The Impact of Delegation Splitting on Effective Decision Making for Decentralized Autonomous Organizations

Justin Mack
Harvard College
B.A. Computer Science
Cambridge, MA
justinmack@college.harvard.edu

Mohib Jafri
Harvard School of Engineering and Applied Sciences
M.S. Electrical and Computer Engineering
Cambridge, MA
mjafri@g.harvard.edu

Sherry Xie
Harvard College
B.A. Applied Math
Cambridge, MA
sherryxie@college.harvard.edu

Hunter Siebel
Harvard College
B.A. Computer Science
Cambridge, MA
hsiebel@college.harvard.edu

***Abstract*—We present a theoretical model for the voting system used by many decentralized autonomous organizations (DAOs) to make decisions. Within this general framework, we define several distributions representing potential allocations of voting weight across different voters, competence levels for each voter, voter perceptions of one another, and how voters choose who to delegate to. We then simulate the voting process many times to determine how the maximum number of voters one is allowed to delegate to influences (i) the group's probability of selecting the correct alternative, (ii) the expected vote share of the correct alternative, and (iii) the distribution of individual voters' influence on the final outcome. Our results show that DAOs should increase the maximum number of allowed delegations from 1 to 10 as this increases the decision success rate and expected vote share of the correct alternative or makes voter influence more equitable in all real-world scenarios.**

I. Introduction

A. Background - Liquid Democracy

Liquid democracy describes a voting model where agents are permitted to delegate their voting power to other agents, who vote on their behalf. Standard liquid democracy assumes that every agent has one vote to cast, and can choose to delegate that vote to a single agent (including themselves), as long as that delegation does not create a cycle.

Standard liquid democracy is usually modeled via a labeled, directed graph. Each vertex is a voter and its directed edges represent delegations from one voter to another. In simplistic models, one assumes two alternatives: "correct" and "incorrect". Each voter has a competence level, which is their probability of voting correctly. In standard liquid democracy, voters will delegate to other voters with higher competency than their own.

Liquid democracy outperforms direct democracy in many cases [9]. When low competency voters delegate to higher competency voters, the probability of group accuracy ought to increase. However, it can be worse than direct democracy when the model extends to include voters' inaccurate guesses of competence and results in a strong consolidation of power to a low competency voter.

Finding a mechanism for optimal delegation under standard liquid democracy is a complex problem. It is known that approximating an optimal delegation within an additive term of $\frac{1}{16}$ under standard liquid democracy is an NP-hard problem [10]. This problem becomes more complex as standard liquid democracy is modified to accept partial delegations, weighted voters, local delegation mechanisms, and other variations.

B. Background - DAOs

A particularly interesting application of liquid democracy can be observed in the recent rise of decentralized autonomous organizations (DAOs). While more academic work is required to offer an operational definition for a DAO, we strictly explore its mechanisms with respect to voting. DAOs use tokens– governance ERC-20's, transferable ERC-721's, among other technologies– held by a set of addresses on a blockchain network to vote on proposals [11]. Votes cast onto the blockchain are verified by validators via zero-knowledge proofing techniques to ensure trustless verification of voting preferences [7], and these votes can be computed through any voting strategy to trigger on-chain events such as a movement of tokens or a change in protocol parameters [11]. It is of high interest to explore these mechanisms deeply, considering the market size of 10.4 billion in DAO treasury funds, 1.7 million governance token holders, and thousands of voting events occurring every week across the 4,500+ DAOs in existance [4].

Unique dynamics appear in DAOs when evaluating the token utilization for proposals, particularly with respect to voter participation and respective weight distributions. One of

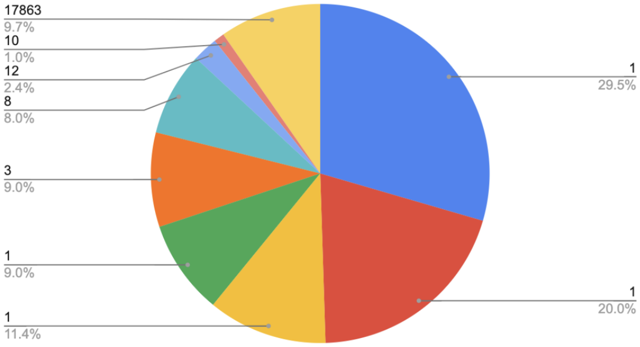

Fig. 1. BitDAO's effective voting power distribution, as measured by token holder count.

the largest DAOs by treasury size–BitDAO, a DAO-directed treasury managing over two billion USD worth of tokens–has over 17,000 stakeholders, and yet only 34 actively voting members according to DAO analytics platform DeepDAO [3]. Moreover, of the total stakeholders, 90.3% of the governance power of the DAO is held by 0.2% of the stakeholders; as shown in Figure 1, 37 total addresses largely dictate the direction of the DAO.

While this exponential difference in token ownership is to be expected when both large venture firms and retail individuals converge to vote together, DAOs employ the practice of delegation to increase token-participation. A popular voting framework, Moloch, has been adopted by 600+ DAOs, including investment DAOs like MetaCartel Ventures and Raid Guild [1]. Under Moloch's protocol, each proposal sent to a DAO has two alternatives: approve or reject. Token holders may allocate all of their tokens to a binary alternative or delegate their decision to another member [6]. Little research has been conducted to exhaustively highlight this model as an optimal voting strategy. Therefore, our paper takes an empirical approach to finding optimal delegation splitting rules that DAOs ought to employ to robustly approach good decisionmaking under a wide range of model parameters.

C. Literature Review

Since it was introduced in 2009 [9], there have been many papers exploring various aspects of liquid democracy. Such literature tends to fall into a few categories: identifying the drawbacks of liquid democracy [2], developing better mechanisms for delegation [8] [12], and analyzing the computational aspects of these proposed delegations [5].

Reference [12] in particular studies an extended model of liquid democracy in which agents are not restricted to delegating their vote to a single agent but can instead employ a mixed delegation strategy – a probability distribution demarcating the agent they will delegate their vote to. The paper finds that with mixed delegations, it is possible to identify delegation structures that optimize the accuracy of the group.

The paper finds and defends an algorithm for calculating optimal weights for each voter's delegation distribution that will maximize group accuracy. The authors run various simulations using this algorithm to compare it to standard liquid democracy and greedy algorithms that attempt to maximize individual accuracy (as opposed to group accuracy).

The authors conclude by summarizing their findings to claim:

1) Weighted delegations enable optimal group accuracy and better equilibria than standard liquid democracy
2) Nash Equilibria (NE) in greedy delegation (GD) games with weighted profiles are shown to never be worse than NE in GD games with pure delegations (and are sometimes better)

The final aspect of the paper discussed three future areas of exploration. The first serves as the primary inspiration for our project– weighted delegations as partial allocations as opposed to probability of delegation. The second explores scenarios in which voters cannot freely delegate to all other voters. The final topic concerned testing weighted delegations under different voting mechanisms, such as random dictatorship.

In our paper, we build off of this model and extend it further to reflect the voting processes of decentralized autonomous organizations by introducing agents with varying numbers of initial votes, as well as by explicitly allowing agents to partially delegate to multiple other agents, up to a global delegation cap.

D. Goals

We aim to determine the relationship between the global delegation limit k and

1) The probability of selecting the correct alternative.
2) The expected value of the vote share of the correct alternative.
3) The distribution of voters' influence on the final outcome.

II. Modeling the DAO Voting System

A. Model Specifications

There is a set of two **alternatives** $A = \{a_1, a_2\}$, a set of n **agents** $N = \{1, 2, \dots, n\}$ indexed by i, and a global delegation cap $k \in \{1, 2, \dots, n\}$.

Each agent can be fully described by the following set of properties:

- **Weight** w_i which specifies the number of tokens (i.e. votes) they start with.
- **Competence** $c_i \in [0, 1]$ which is their probability of voting for the correct alternative.
- **Perceptions** $\mathbf{p_i} \in [0, 1]^n$ which is their estimate of every agent's competence.
- **Limit** $\ell_i \leq k$ which is the number of agents they will attempt to delegate to.

Given these properties, each agent in succession selects a delegation set $D_i \subseteq N$. To do so, they check if adding the agent with the highest perception score would cause a delegation cycle. If not, this agent is added to the delegation set. This agent then has their perception set to zero, and the process is repeated until the delegation set contains ℓ_i agents, or all agents have a perception score of zero. An agent can always delegate to themselves without creating a delegation cycle, so every delegation set is guaranteed to be nonempty.

This method of ensuring the delegation graph is realistic for the case of decentralized autonomous organizations, because whenever an agent in the DAO makes delegation decisions they do so effectively on a fixed delegation graph simply due to the frequency of delegations in real time. Some DAOs disallow delegation decisions that would create cycles, while others instead remove voters in cycles from the graph. We arbitrarily chose the first approach, but find it unlikely that this choice would alter our findings.

After picking a delegation set, each voter then splits their weight w_i among the agents in the set according to a global **split** distribution.

Finally, voting weight is delegated according to the delegation sets and splits, and votes are cast according to the competence values. The winner is selected by simple majority.

B. Parameters

The model is defined over the parameters n and k, as well as 5 distributions: weights, competences, perceptions, limits, and splits. We define a number of different options for each of these initial conditions.

1) Weights: The number of votes each agent starts with. This distribution represents the power of each agent in the DAO.

- **Exponential**: Weights are distributed according to an exponential distribution with $\lambda = 0.7$ and normalized using the softmax. This reflects the high degree of token concentration usually seen in large DAOs [3].
- **Even**: Each agent starts with a fraction $\frac{1}{n}$ of the voting power. This better reflects traditional voting systems such as regular democratic elections.

2) Competences: The probability that each voter chooses the correct alternative.

- **Normal** (μ, σ): Perceptions are distributed randomly between 0 and 1 (inclusive) according to a truncated normal distribution with mean μ and standard deviation σ.
- **Weighted**: Higher weight agents have higher competences.

3) Perceptions: The metric that agents use to choose who to delegate to. Intuitively, this is each agent's guess for the true competence of every other agent. However, the level of generality allows it to also capture any other factors that an agent may use in determining their delegation decision. This distribution can also be used to model a neighborhood graph by having agents set the perception score to zero of agents they do not know.

- **Uniform**: Perceptions are distributed randomly between 0 and 1 (inclusive) according to a uniform distribution.
- **Normal**(μ, σ): Perceptions are distributed randomly between 0 and 1 (inclusive) according to a truncated normal distribution with mean μ and standard deviation σ.
- **True**: Perceptions are equivalent to the true competences.
- **Noisy True**: Perceptions are equivalent to true competences with added noise distributed normally between -0.2 and 0.2.
- **Inverse**: Perceptions are the inverse of the true competences.
- **Weighted**: Agents assign higher perception scores to agents with higher weight.

4) Limits: : The upper bound on the number of delegations is k, but agents can freely choose to delegate to fewer people. This distribution defines how many people each agent chooses to delegate to.

- **Uniform**: Limits are distributed randomly between 1 and k (inclusive) according to a uniform distribution.
- **Lazy**: Limits are distributed randomly between 1 and k (inclusive) according to a truncated normal distribution with $\mu = 1, \sigma = 1$.

5) Splits: : the distribution by which to split votes to k voters.

- **Weighted**: Agents distribute weight proportionally to other agents in their delegation set depending on perception values.

C. Simulation Details

We simulated this model for $n = 50$ and $k = \{1, 2, \dots, 50\}$ under a variety of different distributions. For each value of k, we ran 1,000 trials to account for the randomness introduced by the distributions.

III. Experimental Setup

For each set of distributions simulated, we collected the following data:

1) Accuracy Results:
   For each trial, we recorded the total votes allocated to the correct alternative. This permits the calculation global delegation impact on the expected vote share, the expected success probability, and the variance of both.
2) Concentration Results:
   For each trial, we also recorded the post-delegation vote share of the agent with the largest such value. This data allowed us to determine the impact the global delegation limit has on the concentration of power within the DAO.

We simulated the model for the following choices of distributions, each of which was assigned a label which appears in the plots below. The distributions are listed in the order weights, competences, perceptions, limits, splits.

- **ENNUW**: Exponential, Normal, Normal, Uniform, Weighted
- **25ENNUW**: Exponential, Normal, Normal, Uniform, Weighted (with a decision threshold of 0.25)
- **ENTUW**: Exponential, Normal, True, Uniform, Weighted
- **ENIUW**: Exponential, Normal, Inverse, Uniform, Weighted

We have included results for the aforementioned distribution sets as we believe them to be most realistic for DAO settings. For weights, an exponential distribution reflects the highly concentrated distribution of token ownership in most DAOs. For competences, the distribution is normal to reflect that most members of a DAO are similar, but there are a few outliers in both directions. For splits, we believe that distributing tokens proportionally based on perceived competence is a reasonable approximation of human behavior.

IV. Results

In [12], prior to the introduction of a centralized solution, agents greedily maximize their own accuracies. Though the expected accuracy for any given decision (tokens cast for the correct alternative / total tokens delegated) may be maximized by delegating to the single voter with the greatest perceived accuracy, diversification is a widely employed investment strategy that an outcome-optimizing (number of times correct / number of decisions made) agent could be expected to employ. Indeed, beyond simply looking for the optimal voting rule for DAOs, this behavior further motivates our consideration of various values of k, or the maximum number of voters an individual may delegate to. The remainder of this paper will refer to these metrics, for the entire group, as **"accuracy maximization"** and **"outcome optimization"**.

A. Accuracy Maximization

We first consider the expected accuracy for any given decision (number of votes for the correct alternative / total number of votes cast) at the group level. As shown in Figure 2, when agents are good at identifying the competence of other agents, increasing k has a negative effect on the expected accuracy for each decision; this makes sense because increasing k is akin to diluting D_i with less competent voters. As shown in Figure 3, when perceptions are the inverse of true competence levels (i.e. voters are inept at discerning one another's competence levels), increasing k improves the expected accuracy for any given decision. Here, delegating to more people increases both the maximum and weighted average competence level of voters in D_i. Finally, as shown in Figure 4, when perceptions are random (e.g. uniform or normal), increasing k does not change the expected accuracy for any given decision. This also aligns with our intuition; when perceptions are random, each additional voter one may delegate to has equal probabilities of being more or less likely to choose the correct alternative than voters already delegated to.

Though increasing k does not consistently improve expected accuracy, it is worth noting that small increases in k significantly lower the variance across independent decisions for all of these distributional choices. This may be an optimal property depending on how expected accuracy is distributed across trials; for example, given two normal distributions with $\mu = 0.6$, drawing from one with $\sigma^2 = 0.1$ will give more values greater than 0.5 (i.e. the threshold for a majority rules vote) than one with $\sigma^2 = 0.2$. This is explored further in the following section.

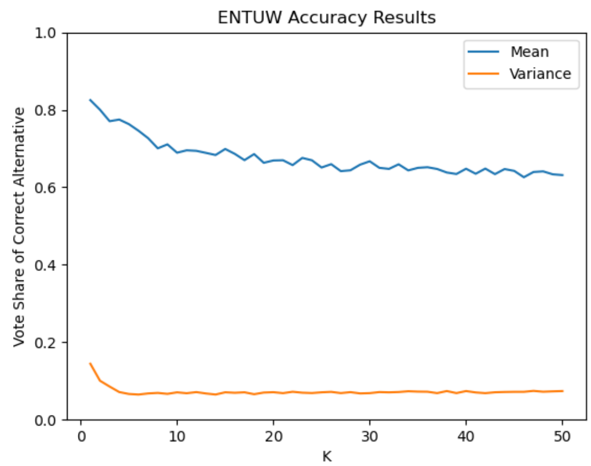

Fig. 2. Perception: True. Competence: $\mathcal{N}(0.6, 0.1)$.

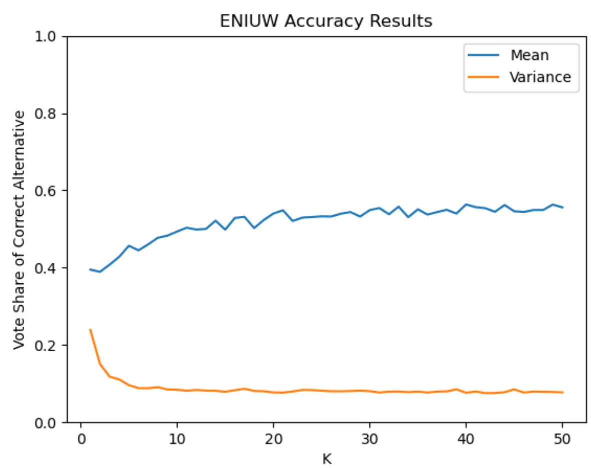

Fig. 3. Perception: Inverse. Competence: $\mathcal{N}(0.6, 0.1)$.

B. Outcome Optimization

By mapping the percentage of correct votes out of total votes cast to a binary value (1 if at least 50% of votes cast are correct, 0 otherwise), we can investigate how increasing k affects actual decision outcomes. We do not include plots for the True (or Inverse) perception distributions as the results are fairly intuitive; we saw previously that expected accuracy per decision is decreases (or increases), and the decision success rate decreases (or increases) as well. We instead focus on the random (Normal) perception distribution. The decision success rate for the distribution from Figure 4 is shown in Figure 5.

At first, it may seem surprising that the decision success rate does not improve, even though the variance in expected accuracy per decision decreases. However, looking at the distributions of expected accuracy across trials explains this phenomena; variance decreases as k increases because our expected accuracy per decision becomes less bimodal, but the

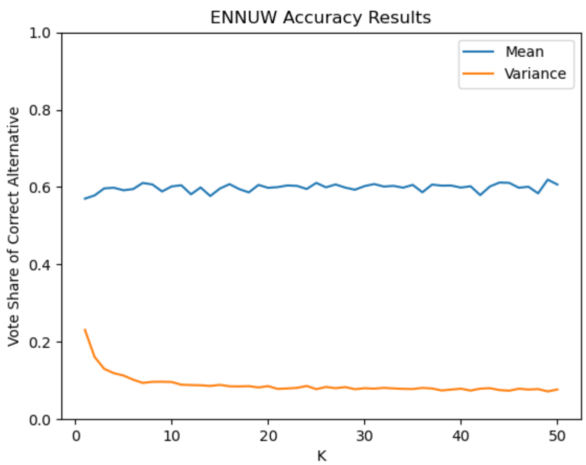

Fig. 4. Perception: $\mathcal{N}(0.5, 0.1)$. Competence: $\mathcal{N}(0.6, 0.1)$.

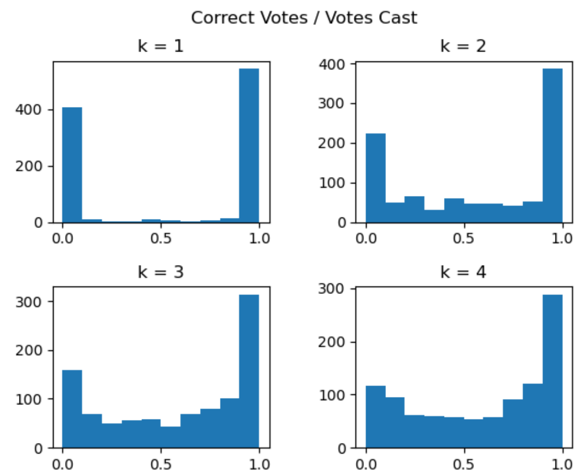

Fig. 6. Perception: $\mathcal{N}(0.5, 0.1)$. Competence: $\mathcal{N}(0.6, 0.1)$.

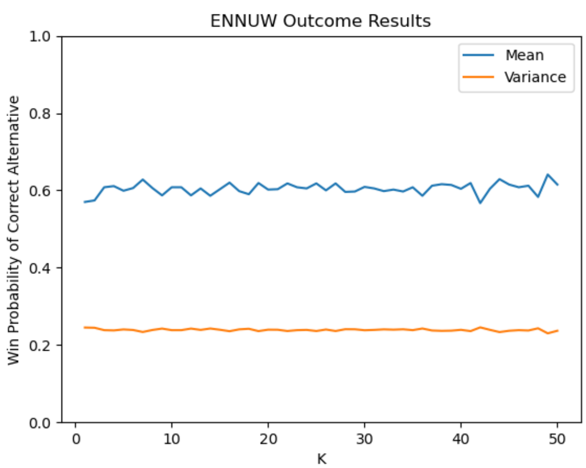

Fig. 5. Perception: $\mathcal{N}(0.5, 0.1)$. Competence: $\mathcal{N}(0.6, 0.1)$.

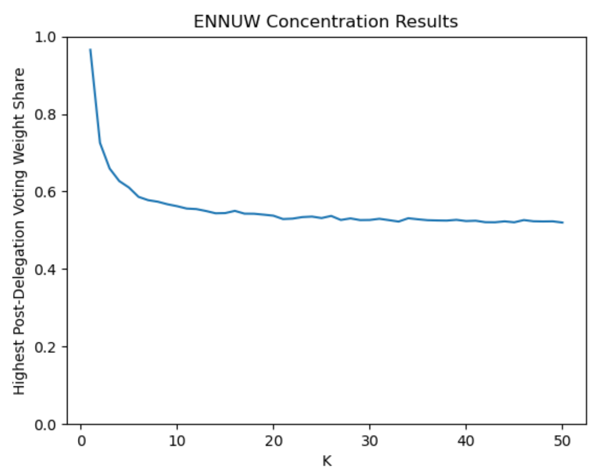

Fig. 7. Perception: $\mathcal{N}(0.5, 0.1)$. Competence: $\mathcal{N}(0.6, 0.1)$.

number of instances on each side of the 0.5 threshold does not change. Specifically, the bimodal distribution in the top left of Figure 6 is the result of the formation of a long delegation chain, in which the final vote is cast by a single voter whose competence was drawn from $\mathcal{N}(0.6, 0.1)$.

Thus, we find that allowing delegation to more voters has the effect of increasing voter participation at no cost to expected accuracy for a given decision or the success rate across multiple decisions. This is shown more explicitly in Figure 7. Additionally, we can see from Figure 8 that for different thresholds, the decision success rate does in fact increase with k.

V. Discussion

Several interesting discussions arise from the results of our simulations.

First, under the accuracy maximizing lens we observed that for perceptions inversely correlated with true competence levels, increasing k does increase the group's expected accuracy for each decision. Inverse perceptions (highest willingness to delegate to those with the lowest true competence levels) may occur in the real world. Suppose certain members of a DAO have low true competence but maintain disproportionately large platform audiences on social media. These members (a) appear knowledgeable enough for the many people who see them to delegate to them and/or (b) gain a cult following where people delegate to them regardless of their actual or portrayed competence. Regardless of the reason for delegation, both (a) and (b) are reflected in our perception distribution. Thus, when DAO membership is subject to inverse perceptions, encouraging voters to delegate to more people improves expected accuracy by increasing the likelihood of delegating to a more competent voter.

Second, across a variety of distributional choices one robust result is that increasing k reduces the maximum post-delegation vote share. We see that for small values of k, long

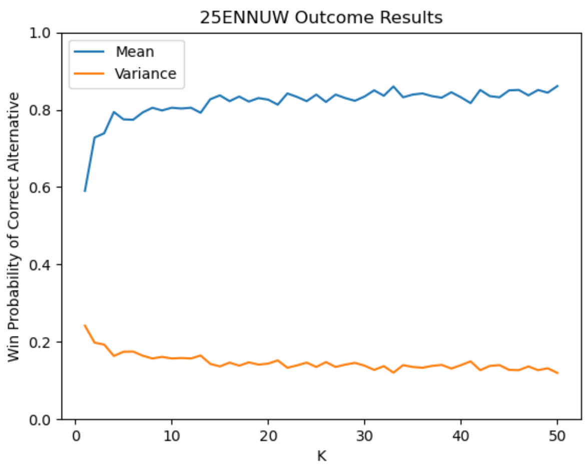

Fig. 8. Perception: $\mathcal{N}(0.5, 0.1)$. Competence: $\mathcal{N}(0.6, 0.2)$.

delegation chains form because few agents choose to delegate to themselves, resulting in extreme concentration of voting power. This dynamic holds truer in larger DAOs like BitDAO [3]. However, as k increases, each agent becomes more likely to delegate at least some weight to themselves, and a more equitable distribution emerges. The extreme diminishing returns we see for this effect suggests that DAOs could achieve a more even distribution of power by allowing a very small amount of vote splitting.

Lastly, we see that with random perception distributions and voting thresholds below 0.5, increasing k can increase the decision success rate. Though a DAO is unlikely to make a decision where the "correct" alternative is known prior to the vote, consider the following illustrative scenario in which this result could be useful. Suppose a referendum is up for voting, and the "correct" alternative is "No". If "No" receives at least 25% of votes, then the referendum will be vetoed. Thus, increasing k will increase the likelihood that the referendum is vetoed.

VI. Conclusion

Overall, we recommend that DAOs increase the maximum number of allowed delegations from 1 to 10. Our simulations shows that this k value is optimal when one cannot make an informed choice on various voter perception distributions. Additionally, it is socially more attainable as as larger k will overwhelm voters and further decrease voter participation. If voter perceptions are inversely related to true competence levels, increasing the number of allowed delegations has a clearly positive effect on the group's expected accuracy as well as the decision success rate. If voters have limited perception accuracy (but not inversely so), increasing the number of allowed delegations will increase voter participation without sacrificing expected accuracy or decision success rate. Only when voters' perceptions are highly correlated with true competence levels will increasing the number of allowed delegations decrease the vote share of the correct alternative as well as the decision success rate. However, in the real world, voters are unlikely to have very accurate assessments of one another's competence levels, especially when DAO members may be anonymous.

References

[1] Gabriel Shapiro Ameen Soleimani, Peter 'pet3rpan'. *MetaCartel Ventures Whitepaper*, 2019 [Online].

[2] Ioannis Caragiannis and Evi Micha. A contribution to the critique of liquid democracy. In *IJCAI*, pages 116–122, 2019.

[3] DeepDAO. Bitdao: Organization 30d12da9-5fd5-411a-8ca8-1422a4ce1373, Apr. 1, 2022 [Online].

[4] DeepDAO. Deepdao home page, Apr. 1, 2022 [Online].

[5] Bruno Escoffier, Hugo Gilbert, and Adele Pass-Lanneau. Iterative delegations in liquid democracy with restricted preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1926–1933, 2020.

[6] Ameen Soleimani et. al. *Moloch v2*, 2019 [Online].

[7] Youssef Faqir-Rhazoui, Javier Arroyo, and Samer Hassan. A comparative analysis of the platforms for decentralized autonomous organizations in the ethereum blockchain. *Journal of Internet Services and Applications*, 12(1), 2021.

[8] Paul Gölz, Anson Kahng, Simon Mackenzie, and Ariel D Procaccia. The fluid mechanics of liquid democracy. In *International Conference on Web and Internet Economics*, pages 188–202. Springer, 2018.

[9] Alois Paulin. An overview of ten years of liquid democracy research. In *The 21st Annual International Conference on Digital Government Research*, dg.o '20, page 116–121, New York, NY, USA, 2020. Association for Computing Machinery.

[10] Ariel Procaccia. Liquid democracy. http://procaccia.info/courses/CS238-S22/cs238s22-7.pdf, 2022. Accessed: 2022–27-05.

[11] Alexandra Sims. Blockchain and decentralised autonomous organisations (daos): The evolution of companies? *New Zealand universities law review*, 28(3):423–458, 2019.

[12] Yuzhe Zhang and Davide Grossi. Tracking truth by weighting proxies in liquid democracy, 2021.