

# Characterizing Properties and Trade-offs of Centralized Delegation Mechanisms in Liquid Democracy

Brian Brubach  
bb100@wellesley.edu  
Wellesley College  
Wellesley, MA, USA

Audrey Ballarin  
ab14@wellesley.edu  
Wellesley College  
Wellesley, MA, USA

Heeba Nazeer  
hn1@wellesley.edu  
Wellesley College  
Wellesley, MA, USA

## ABSTRACT

Liquid democracy is a form of transitive delegative democracy that has received a flurry of scholarly attention from the computer science community in recent years. In its simplest form, every agent starts with one vote and may have other votes assigned to them via delegation from other agents. They can choose to delegate all votes assigned to them to another agent or vote directly with all votes assigned to them. However, many proposed realizations of liquid democracy allow for agents to express their delegation/voting preferences in more complex ways (e.g., a ranked list of potential delegates) and employ a centralized delegation mechanism to compute the final vote tally. In doing so, centralized delegation mechanisms can make decisions that affect the outcome of a vote and where/whether agents are able to delegate their votes. Much of the analysis thus far has focused on the ability of these mechanisms to make a correct choice. We extend this analysis by introducing and formalizing other important properties of a centralized delegation mechanism in liquid democracy with respect to crucial features such as accountability, transparency, explainability, fairness, and user agency. In addition, we evaluate existing methods in terms of these properties, show how some prior work can be augmented to achieve desirable properties, prove impossibility results for achieving certain sets of properties simultaneously, and highlight directions for future work.

## CCS CONCEPTS

• **Applied computing** → **Voting / election technologies; E-government**; • **Theory of computation** → *Algorithmic game theory and mechanism design*.

## KEYWORDS

liquid democracy, voting, computational social choice, fairness, accountability, transparency

## ACM Reference Format:

Brian Brubach, Audrey Ballarin, and Heeba Nazeer. 2022. Characterizing Properties and Trade-offs of Centralized Delegation Mechanisms in Liquid

Democracy. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, June 21–24, 2022, Seoul, Republic of Korea. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3531146.3533219>

## 1 INTRODUCTION

Originally championed by Charles Dodgson (AKA Lewis Carroll) in 1884 [8, 25], liquid democracy is an old idea that has attracted renewed interest from both computer science [4, 5, 7, 14, 19, 21, 26] and political philosophy [3] perspectives. The basic concept is simple: it is a voting system in which each agent can choose to either vote directly or delegate their vote(s) to another agent. However, unlike the more common and restricted system of proxy voting, delegation in liquid democracy is transitive: If some votes are delegated to an agent, that agent may choose to delegate all of those votes (and their own vote) to another agent.

As liquid democracy has become a topic of serious study in the computer science and artificial intelligence communities, works have studied its ability to make a correct choice [7, 19] as well as how best to collect and aggregate agents' delegation preferences [11, 14, 26]. It was implemented internally at Google [17], by Germany's Pirate Party [20], and in several other cases [24]. The mechanisms proposed and studied in these works allow agents to specify delegation preferences in a variety of ways (e.g., providing a ranked or unranked list of acceptable delegates) and sometimes include additional objectives/constraints such as limiting the total number of votes delegated to any single individual. Although still far from widespread adoption, these real-world applications and increased study illustrate the promise of liquid democracy and motivate deeper examination.

The essential idea of liquid democracy may be straightforward to state, but there are many ways to translate it into practice. The seminal work of Blum and Zuber [3] listed four key components:

- (1) *Direct democracy*: Every agent has the option to vote directly on an issue.
- (2) *Flexible delegation*: Agents can delegate their votes to others for a single policy issue, all issues in a policy area, or all issues.
- (3) *Meta-delegation*: Agents can also delegate any votes that have been delegated to them (delegation is transitive).
- (4) *Instant recall*: An agent can withdraw a delegation at any time.

Those four components suffice to establish the basic model of liquid democracy. However, they also leave several ambiguities which should be resolved and made explicit in any transparent implementation. We refer the reader to Blum and Zuber [3] for a discussion of many such challenges that can arise from a political science

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

FAccT '22, June 21–24, 2022, Seoul, Republic of Korea

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9352-2/22/06...\$15.00

<https://doi.org/10.1145/3531146.3533219>

perspective as well as which applications are best-suited to this form of governance and why.

Our work seeks to bring a fairness, accountability, and transparency lens to the computational study of mechanisms for translating agent preferences into outcomes within realizations of liquid democracy. Now that there are many proposed (and in some cases implemented) mechanisms for liquid democracy systems, how do we choose between them? Individual mechanisms have been analyzed in terms of certain key features and the qualities they were designed to support, but not thoroughly compared among a broad range of other important criteria. In particular, many questions arise which are not adequately addressed in prior work such as,

- (1) What rights should agents have? Should they have the right to delegate instead of voting directly, and if so, can the mechanism guarantee that? (Section 4.2.2)
- (2) Can the mechanism itself arbitrarily choose the outcome of the vote and what governs its decision? (Section 4.2.5)
- (3) How transparent and interpretable is the mechanism’s explanation of what happened to an agent’s vote in a previous election, and how easily can agents predict the effect of changing their delegation preferences? (Sections 4.3.1 and 4.3.2)

We formalize these types of questions by characterizing several properties of liquid democracy mechanisms and exploring how these properties interact with each other.

## 2 OUTLINE

Section 3 provides some basic definitions and discusses related work not covered elsewhere in the paper. Section 4 lists and defines several properties which we use to analyze and compare mechanisms for liquid democracy. Section 5 analyzes five prominent existing mechanisms in terms of the properties we propose. Finally, Section 6 provides a conclusion and future directions.

## 3 PRELIMINARIES AND RELATED WORK

### 3.1 Basic definitions and notation

In this section, we provide some basic terminology that will be used throughout the paper. We refer to all participants in a liquid democracy system as *agents*. Any agent who prefers to cast a direct vote as opposed to delegating or abstaining is called a *voter* and we limit our scope to mechanisms which always allow any such agent to cast a direct vote in keeping with that basic tenet of liquid democracy. We largely ignore the agents who abstain from both voting and delegation except when considering a switch from abstaining to delegating or voting over the course of multiple elections. We additionally limit the discussion in this paper to simple plurality votes on binary issues.

While we consider multiple ways in which agents can express their delegation and voting preferences, agent preference in this work are represented by some form of directed graph which we call a *preference graph*. A directed edge from agent  $u$  to agent  $v$  in the preference graph means that  $u$  is willing to delegate their vote to  $v$  and prefers this delegation over voting directly. Edges in the preference graph may additionally be assigned ranks or weights as described in Section 4.1. To determine how votes are actually delegated and cast, a *tally algorithm* takes a preference graph as

input and outputs a *delegation graph* which explicitly describes how each agent’s vote(s) are delegated or cast. This delegation graph determines how many will be assigned to each outcome.

### 3.2 Modeling assumptions and limitations

Regarding agents decisions to vote directly, we make two minor simplifying assumptions. First, we assume that any agent who lists at least one approved delegate would prefer to delegate rather than vote directly. Second, while we explore the possibility that a mechanism cannot (or chooses not to) transfer an agent’s vote(s) to any approved delegate, we do not discuss what to do with such agents. We view the following as the most natural options for handling this situation of agents whose votes are not delegated by the mechanism: request a direct vote from these agents, discard these agents, or allow all agents to specify a backup direct vote that can be applied in this case. Because these choices can be applied to any of the mechanisms we discuss, we do not see them as inherent properties of a given mechanism and assume that a mechanism designer will address this issue however they see fit. Our primary focus in this realm will instead be on whether an agent who prefers to delegate is allowed to delegate by the mechanism.

Another aspect we do not engage with is the ability of agents to specify preferences across multiple elections or issues at once. For example, LiquidFeedback [2] allows an agent to delegate to another agent globally (all issues), only for certain categories of issues, or only for specific issues. Google Votes [16, 17] similarly allows agents to tag potential delegates with categories. However, while these mechanisms collect such complex preferences for the sake of efficiency, they generally resolve preferences for a single issue at a time using the type of preference graph we study (a simple, directed graph with ranked or unranked edges).

### 3.3 Epistemic versus equality-based accounts

Blum and Zuber [3] suggests two distinct functions through which we can evaluate liquid democracy. The epistemic account assumes there is some ground truth correct choice and values the ability of a mechanism to find that correct choice. Equality-based accounts on the other hand emphasize the rights of all participants to have a voice and promote their own individual, subjective choice of preferred outcome.

In some ways, these accounts complement each other. Agents want to select a delegate that represents their needs (equality-based) and makes competent decisions in support of those needs (epistemic). However, we will see that these values can also be somewhat competing. Epistemic approaches may favor mechanisms that exert heavier influence over agents in order to ideally increase the chance of making the “correct” choice. For example, Gözl et al. [14] suggests, “If agents are considering multiple potential delegates, all of whom they trust, they are encouraged to leave the decision for one of them to a centralized mechanism.” As such the relevance of our proposed properties can vary depending on which function one favors. For example, giving all agents a right to delegate may be valued more strongly from an equality-based perspective than an epistemic one. A further complication we explore is that

mechanisms exerting a larger influence may obscure the information needed for agents to make intelligent delegation choices, thus potentially harming chances of correctness.

### 3.4 Other related work

Several prior works which propose centralized mechanisms are discussed at length in later sections. Due to our focus on simple binary votes, we do not analyze some proposals that interpret more complex preference expressions [11] or ranked outcomes [5]. Other works that give broad overviews of liquid democracy and compare approaches include Brill [4] and Paulin [24] with the latter focusing especially on the real world implementations.

One prior work, analyzed in Section 5.3, investigates the concept of “incentivising participation” [21] which has a similar flavor to some properties studied in the present work. In short, they argue that participation among voters (which they call *guru participation*) is disincentivised when an agent can change their delegation preferences in such a way that the voter who ultimately casts that agent’s delegated vote is unhappy with the change. However, the voter in this case can still be given the incentive to vote (which they call *cast participation*) and cannot benefit directly from rejecting the vote delegated by the offending agent. Because the voter in this scenario still has an incentive to vote and we see the ability of agents to affect voter power as fundamental to liquid democracy, we do not include this property in our assessment. Further, we note that the voter who is considered “disincentivised” within this model may nevertheless prefer a mechanism which does not satisfy *guru participation* over a mechanism which does. Thus, while the observations in Kotsialou and Riley [21] are surprising and interesting within the context of that work, it could be misleading for a practitioner to use this property in choosing between mechanisms.

## 4 PROPERTIES OF A CENTRAL DELEGATION MECHANISM (CDM)

This section outlines several important properties of a central delegation mechanism. Section 4.1 covers the most common ways that agents are able to express their preferences to the mechanism. Section 4.2 addresses key issues regarding how a mechanism can influence the outcome of a vote or limit an agent’s ability to delegate. Section 4.3 defines properties affecting the related goals of transparency, explainability, and accountability. Section 4.4 briefly discuss the effects of running time on the ability to satisfy other desirable properties. Finally, Section 4.5 helps to clarify which properties are novel to this work.

### 4.1 Expression of delegation preferences

One of the most obvious ways that mechanisms differ is in how they allow agents to express their delegation preferences. In our analysis, we will consider four common ways in which agents can specify their delegation preferences.

- *One neighbor preferences (ONP)*: Agents may only choose one neighbor to delegate to. The preference graph is unweighted with out-degree 1.
- *Multiple ranked preferences (MRP)*: Agents provide a ranked list of neighbors they are willing to delegate to. In the preference graph, the outgoing edges of each vertex are ranked.

- *Multiple unranked preferences (MUP)*: Agents provide an unranked list of neighbors they are willing to delegate to. The preference graph is unweighted.

In all cases, an agent who wishes to vote directly will have a single outgoing edge to one of the two binary choices. We note that the mechanism may bound the length of the list given by each agent in which case one neighbor preferences is a special case of both multiple ranked preferences and multiple unranked preferences with lists of length at most one.

As stated in Section 3.2, we will not cover how or whether a mechanism solicits votes from agents who prefer delegation over voting, but would like to vote if they can’t delegate.


### 4.2 Influence of the mechanism on outcomes and rights to delegation preferences

Now, we introduce properties that address some of our central fairness questions: How and why can a mechanism influence the result of a vote and govern the agents’ abilities to delegate? We begin by reviewing the common issue of delegation cycles and proposing definitions for an agent’s right to delegate and to their top ranking. We then outline two key factors in mechanisms influencing an agent’s right to delegate and/or the outcome of the vote: the goal of limiting concentrated power and the arbitrary decisions that can occur due to ambiguity in the mechanism objective.

**4.2.1 Resolving delegation cycles.** Delegation cycles are perhaps the most well-known challenge that arises when attempting to implement liquid democracy. As shown in Figure 1, we may receive delegation preferences that form a directed cycle (AKA a *delegation cycle*). In this case, the mechanism must decide which agents (if any) are permitted to delegate. We note that prior work primarily addresses the issue of delegation cycles in one of the following three ways:

- *Assume away cycles (AAC)*: In this case, the mechanism is designed for a model that assumes delegation cycles do not occur. For example, some models assume agents only wish to delegate to someone strictly more competent than themselves regarding the issue at hand, making directed cycles impossible.
- *Discard cycles (DC)*: Directed cycles in the graph are entirely discarded (or ignored). This is a natural solution for mechanisms using one neighbor preferences since a directed cycle in the resulting preference graph will not have any outgoing edges with paths to voters.
- *Break cycles (BC)*: The mechanism has an algorithm for breaking a directed cycle in the preference graph while preserving some non-empty subset of the edges in the cycle. For example, we could randomly choose one of the agents in Figure 1 to vote directly while the others are allowed to delegate.

**4.2.2 Agents’ right to delegate.** A central component of liquid democracy is that all agents are allowed to delegate their vote(s) rather than voting directly. The choice to either vote directly or delegate is essential to the motivation that it serves as a compromise between direct and representative democracy. Further, protecting an agent’s right to delegate reduces the burden of requiring backup voting preferences from agents who prefer to delegate or the risk



**Figure 1: Example of a cycle in a preference graph. These agents’ votes cannot be cast unless at least one of them becomes a voter.**

of disenfranchising agents whose delegation preferences are not honored. Regardless, it may be impossible to satisfy every agent’s desire to delegate (and have their delegated vote eventually cast) due to issues such as cycles in the preference graph as discussed in Section 4.2.1. We can see in Figure 1 that at least one of the three agents must forfeit this right.

However, in defining a formal “right to delegate”, we can sidestep the unavoidable issue of cycles in worst-case analysis by focusing on the agents whose preferences combined with the preferences of others will allow for some delegation that ultimately results in votes being cast. We say that an agent has *delegable preferences* if there is a directed path in the preference graph from that agent to some voter. Thus, an agent with delegable preferences has at least one simple delegation path and all such agents could simultaneously be allowed to delegate according to their preferences. This leads to the following property which can apply to unranked or ranked preferences.

- *Right to delegate (RTD)*: A mechanism satisfies the *right to delegate* property if every agent with delegable preferences has all of the votes assigned to them delegated to one or more of their approved delegates.

**4.2.3 Agents’ right to top rank.** When agents specify a ranking, we may also wish to guarantee that a mechanism respects those ranks to some reasonable extent. In theory, a mechanism could completely ignore ranks, although we are not aware of such an audacious approach. More commonly, a mechanism may make some sacrifice in respecting ranks in the service of other constraints or objectives. For example, if an agent’s top-ranked delegate abstains or is part of a directed cycle, using a lower-ranked delegate would be preferable to not delegating/voting at all (See Figure 2 in Section 5.1 for a detailed illustration). Keeping this in mind, we establish a *right to top rank* property which essentially guarantees that agents can delegate along a path using only top-ranked edges if such a path exists after removing some unusable edges from the graph.

We start by assuming some minor preprocessing of the preference graph to remove all agents which have no path to a voter. Let the *delegable preference graph* be a graph constructed by removing all agents without delegable preferences from the preference graph<sup>1</sup>. An agent has a *top rank delegation path* in the delegable preference graph if there is a path from the agent to a voter in the delegable preference graph using only the top-ranked outgoing edges of each agent. Note that this path may use edges which were

<sup>1</sup>An algorithm is not required to explicitly construct a delegable preference graph to satisfy right to top rank, although such a graph is easy to construct.

not top-ranked outgoing edges in the original preference graph and that this path is guaranteed to be simple. This yields a natural property for ranked preferences.

- *Right to top rank (RTTR)*: A mechanism satisfies the *right to top rank* property if every agent with a top rank delegation path has all of the votes assigned to them delegated along that path.

We note that the right to top rank and right to delegate properties do not imply each other.

**4.2.4 Limiting concentrated power.** Analysis of liquid democracy in practice has revealed that some individual agents will accumulate a large number of votes, becoming *super-voters*<sup>2</sup> [20]. Several works have argued that these super-voters are or could be detrimental. Within the theoretical model of voter competence in Kahng et al. [19], super-voters can make liquid democracy less likely than direct democracy to choose the correct outcome in the worst-case due to a wisdom of the crowd effect. On the practical side, Kling et al. [20] observed that the presence of super-voters may have discouraged other voters from participating, but ultimately concluded that, “While we find that the theoretical and potential power of super-voters is indeed high, we also observe that they stabilise the voting system and prevent stagnation while they use their power wisely. Super-voters do not fully act on their power to change the outcome of votes, and they vote in favour of proposals with the majority of voters in many cases.”

Concerns regarding super-voters have led to the creation of mechanisms that restrict or minimize vote accumulation [14, 19]. We do not take a position on the merits of limiting concentrated power, but rather explore how approaches to this goal interact with other properties we study.

For our purpose, we define two broad categories of approaches to limiting power:

- *Capped power (CP)*: A mechanism that caps power will place a hard upper bound on the weight or number of votes than can be delegated to a single individual.
- *Minimized power (MP)*: A mechanism that minimizes power will seek to minimize the weight or number of votes than are delegated to a single individual while respecting all constraints imposed by the mechanism.

We will see in Section 5.5 that minimizing power can lead to an opaque system with some deficiencies in transparency and accountability. On the other hand, capping power can limit the ability of agents to delegate when they prefer to do so and is incompatible with guaranteeing the right to delegate. In other words, a mechanism cannot have both capped power and a right to delegate.

**THEOREM 4.1.** *No mechanism using one neighbor preferences, multiple ranked preferences, or multiple unranked preferences can satisfy both the right to delegate and capped power for any cap less than the number of agents  $n$ .*

**PROOF.** This can easily be seen with just two agents. A cap of less than two will not allow either of them to delegate to the other. □

<sup>2</sup>The term *super-voter* is used in prior work to denote any agent accumulating a large number of votes, regardless of whether they choose to cast those votes directly or delegate.

**4.2.5 Avoiding arbitrary decisions by the mechanism.** Situations can arise in which the mechanism must choose arbitrarily between multiple valid delegations that would result in different outcomes for the vote or different subsets of agents being allowed to delegate. In other words, there can be a preference graph input such that a given mechanism must make a completely arbitrary decision which decides the outcome of the vote or decides who can or cannot delegate. This occurs most naturally in mechanisms that allow multiple unranked preferences, but can occur with other types of delegation preferences as well.

We view arbitrary decision-making as a negative characteristic. As such, we frame three desirable properties in this area followed by a catch-all less desirable property:

- **Single delegation (SD):** A mechanism is a single delegation mechanism if every set of agent preferences has exactly one allowable delegation which satisfies the constraints and objective of the mechanism. I.e., there is a many-to-one mapping from preference graphs to delegations.
- **Single distribution over delegations (SDOD):** A randomized mechanism satisfies SDOD if every set of agent preferences has exactly one allowable distribution over delegations which satisfies the constraints and objective of the mechanism. I.e., there is a many-to-one mapping from preference graphs to distributions on delegations.
- **No arbitrary decisions (NAD):** If there exists a set of agent preferences such that 1) the constraints and objective of the mechanism allow for more than one valid delegation (or more than one valid distribution over delegations for a randomized algorithm); 2) the mechanism deterministically chooses a delegation (or distribution); and 3) the choice of delegations affects the outcome, then we say the mechanism makes an *arbitrary decision*. A mechanism satisfying the NAD property must not make an arbitrary decision, possibly by satisfying one of the previous two properties.
- **Arbitrary:** We call a mechanism arbitrary if there exists a preference graph input on which it makes an arbitrary decision as defined in the previous property.

Clearly, single delegation is a special case of single distribution over delegations. Both satisfy the no arbitrary decisions property assuming that the problem of finding the delegation or distribution over delegations for a given preference graph input is tractable.

Note that these properties are defined in terms of the objective and constraints of the mechanism rather than the tally algorithm. In other words, a given tally algorithm may create a many-to-one mapping from preference graphs to delegations (e.g., due to breaking ties based on some arbitrary ordering of the agents), but this is not enough to satisfy the single delegation property.

### 4.3 Transparency, explainability, and accountability

In the context of liquid democracy, we explore transparency, explainability, and accountability by asking how well agents can understand what has happened with their vote(s), why it happened, and what will happen if they make specific choices in the future. This information allows agents to make informed decisions about their preferences.

**4.3.1 Post hoc explainability.** Properties of post hoc explainability capture the ideal in liquid democracy that agents are able to understand the outcome of their preference choices in order to inform their future preferences and delegation. Knowing where their vote(s) went is one piece of information that an agent needs to refine their preferences between voting rounds and allows them to hold their delegates accountable for past votes.

We consider two natural ways in which a mechanism can give agents feedback on what happened to their vote:

- **$\psi$ -path explainability ( $\psi$ -PE):** We define a *vote path* to be a simple path in a preference graph from an agent to an issue with a label specifying how the final agent in the path voted. A mechanism satisfies  $\psi$ -PE if it can provide every agent with a list of at most  $\psi$  total vote paths and the fraction of that agent's vote which was delegated to each path. We will focus exclusively 1-PE here, but use this more general definition in case it is relevant for future work. We are not aware of any mechanism satisfying  $\psi$ -path explainability for a bounded  $\psi > 1$ .
- **Golden rule explainability (GRE):** The “Golden Rule” as described in Hardt [16] and Hardt and Lopes [17] for the Google Votes system is that an agent should be given a single path explanation for where their own vote went. However, as we show in Section 5.1, the votes delegated to that agent may travel along a different path, violating the 1-path explainability property described above.
- **Local feedback explainability (LFE):** A mechanism satisfies LFE if it can provide every agent with the fraction of that agent's vote which was delegated to each neighbor and what percent of the vote delegated to each neighbor supported each outcome.

We note that any mechanism satisfying 1-path explainability will satisfy Golden Rule explainability and any mechanism satisfying Golden Rule explainability will satisfy local feedback explainability, but the reverse is not true.

**4.3.2 Transparency and accountability.** A crucial pillar of liquid democracy is accountability which is enforced partly by the ability of an agent to change their delegation preferences at any time. At the same time, it is important to consider transparency; i.e., what information an agent needs or can use to make this decision and how well they can anticipate the results of a change in their delegation preferences. This description from German news website, Der Spiegel, of the fluctuating support for super-voter Martin Haase within the German Pirate Party illustrates how agents engage with transparency and accountability in practice [1].

Many pirates gave [Haase] their votes in 2010, when the party was seeking to define its position on the concept of an unconditional basic income guarantee. They were confident that Haase would support it. To their surprise, however, he transferred his votes to another party member, who voted against the motion, and was defeated.

As a linguist, says Haase, he was also opposed to the initiative because of linguistic weaknesses. But this didn't convince his supporters, and about 50 Pirates

promptly withdrew the votes they had assigned to him. Haase eventually approved a revised motion, and since then the number of members supporting him has gone up again.

Often, the main factor conflicting with transparency and accountability in liquid democracy is the competing goal of privacy. An agent may wish to see the voting/delegation history of another agent before choosing them as a possible delegate, but this history may be hidden to respect privacy. While interesting, this tension between transparency and privacy tends to be independent of the mechanism choice and therefore outside the scope of this paper.

Along those lines, we would prefer to identify a transparency issue inherent in a given mechanism. To do so, we consider a model in which an agent knows the voting/delegation history of other agents, but does not know or cannot run the tally algorithm. In this model, an agent chooses potential delegates based on their history and should be able to assume that their future voting/delegation will be similar.

More formally, we setup the following scenario. Suppose there are two rounds of votes in sequence. After the first vote, let each agent have a rating which determines what percentage of their vote went to each of the two outcomes. Let exactly one agent  $a_c$  be the only agent to change their delegation or voting preferences between the first and second votes. Let  $P_1$  and  $P_2$  be the preferences of  $a_c$  for the first and second votes, respectively. We say that  $P_2$  *strictly favors* an outcome  $O$  over  $P_1$  if every agent in  $P_2$  had a higher expected percentage of their vote(s) cast for  $O$  during the first vote than every agent in  $P_1$ . Using this scenario, we establish the following definition of *local predictability* which captures the ability of an agent to predict the effect of changing their delegation preference. We note that for deterministic tally algorithms, local feedback explainability suffices to provide the expected percentage of an agent's vote(s) cast for  $O$  while a randomized algorithm would need to give the *ex ante* expected percentages in the preference graphs of previous votes.

- *Locally predictable*: We say that a mechanism is locally predictable if this change in the preferences of  $a_c$  cannot reduce the expected vote share assigned to outcome  $O$  in the second vote of the scenario. The expectation is taken over any randomness in the tally algorithm.

It may seem surprising that mechanisms could violate the local predictability, and one could ask if it is too low of a bar to be useful. However, we show in Section 5.5 that a reasonable proposed mechanism does not satisfy this property.

#### 4.4 Running time and tractability

Another important consideration is the running time of computing the tally. For different mechanisms, the tractability of computing the tally can range from solving an NP-hard problem to running a simple linear time algorithm. In addition to the usual concerns of efficiency, tractability of the tally problem can indirectly affect other properties. If a heuristic or approximation algorithm is required to tally votes and resolve delegations, then the choice of heuristic/algorithm or the larger space of acceptable solutions can potentially influence the outcome and/or reduce transparency. For example, two different approximation algorithms with the same

approximation guarantees may lead to different delegations with different election outcomes, making the choice of which algorithm to use an arbitrary decision affecting the outcome.

#### 4.5 Properties contributed by this work

Many of these properties are explicitly or implicitly defined in prior work. However, they have not been collected in this way or compared through the lenses of accountability, transparency, explainability, fairness, and user agency. In addition, right to delegate (Section 4.2.2), right to top rank (Section 4.2.3), avoiding arbitrariness of decisions (Section 4.2.5), and local predictability (Section 4.3.2) present novel properties which we are not aware of in any prior work.

### 5 ANALYSIS OF EXISTING MECHANISMS

Here, we analyze five mechanisms from prior work using the properties defined in Section 4 with several goals in mind. First, we aim to illustrate and justify the properties through example. Second, although this is far from an exhaustive list of prior work, it gives a representative snapshot of which combinations of properties have been achieved, highlighting future directions. Third, the analysis suggests additional impossibility results which are provided throughout.

#### 5.1 Google Votes

The Google Votes mechanism, mentioned in Hardt and Lopes [17] and described in more detail in a Google TechTalk [16], is a liquid democracy system that was used internally within Google for Google employees to vote on issues such as meals, snacks, and project logos. It was implemented as a web app on the internal corporate Google+ network, and was in use from March 2012 to February 2015. In practice, Google Votes allowed users to express preferences for multiple issues and categories of issues all at once. However, we focus on the special case of a single binary issue for simplicity and Google Votes clearly matches the multiple ranked preferences property in this case. It is also easy to see that Google Vote does not attempt to limit concentrations of power.

The Google Votes paper [17] defines the “Golden Rule of Liquid Democracy” as one of the key properties of the Google Votes system. Similar to 1-path explainability, this property states that an agent who delegates their vote must be able to see the entire path that their vote travels on, including the final vote it was cast for. The Google Votes tally algorithm achieves this by selecting one final vote path per agent. However, there are some ambiguities in the algorithm definition which make it difficult to determine whether Google Votes is actually guaranteed to satisfy 1-path explainability. It appears that their “Golden Rule of Liquid Democracy” applies only to the single vote that an agent starts with, rather than all votes delegated to an agent.

Google Votes resolves preferences with a tally algorithm which first enumerates all simple paths from every agent to some voter. As such, the worst case running time is lower bounded by the number of paths which is exponential in the number of agents  $n$ . Then, for each agent  $a$ , it identifies that agent's highest ranking neighbor  $a'$  for which there exists at least one delegation path

## Properties of Previously Proposed Mechanisms


Mechanism	Preferences	Limit Power	Outcomes	Cycles	Right to Delegate	Right to Top Rank	Explainability	Locally Predictable	Running Time
Google votes	MRP	No	SD	BC	Yes	Yes**	Golden Rule	Unclear	$\Omega(2^n)^*$
LiquidFeedback	ONP	No	SD	DC	Yes	N/A	1-PE	Yes	$O(n)$
Breadth-first	MRP	No	SD	BC	Yes	No	1-PE	Yes	$O(n + m)^*$
GreedyCap	MUP	CP	SDOD**	AAC	No	N/A	1-PE	Yes	$O(n \lg n)^*$
Fluid mechanics	MUP	MP	Arbitrary	BC	Yes	N/A	1-PE	No	Varies

**Table 1: Summary of properties for five prior mechanisms analyzed in this section. \* indicates a running time not stated in prior work, but easily derived. The Google votes running time is a lower bound on the worst case performance. \*\* indicates a property which can be achieved with a minor extension or resolution to an ambiguous step described in the section covering that algorithm.**

starting with  $a'$ . From this point in the algorithm, it is not possible to determine from existing documentation how Google Votes chooses among multiple potential paths. Nevertheless, we will analyze properties based on two potential implementations of the ambiguous step and show an impossibility result that holds for any possible resolution of this step. We further note that for several possible resolutions of the ambiguous step, a fairly straightforward polynomial-time algorithm can be devised, but leave that outside the scope of this work.

To explain the issues with specifying additional properties for Google Votes, we refer to Figure 2. This example shows three agents ( $a_1, a_2,$  and  $a_3$ ) who all prefer to delegate to each other in a directed cycle while also having second choices ( $a_4, a_6,$  and  $a_5,$  respectively) who are voters. According to the limited description in Hardt [16],  $a_1$  will either be given the path  $(a_1, a_2, a_6)$  or  $(a_1, a_2, a_3, a_5)$  based on two natural approaches that yield different delegations. Approach 1 would be to pick the path with highest ranked prefix (or highest lexicographical ordering). In other words, first choose the highest-ranked first step, then the highest-ranked second step, then the highest-ranked third step, and so on. This approach would choose the path  $(a_1, a_2, a_3, a_5)$  for  $a_1$  since  $a_2$  prefers  $a_3$  over  $a_6$ , and it clearly satisfies right to top rank. Approach 2 would be to choose the shortest path from an agent which starts with its highest ranked neighbor who has some valid path that does not include the agent itself (breaking ties in length by highest ranked prefix). By contrast, this approach would choose the shorter path  $(a_1, a_2, a_6)$  for  $a_1$ , and it does not satisfy right to top rank since the shortest path is not guaranteed to choose top-ranked edges. Approach 1 appears to be the most likely according to Hardt [16]. However, neither approach nor any possible resolution to this ambiguous step can satisfy 1-path explainability. We formalize this with a theorem and brief proof sketch.

**THEOREM 5.1.** *Any tally algorithm which must delegate each agent’s vote to its highest ranked neighbor through which a delegation path exists, only allows simple paths, and guarantees that the vote will eventually reach a voter, cannot satisfy 1-path explainability.*



**Figure 2: A preference graph that illustrates ambiguity in the the Google Votes tally algorithm.**

**PROOF.** Consider the example in Figure 2. Each of  $a_1, a_2,$  and  $a_3$  will delegate their own vote to their top-ranked neighbor, and that vote can never return to them due to the simple path requirement. However, for votes to be cast, at least one of them must delegate some votes they receive through delegation to one of the voters ( $a_4, a_5,$  or  $a_6$ ). Thus, at least one agent will send votes on more than one path, violating 1-path explainability.  $\square$

We note that the ambiguous tally algorithm of Google Votes is also discussed in Kotsialou and Riley [21] which considers a similar ranked preference model with a focus on how participation is incentivised (or not). They characterize Approach 1 above as a *depth-first delegation* rule and use an example similar to our Figure 2 to illustrate how an agent adding delegation preferences can alter the weights assigned to voters. They propose an alternative *breadth-first delegation* rule which is almost identical to Approach 2, except that it prioritizes the shortest path starting from the agent rather than from their highest-ranked neighbor. Applying the breadth-first delegation rule [21] to Figure 2 would simply give the path  $(a_1, a_4)$  to  $a_1$ . See Section 5.3 for further details.

### 5.2 LiquidFeedback

LiquidFeedback [2] is possibly the most prominent real world implementation of liquid democracy and also one of the most simple approaches in terms of its mechanism. It has been used in practice

(most notably by the German Pirate Party) and is available as open-source software<sup>3</sup> which includes many features beyond simply implementing liquid democracy. A hallmark of LiquidFeedback’s take on liquid democracy is a light touch approach that favors minimal influence of a centralized mechanism. Agents specify one neighbor preferences and delegation cycles are simply ignored by the tally algorithm (agents are expected to resolve cycles themselves by at least one member of the cycle becoming a voter). Similarly, it satisfies the right to delegate and there is no attempt to limit concentration of power. After cycles are removed, the preference graph becomes a directed tree with a maximum out-degree of one. This easily implies 1-path explainability, single delegation, and no arbitrary decisions.

### 5.3 Breadth-first delegation

As noted in Section 5.1, Kotsialou and Riley [21] proposed *breadth-first delegation* as an alternative to “depth-first” approaches such as Google Votes for tallying multiple ranked preferences. A breadth-first delegation chooses the shortest path from an agent to some voter and only uses edge ranks to break ties amongst multiple paths of the same length. In motivating this approach, Kotsialou and Riley [21] argue, for example, that an agent may prefer to delegate to their second choice over their top choice when their second choice is a voter, but their top choice will transitively delegate. Thus, breadth-first delegation explicitly chooses not to satisfy the right to top rank. In fact, if an agent’s worst-ranked delegate is a voter while all better ranked delegates are agents who delegate to voters, then a breadth-first tally will choose the worst-ranked delegate. We acknowledge that this preference may exist in some settings and describe some advantages offered by a breadth-first tally, but note that it goes against the traditional model of liquid democracy and violates the explicit ranking preferences of the agent. Delegating to someone else who then further chooses whom to delegate to (an expert on selecting experts) is fundamental to liquid democracy and a major feature distinguishing it from proxy voting.

Despite this key philosophical difference of prioritizing path length over ranks, breadth-first delegation shares many properties in common with Google Votes. By choosing the unique shortest path which has the highest lexicographical ordering of ranks for each agent, it ensures that each preference graph implies a single delegation with the right to delegate and break cycles properties. It also does not limit power in the worst case. However, Kotsialou and Riley [21] suggest the open question of whether the shorter paths will accumulate fewer votes on real world input and have the effect of limiting concentrated power in practice.

On the other hand, breadth-first delegation differs from Google Votes by offering 1-path explainability and local predictability. Lemma 2 of Kotsialou and Riley [21] implies 1-path explainability by showing that if one agent is on the chosen delegation path of another agent, then both of them are delegating to the same voter. To briefly sketch why local predictability holds, consider an agent  $a_c$  with distinct sets of preferences  $P_1$  and  $P_2$  for the first and second votes, respectively. Let  $O$  be the outcome strictly favored by  $P_2$  and  $a_d \in P_2$  be the agent whom  $a_c$  ultimately delegates to in the second vote. By the single delegation property and 1-path explainability,  $a_c$

and  $a_d$  and must have had disjoint delegation paths that delegated all of their assigned votes to different outcomes in the first vote. It follows that  $a_d$  will have the same delegation path in the second vote as well and will not lose any of the votes delegated to them in the first vote although some votes may now travel to  $a_d$  through  $a_c$ . Thus, in the second vote, outcome  $O$  will receive all of the votes it received in the first vote plus at least one additional vote (from  $a_c$ ).

### 5.4 The GreedyCap algorithm

The GreedyCap algorithm was introduced in Kahng et al. [19] to prove that there exists a non-local delegation mechanism which satisfies their properties of positive gain and do no harm with respect to direct democracy. Roughly speaking, *positive gains* means there is some preference graph on which a given liquid democracy mechanism outperforms direct voting at finding the ground truth in a model in which there is one correct outcome. *Do no harm* is the complementary property stating that the liquid democracy mechanism does not perform worse than direct democracy.

The input to GreedyCap is an unweighted, directed graph with vertices representing agents and a cap  $C$  on the number of votes that can be delegated to a single agent. A directed edge  $(u, v)$  signifies that agent  $u$  approves of delegating their vote to agent  $v$ . This simple greedy algorithm proceeds by iteratively choosing the remaining agent  $v$  with the most approvals (incoming edges) and at most  $C - 1$  agents that approve of  $v$ . These agents (including  $v$ ) are removed from the graph with their votes all being delegated to  $v$ . At the end of this process, all agents whose votes have not been delegated to someone else are instructed to vote.

We can now take a look at which properties are satisfied by GreedyCap. It is easy to see that the input graph here satisfies the multiple unranked preference property, and power is strictly capped by the parameter  $C$ . Cycles are not considered because Kahng et al. [19] assumes a model in which each agent has a competence value and will only delegate to someone strictly more competent, making this a mechanism that assumes away cycles. Explainability is not considered in Kahng et al. [19]. However, because this mechanism only applies “one hop” of delegation, it is technically a form of proxy voting and clearly offers 1-path explainability. Similarly it also satisfies local predictability since all agents are delegating directly to voters. However, we note that generalizing this approach of capping power to true liquid democracy may result in the loss of local predictability. In fact, it is easy to see that some natural mechanisms which cap voter power would not have local predictability, and we leave this as an exercise for the reader.

Implementing this algorithm as defined in Kahng et al. [19] could make arbitrary decisions that affect the outcome. There is no explanation in that paper for how ties are broken at various stages (it isn’t relevant to the theorem they’re proving). However, if these ties are broken uniformly at random, we get a single distribution over delegations.

### 5.5 The fluid mechanics approach

The fluid mechanics approach [14] takes inspiration from fluid mechanics to address the issue of concentrated power. They seek to minimize the maximum power of any voter using a flow-based tally

<sup>3</sup><https://liquidfeedback.com/en/>



algorithm with the following high-level idea: “Put another (more whimsical) way, we wish to design liquid democracy systems that emulate the *law of communicating vessels*, which asserts that liquid will find an equal level in connected containers”

The input is an unweighted, directed graph (multiple unranked preferences) with each agent being a *source* of one vote and all voters being *sinks*. The tally algorithm searches for a confluent flow in this graph to establish the delegation. In addition to the common *flow conservation* property, a confluent flow has the *confluence* property that at most one outgoing edge from each vertex may carry nonzero flow. Finding such a flow is NP-hard leading to three potential tally options described in Gözl et al. [14]: finding an exact solution using a MILP on sufficiently small instances, a heuristic, or an approximation algorithm. As such, the running time varies greatly depending on specific implementation choices.

The reasoning for tackling the NP-hard confluent flow problem rather than an easier splittable flow variant is to allow for 1-path explainability. Although Gözl et al. [14] does not explicitly discuss giving agents single-path explanations, they note that confluent flow is more intuitive and transparent leading to higher accountability. Further, their experimental results do not show a significant decrease in maximum power when switching to splittable flow.


In addition to 1-path explainability, it is easy to see that this method provides a right to delegate. Any agent with a path to some voters, will send flow to one of those voters. Similarly, cycles are broken if at least one agent in the cycle has some other delegation path, while agents connected components with no path to any voter will be discarded (alternately, they could be allowed to vote directly, but this is not discussed in the paper).

One potential drawback to this system and using confluent flows is that some preference graphs have multiple distinct delegation graphs which optimize the minimum confluent flow objective. Figure 3 illustrates a simple example of this. The vote of  $a_1$  must be delegated to either  $a_2$  or  $a_3$ . Both options are optimal with respect to the objective and confluence constraint and this arbitrary choice determines whether the outcome will be “yes” or “no”. Thus, this method can potentially make arbitrary decisions that affect the outcome of a vote. We can formalize and generalize this observation with the following theorem which states that a deterministic mechanism cannot guarantee multiple unranked preferences, the right to delegate, single-path explainability, and no arbitrary decisions.

**THEOREM 5.2.** *No deterministic mechanism with multiple unranked preferences can simultaneously satisfy the right to delegate, single-path explainability, and no arbitrary decisions.*

**PROOF.** The proof follows from the above discussion of Figure 3.  $\square$

Another issue is the potential to produce outcomes that are counter to an agent’s intention and violate the locally predictable property. The preference graphs in Figure 4 and their minimum confluent flows illustrate this effect. In (A), we see the first vote in which  $a_1$  prefers to delegate only to  $a_2$  who votes for “yes”, but “no” ultimately wins (3 votes for “yes” and 4 for “no”). Now, suppose  $a_1$  decides that “no” was truly the better choice and decides they would rather delegate to  $a_3$  who’s vote previously followed a path to “no” in the first vote (A). Then, in (B), we see a second vote in



**Figure 3: Illustrating a preference graph which forces the fluid mechanics approach to arbitrarily choose the outcome of the vote. The vote of  $a_1$  must be delegated to either  $a_2$  or  $a_3$ . Both options are optimal with respect to the objective and confluence constraint and this arbitrary choice determines whether the outcome will be “yes” or “no”.**

which  $a_1$  has changed their preference to  $a_3$ , but this change causes “yes” to win (4 votes for “yes” and 3 for “no”).

## 6 CONCLUSION AND FUTURE DIRECTIONS


The analysis in Section 5 suggests that ideals of fairness, explainability, transparency, and accountability can be at odds with goals of limiting concentrated power or selecting the ground truth correct outcome. A similar divide is present between practical implementations [2, 17] and theoretical work [14, 19] with the former giving less influence to a centralized mechanism while the latter type tend to explore how a centralized mechanism could exert a beneficial influence on the system. Given these trade-offs between desirable traits, it would be useful to better understand what other combinations of properties can or cannot be achieved.

Applications similar to those which have used Google Votes or LiquidFeedback might prefer a mechanism with certain desirable properties of bread-first delegation, but prioritizing edge ranks over shortest paths. In addition to multiple ranked preferences, SD, BC, right to delegate, right to top rank, 1-PE, local predictability, and polynomial running time, it would satisfy right to top rank and/or some other guarantee relating to edge ranks. This implies the additional open question of what other properties we could guarantee for giving agents their best ranked paths. E.g., one might incorporate an egalitarian objective which maximizes the minimum rank of any edge in the delegation graph.

When minimizing power is a high priority, we can ask if it’s possible to minimize power in some way while ensuring local predictability and SD or SDOD. This might involve using multiple ranked preferences or abandoning 1-PE to satisfy SD to sidestep the impossibility result of Theorem 5.2. Moreover, as suggested by Kotsialou and Riley [21], an empirical study could reveal which mechanisms tend to yield lower concentrations of power on real or simulated data.

On the epistemic side, there is the question of whether a liquid democracy mechanism can provide strict improvement over direct democracy while achieving features not possessed by GreedyCap, such as a right to delegate. Due to the negative result of Kahng et al. [19], the most promising direction might involve developing and justifying a new model of voter competence that is not subject to the same impossibility.

Finally, the approach of this paper could be extended to other democratic systems in which computation plays a role. For example, the computer science and mathematics communities have given



**Figure 4: A pair of preference graphs showing how the Fluid Mechanics approach can produce an output contrary to an agent's intention. The numbers above each edge indicate the amount of flow traveling through the edge after minimizing confluent flow. (A) We see the first vote in which  $a_1$  prefers to delegate only to  $a_2$  who votes for "yes", but "no" wins. (B) We see the second vote in which  $a_1$  changes their preference to  $a_3$  who had voted for "no" in the first vote, but this change causes "yes" to win.**

increasing attention to redistricting and gerrymandering in district-based representative democracies. Algorithms have been proposed to draw congressional district maps, thereby controlling the delegation of citizens' votes to representatives and affecting election outcomes [10, 13, 15, 22]. They are also used to measure gerrymandering in existing maps [9, 12, 18, 23] and have been shown to affect theoretical guarantees on voter incentives [6].

## REFERENCES

- [1] Sven Becker. 2012. "Web Platform Makes Professor Most Powerful Pirate". *Der Spiegel* (2012). <https://www.spiegel.de/international/germany/liquid-democracy-web-platform-makes-professor-most-powerful-pirate-a-818683.html> [Online; accessed 24-April-2022].
- [2] Jan Behrens, Axel Kistner, Andreas Nitsche, and Björn Swierczek. 2014. *The principles of LiquidFeedback*. Interaktive Demokratie e. V. Berlin.
- [3] Christian Blum and Christina Isabel Zuber. 2016. Liquid democracy: Potentials, problems, and perspectives. *Journal of Political Philosophy* 24, 2 (2016), 162–182.
- [4] Markus Brill. 2018. Interactive democracy. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. 1183–1187.
- [5] Markus Brill and Nimrod Talmon. 2018. Pairwise Liquid Democracy. In *IJCAI*, Vol. 18. 137–143.
- [6] Brian Brubach, Aravind Srinivasan, and Shawn Zhao. 2020. Meddling metrics: the effects of measuring and constraining partisan gerrymandering on voter incentives. In *Proceedings of the 21st ACM Conference on Economics and Computation*. 815–833.
- [7] Ioannis Caragiannis and Evi Micha. 2019. A Contribution to the Critique of Liquid Democracy. In *IJCAI* 116–122.
- [8] Lewis Carroll. 1884. *The principles of parliamentary representation*. Harrison and Sons.
- [9] Maria Chikina, Alan Frieze, and Wesley Pegden. 2017. Assessing significance in a Markov chain without mixing. *Proceedings of the National Academy of Sciences* 114, 11 (2017), 2860–2864.
- [10] Vincent Cohen-Addad, Philip N Klein, and Neal E Young. 2018. Balanced centroidal power diagrams for redistricting. In *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. 389–396.
- [11] Rachael Colley, Umberto Grandi, and Arianna Novaro. 2021. Smart voting. In *Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI 2020)*, International Joint Conferences on Artificial Intelligence (IJCAI), 1734–1740.
- [12] Daryl DeFord, Moon Duchin, and Justin Solomon. 2019. Recombination: A family of Markov chains for redistricting. *arXiv preprint arXiv:1911.05725* (2019).
- [13] Seyed A Esmaeli, Hayley Grape, and Brian Brubach. 2022. Centralized Fairness for Redistricting. *arXiv preprint arXiv:2203.00872* (2022).
- [14] Paul Gözl, Anson Kahng, Simon Mackenzie, and Ariel D Procaccia. 2018. The fluid mechanics of liquid democracy. In *International Conference on Web and Internet Economics*. Springer, 188–202.
- [15] Wes Gurnee and David B Shmoys. 2021. Fairmandering: A column generation heuristic for fairness-optimized political districting. In *SIAM Conference on Applied and Computational Discrete Algorithms (ACDA21)*. SIAM, 88–99.
- [16] Steve Hardt. 2014. Liquid Democracy with Google Votes. <https://www.youtube.com/watch?v=F4lkCECSBFw>
- [17] Steve Hardt and Lia CR Lopes. 2015. Google votes: A liquid democracy experiment on a corporate social network. (2015).
- [18] Gregory Herschlag, Han Sung Kang, Justin Luo, Christy Vaughn Graves, Sachet Bangia, Robert Ravier, and Jonathan C Mattingly. 2020. Quantifying gerrymandering in north carolina. *Statistics and Public Policy* 7, 1 (2020), 30–38.
- [19] Anson Kahng, Simon Mackenzie, and Ariel Procaccia. 2021. Liquid democracy: An algorithmic perspective. *Journal of Artificial Intelligence Research* 70 (2021), 1223–1252.
- [20] Christoph Carl Kling, Jérôme Kunegis, Heinrich Hartmann, Markus Strohmaier, and Steffen Staab. 2015. Voting behaviour and power in online democracy: A study of LiquidFeedback in Germany's Pirate Party. In *Ninth International AAAI Conference on Web and Social Media*.
- [21] Grammateia Kotsialou and Luke Riley. 2018. Incentivising participation in liquid democracy with breadth-first delegation. *arXiv preprint arXiv:1811.03710* (2018).
- [22] Yan Y. Liu, Wendy K. Tam Cho, and Shaowen Wang. 2016. PEAR: a massively parallel evolutionary computation approach for political redistricting optimization and analysis. *Swarm and Evolutionary Computation* 30 (2016), 78–92.
- [23] Jonathan C Mattingly and Christy Vaughn. 2014. Redistricting and the Will of the People. *arXiv preprint arXiv:1410.8796* (2014).
- [24] Alois Paulin. 2020. An overview of ten years of liquid democracy research. In *The 21st Annual International Conference on Digital Government Research*. 116–121.
- [25] José Ramos. 2015. Liquid democracy and the futures of governance. In *The Future Internet*. Springer, 173–191.
- [26] Yuzhe Zhang and Davide Grossi. 2021. Tracking Truth by Weighting Proxies in Liquid Democracy. *arXiv preprint arXiv:2103.09081* (2021).