Liquid Democracy in Extended Settings

Aditya Khambete
IIT Bombay
Mumbai, India
khambete@iitb.ac.in

Mayank Sandh
IIT Bombay
Mumbai, India
22b0962@iitb.ac.in

Aditya Sanapala
IIT Bombay
Mumbai, India
23b0912@iitb.ac.in

ABSTRACT

We study the setting of liquid democracy, where the voters are allowed to vote directly or delegate there votes to those more informed. In this paper we propose some extension to existing local delegation mechanism of `GreedyCap`, and then extend the setting to from binary alternatives (1 correct, 1 incorrect) to multiple alternatives.

KEYWORDS

Liquid Democracy, Delegation Mechanism, Positive Gain, Do No Harm, GREEDYCAP, Local Mechanisms, Multi-option Voting

ACM Reference Format:
Aditya Khambete, Mayank Sandh, and Aditya Sanapala. 2023. Liquid Democracy in Extended Settings. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 4 pages.

1 INTRODUCTION

In a Direct Democracy, the agents can vote on every issue by themselves. In Liquid Democracy, in addition to this, agents can delegate their votes transitively. For example if A delegates her vote to B, and B does so to C. In this case, agent C will bear the weight of votes of all three agents, including herself. The liquid metaphor could be seen through, how the vote of A flows to B, and hence to C. More generally the votes flow freely in a delegation graph until they reach a sink. The roots of liquid democracy could be traced back as long as Miller III [7], the key difference is here the proxies who are selected by the voters to vote on their behalf, may delegate their own vote to a proxy. In recent years, this approach has been implemented and been used on a large scale, particularly by eclectic political parties such as German Pirate Party in Germany, and Demoex in Sweden. One of the main reasons for the success of liquid democracy is it provides the best of both worlds of direct democracy and representative democracy.

The existing work in liquid democracy focuses majorly on binary outcomes, where one choice is correct and one is incorrect, and there are some papers such as Bersetche [1] focus on multi-agent delegation mechanism, but none go into generalising it for more choices.

In this paper, we study some existing local mechanisms, mainly `GreedyCap` from Kahng et al. [4] try to optimize them a bit more, and also extend them to the general setting of non-binary outcomes with single-agent delegation.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

2 CONTRIBUTIONS

This project builds upon the foundational work of Kahng et al. [4], which introduces the concepts of Positive Gain (PG), Do No Harm (DNH), and the non-local GREEDYCAP mechanism, which satisfy both DNH and PG in a restricted setting.

- We propose two variations of the `GreedyCap` Algorithm-
  - `Modified GREEDYCAP`: Prioritizes delegation from less competent voters when applying the cap.
  - `DOUBLEGREEDYCAP`: Processes the least competent voters first and delegates to the most competent available neighbor under the cap. Our simulations suggest this approach yields improved empirical performance compared to the original GREEDYCAP.
- We propose the extension of LD to non-binary models, based on the idea of Mallows [6]. This utilises the kendall tau distance Kendall [5] and the defined dispersion parameter ϕ_i for each agent. This allows us to find the probability distribution of a person over the alternatives. This naturally extends, to allow the notions of
  - PG, DNH
  - Local/Non-Local Mechanisms
  - Specific Mechanisms, such as `GreedyCap`

  in the general setting as well.

3 RELATED WORK

Kahng et al. [4] is a foundational work, and establishes a clear distinction between local and non-local delegation mechanisms. Establishes that local mechanisms can't guarantee better results than direct voting, and propose some non-local mechanisms which perform better under a certain set of assumptions.

Gölz et al. [3] proposes allowing agents to specify multiple delegation options instead of just one to avoid the weakness of liquid democracy that a small subset of agents may gain massive influence

Brill et al. [2] introduce multi-agent ranked delegations, enabling agents to specify backup delegates through preference rankings.

Tyrovolas et al. [8] extends delegation through monotonic Boolean functions, allowing conditions like "vote yes if a majority of friends do." They analyze MinSum and MinMax unravelling procedures, proving computational hardness for MinSum with non-trivial functions while showing polynomial-time solutions for restricted cases like OR/AND operations.

4 BINARY OUTCOME SETTING

4.1 Model

We represent the first instance of our problem, which is in binary, very similarly to that in Kahng et al. [4]. Let

$$G = (V, E, \vec{p}) \text{ ,where } V = \{1, 2 \cdots n\}$$

denote the graph of voters, where E represents the pair of voters which know each other. Each voter is labeled by his competence level $p_i \in [0, 1]$, which is the probability of him currently voting on some issue.

Further, we define a constant α which allows us to define the set of people a voter might delegate their vote to.

We assume i approves j iff and only if-

$$p_j > p_i + \alpha$$

This helps us to define naturally the set of all approved voters by i

$$A_G(i) = \{j \in V \mid (i, j) \in E, i \text{ approves } j\}$$

Each agent delegates his vote using some delegation mechanism M, which takes inputs the graph G as defined above, and outputs a probability distribution over $A_G(i) \cup \{i\}$. Once we have this probability distribution we apply M to G, and sample the resulting probability distribution to get a delegation graph, with each sink (vertex with no outgoing edge) having weight equal to the number of vertices with directed path leading to that sink. Each sink votes with his probability of p_i, and based on the weighted majority we make a decision.

The probability of this process yielding the correct decision is denoted by $P_M(G)$. Notice, there is always a trivial mechanism where no delegation happen, we call that the 'direct mechanism' and denote it by D, and hence we denote the probability of correct decision here as $P_D(G)$

A mechanism is called local, where each voter makes a independent delegation delegation, without any central coordination.

Since we want to compare the difference between direct voting and given mechanism, we define the gain of a mechanism M over G as-

$$Gain(M, G) = P_M(G) - P_D(G)$$

Based on this we define some properties of a mechanism,

Definition 4.1 (Do No Harm (DNH)). A mechanism M satisfies DNH if $\forall \epsilon > 0, \exists n_1$ such that for all graphs G_n with $n \geq n_1$ vertices, $Gain(M, G_n) \geq -\epsilon$

Definition 4.2 (Positive Gain (PG)). A mechanism M satisfies PG if $\exists \gamma > 0, n_0 \in \mathbb{N}$, such that for all $n \geq n_0, \exists G_n$ with n vertices such that, $Gain(M, G_n) \geq \gamma$

Intuitively, we can see PG says there is some graph where gain is positive, while DNG says that after a certain point, gain can't be lower than a certain threshold.

4.2 Main Results

We state the following result without proof, from Kahng et al. [4]

Theorem 4.3. *For any $\alpha_0 \in [0, 1)$, there is no local mechanism that satisfies the PG and DNH properties.*

The main idea behind proof is after a certain point, liquid democracy can lead to a point where mistakes of few popular voter might tip the scales in wrong direction, so we need to explore the possibility of non-local mechanisms.

An example non-local mechanism is explored in [4] which is the `GreedyCap` algorithm. Under a particular setting, this algorithm in fact satisfies both DNH and PG properties as stated by the theorem below-

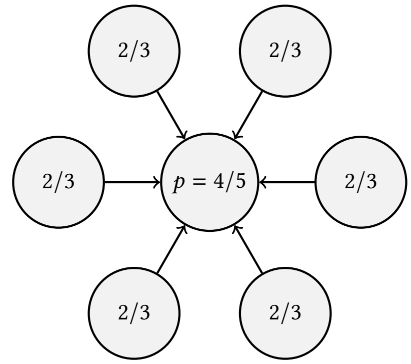

Figure 1: Notice if everyone delegates to the center here, the P_M is 4/5, while the P_D goes to 1 as number of leaves increase

Algorithm 1 GreedyCap

Require: labeled graph G, cap $C : \mathbb{N} \to \mathbb{N}$
$V' \leftarrow V$
while $V' \neq \emptyset$ **do**
  let $i \in \text{argmax}_{j \in V'} |A_G^{-1}(j) \cap V'|$
  $J \leftarrow A_G(i) \cap V'$
  if $|J| \leq C(n) - 1$ **then**
    $J' \leftarrow J$
  else
    let $J' \subseteq J$ such that $|J'| = C(n) - 1$
  end if
  vertices in J' delegate to i
  $V' \leftarrow V' \setminus (\{i\} \cup J')$
end while

Theorem 4.4. *Assume that there exists $\beta \in (0, \frac{1}{2})$ such that all competence levels are in $[\beta, 1 - \beta]$. Then for any $\alpha \in (0, 1 - 2\beta)$, GreedyCap with cap $C : \mathbb{N} \to \mathbb{N}$ such that $C(n) \in \omega(1)$ and $C(n) \in o(\sqrt{\log n})$ satisfies the PG and DNH properties.*

But this mechanism, still has some shortcomings, notice that we are randomly chosing $C(n) - 1$ delegators, we can make an easy modification as follows.

Algorithm 2 ModifiedGreedyCap

Require: labeled graph G with n vertices, cap $C : \mathbb{N} \to \mathbb{N}$
$V' \leftarrow V$
while $V' \neq \emptyset$ **do**
  let $i \in \text{argmax}_{j \in V'} |A_G^{-1}(j) \cap V'|$
  $J \leftarrow A_G^{-1}(i) \cap V'$
  if $|J| \leq C(n) - 1$ **then**
    $J' \leftarrow J$
  else
    let $J' \subseteq J$ be the $C(n) - 1$ nodes with the lowest competence
  end if
  vertices in J' delegate to i
  $V' \leftarrow V' \setminus (\{i\} \cup J')$
end while

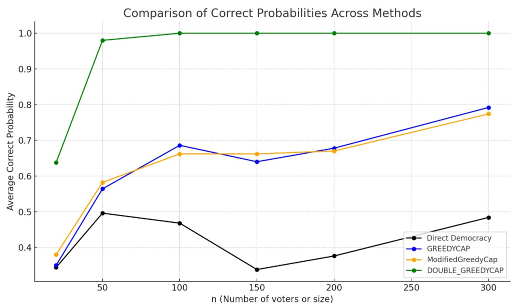

Figure 2: Comparision of the Mechanisms

This little modification ensures more competent voters remain available to receive delegations in later iterations, potentially improving overall decision quality, which aligns with the basic principle of liquid democracy of rewarding the higher competent voters. Also since, we are changing nothing except taking $C(n) - 1$ least competent voters instead of random, so the following theorem which is equivalent of Theorem 4.4 is immediate-

Corollary 4.5. *Assume that there exists $\beta \in (0, \frac{1}{2})$ such that all competence levels are in $[\beta, 1-\beta]$. Then for any $\alpha \in (0, 1-2\beta)$, ModifiedGreedyCap with cap $C : \mathbb{N} \to \mathbb{N}$ such that $C(n) \in \omega(1)$ and $C(n) \in o(\sqrt{\log n})$ satisfies the PG and DNH properties.*

Now we introduce another mechanism, which does better than both GreedyCap and ModifiedGreedyCap in our simulations. We take an opposite approach, and greedily delegate the least competent voter to their most competent neighbour, who hasn't reached the cap of $C(n)$ Given the results of the simulation, we clearly expect DoubleGreedyCap to perform as good as GreedyCap, so naturally we expect them to follow the equivalent of Theorem 4.4, as stated below

Theorem 4.6. *Assume that there exists $\beta \in (0, \frac{1}{2})$ such that all competence levels are in $[\beta, 1-\beta]$. Then for any $\alpha \in (0, 1-2\beta)$, DoubleGreedyCap with cap $C : \mathbb{N} \to \mathbb{N}$ such that $C(n) \in \omega(1)$ and $C(n) \in o(\sqrt{\log n})$ satisfies the PG and DNH properties.*

Proof Sketch. The proof follows the same structure as the proof for Theorem 4.4 stated in Kahng et al. [4] The argument mainly rely on the bound $[\beta, 1-\beta]$ of p_i, and $(0, 1-\beta)$ of α, and the cap condition $C(n) \in o(\sqrt{\log n})$ to bound votes and ensure DNH. The condition $C(n) \in \omega(1)$ ensures sufficient delegation occurs for PG. Since DoubleGreedyCap operates under the same bounds for p_i, α and same cap $C(n)$, the essential components of the original proof apply directly. The change in the order of dealing with nodes does not change the validity of the bounds derived from the cap size. The mechanism of delegating to more competent approved neighbors ensures positive gain remains. □

5 NON-BINARY OUTCOMES

5.1 Model

We now extend the model to handle scenarios with $n \geq 3$ possible outcomes or alternatives, building upon the above framework.

Algorithm 3 DoubleGreedyCap

Require: labeled graph G with n vertices, cap $C : \mathbb{N} \to \mathbb{N}$
$V' \leftarrow V$
Initialize weights $w_i \leftarrow 1$ for all $i \in V$
while $V' \neq \emptyset$ **do**
  let $i \in \text{argmin}_{j \in V'} p_j$ ▷ Select least competent voter
  if $w_i \geq C(n)$ **then**
    i votes with weight w_i
  else
    let S be approved neighbors of i in V', sorted by p_j descending
    for each $j \in S$ **do**
      if $w_j + w_i \leq C(n)$ **then**
        Delegate w_i to j
        $w_j \leftarrow w_j + w_i$
        $w_i \leftarrow 0$
        break
      end if
    end for
    if no delegation occurred **then**
      i votes with w_i
    end if
  end if
  $V' \leftarrow V' \setminus \{i\}$
end while

The underlying graph structure $G = (V, E)$ with voters $V = \{1, 2, \ldots, n\}$ and edges E remains the same as in the binary case. However, instead of a simple binary competence, we need to define voter quality and preferences are defined differently.

Let $\mathcal{A}$ be the set of m distinct alternatives ($m \geq 3$). Voters' preferences are represented by permutations(rankings) of these alternatives. Let $\mathcal{S}$ be the set of all $m!$ possible rankings. We assume there exists a ground truth ranking $\sigma^* \in \mathcal{S}$.

Each voter i is characterized by a dispersion parameter $\phi_i \in (0, \infty)$. This parameter is inversely related to competence or accuracy; a smaller ϕ_i indicates the voter's sampled rankings are more likely to be close to the ground truth σ^*. $\phi_i = 0$ means the voter is competent, and always will choose the correct outcome, while as $\phi_i \to \infty$ the voter will almost always choose the worst ranking, which is the exact reverse of the ground truth.

Now we need a measure to quantify the difference between two rankings $\sigma_1, \sigma_2 \in \mathcal{S}$, we use the Kendall tau distance, denoted by $d(\sigma_1, \sigma_2)$, which counts the number of pairs of alternatives ranked discordantly between σ_1 and σ_2.

The probability that a voter i internally samples a specific ranking $\sigma \in \mathcal{S}$, given the ground truth σ^*, is proportional to their dispersion parameter raised to the power of the Kendall tau distance between σ and σ^*:

$$P_i(\sigma \mid \sigma^*) = \frac{\phi_i^{d(\sigma, \sigma^*)}}{\sum_{\sigma' \in \mathcal{S}} \phi_i^{d(\sigma', \sigma^*)}}$$

The summation in the denominator normalises, so the overall summation is 1.

So clearly, delegation decisions are now based on comparing dispersion parameters. Let $0 < \lambda' < 1$ be a constant threshold parameter (analogous to α in the binary model, but used multiplicatively this time). Voter i approves voter j if $(i, j) \in E$ and j is sufficiently less dispersed (more competent) than i, specifically

$$A_G(i) = \{j \in V \mid (i, j) \in E \text{ and } \phi_j < \phi_i \cdot \lambda'\}$$

This naturally defines the set $A_G(i)$ of voters approved by i for delegation, very similar to the binary case.

Given this setup, the concept of a delegation mechanism M outputting a probability distribution over $A_G(i) \cup \{i\}$ remains identical. Similarly, sampling these distributions outputs a delegation graph with sinks s and corresponding weights $w(s)$ representing the number of voters whose delegation path terminates at s.

Another crucial difference lies in how the final outcome is chosen.

(1) Each sink node s samples a ranking $\sigma_s \in \mathcal{S}$ according to its own probability distribution $P_s(\cdot \mid \sigma^*)$.
(2) Sink s then casts its entire weight $w(s)$ for the single alternative that is ranked first in its sampled ranking, i.e., for $top(\sigma_s)$.
(3) The votes for each alternative across all sinks are summed.
(4) The final winning alternative is the one that receives the highest total weight (First-Past-The-Post).

Success in this model is defined as the winning alternative matching the top-ranked alternative in the ground truth ranking, $top(\sigma^*)$. Similar to the earlier case, we denote the probability of success under mechanism M as $P_M(G)$ and under the direct mechanism (where everyone votes based on their own sampled top choice) as $P_D(G)$.

With these probabilities defined based on top-rank accuracy, the definitions for $Gain(M, G)$, Do No Harm (DNH), and Positive Gain (PG) remain syntactically identical to those presented in the binary case definitions (Definition 1 and 2), but are now interpreted in the context of this n-outcome model and the First-Past-The-Post aggregation rule.

5.2 Main Results

Notice our new framework allows us to retain all of the earlier ideas and notions, just the main difference being the number of possible outcomes, and the use of dispersion parameters. In fact, we can extend all our results by a very similar argument to this setting as well.

Particularly, the definition of Local and Non-Local mechanism stays the same, so based on that, we state the following result-

THEOREM 5.1. *Consider the extended model for non-binary outcomes ($m \geq 3$) with success defined as top-rank accuracy under First-Past-The-Post aggregation. For any approval threshold parameter $\lambda'_0 \in (0, 1)$, there is no local delegation mechanism that satisfies both the PG and DNH properties.*

We can show this using a very similar argument to that used for proving the binary counterpart to this theorem.

Roughly, we construct a specific 'star graph' where less informed 'leaf' voters see only a few 'center' voter, locality ensures the leafs independently delegate causing a heavy concentration in the center. In direct voting, this succeeds with probability going to 1 averaging out individual errors. This ensures the mechanism's gain eventually becomes permanently negative, violating DNH.

Notice we can extend our algorithms (namely GREEDYCAP, MODIFIEDGREEDYCAP and DOUBLEGREEDYCAP) to this setting as well. So clearly, it is very natural to expect an equivalent of Theorem 4.3 for this setting as well. However for that, we face a small problem, the earlier version requires the competence values $p_i \in [0, 1]$, but due to the framing of this setting, we can't directly define p_i in this case, but we do have dispersion parameter, roughly determining the competence of each voter, so we can write the equivalent theorem as-

THEOREM 5.2. *If*

- *Dispersion parameters ϕ_i are in $[\phi_{min\_bound}, \phi_{max\_bound}]$, for some $0 < \phi_{min\_bound} \leq 1 \leq \phi_{max\_bound} < \infty$,*
- *The approval parameter $\lambda' \in (\phi_{min\_bound}, 1)$,*
- *The cap $C(n) \in \omega(1) \cap o(\sqrt{\log n})$,*

then GREEDYCAP *satisfies both PG and DNH under First Past the Post aggregation.*

The reasoning behind the proof is also not very hard. Delegation using the ratio rule directs votes towards competent sinks ($\phi < 1$), and the cap $C(n) \in o\sqrt{\log n}$ prevents any single sink's vote (even from many anti-competent delegators) from dominating due to excessive weight. So, in favorable structures, concentrating votes on competent sinks achieves PG. Similarly, on random graphs, the cap limits harmful concentration, ensuring DNH.

6 CONCLUSION

So we have successfully extended the binary choice model, to a model which works for any general n outcomes. We rely on existence of Ground Truth. The model described here relies on plurality (i.e. First-Past-The-Post) for aggregation, however for other voting models such as approval voting, it should not be very hard to generalise the results. We can also restrict the permutation space, to say single peaked sets, the normalisation term in the expression of probability distribution changes, keeping the main ideas same.

REFERENCES

[1] Francisco M Bersetche. 2022. Generalizing Liquid Democracy to multi-agent delegation: A Voting Power Measure and Equilibrium Analysis. *arXiv preprint arXiv:2209.14128* (2022).
[2] Markus Brill, Théo Delemazure, Anne-Marie George, Martin Lackner, and Ulrike Schmidt-Kraepelin. 2022. Liquid democracy with ranked delegations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 4884–4891.
[3] Paul Gölz, Anson Kahng, Simon Mackenzie, and Ariel D Procaccia. 2021. The fluid mechanics of liquid democracy. *ACM Transactions on Economics and Computation* 9, 4 (2021), 1–39.
[4] Anson Kahng, Simon Mackenzie, and Ariel Procaccia. 2021. Liquid democracy: An algorithmic perspective. *Journal of Artificial Intelligence Research* 70 (2021), 1223–1252.
[5] M. G. Kendall. 1938. A New Measure of Rank Correlation. *Biometrika* 30, 1/2 (1938), 81–93. http://www.jstor.org/stable/2332226
[6] C. L. Mallows. 1957. Non-Null Ranking Models. I. *Biometrika* 44, 1/2 (1957), 114–130. http://www.jstor.org/stable/2333244
[7] James C Miller III. 1969. A program for direct and proxy voting in the legislative process. *Public choice* 7, 1 (1969), 107–113.
[8] Giannis Tyrovolas, Andrei Constantinescu, and Edith Elkind. 2024. Unravelling expressive delegations: complexity and normative analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 9918–9925.