Summary_Judgment_Reply___Response.pdf.
14. Curve. Improvement proposal: reduction to crv token inflation, 2020. Available at: <https://gov.curve.fi/t/discussion-reduction-to-crv-inflation/851>.

15. Decrypt. Steem vs tron: The rebellion against a cryptocurrency empire, 2020. Available at: <https://decrypt.co/38050/steem-steemit-tron-justin-sun-cryptocurrency-war>.
16. Maya Dotan, Aviv Yaish, Hsin-Chu Yin, Eytan Tsytkin, and Aviv Zohar. The Vulnerable Nature of Decentralized Governance in DeFi. Technical report, The Hebrew University, August 2023. arXiv:2308.04267 [cs] type: article. URL: <http://arxiv.org/abs/2308.04267>, doi:10.48550/arXiv.2308.04267.
17. Rainer Feichtinger, Robin Fritsch, Yann Vonlanthen, and Roger Wattenhofer. The hidden shortcomings of (d) aos—an empirical study of on-chain governance. *arXiv preprint arXiv:2302.12125*, 2023.
18. FIOD. Arrest of suspected developer of tornado cash, 2022. Available at: <https://www.fiod.nl/arrest-of-suspected-developer-of-tornado-cash/>.
19. Bryan Ford and Rainer Böhme. Rationality is self-defeating in permissionless systems. *arXiv preprint arXiv:1910.08820*, 2019.
20. Robin Fritsch, Marino Müller, and Roger Wattenhofer. Analyzing voting power in decentralized governance: Who controls daos?, 2022. arXiv:2204.01176.
21. Michael Fröwis and Rainer Böhme. The operational cost of ethereum airdrops. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology: ESORICS 2019 International Workshops, DPM 2019 and CBT 2019, Luxembourg, September 26–27, 2019, Proceedings 14*, pages 255–270. Springer, 2019.
22. Zhiguo He, Jiasun Li, and Zhengxun Wu. Don’t trust, verify: The case of slashing from a popular ethereum explorer. In *Companion Proceedings of the ACM Web Conference 2023*, WWW ’23 Companion, page 1078–1084, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3543873.3587555.
23. 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, 2021. arXiv:2102.10096.
24. Valerie Laturmus. The economics of decentralized autonomous organizations. *Available at SSRN 4320196*, 2023.
25. Robert Leshner and Geoffrey Hayes. Compound: The money market protocol. Technical report, Compound, 2019. Available at <https://bit.ly/3ioW0jW>.
26. MakerDAO. The maker protocol: Makerdao’s multi-collateral dai (mcd) system. Technical report, MakerDAO, 2020. Available at <https://makerdao.com/en/whitepaper>.
27. Messari. Crypto theses for 2022, 2022. Available at: <https://messari.io/pdf/messari-report-crypto-theses-for-2022.pdf>.
28. Matthias Nadler and Fabian Schär. Decentralized Finance, Centralized Ownership? An Iterative Mapping Process to Measure Protocol Token Distribution. *arXiv:2012.09306 [cs, econ, q-fin]*, December 2020. arXiv: 2012.09306. URL: <http://arxiv.org/abs/2012.09306>.
29. Rekt news. Sushiswap scandal, 2021. Available at: <https://rekt.news/sushiswap-scandal/>.
30. Arlind Nocaj, Mark Ortmann, and Ulrik Brandes. Untangling the hairballs of multi-centered, small-world online social media networks. *Journal of Graph Algorithms and Applications*, 19(2):595–618, 2015.
31. The Board of the International Organization of Securities Commissions. Policy recommendations for decentralized finance (defi) consultation report, 2023. Available at: <https://www.iosco.org/library/pubdocs/pdf/IOSCOPD744.pdf>.
32. Evis Sinani, Anna Stafsudd, Steen Thomsen, Christofer Edling, and Trond Randøy. Corporate governance in scandinavia: Comparing networks and formal institutions. *European Management Review*, 5(1):27–40, 2008.

33. Snapshot. Documentation. Technical report, Snapshot, 2023. Available at <https://docs.snapshot.org/>.
34. Katerina Stroponiati, Ilya Abugov, Yiannis Varelas, Kostas Stroponiatis, Modesta Jurgeleviciene, and Yashoda Savanth. Decentralized governance in defi: Examples and pitfalls. Technical report, DappRadar, 2020.
35. Xiaotong Sun, Xi Chen, Charalampos Stasinakis, and Georgios Sermpinis. Voter coalitions and democracy in decentralized finance: Evidence from makerdao, 2023. [arXiv:2210.11203](https://arxiv.org/abs/2210.11203).
36. Xiaotong Sun, Charalampos Stasinakis, and Georgios Sermpinis. Decentralization illusion in decentralized finance: Evidence from tokenized voting in makerdao polls, 2023. [arXiv:2203.16612](https://arxiv.org/abs/2203.16612).
37. Sushiswap. Documentation. Tech. rep., Sushiswap, 2022. Available at <https://docs.sushi.com/>.
38. UnchainedCrypto. Arbitrum backtracks on aip-1 after community backlash, 2023. Available at: <https://unchainedcrypto.com/arbitrum-backtracks-on-aip-1-after-community-backlash/>.
39. Uniswap. Improvement proposal: Making protocol fees operational, 202. Available at: <https://gov.uniswap.org/t/making-protocol-fees-operational/21198>.
40. Uniswap. A user posts advances to try and buy votes, 2021. Available at: <https://gov.uniswap.org/t/consensus-check-abolish-delegates-and-change-the-uni-governance-voting-system/13458>.
41. Uniswap. Improvement proposal: deploy uniswap v3 on bnb chain, 2022. Available at: <https://gov.uniswap.org/t/rfc-update-deploy-uniswap-v3-1-0-3-0-0-5-0-01-on-bnb-chain-binance/19734>.
42. Uniswap. Improvement proposal: Uniswap dao can help 80k uni to turkiye to relief after the big earthquake disaster, 2023. Available at: <https://gov.uniswap.org/t/governance-proposal-uniswap-dao-can-help-80k-uni-to-turkiye-to-relief-after-the-big-earthquake-disaster/20758>.
43. Qin Wang, Guangsheng Yu, Yilin Sai, Caijun Sun, Lam Duc Nguyen, Sherry Xu, and Shiping Chen. An empirical study on snapshot daos. *ArXiv*, abs/2211.15993, 2022. URL: <https://api.semanticscholar.org/CorpusID:254069759>.
44. Duncan J Watts and Steven H Strogatz. Collective dynamics of ‘small-world’ networks. *nature*, 393(6684):440–442, 1998.
45. Weiqin Zou, David Lo, Pavneet Singh Kochhar, Xuan-Bach Dinh Le, Xin Xia, Yang Feng, Zhenyu Chen, and Baowen Xu. Smart contract development: Challenges and opportunities. *IEEE Transactions on Software Engineering*, 47(10):2084–2106, 2019.

A Data

This Annex provides additional material to Section 3. We start by listing more detailed information on the data sources (A.1) and how we extracted the contributor roles (A.2). Additionally, we include how we gathered information on Total Value Locked (TVL) of DAOs and describe the scrape and join procedure in (A.3). We describe the cleaning procedure of the *raw* data in (A.4), add more information on the voting strategies (A.5) and finally outline the data validation procedure (A.6).

A.1 Sources

We gathered data from the following sources:

- **Snapshot:** We received a comprehensive database dump from the Snapshot off-chain voting platform, provided by the platform’s developers. It encompasses all information about the DAO spaces, associated proposals, and votes on each proposal, spanning from the platform’s inception in November 2020 up until the 7th of December 2022.
- **Ethereum blockchain:** We run a full Erigon Ethereum archive node to acquire further information on code account creators and cryptoasset balances.
- **Ethereum Name Service (ENS):** Taking advantage of Snapshot’s data structure, all DAO spaces are linked to their corresponding ENS domains⁸, which reveal information about addresses associated with a DAO and the roles of users in charge.
- **The Graph:** This platform makes DLT data accessible through GraphQL and supplies additional information about ENS domains and their sub-domains⁹, which we are able to utilize and interlink.

A.2 Contribution roles

We extract contributions with different roles using the following sources:

- *Owner:* we retrieve the domain owner address from the ENS service of the Ethereum node and The Graph.
- *Administrator:* we retrieve the addresses of administrators from Snapshot.
- *Developer:* from Snapshot we obtain all contract account (CA) addresses associated with a certain space. If a CA represents a governance token, we keep it only if it can unambiguously be associated with that space¹⁰. Then we retrieve the addresses of contract creators from executed on-chain *create* transactions. Adopting a conservative approach, we only consider direct creations by users initiated via their externally owned accounts (EOAs). We discern indirect creations via contract accounts (CAs).

⁸ Refer to <https://docs.snapshot.org/user-guides/spaces/space-roles>

⁹ <https://thegraph.com/hosted-service/subgraph/ensdomains/ens>

¹⁰ Some spaces might utilize external governance tokens or prominent tokens, e.g., stablecoins, for voting.

A.3 Total Value Locked

We use the Total Value Locked (TVL) as the benchmark of the financial value of a DAO, and therefore as an indicator of its relevance. We collect the TVL by scraping available dApps from the DeFiLlama API¹¹. We focus on those with TVL available on Ethereum, and exclude centralized entities such as centralized exchanges (CEX).

To link the TVL data to snapshot DAO spaces, we directly included all scraped protocols that had governance identifiers in Snapshot. For those without such identifiers, we employed a name-matching approach based on Jaccard similarity. We accomplished this by utilizing an algorithm that compares names and identifier strings between Snapshot and DeFiLlama entries. We introduced a threshold of 0.5 for the Jaccard similarity of the names and 0.6 for the identifiers to link similar entries. To ensure consistency, we conducted a manual check of the entries and removed those with uncertain matches.

For the remaining matches that passed both the automatic and manual checks, we collected the TVL on Ethereum at the latest available time point in 2022. In cases where DeFi protocols had multiple versions, resulting in a $1:n$ relationship between Snapshot and DeFiLlama, we aggregated all versions of the same snapshot identifier to establish a $1:1$ relationship.

A.4 Data cleaning procedure

We now provide additional details on the data cleaning procedure described in Section 3. To identify DAO spaces with mature governance structures, we computed four features for each space: the total number of proposals; the number of proposals with more than ten votes; the number of users following a space; and the TVL obtained from DeFiLlama. We then applied the following cleaning steps: first, we removed DAOs with TVL smaller than 100k USD and selected the set of DAOs in the first 500 positions for at least one of the aforementioned measures.

Next, we removed 1806 non-final proposals with *pending* or *invalid* status and only took into account those with positive *final* status. We also restricted our dataset to proposals using the *single-choice* voting type, which reflects the majority of proposals (86.97%). In a single-choice vote, users can choose only one option, which simplifies the evaluation of voting outcomes. Finally, we verify the consistency of the scoring information of each proposal, i.e., we measure if the sum of individual votes and aggregate Snapshot data correspond and remove 17 inconsistent proposals.

A.5 Voting strategies

There are several ways to enable voting rights on proposals. We denote the methods to determine the power of a vote as *strategies*, following the terminology of Snapshot.

¹¹ <https://api.llama.fi/protocols>

Name	Number of proposals	Percentage
erc20-balance-of	24 359	53.85
delegation	5815	12.86
multichain	3842	8.49
erc721	3806	8.41
ticket	3543	7.83
contract-call	3514	7.77
erc721-with-multiplier	2549	5.64
decentraland-estate-size	2260	5.0
pancake	2231	4.93
cake	1350	2.98
erc20-balance-of-delegation	1111	2.46
erc20-with-balance	700	1.55
uniswap	699	1.55
balance-of-with-min	660	1.46
pagination	637	1.41
eth-balance	625	1.38
whitelist	566	1.25
sushiswap	460	1.02
erc1155-balance-of	447	0.99
masterchef-pool-balance	377	0.83

Table A.1. Top 20 Strategies by number of proposals. Strategies are used to determine the voting power of votes. We list most frequently used ones by number and percentage of proposals.

The Snapshot *raw* data set introduced in Section 3 consists of 8 365 707 votes on 76 851 proposals. In these proposals, we identified 208 distinct voting strategies, some of which are specific to individual DAOs. The most prominent one is *erc20-balance-of* (f^{erc20}), which defines the voting power to be proportional to the balance of a specific token, typically the DAO governance token. It is used in 53.85% of proposals. Similarly, the strategy *erc721* (f^{erc721}) computes voting power proportionally to the amount of some specific NFTs held by each user and *eth-balance* from the ETH balance. Also, the *ticket* strategy is relevant, implementing the “one person, one vote” approach. We note that strategy functions can be implemented arbitrarily and be more complex and also exploit holdings of other cryptoassets such as stablecoins (USDT) or governance tokens of other DAOs. In Table A.1, we show the top 20 strategies by number of proposals that use them. Note that strategies can vary significantly, e.g., from using tokens of multiple chains to being tailored to a specific DAO. The retrieval of this information necessitates a substantial investment in blockchain infrastructure and ongoing maintenance, concurrently making more complicated for outsiders to validate the correctness.

A.6 Data validation procedure

DAO space	Token	Error Count
bgansv2.eth	0x31385d3520bcd94f77aae104b406994d8f2168c [1]	3003
theopendao.eth	0x3b484b82567a09e2588a13d54d032153f0c0aee0 [1]	2312
purrnelopes-countryclub.eth	0x9759226b2f8ddef81583e244ef3bd13aaa7e4a1 [1]	2168
goopsnapshot.eth	0x15a2d6c2b4b9903c27f50cb8b32160ab17f186e2 [1]	1677
9x9x9.eth	0x5219c2f6f8ed1e76c937ed1269eda2658ba3c721 [1]	1296
cryptohoots.eth	0x5754f44bc96f9f0fe1a568253452a3f40f5e9f59 [1]	1168
pandaparadise.eth	0x24998f0a028d197413ef57c7810f7a5ef8b9fa55 [1]	202
cryptohoots.eth	0x196f7e9c5769fc777909a3f1b9bd65959f3f64fb [1]	102
gfanchan11.eth	0xa0b86991c6218b36c1d19d4a2e9eb0ce3606eb48 [4]	14
shibvinci.eth	0xe9615071341c6f0392a5dfde1645ad01b810cb43 [3]	9
gfanchan11.eth	0x6b175474e89094c44da98b954eedeac495271d0f [1]	7
ilvgov.eth	0xaebd9bd588f044cbdec8f3cf1e80277a7a52dc69 [2]	4
gfanchan11.eth	0xdac17f958d2ee523a2206206994597c13d831ec7 [4]	3

Table A.2: **Report of DAO spaces and patterns of mismatches.** This table contains the DAO spaces with inconsistent values between Snapshot and on-chain data. Patterns of errors can be found and categorized in classes.

We aim to validate the voting power w of the *cleaned* data set with on-chain data wherever possible. Therefore, we focus on the on-chain Ethereum blockchain and on proposals covered by the selected strategies $F' \subseteq \mathbb{P}(\{f^{erc20}, f^{erc721}, f^{eth}\})$ ¹². The reasoning for the restriction to F' is that strategies must be re-implemented manually, and F' contains the three most widely adopted strategies (which define the voting power to be proportional to the balance of a specific token, typically the DAO governance token or to the ether balance).

We start by coding the aforementioned strategies, i.e., we gather all asset holdings of the voters and recompute their voting power from on-chain Ethereum data. We apply the strategy information given in the Snapshot data. For governance tokens, this includes the token address as well as the decimal, i.e., the decimal floating point value necessary to convert on-chain data to a floating point number. That is, a token balance is reported on-chain as an integer and must be converted to its actual value. For instance, an illustrative token TKN can be reported on-chain with value 10,000,000 but its actual value is 10 because its decimal is 6. If no decimal is given, we assume it to be zero.

When we compare the queried voting power to the Snapshot data, discrepancies arise. A closer inspection revealed that some Snapshot proposals report incorrect values for the decimal value, e.g., there is no correspondence with the value reported in the smart contract of the token, or there is no "BalanceOf"

¹² Note, however, that in principle the same approach can be applied to multiple chains to retrieve additional voting balances.

function and therefore it is not possible to convert the on-chain value to the floating point. To avoid considering rounding errors as mismatches, we introduce a tolerance interval $\Delta\epsilon = 10^{-3}$ and consider the on-chain and off-chain values as equal if the difference is smaller than the threshold. We compare the voting power computed with that reported by Snapshot, and found high consistency: with this approach, we could verify that 461 402 (97.48%) of 473 306 Ethereum Snapshot weights are correct. We documented, however, larger deviations for some DAO spaces. The list in Table A.2 reports the spaces and the number of mismatches for each space. We find recurring patterns of errors, which we can categorize and also provide solutions to get also consistent values:

- Solution [1]: Instead of the decimal value given by the Snapshot strategy, we apply the decimal contained in the token smart contract. If none is given, we consider the decimal equal to 0.
- Solution [2]: We repeat the approach of Solution [1] but consider the decimal equal to 18. 18 is the conversion factor from wei to ether, and some currencies adopted it as well.
- Solution [3]: We apply the \log_{10} -function for both Snapshot and on-chain data before comparing them because the on-chain data reported as integers are very large, and therefore the deviation exceeds $\Delta\epsilon$, even though the actual values after the conversion would be small.
- Solution [4]: We substitute the decimal with 18 to convert the floating value.

These DAO spaces leave doubts about the reliability of the reported values. Thus, we selected only 357 out of 367 DAOs where all votes have been fully validated. The resulting *validated dataset*, used to conduct the analyses in Section 6, comprises 119 413 users, who voted 438 668 times on 8116 proposals of 357 spaces. These include 22 878 contributor votes out of 3927 contributions.

A.7 Snapshot activity

The voting platform Snapshot gained popularity in recent years, and especially after 2020. To give an overview, Figure A.1 shows the evolution over time of the monthly number of proposals and the average votes from July 2020 to December 2022. We observe a trend of increasing proposals and average votes, with peaks in November 2021 for the former and in November 2022 for the latter.




Fig. A.1. Evolution of Snapshot activity in time. Number of proposals created each month on the Snapshot platform (top) and the average votes per proposal (bottom). The panel on top shows that the number of proposals created monthly on Snapshot has grown significantly in 2021, with a peak of more than 5,000 proposals executed in November 2021. The bottom Panel shows that the average number of votes cast per proposal was around 100 until the beginning of 2022. Since then, this number is followed an increasing trend.

B Influence of Contributors on DAO Governance

We provide more details on the involvement of individual DAOs (Section B.1) and on self-decisions (Section B.2). First, we provide additional statistics on the contributor involvement and second, we measure to what extent other-space contributors overrule the same-space contributors in making decisions.

B.1 Contributor involvement

We complement the analysis in Section 4.1 by providing additional statistics on the top spaces by TVL and their contributor involvement. Table B.1 lists DAOs with contributor involvement above 50% and additional statistical measures. Figure B.1 shows the distributions across proposals for the top five spaces by TVL. In all these plots, the median is less than the average, but we also find outliers even above 50%, meaning that contributors had majorities in certain proposals.

DAOs	# Proposals	Contributor Involvement				
		Mean	Max	Std	Min	Median
levidao.eth	185	100.0%	100.0%	0.0%	100.0%	100.0%
cryptodigger.eth	29	100.0%	100.0%	0.0%	100.0%	100.0%
zingan.eth	4	100.0%	100.0%	0.0%	100.0%	100.0%
651792.eth	22	100.0%	100.0%	0.0%	100.0%	100.0%
strikeorg.eth	1	100.0%	100.0%	-	100.0%	100.0%
alongdomainsacheapdomain.eth	23	100.0%	100.0%	0.0%	100.0%	100.0%
oligamy.eth	2	100.0%	100.0%	0.0%	100.0%	100.0%
atato.eth	22	100.0%	100.0%	0.0%	100.0%	100.0%
loganzc.eth	28	100.0%	100.0%	0.0%	100.0%	100.0%
test4.shot.eth	28	50.2%	100.0%	42.0%	0.0%	50.0%
longbtc20090103.eth	20	50.23%	100.0%	23.02%	30.77%	41.43%
castledao.eth	3	50.26%	56.98%	6.04%	45.28%	48.51%
mcv.eth	59	50.48%	95.01%	16.08%	0.0%	47.65%
comp-vote.eth	5	51.15%	100.0%	36.55%	0.02%	49.62%
coinpirates.eth	2	51.23%	70.08%	26.65%	32.38%	51.23%
seedifyfund.eth	22	51.3%	100.0%	42.89%	0.0%	28.57%
apeswap-finance.eth	24	52.02%	99.37%	40.76%	0.0%	62.19%
5021314.eth	17	52.94%	100.0%	43.79%	0.0%	66.67%
sarcophagus-ambassadors.eth	41	54.81%	100.0%	28.86%	0.13%	59.03%
rabbitholes.eth	6	56.95%	100.0%	45.48%	0.0%	71.27%
lxdao.eth	37	57.5%	100.0%	32.53%	18.18%	42.86%
fdao.eth	25	57.69%	86.59%	17.32%	19.03%	58.99%
usdaogov.eth	21	57.9%	100.0%	43.55%	0.0%	77.78%
sdbal.eth	31	58.02%	100.0%	50.12%	0.0%	100.0%
layer2finance.eth	2	58.73%	84.87%	36.96%	32.59%	58.73%
meebitsdaopool.szns.eth	22	59.02%	100.0%	41.14%	0.0%	74.94%
metrox.eth	34	59.54%	100.0%	46.08%	0.0%	99.93%
ginocf.eth	64	60.99%	100.0%	47.5%	0.0%	100.0%
weenus	44	61.61%	100.0%	47.63%	0.0%	100.0%
dorg.eth	161	62.81%	100.0%	27.54%	0.0%	67.56%
al409.eth	17	65.11%	100.0%	36.51%	0.0%	66.67%
dakshow.eth	25	66.67%	100.0%	43.03%	0.0%	100.0%
alexec.eth	3	66.69%	100.0%	57.7%	0.06%	100.0%
wintersun.eth	11	68.79%	76.04%	13.53%	28.55%	73.37%
cryengine.eth	13	69.01%	85.03%	17.13%	31.0%	75.4%
pharo.eth	1	69.37%	69.37%	-	69.37%	69.37%
insuretoken.eth	36	69.4%	100.0%	46.69%	0.0%	100.0%
wdefi.eth	23	70.94%	100.0%	42.61%	0.0%	99.54%
thanku.eth	21	71.43%	100.0%	46.29%	0.0%	100.0%
primerating.eth	249	72.06%	100.0%	34.75%	0.0%	100.0%
cvx.eth	265	72.44%	99.58%	26.34%	0.0%	77.57%
testbsw.eth	17	76.4%	100.0%	43.68%	0.0%	100.0%
venus-xvs.eth	17	77.47%	99.88%	36.07%	0.0%	91.47%
lemu.dcl.eth	42	78.16%	100.0%	41.4%	0.0%	100.0%
huangwenchao.eth	72	78.24%	100.0%	37.7%	0.0%	100.0%
melson.eth	14	78.82%	100.0%	19.84%	27.28%	82.33%
cabindao.eth	4	79.53%	97.16%	20.28%	56.86%	82.05%
grantsdao.eth	36	80.42%	100.0%	20.64%	33.33%	77.5%
polywrap.eth	132	80.58%	100.0%	20.03%	0.0%	83.57%
cgpool.eth	1	81.04%	81.04%	-	81.04%	81.04%
4.spaceshot.eth	9	81.19%	100.0%	34.68%	0.24%	100.0%
xunfa.eth	35	81.86%	100.0%	5.91%	79.83%	79.83%
tommyg.eth	14	83.33%	100.0%	36.4%	0.0%	100.0%
retokendao.eth	36	84.46%	100.0%	10.09%	54.55%	85.43%
frami.eth	206	85.17%	100.0%	34.98%	0.0%	100.0%
adaocompany.eth	20	85.85%	100.0%	32.06%	0.0%	100.0%
3.spaceshot.eth	69	86.22%	100.0%	23.12%	0.0%	100.0%
mycontext.eth	9	87.04%	100.0%	26.06%	33.33%	100.0%
*	15	87.78%	100.0%	28.5%	0.0%	100.0%
btc1.eth	2	89.49%	100.0%	14.87%	78.97%	89.49%
szns.shean.eth	61	91.15%	100.0%	27.93%	0.0%	100.0%
bridgeswap.eth	25	91.73%	98.7%	5.11%	83.9%	94.69%
oracles.opiumprotocol.eth	25	94.77%	100.0%	18.74%	8.77%	100.0%
ecashxec.eth	1	95.82%	95.82%	-	95.82%	95.82%
halodao-kovan.eth	41	97.56%	100.0%	15.62%	0.0%	100.0%
alexjyoung.eth	24	98.75%	100.0%	6.11%	70.05%	100.0%

Table B.1. Statistics of contributor voting power. Additional statistics on voting power for the 66 DAOs with mean contributor involvement above 50%.

* DAO name contains character that cannot be displayed.




Fig. B.1. Contributors’ involvement in top-TVL DAOs. Box- and violin plots of relative voting power of contributors across top TVL DAOs. We find the mean values (X) above the median, for some proposals but outliers even above 50% relative voting power. Thus, also in high TVL DAOs, contributors had majorities in some proposals.

B.2 Contributor self-decisions

Decisions by other-space contributors To provide a broader picture, we complement the measure of *contributor self-decisions*, introduced in Section 4, with the *contributor other-decisions*, that is, we measure to what extent other-space contributors overrule the same-space contributors in making decisions. We define three joint conditions for proposals to be *decisive other-votes*. The first two conditions are inspired by *decisive self-votes* applied to votes of other-space contributors, which we denote as “other-contributor votes”. First, other-contributor votes have a relative majority within the voting power for the decision. Second, we consider also the second-ranked option and select those proposals where the weight of other-space contributors in the decision is higher than the weight of the second-ranked option. Third, we take into account only proposals where same-space contributors are involved in the voting for the second ranked option.

Formally, we define the set of *decisive other votes* $V_{DO}^p = V_{o1}^p \cap V_{oS}^p$ and also the complement set $V_{CO}^p = V_{o1}^p \setminus V_{DO}^p$; we identify the fractions of relative voting power for decisive other votes (7), for the complement set (8) and the same-space contributors for the second option $V_{SS2}^p = V_{o2}^p \cap V_{SS}^p$ as

$$\tilde{w}_{DO}^p = \sum_{v_i \in V_{DO}^p} \tilde{w}_i, \quad (7) \quad \tilde{w}_{CO}^p = \sum_{v_i \in V_{CO}^p} \tilde{w}_i, \quad (8) \quad \tilde{w}_{SS2}^p = \sum_{v_i \in V_{SS2}^p} \tilde{w}_i, \quad (9)$$

We consider proposals to be decided by other-contributors, when all three conditions are fulfilled:

$$[(\tilde{w}_{DO}^p > \tilde{w}_{CO}^p) \wedge (\tilde{w}_{DO}^p > \tilde{w}_{SS2}^p) \wedge (\tilde{w}_{SS2}^p > 0)], \quad (10)$$

In rare cases, we find that the outcome of a proposal has been dominated by contributors of other spaces over the self-contributors of the proposal’s space. We show the relative voting power for these proposals in Figure B.2, by a stacked bar plot of votes on different choices. Each proposal is illustrated by three bars, respectively corresponding to the relative voting power of the winning choice, the second choice by voting power and all other choices aggregated.




Fig. B.2. Decisions of other-space against same-space contributors. In 30 proposals, other-space contributors (■) had the relative majority in the winning choice and overruled the same-space contributors (■), voting for the second choice.

C Network Analysis

We present additional material about network descriptive statistics (C.1), an analysis of the small-world properties of the networks (C.2), an investigation

Network	G_{AA}	G_{AW}	G_{TA}	G_{TW}
Daos	All	All	Top-100	Top-100
Votes	All	Winning	All	Winning
Num Nodes	104 863	75 879	20 401	14 494
Num Edges	739 813 062	107 374 710	19 917 792	6 045 065
Max Degree	57 304	31 664	14 456	9 039
Avg. Degree	14 110.09	2830.16	1952.63	834.15
Contr. Nodes	1.29%	1.45%	3.25%	4.5%
Contr. Edges	1.61%	1.76%	3.4%	8.0%
Contr. Max. Degree	57 303	24 346	14 196	7 755
Contr. Avg. Degree	8795.3	3434.12	2079.8	1490.86
Giant Component	98.1%	98.8%	99.9%	99.4%
Assortativity	0.02	0.05	0.11	0.07
Avg. Path Length	2.24	2.58	2.21	2.37
Diameter	32	21	7	9
Clustering	1.55	1.59	1.99	2.06
Small Worldliness	x	3.85	4.31	6.41
Louvain Communities	x	330	28	12
Largest Community	x	42%	43%	34%

Table C.1. Network statistics of four co-voting networks. The top of the table defines the four networks as a unique combination of two features: DAOs and Votes. *Top-100* DAOs are ranked by total value locked (TVL); *Winning* votes are votes for the choice that ultimately won the majority of voting power. The rest of the table contains network statistics for the four graphs, the central part referring to all nodes and edges, and the bottom parts only to the giant component. *Some statistics could not be computed for the largest network due to memory limitations.*

of the network of contributors (C.3), additional plots for the centrality and k-core analysis (C.4), and finally, a sensitivity analysis for the choice of the edge threshold (C.5).

C.1 Co-voting network statistics

In Section 5 we built the four co-voting networks G_{AA} , G_{AW} , G_{TA} , G_{TW} . Their exhaustive statistics are presented in Table C.1. The table is split in four segments, starting on top with the distinction of the networks by combination of their features. Then, it provides descriptive statistics of the network nodes and edges, of the contributor nodes only, and, finally, insights on network measures computed on the giant component.

C.2 Small-world properties

The co-voting network with all DAOs and all votes (G_{AA}) counts 104 863 nodes and is about 5 times larger than the corresponding network with the top-100

DAOs by TVL (G_{TA}). Despite the different sizes, the four networks present similar topological features: a very large giant component (about 99% of all nodes) with clear small world features. Following the procedure in [32], we computed the “small-worldliness” coefficient of each graph, comparing the ratio of average path length (APL) and clustering coefficient (CC) normalized by the values of those statistics in a random network (rnd) with the same number of nodes and degree distribution (test):

$$SW = (APL_{test}/APL_{rnd})/(CC_{test}/CC_{rnd}) \quad (11)$$

Intuitively, to be a small world, a network needs APL to be similar to that of a random network, but CC should be much larger, leading to a “small-worldliness” coefficient SW considerably larger than one. This is indeed the case for all our networks (SW between 3.85 and 6.41). The non-random structure is also confirmed by a small, yet positive assortativity coefficient, indicating a certain degree of homophily.

C.3 The contributors network

Users might connect spaces through their contributions. We construct the *contributors network* G_{SU} as a bipartite network of the contribution relation \mathcal{C} , using spaces \mathcal{S} and users \mathcal{U} as two node types without taking into account their specific roles \mathcal{T} .

We utilized the raw data set from Section 3, containing 12 294 DAO spaces, and extracted the contribution roles with sources defined in Section A.2. Using these sources, we found a total of 12 272 DAOs 17 561 addresses. 11 949 of these addresses are voters in the raw data, which would strictly interpreted only be identified as contributions. However, we are interested in the most comprehensive network of contributors to DAOs to find all possible relations. We construct the network G_{SU} linking 12 272 DAOs to all 17 561 contribution addresses. Thus, on average, each DAO has 1.43 contributors. This network appears fragmented, as it separates into 10 651 small components, mostly with ten or fewer nodes and with rather simple relations between addresses and a few larger ones. In Figure C.1 we plot the distribution of the connected component. We limit the network to components size larger than 10 and compute the bipartite projection of G_{SU} on DAO spaces. The resulting network in Figure C.2 links DAO nodes when they have contributors in common. We observe several star-like structures, but there also exist larger structures created by contributors bridging multiple DAOs together. However, keeping the distribution of Figure C.1 in mind, most of the relation consist only of rather simple relations between DAOs.




Fig. C.1. Number of connected components of the G_{CD} network. The network mainly consists of small components and a larger component with above 100 nodes.




Fig. C.2. Network of DAOs connected through contributions. Bipartite network projection of contributors network G_{SU} with more than 10 members. The network illustrates nodes that are linked when they have a common contribution. It appears fragmented, with star-like structures and a view larger bridged DAO communities.

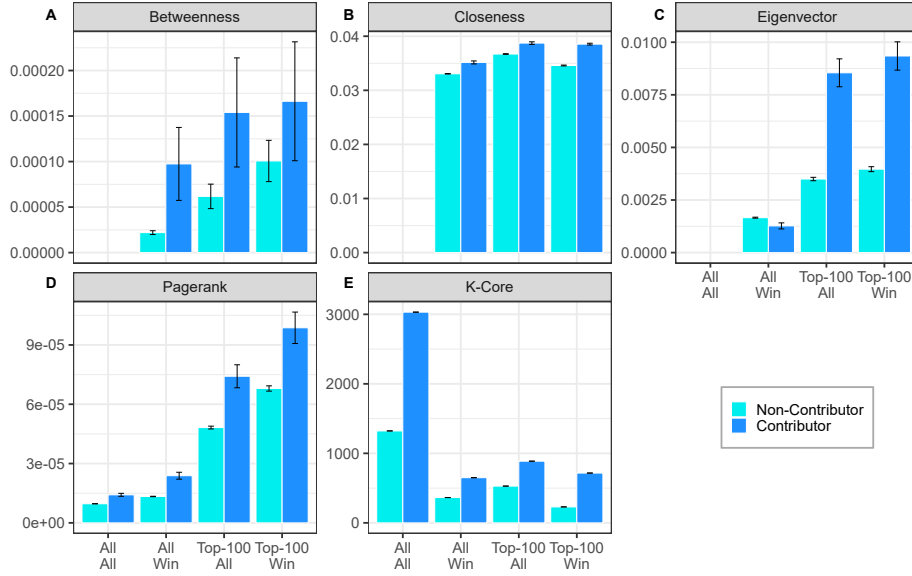


Fig. C.3. Centrality and k-core statistics in the four co-voting networks. Contributors tend to have higher centrality and k-core across networks and measures. Overall, it suggests that contributors have influence over the entire network. **A-D.** Centrality measures (some measures for the largest network could not be computed due to memory limitations). All centrality measures make use of edge weights and are applied on the giant connected component. **E.** Geometric mean of K-core. Error bars are 95% confidence intervals of the means.

C.4 Centrality and k-core

We report in Figure C.3 all the plots of all centrality measures mentioned in the main text. Contributors tend to score higher in all centrality measures, but one (eigenvector centrality in the network with all DAOs and only winning votes).

In addition, in Figure C.4, we report the density plots for the distributions of k-core across all networks. Contributors are less frequent in the portion of the distributions with lower k-core; however, in two networks (G_{AW} , G_{TA}) an outlier cluster of non-contributors with very high k-core is detected. For this reason, we opted to report the geometric mean in the main text. Econometric tests with the arithmetic means are mixed: contributors have on average significantly higher k-core (t-test $p < 0.001$) in two out of four networks.

C.5 Edge-weight threshold sensitivity analysis

In the main text, we report results for networks built with an edge-weight threshold equal to ten, so that an edge between two nodes is added if and only they have have co-voted in more than $T=10$ proposals. The value of ten has been

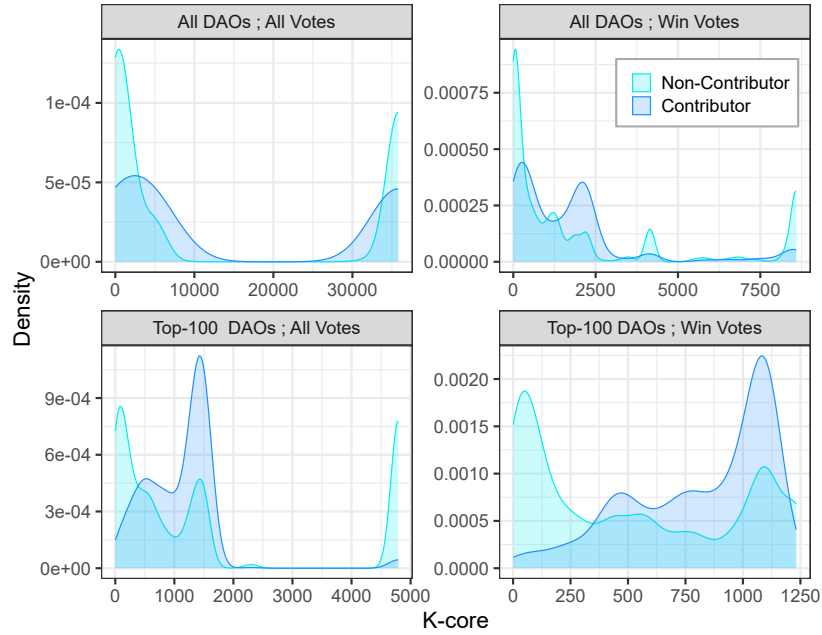


Fig. C.4. K-core density plots for contributor and non-contributor nodes in the four co-voting networks. Contributors are less frequent in the lower portion of the distribution in all plots. In two plots (G_{AW} , G_{TA}) an outlier cluster of non-contributor with very high k-core is detected.

conservatively chosen to be low enough to keep as much of the original network structure intact, but at the same time, allowing us to reduce both data noise and the complexity of computation. We tested how the choice of a threshold affected the main network statistics, and the results are reported in Figure C.5. For most statistics—i.e., number of edges, number of nodes, average path length—the curve varies smoothly with T ; however, assortativity and the pseudo-diameter seem more sensitive to the choice of the threshold. Further analyses may precisely quantify the impact of the choice of threshold on our results.

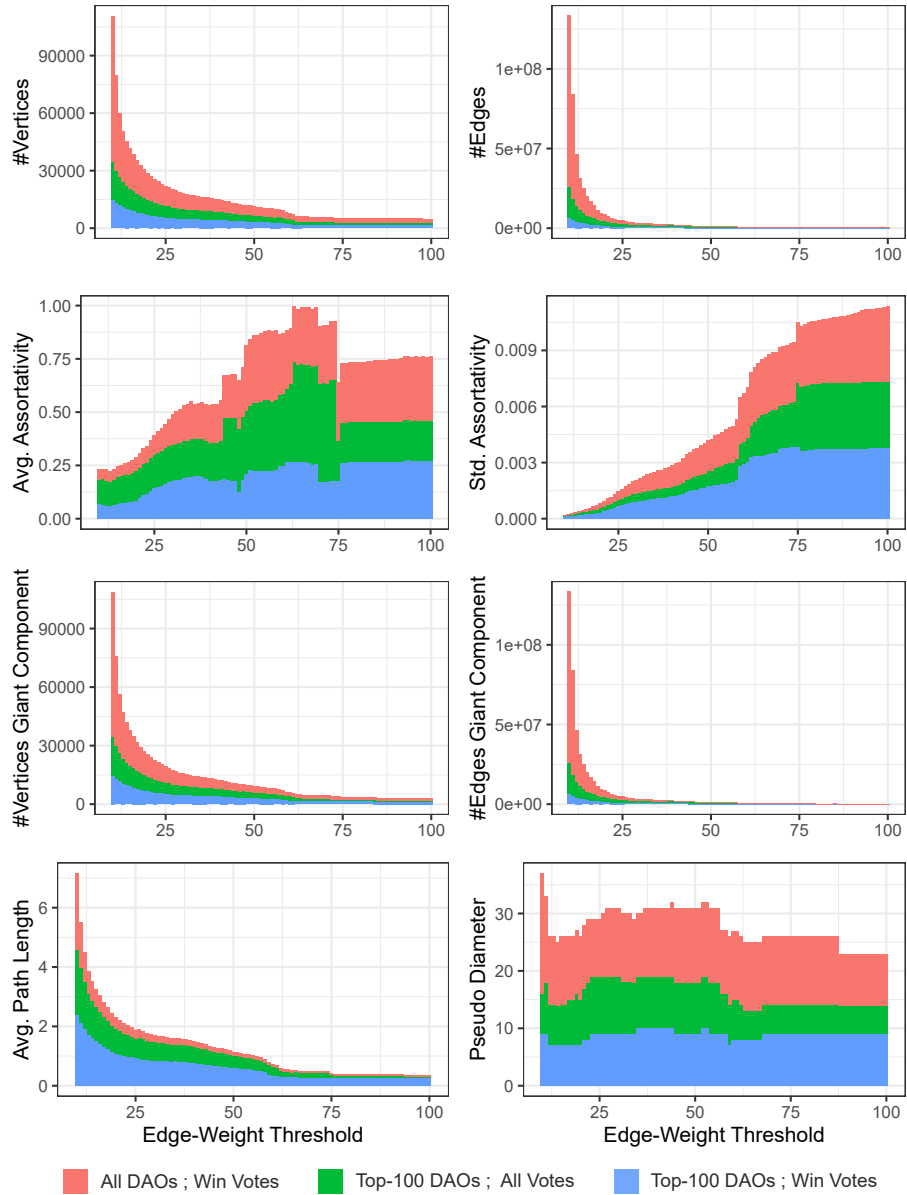


Fig. C.5. Overview of network measures for different edge-weight thresholds. Two nodes must have co-voted in more than T proposals, where T is the edge-weight threshold. The value $T = 10$ used in the main text corresponds to the left-most point on the x-axis. All plots are stacked.