Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

# Democracy on the Blockchain

Bachelor's Thesis

Marino Müller
muemarin@ethz.ch

Distributed Computing Group
Computer Engineering and Networks Laboratory
ETH Zürich

**Supervisors:**
Robin Fritsch
Prof. Dr. Roger Wattenhofer

February 10, 2022

# Acknowledgements

I would like to thank my supervisor, Robin Fritsch, for making this work possible. His advice and feedback in the weekly meetings helped me immensely in completing this project.

I would also like to thank Frederik Semmel who provided his expertise for some mathematical problems I encountered and Maximilian Winzely who I could always consult for questions regarding form and layout of the work.

Furthermore I want to thank all my friends and family for their ongoing support.

# Abstract

With the steep rise of blockchain projects, the need for governance systems has also become apparent. Some projects try to let their community decide important decisions with voting schemes on the blockchain. To do this, they use delegative democracy, where the mechanism requires you to first delegate your voting power before it can be used.
Using different techniques we try to get an insight into voting behavior, voting power and fairness of projects implementing such a kind of governance system. Mainly, we want to answer whether the voting behavior of small token holders differs from that of large token holders, how governance tokens are distributed, and who holds power in these systems.
To accomplish this, we use publicly available blockchain data of the prominent blockchain-based applications *Uniswap*, *Compound* and *ENS*. In the end we compare our findings for the conventional on-chain voting mechanism with *Snapshot*, a no gas cost off-chain voting alternative.

# Contents

Chapter 1

# Introduction Blockchain Voting

Blockchain technology has become increasingly important in recent years. Later iterations of these technologies are making use of immutable deterministic computer programs commonly referred to as smart contracts [1].
These smart contracts enable the implementation of a variety of decentralized projects on the blockchain. The best-known example is certainly Decentralized Finance (DeFi) with projects such as *Uniswap* [2] and *Compound* [3]. Other promising projects, such as the name and lookup service *Ethereum Name Service (ENS)*[1], are also emerging with increasing frequency.

These blockchain projects seek to make important decisions through voting schemes built on smart contract. These decisions take the form of votes on proposals and can range from simple parameter changes in the protocol to the use of the project's own governance treasury.
A rather controversial example of the second case would be the fifth proposal on *Uniswap*. This proposal was to allocate one million UNI tokens (roughly 19.5 million CHF at the time of the proposal July 6, 2021) from the governance treasury to create the non-profit *DeFi Education Fund*, which provides grants for political, educational and legal engagement [2].

The governance process involves multiple steps, only the last of which occurs on the blockchain. For example, in *Uniswap*, there is first a *temperature check* and a *consensus check*, both of which are off-chain and require a certain threshold of yes-votes. Only when these two thresholds are reached is a proposal made on-chain. In this work, we mainly focus on on-chain voting. Off-chain voting is briefly examined at the end.

The creation and issuance of governance tokens through these blockchain-based governance projects is intended to achieve a decentralized distribution of voting rights [4].

[1] `https://ens.domains`
[2] `https://app.uniswap.org/#/vote/1/1`

## 1.1 Governance Tokens

Governance tokens are fungible tokens that establish a voting logic by allowing holders to express their intent for the protocol development in majority-voting schemes, with each token corresponding to one vote [4].
The initial distribution of these governance tokens varies widely for each project, but almost all of them distribute the tokens among the founders, the development team, potential investors, the ecosystem treasury and some external agents as an incentive for future use of the application [4].

In order to participate each governance token holder has to delegate their tokens, and therefore their voting rights to another address, the delegate. Each holder can also delegate to their own address, which makes them their own delegate. However, each holder can only delegate to one address and only with their entire balance of governor tokens. Therefore, anyone who wants to delegate their tokens to multiple delegates would first have to distribute their tokens to multiple addresses.

Delegates can use the governance token balance delegated to them to vote for or against a proposal in a binary voting scheme [5]. The more governance token are delegated to them, the more weight their vote on a proposal holds.

Like traditional equities, governance tokens are traded on secondary markets, giving them a monetary value in addition to voting rights.

## 1.2 Dataset

We use the indexing protocol *The Graph*[3] to query data from the Ethereum blockchain for governance tokens from the renowned DeFi applications *Uniswap*, *Compound* and the relative new but promising project *Ethereum Name Service* (ENS). With the help of the *GraphQL Query* API, we are able to retrieve the complete set of holders, delegates, and proposals.
Furthermore, it is also possible to get important information such as number of tokens held, number of tokens delegated, votes and details about proposal outcomes.
Table 1.1 shows an overview of the chosen projects as of January 13, 2022.

[3] `https://thegraph.com`

| **Protocol** | **Holders** | **Delegates** | **Supply (max)** | **delegated** | **Proposals** |
|---|---|---|---|---|---|
| Compound | 184,350 | 1,759 | 6,324,007 ($10^7$) | 27.72% | 80 |
| Uniswap | 296,780 | 4,360 | 627,287,936 ($10^9$) | 20.56% | 10 |
| ENS | 59,380 | 12,200 | 20,244,862 ($10^8$) | 7.25% | 0 |

Table 1.1: Overview of the dataset. **Supply (max)** corresponds to the current circulating supply and the maximum supply that can ever exist. **delegated** means the ratio of delegated tokens to the circulation supply.

## 1.3 Visualization of Delegation Network

We model the holders and delegates from the dataset as a network to determine their connections to each other. Using *NetworkX* [6], we create two sets of nodes for holders and delegates, as well as edges between nodes from different sets when there is a delegation. We also assign a weight to each node and edge corresponding to the number of tokens delegated or held.
To create the visualization of the networks shown in Fig 1.1, we use *Pyvis*[4] The red vertices symbolize delegates and the blue vertices holders. An edge connecting them represents a delegation from the holder to the delegate. The size of the vertices and edges is proportional to the number of governance tokens held or delegated.
Since each holder can only delegate to one delegate, the network consists of many structures in which holders are arranged circular around a delegate.

In Fig 1.1, we see that most delegates in *Uniswap* and *Compound* have one main delegator from which most delegated tokens originate. This could be due to holders creating new addresses and delegating to themselves, or an entity creating multiple addresses to fund numerous delegates.
With *ENS*, the network looks quite different. Here we see that almost every delegate represents many holders and that the distribution of the delegated token amounts is much more evenly distributed among the delegates. This could be because *ENS* is distributing a large portion of its issued governance tokens (25%) to over 137,000 users via an ongoing airdrop (not all have been claimed yet).

[4] https://pyvis.readthedocs.io

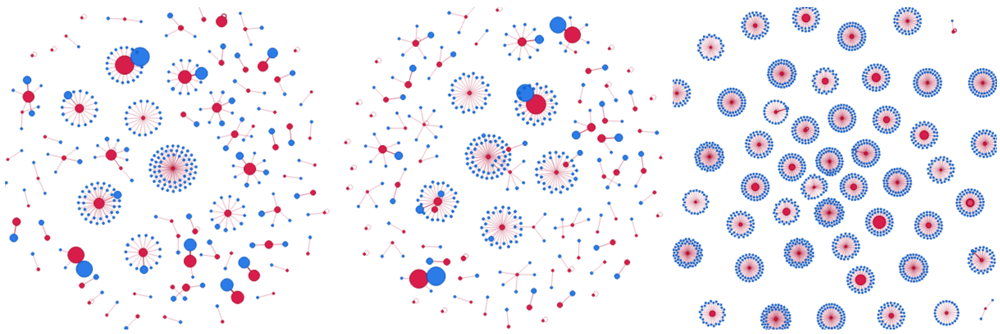

(a) (b) (c)

Figure 1.1: (a) Uniswap (b) Compound (c) ENS
Visualization of the network of holders (blue) and delegates (red) for each project. An edge means delegation and the size of the vertices corresponds to the amount of held and thus delegated tokens.

## 1.4 Related Work

The paper *How Decentralized is the Governance of Blockchain-based Finance* [4] addresses the decentralization of governance token distribution for DeFi projects. It puts great emphasis on Gini coefficients and the Nakamoto coefficient. In contrast to this work, the distribution of governance tokens rather than the distribution of delegations of governance tokens is examined there.

In the paper *Decentralized governance in DeFi: Examples and pitfalls* [7] the structure and centralization of prominent DeFi DAOs is analyzed.

The report *Uniswap Research Report: Discord, Governance, Community* [8] examines *Uniswap* governance qualitatively, while we take a quantitative approach.

While our research is limited to *Uniswap*, *Compound* and *ENS*, the 2021 research paper *Centralized Governance in Decentralized Finance (DeFi): A Case Study of MakerDAO* [9] analyzes the centralization of governance in the *Maker* protocol, a blockchain-based DeFi project.

Chapter 2

# Analysis

After downloading and processing the data, we examine delegates' voting behavior and voting power using various tools and metrics. We want to answer if the voting behavior of small token holders differs from that of large token holders, how governance tokens are distributed, and who holds power in these systems.

## 2.1 Voting Behaviour

In this section, we investigate how voting behavior differs between small and large token holders. For this purpose, we examine the approval rate and the voting frequency and try to cluster the behavior using *Principal Component Analysis*.

### 2.1.1 Approval Rate

The approval rate is the percentage of positive votes for each delegate [10]. In Fig 2.1 we see the approval rate for both *Compound* and *Uniswap* plotted against

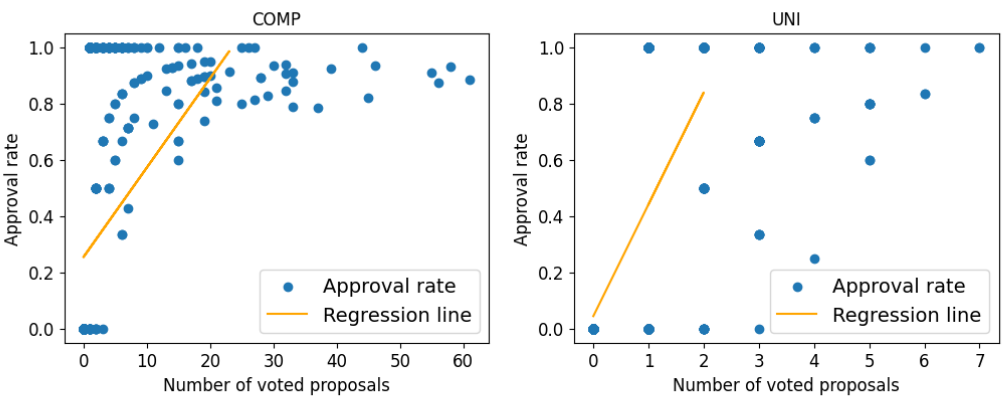


Figure 2.1: Approval rate

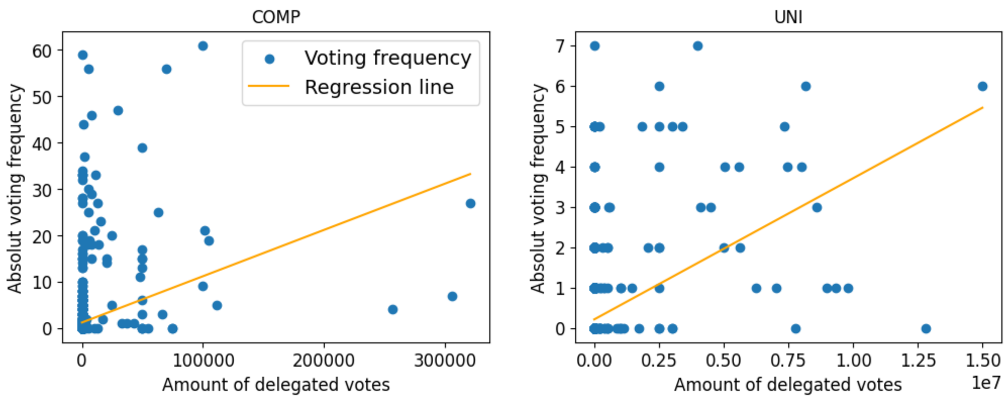


Figure 2.2: Voting frequency

the number of proposals voted by each delegate. Using a regression line we can see that the approval rate increases rapidly with the number of voted proposals. This correspondence could be explained by the fact that some delegates only vote when they have a strong opinion on a proposal. These delegates would then have a lower approval rate because there are many proposals with almost 100% approval in which they did not participate.

### 2.1.2 Voting Frequency

In Fig 2.2 we plot the absolute voting frequency against the amount of delegated votes for each delegate.
Although the upward trend indicated by the regression lines is difficult to discern, we can at least note that for both protocols, above a certain number of delegated tokens, the majority of delegates voted at least once. The only real exception is the delegate in the *Uniswap* network who received the second most delegations. This delegate did not vote once, which means that all tokens delegated to him (roughly 12.8 million UNI $\approx$ 188 million CHF as of January 13, 2022) never participated in any votes on proposals.

### 2.1.3 Behaviour Clustering

We can represent the voting behaviour of each delegate with a vector $v$ whose dimension corresponds to the number of proposals executed. We set each entry $v_i \in \{-1, 0, +1\}$, with -1 if the delegate voted negatively on proposal $i$, 0 if he did not vote and +1 if he voted positively. In the case of *Compound*, we now have a vector in 80-dimensional space that is impossible to visualize. Therefore, we use *Principal Component Analysis* (PCA) to reduce the dimensionality.
PCA accomplishes this by transforming to a new coordinate system whose basis

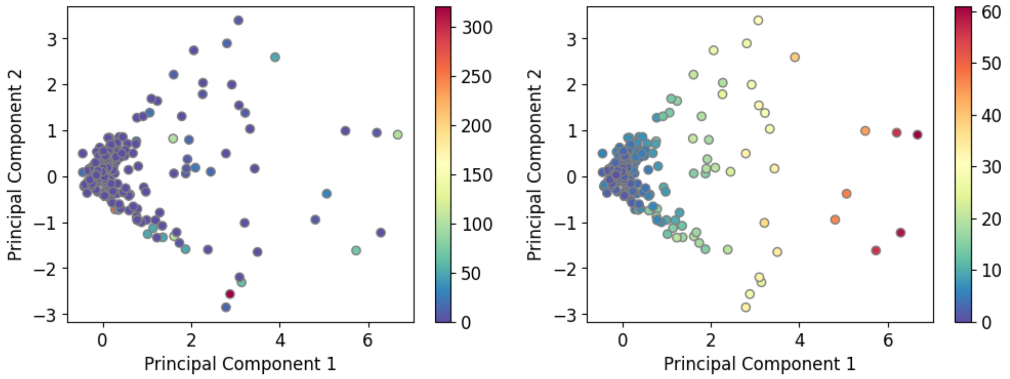


Figure 2.3: PCA Compound. On the left, each delegate is colored according to the number of tokens delegated ($\times 10^3$) to him. On the right according to how often they have voted.

vectors, the principal components (PCs), are ordered in a way that the first few contain most of the variation present in all of the original variables [11].
After this transformation, we can plot the data points along the first two principal components. This way, we can represent the voting behaviour in a two-dimensional space while losing as little information as possible.

With the help of *Scikit-learn* [12], we are able to apply PCA and calculate the PCs. Fig 2.3 shows the first two PCs of the transformed voting behavior of all delegates in the new coordinate system. On the left, each delegate is colored according to the number of tokens delegated to him. On the right side, however, the delegates are colored according to how often they have voted. In the figure to the right, it is clear that delegates who voted a similar number of times have a similar first principal component, suggesting that the voting decisions of these delegates are somewhat similar. On the other hand, delegates who voted almost all the time or almost never have a lower variance in the second PC than delegates who voted about half the time. This indicates that delegates who voted about 25 to 35 times show greater differences in their voting decisions than always-voters or never-voters.

The first PC explains 22% of the total variance in the data, while the second PC explains only 8%, the third 6%, and the fourth 4%. This means that a meaningful analysis of the other PCs is probably not possible.

We did not perform *PCA* for *Uniswap* and *ENS* because there are not enough proposals for both at this time.

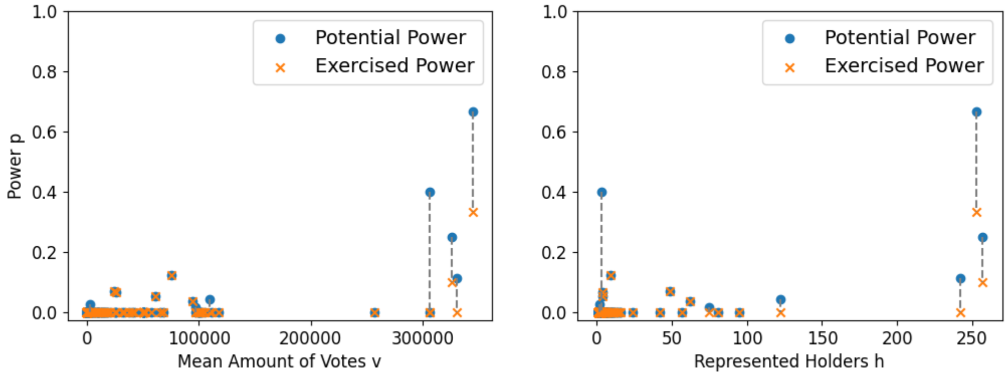


Figure 2.4: Average potential power and exercised power for the amount of average votes respectively given number of delegations COMPOUND.

## 2.2 Voting Power

In this part, we examine the issue of the distribution and exercise of voting power.

### 2.2.1 Potential and Exercised Power

Using the metrics of potential and exercised voting power, we investigate whether someone had the power to decide proposals and whether that power was actually used to change the outcome.

**Definition 2.1** (Potential Power)**.** The ability to decide a vote is calculated with the sum of weights of positive $W_p^m$ and negative $W_n^m$ votes in a voting $m$, testing if the weight $w_{im}$ of voter $i$ is bigger than the distance to quorum $q_m$ without $i$:

$$\gamma_{im}^{p} = \begin{cases} 1 & , w_{im} > q_m(W_m^p + W_m^n) - W_m^p + w_{im}\bar{v}_{im} > 0 \\ 0 & , \text{else} \end{cases} \tag{2.1}$$

where $\bar{v}_{im} \in \{0,1\}$ indicates the decision of voter $i$ in voting $m$ [10].

**Definition 2.2** (Exercised Power)**.** Similarly, we can look at the actual vote of voters and see whether the power actually was used to reverse the voting result [10].

$$\gamma_{i}^{e} = \begin{cases} 1 & , \left(\frac{W_m^p - w_{im}\bar{v}_{im}}{W_m^p + W_m^m - w_{im}} > q_m\right) \oplus \left(\frac{W_m^p}{W_m^p + W_m^n} > q_m\right) \\ 0 & , \text{else} \end{cases} \tag{2.2}$$

In Fig 2.4, we see the average potential and exercised voting power for *Compound* plotted against the average number of votes on the left and against the number of represented holders on the right.
We see that the potential voting power per delegate increases with the average

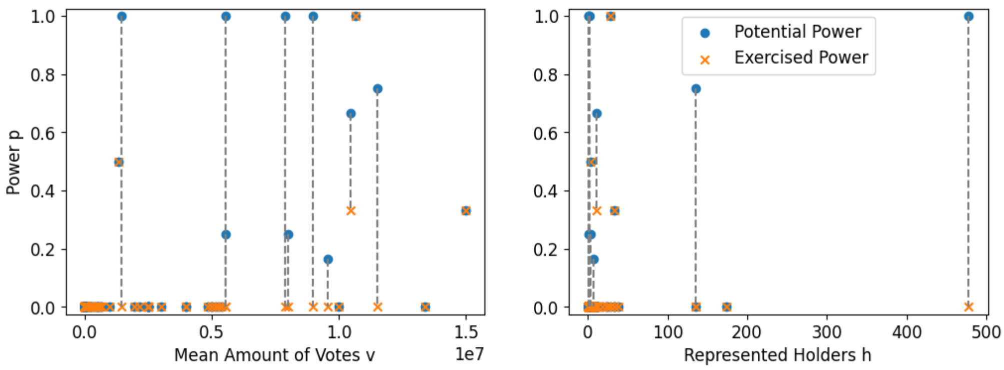


Figure 2.5: Average potential power and exercised power for the amount of average votes respectively given number of delegations UNISWAP.

number of votes per proposal. Except for a few delegates with very few holders represented, the same is true for the relationship between potential voting power and number of represented holders. It is also noticeable that most delegates do not always use their full voting power, so that in most cases the exercised voting power is lower than the potential voting power.

In the case of *Uniswap*, many delegates have a very high potential voting power. As can be seen in Fig 2.5, there are even some delegates who have an average potential voting power of 1, which means that they would have had the power to change the outcome of any proposal on which they voted. However, the voting power exercised is also mostly less than the potential voting power, which means that this power was rarely used.

### 2.2.2 Gini coefficient and Nakamoto coefficient

We measure the inequality of the distribution of delegated governance tokens using the Gini coefficient and the Nakamoto coefficient.

**Definition 2.3** (Gini coefficient)**.** The *Gini coefficient* $G$ is calculated as follows:

$$G = \frac{\sum_{i=1}^{n}(2i - n - 1)x_i}{n\sum_{i=1}^{n} x_i} \tag{2.3}$$

where the list of delegate addresses $[x_i|i \in [1, ..., n]]$ is sorted in ascending order of delegation quantity, such that $x$ has rank $i$. Here, $n$ is the number of total delegate addresses. [4]

We interpret the *Gini coefficient* as a measure of inequality that essentially calculates the closeness of the list of delegate addresses to a uniform distribution [4].

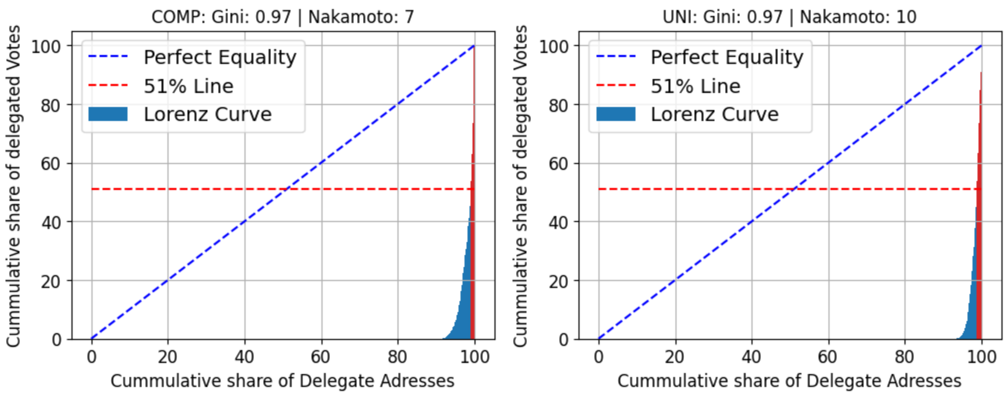


Figure 2.6: Lorenz Curve

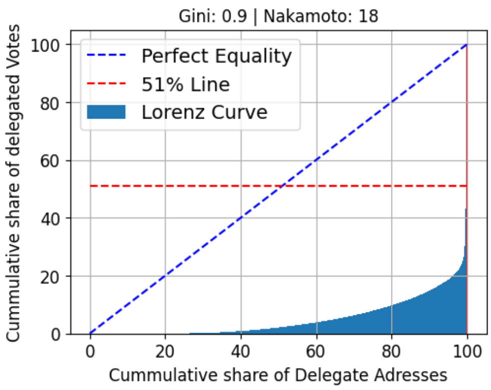


Figure 2.7: Lorenz Curve ENS

**Definition 2.4** (Nakamoto coefficient)**.** Given a distribution of delegate addresses $d$ with $K$ entities where $p_1 > ... > p_K$ are the proportions of governance tokens delegated to each of the delegate addresses $a_1, ..., a_k$ such that $\sum_i^K = 1$. Then we define the *Nakamoto coefficient* $N$ as [13]:

$$N = \min_{k \in [1,...,K]} \sum_{i=1}^{k} a_i \geq 0.51 \tag{2.4}$$

This means that the *Nakamoto coefficient* for a distribution of delegate addresses $d$ is the smallest number of entities whose shares sum up to $> 51\%$ of the delegated governance tokens. Therefore these delegates could decide every vote if they would vote the same way.

Fig 2.6 shows the result of these calculations. The *Lorenz curve*, is the cumulative distribution of the amount of incoming delegations for each of delegate. The area between perfect equality and *Lorenz curve* is the *Gini coefficient*. The

*Nakamoto coefficient* is represented by the red dotted line and the bars passing through this line indicate the smallest number of entities that together control $> 51\%$ of all delegated votes [4].

A *Gini coefficient* of 0.9, as opposed to 0.97 for *Uniswap* and *Compound*, shows that the distribution of delegations in *ENS* is much closer to an equal distribution. This is also evident from the higher *Lorenz* curve in Fig 2.7.

## 2.3 Clusters

The goal of this section is to group delegates who received the majority of their tokens from the same institution.

It is known that for each delegate there are holders who delegate their tokens to him. However, some delegates got most of their delegations from one holder, let's call him the *main delegator*.

Using the *Etherscan* API[1], we try to find out where most of the main delegator's tokens come from to see if the main delegators of different delegates receive their tokens from the same source.
To this end, we query the *Etherscan API* for incoming governance token transactions to the main delegator's address. If we find a transaction that was responsible for more than 50% of the main delegator's delegation balance, we put the sender on a list of possible sources of this main delegator. Then we iterate through this list of possible sources and query the *Etherscan API* for incoming transactions to those addresses, and if we again find one that was responsible for more than 50% of the balance, we also put this sender on the same list of possible sources corresponding to that same main delegator. In the end, we compare the lists of all the main delegators, and if an address appears in multiple, it means that these main delegators, and thus the associated delegate, have obtained most of their tokens from the same source.

Indeed, in the case of *Compound*, we are able to identify two such clusters of delegates where the *main delegators* obtained most of their tokens either directly or indirectly from the same address. This suggests that an entity controls all delegates in a cluster. However, this is not really provable as these addresses are anonymous and we cannot tell who they belong to.
The first cluster found (*Cluster 1*) contains 9 of the 20 largest delegates in terms of delegated tokens, while the second cluster (*Cluster 2*) contains 4 delegates.

For *Uniswap*, on the other hand, we find only one such cluster. This cluster contains 8 of the 50 largest delegates in terms of delegated tokens.

[1] https://docs.etherscan.io

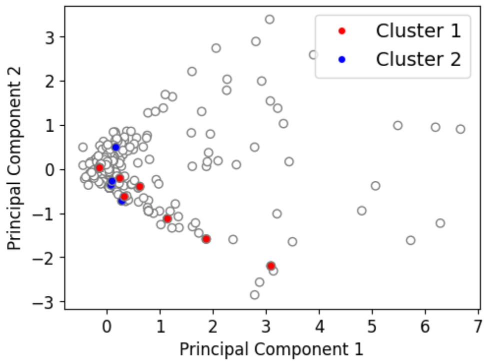


Figure 2.8: PCA Compound. Delegates which are part of the same cluster are colored accordingly.

### 2.3.1 PCA

Fig 2.8 shows the first two PCs of the transformed voting behavior of *Compound* (see chapter 2.1.3), but now color-coded according to which cluster that delegate belongs to.
At first glance, it looks like all delegates in *Cluster 2* voted very similarly. However, a closer look shows that all delegates in *Cluster 2* voted very infrequently, which could explain the closeness in the PCA plot.

To further examine the voting behavior of these delegates, we look at the average approval rates of all delegates in the *Compound* clusters. We calculate that delegates in *Cluster 1* have an mean approval rate of 0.9, excluding 4 never voters. Delegates in *Cluster 2* have an average approval rate of 0.86. Both are somewhat similar to the mean approval rate of the top 20 delegates, which is 0.9 if we exclude 4 never voters. The overall mean approval rate of all delegates is 0.92, which is also very close. From this, we can conclude that delegates in *Cluster 2* vote negatively more often than the average delegate, although the deviation is not very large.

We did not perform PCA or further behavioral analysis for the cluster found in *Uniswap* because there are not enough proposals so far.

### 2.3.2 Visualization

Using the same technique as in Chapter 1.3, we attempt to visualize the extent of the clusters for *Compound* and *Uniswap*.

For *Compound* Fig 2.9 shows, on the left, the network of the 20 largest delegates (red), as measured by delegation amount, and the holders delegated to each of these delegates (blue). An edge in this network indicates a delegation of governance tokens from the holder to the delegate. The size of the vertices

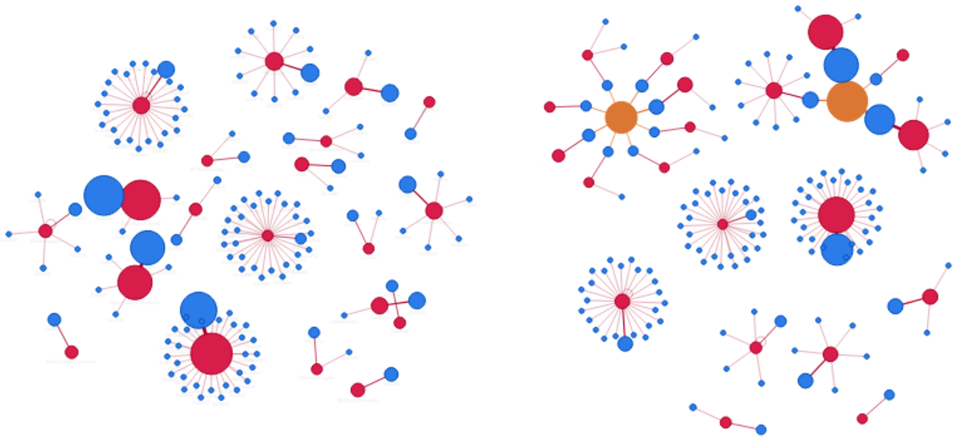

(a) (b)

Figure 2.9: Compound Network with holders in blue and delegates in red.
(a) Top 20 delegates and their represented holders
(b) Cluster sources are added in orange.

indicates the amount of delegated tokens.
On the right side, the two cluster sources are colored in orange for better illustration. This clearly shows that some *main delegators* receive tokens from the same source (orange) and are thus connected.

Fig 2.10 shows the *Uniswap* network of the 50 largest delegates on the left and on the right with the added cluster in orange. We visualize only the top 50 because the delegates in the cluster are ranked lower in delegation amount than in *Compound*. This also shows that the *main delegators* in this cluster receive tokens from the same source (orange) and are therefore connected.

## 2.4 Fairness

The *Gini coefficient* and the *Nakamoto coefficient* of delegates are good indicators of how fair delegations are in terms of distribution. We assume that a more even distribution of delegations is considered fairer for the majority of users, although it could be argued that it is not per se unfair for a person to have a lot of voting power because they have bought a lot of governance tokens.

We can clearly see that a *Gini coefficient* of 0.97 for *Compound* and *Uniswap* is very far from a fair equal distribution of delegated governance tokens. In Fig 2.6, we also see that the largest 10% of delegate addresses effectively control 100% of governance tokens. This means that the remaining 90% have virtually

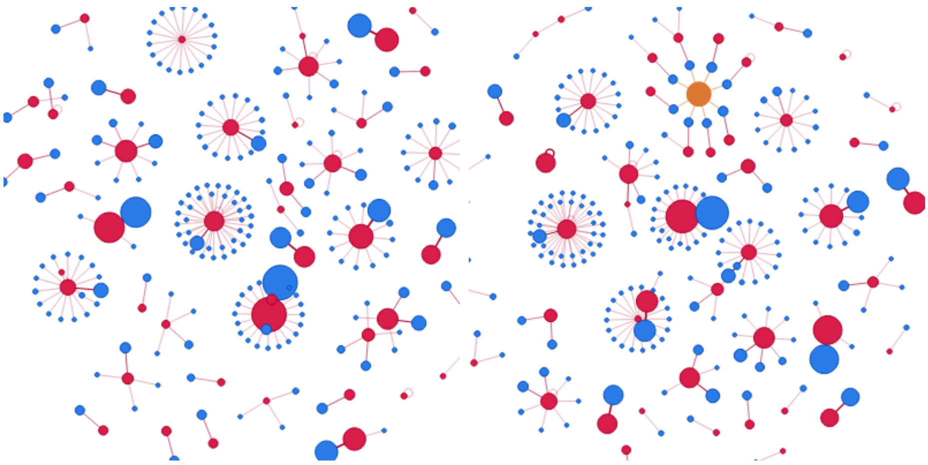

(a) (b)

Figure 2.10: Uniswap Network with holders in blue and delegates in red.
(a) Top 50 delegates and their represented holders
(b) Cluster sources are added in orange.

no impact on the outcome of the proposals.
For *ENS*, on the other hand, the *Gini coefficient* is 0.9, which is also relatively far from a equal distribution, but still much nearer than the other two. As we can see in Fig 2.7, the smallest 90% of delegates still control about 15% of the delegated tokens.

For comparison, the *World Bank*[2] calculates the *Gini coefficient* of income to be 0.63 in South Africa, 0.41 in the United States and 0.33 in Switzerland.

If we look at the visualization of the network of holders and delegates in Fig 1.1, we can see that in *Uniswap* and *Compound*, most delegates have very few holders delegating to them. We also see that most delegates have one main delegate from which most tokens originate.
As shown in Section 2.3, we are even able to identify clusters of delegates for *Compound* and *Uniswap* where the main delegators obtained most of their tokens either directly or indirectly from the same address. This suggests that one entity controls all delegates in a cluster, which in turn may indicate that the distribution of delegations may be even more centralized than previously thought.

The situation is different for *ENS*, where we see that almost all delegates have many holders delegating to them. In fact, no main delegates are identifiable. So, also from this aspect, *ENS* seems to have a more equitable distribution of delegations.
In the absence of current *ENS* proposals, it remains to be seen whether this more fair distributions will be retained.

[2]https://data.worldbank.org/indicator/SI.POV.GINI/

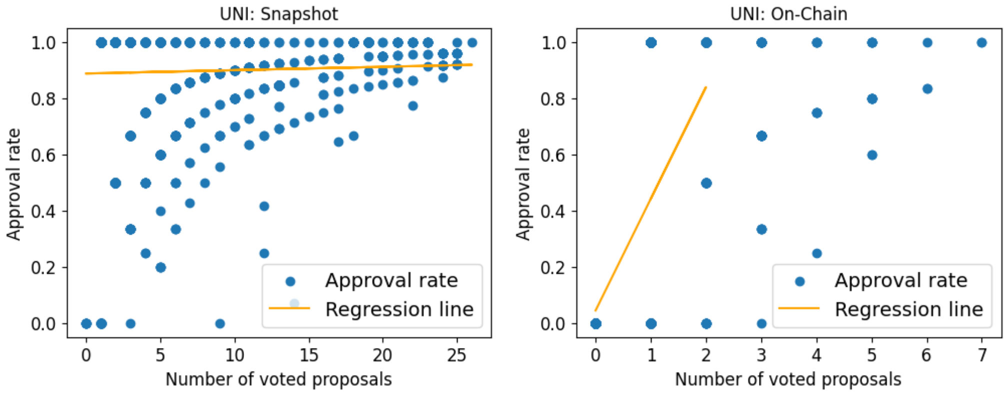


Figure 2.11: Snapshot Approval Rate

## 2.5 Snapshot

*Snapshot*[3] is an off-chain decentralized multi-governance voting system. Unlike voting on the aforementioned governance projects, creating proposals and voting on *Snapshot* does not cost any gas because the process is off-chain.
This gives us the opportunity to compare what is essentially a free off-chain voting system to the standard on-chain voting system where you pay gas fees when you make a proposal or vote on a proposal.

We did not perform analysis for *Snapshot*-voting for *Uniswap* and *ENS* because there are not enough proposals for both at this time.

### 2.5.1 Differences in Voting Behaviour

Comparing the approval rate per number of proposals voted on *Snapshot* and on the *Uniswap* chain in Fig 2.11, we see that the regression line of the approval rate on Snapshot is relatively horizontal compared to on-chain. This suggests that voting frequency is less correlated with approval rate on the *Snapshot* network than on the *Uniswap* network. This could be due to the fact that, because *Snapshot* is free, more of the smaller delegates vote on proposals that were known from the beginning to be almost 100% approved.
This result should be taken with caution because due to the fact that only 8 proposals were executed on the *Uniswap* chain, we cannot rely on the accuracy of the regression line.

[3] https://docs.snapshot.org

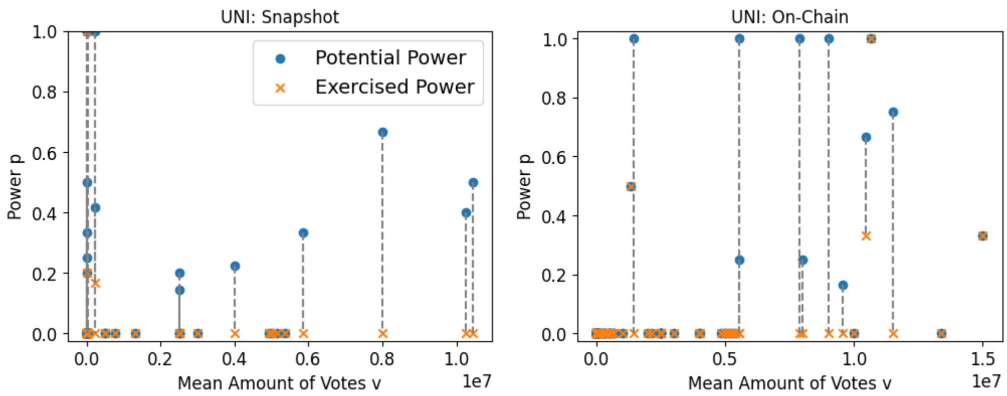


Figure 2.12: Snapshot Voting Power

### 2.5.2 Differences in Voting Power

When we compare the potential and exercised voting power of *Uniswap* on *Snapshot* and on-chain in Fig 2.12, we see that the mean potential power on *Snapshot* is much lower. The relatively high potential and exercised voting power on *Snapshot* among delegates with very few votes can be explained by the fact that proposals do not require a quorum to be reached, so there can be proposals with very low participation.
Apart from that, the mean exercised power on *Snapshot* is almost always 0, which means that the already lower potential power is almost never used.

Both findings point to more fair voting conditions on *Snapshot* then on-chain for *Uniswap*. Whether this is due to the fact that you can vote and create proposals without paying of gas fees remains an open question.

Chapter 3

# Conclusion

In this work, we aimed to examine the relatively new type of delegative governance used in an increasing number of blockchain projects.
With the help of quantitative analysis, we attempted to gain insight in the voting behavior and voting power for different blockchain-based governance projects.

Our research led to the conclusion, that the distribution of power in the case of the projects *Uniswap* and *Compound*, is very centralized to a few big players. We also saw that we can observe a similarity in the voting behavior of delegates, that vote with similar frequency.

In comparison, the situation is different for the relatively new *ENS* project. *ENS* is currently much more evenly distributed in terms of power. However, it remains to be seen whether this will continue, as there are no proposals so far.

We were also able to show the existence of cluster in *Compound* and *Uniswap* where some delegates receive most of their delegations indirectly from the same source. Whether this means that delegates are controlled by the same entity, consequently changing the centralization of the delegation distribution, remains an open question for future research.

In addition, we looked at *Snapshot*, an easy-to-use and free voting platform. In the case of *Uniswap* we were able to identify possible fairer voting conditions. Further research is needed to determine if *Snapshot* could be a possible, fairer alternative to on-chain voting in the future.

# Bibliography


[1] A. Antonopoulos, G. Wood, and G. Wood, *Mastering Ethereum: Building Smart Contracts and DApps.* O'Reilly Media, Incorporated, 2018.

[2] H. Adams, N. Zinsmeister, M. Salem, R. Keefer, and D. Robinson, "Uniswap v3 Core," 2021.

[3] R. Leshner and G. Hayes, "Compound: The Money Market Protocol," 2019.

[4] J. R. Jensen, V. von Wachter, and O. Ross, "How Decentralized is the Governance of Blockchain-based Finance: Empirical Evidence from four Governance Token Distributions," 2021.

[5] R. Leshner, "Compound Governance," 2020, accessed: February 10, 2022. [Online]. Available: https://medium.com/compound-finance/compound-governance-5531f524cf68

[6] A. A. Hagberg, D. A. Schult, and P. J. Swart, "Exploring Network Structure, Dynamics, and Function using NetworkX," in *Proceedings of the 7th Python in Science Conference*, 2008, pp. 11–15.

[7] K. Stroponiati, I. Abugov, Y. Varelas, and K. Stroponiatis, "Decentralized governance in DeFi: Examples and pitfalls." 2020.

[8] X. Sun, "Centralized Governance in Decentralized Finance (DeFi): A Case Study of MakerDAO," 2021.

[9] K. Stroponiati, I. Abugov, Y. Varelas, and K. Stroponiatis, "Decentralized governance in DeFi: Examples and pitfalls." 2021.

[10] C. C. Kling, J. Kunegis, H. Hartmann, M. Strohmaier, and S. Staab, "Voting Behaviour and Power in Online Democracy: A Study of LiquidFeedback in Germany's Pirate Party," 2015.

[11] I. T. Jolliffe, *Principal Component Analysis.* Springer Verlag, 1986.

[12] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, and B. Thirion, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[13] B. S. Srinivasan and L. Lee, "Quantifying Decentralization," 2017, accessed: February 10, 2022. [Online]. Available: https://news.earn.com/quantifying-decentralization-e39db233c28e