

Tracking Truth by Weighting Proxies in Liquid Democracy

Yuzhe Zhang and Davide Grossi

Abstract

We study wisdom of the crowd effects in liquid democracy when agents are allowed to apportion weights to proxies by mixing their delegations. We show that in this setting—unlike in the standard one where votes are always delegated in full to one proxy—it becomes possible to identify delegation structures that optimize the truth-tracking accuracy of the group. We contrast this centralized solution with the group accuracy obtained in equilibrium when agents interact by greedily trying to maximize their own individual accuracy through mixed delegations, and study the price of anarchy of these games. While equilibria with mixed delegations may be as bad as in the standard delegations setting, they are never worse and may sometimes be better.

1 Introduction

Liquid democracy [6] is a form of proxy voting [25, 32, 2, 20, 11] where each proxy is delegable, thereby giving rise to so-called transitive delegations. In such a system each voter may choose to cast her vote directly, or to delegate her vote to a proxy, who may in turn decide whether to vote or delegate, and so pass the votes she has accrued further to yet another proxy. The voters who decide to retain their votes—the so-called gurus—cast their ballots, which now carry the weight given by the number of delegations they accrued. The system has been deployed in online democratic decision-support tools like LiquidFeedback [4]¹ and has been recently studied in political science and eDemocracy, as well as artificial intelligence (see [28] for a recent overview).

Contribution Our paper studies aspects of the truth-tracking properties of liquid democracy when agents are allowed to express delegations consisting of the apportionment of shares of a unit weight (i.e., the agent’s voting weight) to their proxies. We interpret these weights probabilistically, that is, as mixing of pure delegations. The issue we are after is to understand the extent to which weighted delegations could help the truth-tracking behavior of liquid democracy. We make three main contributions. *First*, we show that in this more general setting it is always possible for the agents to achieve maximal group accuracy by centrally coordinating their delegations (Theorem 2). *Second*, we extend the strategic model of liquid democracy developed in [5, 33] to the setting involving weighted delegations, i.e., mixed delegation strategies. In games in which agents greedily try to maximize their individual accuracies we show that weighted delegations enable equilibria that are better in terms of group accuracy, with respect to equilibria with pure delegations (Theorem 3). This, however, comes at the cost of a higher price of anarchy with respect to games with pure delegations. *Third*, we provide an interpretation of weighted delegations alternative to mixing, in which weights are modeled as shares of voting power that agents apportion to their proxies. This model leads to an alternative notion of utility in delegation games, and therefore to different equilibria. We prove the resulting notion of equilibrium to be weaker than the mixed delegations one (Theorem 5).

Related Work In the last couple of years, several papers in the area of computational social choice have focused on liquid democracy. Three lines of research have mainly been pursued, which are broadly intertwined. First, papers have pointed to potential weaknesses of voting with transitive delegations and suggested alternative schemes for liquid democracy, specifically

¹<https://liquidfeedback.org/>

focusing on delegation methods. Problems the literature has focused on include: delegation cycles and the failure of individual rationality in multi-issue voting [10, 8]; poor accuracy of group decisions as compared to those achievable via direct voting in non-strategic settings [22, 9], as well as strategic ones [5]; issues related to power [33]. In response to issues such as the ones highlighted above research has focused on the development of better behaved delegation schemes, e.g.: delegations with some level of centralized control [22]; delegations with preferences over trustees [8] or over gurus [18]; multiple delegations [19]; complex delegations like delegations to a majority of trustees [12]; dampened delegations [7]; breadth-first delegations [24]. Second, papers have focused on computational aspects of some of the themes mentioned above, like the computation of equilibria in delegation games [18] or the resolution of delegation cycles [16]. Finally, research has focused on the empirical study of implementations of liquid democracy in real-world applications [23, 28].

Our paper is a contribution to the first line of research mentioned above and is most directly related to [22, 9], which studied the truth-tracking properties of liquid democracy as opposed to direct voting. In particular [22] showed that no ‘local’ probabilistic procedure to select proxies can guarantee that liquid democracy is, at the same time, never less accurate (in large enough graphs) and sometimes strictly more accurate than direct voting. These negative results have been further strengthened along several lines in [9]. Like these papers we study delegations from a probabilistic perspective, but we use probabilistic delegations to leverage existing generalizations of the Condorcet jury theorem ([26, 31]) in which agents are assigned weights depending on their accuracy. Our starting point is to exploit these results to determine delegation graphs which are optimal in terms of truth-tracking behavior, and which are not feasible in the standard delegation setting. We then compare this centralized approach with a decentralized one in which agents choose their delegations with the sole aim of improving their individual accuracy. Besides the aforementioned [22, 9] the idea of weighted delegations has been proposed and studied in other papers on liquid democracy and proxy voting (see [14, 19, 1]). Our paper further develops the theory of weighted delegation schemes in a truth-tracking context.

Outline Section 2 introduces our model of liquid democracy and Section 3 describes a centralized mechanism to use weighted delegations in order to achieve optimal group accuracy. Sections 4 and 5 studies weighted delegations from a game-theoretic perspective and focuses on the quality of equilibria in those games measured as group accuracy. Finally, Section 6 studies an alternative model of weighted delegations based on the transfer of voting power. Section 7 concludes.

2 Preliminaries

2.1 Basic Model: Binary Voting for Truth-Tracking

Our model is based on the binary voting setting for truth-tracking [13, 21, 17]. The setting has already been applied to the study of liquid democracy by [22, 5, 9].

A finite set of agents $N = \{1, 2, \dots, n\}$ has to vote on whether to accept or reject an issue. The vote is supposed to track the correct state of the world—that is whether it is ‘best’ to accept or reject the issue. The agents’ ability to make the right choice (i.e., the agents’ error model) is represented by the agent’s *accuracy* $q_i \in [0.5, 1]$, for $i \in N$. Each agent is endowed with one vote and the result of such an election is determined by simple majority.

When agent $i \in N$ delegates to agent $j \in N$ we write $d_i = j$. Then $\mathbf{d} = (d_1, d_2, \dots, d_n)$ is called a *delegation profile* (or simply *profile*) and is a vector describing each agent’s delegation. Equivalently, delegation profiles can be usefully thought of as maps $\mathbf{d} : N \rightarrow N$, where $\mathbf{d}(i) = d_i$. When $d_i = i$, agent i votes on her own behalf. We call such an agent a *guru*. On the other hand, any agent who is not a guru, is called a *delegator*. For profile \mathbf{d} , and $C \subseteq N$, let $Gu(C, \mathbf{d})$ denote all gurus in C in profile \mathbf{d} , i.e., $Gu(C, \mathbf{d}) = \{i \in C \mid d_i = i\}$. Moreover we

simply denote by $Gu(\mathbf{d})$ all gurus in the whole agent set N . A delegation profile in which all agents are gurus (i.e., for all $i \in N$ $d_i = i$) is said to be *trivial*.

Any profile \mathbf{d} can also be represented by a directed graph. An edge from agent i to j exists whenever $d_i = j$. Consider then a profile \mathbf{d} where a path exists from i to j , i.e., $i \rightarrow i_1 \rightarrow \dots \rightarrow i_k \rightarrow j$. We call such paths *delegation chains*. When such a chain from i to guru j exists, every agent on this delegation chain (directly or indirectly) delegates to j , and we denote i 's guru by $d_i^* = \mathbf{d}^*(i) = j$. For any profile \mathbf{d} and any guru $i \in Gu(\mathbf{d})$, let $w_{\mathbf{d}}(i)$ denote the weight collected by i , i.e., $w_{\mathbf{d}}(i) = |\{j \in N \mid \mathbf{d}^*(j) = i\}|$. A *delegation cycle* is a chain where the first and last agents coincide. In such a case, no agent in the chain is linked to a guru. Therefore no agent linked via a delegation chain to an agent in a delegation cycle has a guru.

In several models for liquid democracy agents are thought of as nodes in an underlying network (e.g., [22, 5]). An underlying network is a directed graph $R = \langle N, E \rangle$, where E are directed edges. Let $R(i)$ denote the *neighborhood* of agent $i \in N$, i.e., $R(i) = \{i\} \cup \{j \in N \mid (i, j) \in E\}$. In those models, agents are only allowed to delegate agents in their neighborhoods. In this paper, we mainly focus on the setting in which such a network is complete, that is, every agent can delegate to every agent, except for Lemma 1 and Theorem 5 where we establish more general results.

2.2 Weighted Delegations

We are interested in the setting in which, instead of transferring their full vote to a proxy, agents apportion parts of their full weight to different proxies. We work under a probabilistic interpretation of this apportionment, so that the weight attached to a proxy corresponds to the probability of delegating to it (cf. the notion of local delegation mechanism in [22]).² So i 's mixed delegation amounts to a stochastic vector $\mathbf{D}_i = (D_{i1}, \dots, D_{in}) \in \mathbb{R}^n$ with $\sum_{j \in N} D_{ij} = 1$. We call such delegations weight-apportioning delegations or, simply, *weighted delegations*. A profile of weighted delegations (or, *weighted profile*) is an $n \times n$ -dimensional stochastic matrix $\mathbf{D} = (\mathbf{D}_1, \dots, \mathbf{D}_n)$, and \mathbb{D} is the collection of all weighted profiles. Moreover we can generate a directed graph, which is called a delegation graph, given a weighted profile \mathbf{D} . All agents are nodes in the delegation graph, and for any pair of $i, j \in N$, an directed edge from i to j exists if $D_{ij} > 0$. Clearly, a standard delegation $d_i = j$ corresponds to the vector in which the whole of i 's weight is apportioned to j , and so the j^{th} coordinate has value 1 and the remaining coordinates have 0. We will be referring to the standard type of delegations also as *pure delegations*.

Agents' weights Taking a weighted profile \mathbf{D} to denote the probability of all pure delegation strategies of each agent, we can obtain the probability of any standard profile \mathbf{d} as $\Pr(\mathbf{d}) = \prod_{i \in N} D_{id(i)}$. Then, given a weighted profile \mathbf{D} , for any agent $i \in N$, let \mathcal{D}_i denote all possible profiles, in which i is a guru, i.e., $\mathcal{D}_i = \{\mathbf{d} \mid \forall j \in N, D_{jd(j)} > 0, i \in Gu(\mathbf{d})\}$. Hence we denote the *weight* of any agent $i \in N$ under \mathbf{D} as i 's expected weight in all possible standard profiles, that is:

$$w_{\mathbf{D}}(i) = \sum_{\mathbf{d} \in \mathcal{D}_i} \Pr(\mathbf{d})w_{\mathbf{d}}(i). \quad (1)$$

In weighted profiles we call all agents, who have positive weight, gurus and denote their set $Gu(\mathbf{D}) = \{i \in N \mid w_{\mathbf{D}}(i) > 0\}$.

Group accuracy We call group accuracy (or majority accuracy) the probability that the group correctly identifies the state of the world via a weighted majority vote. The majority accuracy $q_{\mathbf{D}}$ of a weighted profile \mathbf{D} is therefore the probability that a randomly sampled

²Later in Section 6 we are going to explore also a different interpretation of these weights.

coalition of gurus $C \subseteq Gu(\mathbf{D})$, with property $\sum_{i \in C} w_{\mathbf{D}}(i) \geq \sum_{i \in Gu(\mathbf{D}) \setminus C} w_{\mathbf{D}}(i)$, contains agents that all vote correctly, while all agents in $Gu(\mathbf{D}) \setminus C$ vote incorrectly. That is:

$$q_{\mathbf{D}} = \sum_{C \in \mathcal{W}(\mathbf{D})} \prod_{i \in C} q_i \prod_{i \in Gu(\mathbf{D}) \setminus C} (1 - q_i) \quad (2)$$

where $\mathcal{W}(\mathbf{D}) = \{C \subseteq Gu(\mathbf{D}) \mid \sum_{i \in C} w_{\mathbf{D}}(i) \geq \sum_{i \in Gu(\mathbf{D}) \setminus C} w_{\mathbf{D}}(i)\}$ is the set of winning coalitions. When dealing with standard delegation profiles \mathbf{d} we will also write $q_{\mathbf{d}}$ for the majority accuracy determined by \mathbf{d} .

Example 1. Consider an example, where a set of agents $N = \{1, 2, 3, 4, 5\}$ with accuracies $(0.9, 0.9, 0.6, 0.6, 0.6)$. Assume the given weighted profile is \mathbf{D} , in which $D_{11} = 1$, $D_{22} = 1$, $D_{33} = 1$, $D_4 = (0.4, 0.4, 0.2, 0, 0)$, and $D_5 = (0.4, 0.3, 0.3, 0, 0)$, that is, agents 1, 2 and 3 always vote on their behalves, and the vectors D_4 and D_5 show the weighted delegation strategy of agent 4 and 5, e.g., agent 4 delegates to agent 1 by probability of 40%, to agent 2 by 40%, and to agent 3 by 20%. Then under \mathbf{D} , we can compute agents' weights by Equation 1 as follows:

$$(w_{\mathbf{D}}(1), w_{\mathbf{D}}(2), w_{\mathbf{D}}(3), w_{\mathbf{D}}(4), w_{\mathbf{D}}(5)) = (1.8, 1.7, 1.5, 0, 0).$$

Therefore the set of gurus are $Gu(\mathbf{D}) = \{1, 2, 3\}$, and by weighted majority rule, all winning coalitions are $\mathcal{W}(\mathbf{D}) = \{\{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$. Then take an instance of the winning coalition $\{1, 2\}$, such that the accuracy of this majority is computed by assuming agent 1 and 2 vote correctly, i.e., with probabilities (or their accuracies) $q_1 = 0.9$ and $q_2 = 0.9$, while agent 3 votes incorrectly with probability $1 - q_3 = 0.4$. That is $q_1 q_2 (1 - q_3) = 0.324$. Hence the majority accuracy is the sum of all accuracies of such majorities, which is as follows,

$$q_{\mathbf{D}} = \sum_{C \in \mathcal{W}(\mathbf{D})} \prod_{i \in C} q_i \prod_{i \in Gu(\mathbf{D}) \setminus C} (1 - q_i) = 0.324 + 0.054 + 0.054 + 0.486 = 0.918.$$

Remark 1. In the paper we are working under the one-voter one-vote assumption. It is important to notice, however, that this assumption can easily be lifted. All results presented can be generalized to the setting of voters with heterogeneous weights. It suffices to multiply each weight that agents apportion to proxies by the original weight of the delegating agent.

3 Centralized Weighted Delegations

3.1 Optimal Truth-Tracking via Weighted Delegations

Our motivation to study weighted delegations comes from known generalizations of the Condorcet jury theorem showing that the chance that the voting outcome of the group is correct is maximized if a weighted majority rule is used, instead of simple majority, with agents' weights proportional to the log of the $\frac{q_i}{1 - q_i}$ ratio. We state the relevant theorem here for ease of reference:

Theorem 1 ([21]). *The group accuracy is maximal if each agent i is assigned a weight proportional to $\log\left(\frac{q_i}{1 - q_i}\right)$.*

Proof sketch. Under the proposed weight allocation, for any winning coalition $N_1 \subseteq N$ (with $N - N_1 = N_2$), it holds that

$$\prod_{i \in N_1} q_i \prod_{i \in N_2} (1 - q_i) > \prod_{i \in N_1} (1 - q_i) \prod_{i \in N_2} q_i.$$

The reason is that if we take the log of both sides of the equivalent formula

$$\prod_{i \in N_1} \frac{q_i}{1 - q_i} > \prod_{i \in N_2} \frac{q_i}{1 - q_i},$$

we obtain

$$\sum_{i \in N_1} \log \left(\frac{q_i}{1 - q_i} \right) > \sum_{i \in N_2} \log \left(\frac{q_i}{1 - q_i} \right),$$

as desired. That is, a coalition with higher probability to make a correct decision (e.g., N_1) than its complement coalition (e.g., corresponding N_2) always has more weight. \square

By exploiting Theorem 1 we can show that weighted delegations offer a viable approach to the optimal delegation problem, that is: given a set of agents with different individual accuracies, what is the weighted delegation graph that maximizes group accuracy?

3.2 An Optimal Centralized Delegation Mechanism

We answer the above question by proposing an algorithm that uses one-hop mixed delegations to reallocate weight from the less accurate to the more accurate voters in the group. Let us define the optimal weight of each agent $i \in N$ by:

$$w_i^* = n \cdot \frac{\log \frac{q_i}{1 - q_i}}{\sum_{j \in N} \log \frac{q_j}{1 - q_j}} \quad (3)$$

Notice that this weight is larger than 1 for the more accurate agents whereas it is smaller for the less accurate ones and, as desired, it is proportional to $\log \frac{q_i}{1 - q_i}$. The idea behind the algorithm (Algorithm 1) is then to have the agents i with $w_i^* > 1$ apportion their full weight to themselves, and have each agent j with $w_j^* < 1$ apportion share $w_j^* - 1$ of the excess weight $1 - w_j^*$ to each agent i , normalized by the total excess weight of the i agents. Notice that if all agents are of the same accuracy ($N = N_3$), the algorithm returns the trivial profile. We illustrate the algorithm with the following example.

Example 2 (Example 1 cont'd). *Let us again use Example 1 to show the dynamics of Algorithm 1. We first compute $\log \frac{q_i}{1 - q_i}$ for all $i \in N$ as $(0.9542, 0.9542, 0.1761, 0.1761, 0.1761)$. Then w_i^* is computed by normalizing the above vector by the entire weight, which is in total 5 votes, and for all $i \in N$, w_i^* are $(1.958, 1.958, 0.3613, 0.3613, 0.3613)$. Hence we can observe that w_1^* and w_2^* are larger than their initial weight 1, and then they do not delegate. For any other agents $i \in \{3, 4, 5\}$, they delegate the surplus weight above $w_i^* = 0.3613$ to agents 1 and 2 equally since $w_1^* = w_2^*$. Therefore the returned profile is \mathbf{D} , in which $D_{11} = D_{22} = 1$, $D_3 = (0.31935, 0.31935, 0.3613, 0, 0)$, $D_4 = (0.31935, 0.31935, 0, 0.3613, 0)$, and $D_5 = (0.31935, 0.31935, 0, 0, 0.3613)$. Then for all $i \in N$, $w_{\mathbf{D}}(i) = w_i^*$.*

Theorem 2. *Algorithm 1 outputs an element of $\arg \max_{\mathbf{D} \in \mathbb{D}} q_{\mathbf{D}}$.*

Proof. First we show that in the weighted profile \mathbf{D} returned by Algorithm 1, the weight of any agent $i \in N$ is w_i^* . It can be observed that for any agent $i \in N$, with $w_i^* < 1$, $D_{ii} = w_i^*$, $D_{ij} = 0$ for all $j \in N_1 \setminus \{i\}$, and $D_{ij'} = (1 - w_i^*) \frac{w_{j'}^* - 1}{w}$ for all $j' \in N_2$. On the other hand, for all $i \in N_2$, i.e., agents with $w_i^* > 1$, $D_{ii} = 1$. Then for any $i \in N_1$, the probability of any $\mathbf{d} \in \mathcal{D}_i$, i.e., i is a guru in \mathbf{d} , or equivalently i delegates to herself, is D_{ii} . Moreover, for all $\mathbf{d} \in \mathcal{D}_i$, $w_{\mathbf{d}}(i) = 1$. Thus we have for all $i \in N_1$,

$$w_{\mathbf{D}}(i) = \sum_{\mathbf{d} \in \mathcal{D}_i} \Pr(\mathbf{d}) w_{\mathbf{d}}(i) = D_{ii} = w_i^*.$$

Then for any $i \in N_2$, i is always a guru since $D_{ii} = 1$, and let \mathcal{D}_{i,j_i} denote the set of profiles $\mathcal{D}_{i,j_i} = \{\mathbf{d} \mid i \in \text{Gu}(\mathbf{d}), \mathbf{d}(j) = i, \forall j' \in N, D_{j'\mathbf{d}(j') > 0}\}$ for any $j \in N_1$. That is an agent $j \in N_1$ delegates to $i \in N_2$ in any profile in \mathcal{D}_{i,j_i} . Then the probability of any profile in \mathcal{D}_{i,j_i} is $(1 - w_j^*) \frac{w_i^* - 1}{w}$, which is the weight-apportioning delegation from j to i . Also note that i always retains her weight of 1 in any $\mathbf{d} \in \mathcal{D}_i$. Therefore we have that

$$w_{\mathbf{D}}(i) = \sum_{\mathbf{d} \in \mathcal{D}_i} \Pr(\mathbf{d}) w_{\mathbf{d}}(i) = 1 + \sum_{j \in N_1} (1 - w_j^*) \frac{w_i^* - 1}{w}.$$

Algorithm 1 Optimal weights

Initialize: $w = 0$, $N_1 = \{i \in N \mid w_i^* < 1\}$, $N_2 = \{i \in N \mid w_i^* > 1\}$, $N_3 = \{i \in N \mid w_i^* = 1\}$.

Delegate:

```
if  $N_3 = N$  then
  for  $i \in N$  do
     $D_{i,i} = 1$ 
  end for
else
  for  $i \in N_2$  do
     $D_{i,i} = 1$ 
     $w = w + (w_i^* - 1)$ 
  end for
  for  $i \in N_1$  do
    for  $j \in N_2$  do
       $D_{i,j} = (1 - w_i^*) \frac{w_j^* - 1}{w}$ 
    end for
     $D_{i,i} = 1 - \sum_{k \in N_2} D_{i,k}$ 
  end for
end if
```

Return: \mathbf{D}

Notice that all delegations received by N_2 should equal the amount from N_1 , i.e., $\sum_{j \in N_1} (1 - w_j^*) = w$. Hence $w_{\mathbf{D}}(i) = w_i^*$ for all $i \in N_2$ and therefore, by [21, Th. 13], $q_{\mathbf{D}}$ is maximal. \square

It is worth observing that since pure delegation profiles are just degenerate weighted delegation profiles (so, they are also elements of \mathbb{D}) the optimal group accuracy reachable by weighted delegations cannot be possibly worse than the optimal group accuracy obtainable by pure delegations. But, crucially, it can be better, as shown in the following example.

Example 3 (Example 2 cnt'd). *Let us continue with Example 2. In that example the optimal pure delegation profile is the one in which only one delegation happens: an agent with accuracy of 0.6 delegates to an agent with accuracy of 0.9. Then the optimal (pure profile) majority accuracy is 0.918, which is lower than the optimal accuracy, 0.92664, of the weighted profile \mathbf{D} in Example 2.*

Intuitively, pure delegations allow for only discrete weights and can therefore only roughly approximate a weight distribution among gurus in which each winning coalition C of agents is more accurate than the corresponding losing coalition $N \setminus C$.

Observe finally that Algorithm 1 is polynomial in N . So it provides us with a tractable centralized mechanism to solve the optimal delegation problem when weighted delegations are available via mixing. In the next two sections we move from this centralized perspective to a decentralized one, exploring the behavior of weighted delegations in a game-theoretic setting.

Remark 2. *It is worth to also briefly discuss Algorithm 1 in the context of the GreedyCap algorithm proposed in [22]. GreedyCap is a local probabilistic delegation algorithm, with a centralized element: a cap on the maximal number of delegations. The cap essentially guarantees that delegations do not 'break' the wisdom of the crowd effect of independent voters by introducing too much correlation. Algorithm 1 implements a fully centralized approach to group accuracy by assuming delegations, and their weights, to be centrally determined.*

4 Decentralized Weighted Delegations

Algorithm 1 described a centralized mechanism to establish optimal weighted delegations within a group involved in a truth-tracking task. We move now to study the delegation structures that would arise if agents were to autonomously decide how to delegate.

4.1 Incentives to Delegate

We assume each agent i holds an evaluation function $u_i : N \rightarrow \mathbb{R}$ that associates a cardinal value to any agent $i \in N$. Intuitively $u_i(j)$, with $i, j \in N$, is the utility that i obtains if j acts as i 's guru. For example, u_i may be such that $u_i(j) = q_j$, that is, i 's utility would amount to the accuracy of j if j were to act as i 's guru (this is for instance the case in the game-theoretic setting of [5]). Note that the utility given by function u is independent of the specific delegation chains linking agents to their gurus: in two profiles \mathbf{d} and \mathbf{d}' , for agent $i \in N$, $\mathbf{d}(i) \neq \mathbf{d}'(i)$ but $\mathbf{d}^*(i) = \mathbf{d}'^*(i)$, then the utility of i is identical in the two profiles. Slightly abusing notations, let \mathbf{u}_i be the n -dimensional vector, of which the j -th ($j \in N$) element is $u_i(j)$.

Given a weighted profile \mathbf{D} , let \mathcal{D} be the set of possible pure profiles with positive probability, i.e., $\mathcal{D} = \{\mathbf{d} \mid \forall i \in N, D_{i\mathbf{d}(i)} > 0\}$. Then $U_i(\mathbf{D})$ is the expected utility of agent $i \in N$ under any weighted profile \mathbf{D} . $U_i(\mathbf{D})$ is a mapping $U : \mathbb{D} \rightarrow (N \rightarrow \mathbb{R})$, and defined as

$$U_i(\mathbf{D}) = \sum_{\mathbf{d} \in \mathcal{D}} u_i(d_i^*) \Pr(\mathbf{d}). \quad (4)$$

That is, $U_i(\mathbf{D})$ corresponds to i 's expectation over the utility it would obtain given the distribution over its possible gurus that is induced by \mathbf{D} .

4.2 Delegation Games

Equipped with the notion of utility we move to define delegation games as structures $G = \langle N, S, U \rangle$, where $N = \{1, 2, \dots, n\}$ is the set of agents, $S = \mathbb{D}$ is the strategy space³ of each agent $i \in N$ (so a strategy is the mixing of pure delegations), and U is as in equation (4). Observe that, since utility function U corresponds to the expectation over u given a profile \mathbf{D} , the corresponding delegation game can be viewed as the mixed-strategy version of the delegation game with pure delegations. By Nash theorem (see [27]) we therefore know that such games always have Nash Equilibria (NE), and we call such NE under weighted strategy U -NE. That is not the case for the pure delegation variant of the game (see [18]).

Let us also fix some auxiliary notation. Given a standard profile \mathbf{d} , a deviation by $i \in N$ from $\mathbf{d}(i)$ to d'_i yields the profile denoted (\mathbf{d}_{-i}, d'_i) , in which any agent in $N \setminus \{i\}$ delegates in the same way as in \mathbf{d} , but i delegates to d'_i . This notion of deviation can be extended to weighted profiles in the natural way where (\mathbf{D}_{-i}, D'_i) is the new profile, and D'_i is an n -dimensional stochastic vector.

Example 4 (A 2-agent delegation game). *Consider a game with $N = \{1, 2\}$, $\mathbf{u}(1) = (x_{11}, x_{12})$ and $\mathbf{u}(2) = (x_{21}, x_{22})$, where $x_{12} > x_{11}$ and $x_{21} > x_{22}$. That is, agent 1 prefers delegating to 2 over being a guru, and agent 2 prefers delegating to 1 over being a guru.*

If the strategy of agent 1 is $\mathbf{D}_1^1 = (1, 0)$, i.e., agent 1 delegates to herself by probability of 1, her utility is

$$U_1(\mathbf{D}^1) = x_{11},$$

since in any possible profile \mathbf{d} , $\mathbf{d}(1) = 1$. However, if the strategy of agent 1 is $\mathbf{D}_1^2 = (0, 1)$, her utility becomes

$$U_1(\mathbf{D}^2) = y_{22} \cdot x_{12},$$

³Here we only consider delegation in complete networks. We can restrict the strategy space to $S \subset \mathbb{D}$ for general networks, in which agents are only allowed to delegate to their neighbors in the network.




Figure 1: Equilibria for the delegation game of Example 4

where the strategy of agent 2 is (y_{21}, y_{22}) . This is the case because $y_{21} > 0$, in the corresponding delegation graph the two agents form a delegation cycle and their utility is therefore 0 in that case.

Therefore agent 1 would choose \mathbf{D}_1^1 if $y_{22} < x_{11}/x_{22}$, otherwise she chooses \mathbf{D}_1^2 . Similarly, assume that the strategy of agent 1 is (y_{11}, y_{12}) . For agent 2, we obtain that agent 2 chooses $\mathbf{D}_2^3 = (0, 1)$, i.e., always being a guru, if $y_{11} < x_{22}/x_{21}$, otherwise she chooses $\mathbf{D}_2^4 = (1, 0)$. The best strategies of both agents are shown in Fig. 1, and the U-NE are $\begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix}$, $\begin{pmatrix} \frac{x_{22}}{x_{21}} & 1 - \frac{x_{22}}{x_{21}} \\ 1 - \frac{x_{11}}{x_{12}} & \frac{x_{11}}{x_{12}} \end{pmatrix}$, and $\begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix}$.

4.3 Greedy Delegation (GD) Games

In the remainder of the paper we focus on a specific class of delegation games in which the utility u_i of agent i is set to equal the accuracy of i 's guru. So agents delegate by greedily trying to optimize their own individual accuracy. A *greedy delegation game* (GD game) is a delegation game as defined above where, however, the utility function U is such that:

$$U_i(\mathbf{D}) = \sum_{\mathbf{d} \in \mathcal{D}} q_{\mathbf{d}^*(i)} \Pr(\mathbf{d}). \quad (5)$$

So, in a GD game all agents evaluate gurus in the same way and i 's utility corresponds to the expected accuracy of her gurus. The pure delegation variant of these games has been studied in [5], where it has been proven that pure strategy NE always exist for them.

We observe that NE in GD games never contain weighted delegation cycles. This observation generalizes results for pure delegation games from [5] and it will be of use later in Section 6. In what follows, let $[k] = \{1, 2, \dots, k\}$ for any $k \in \mathbb{Z}_{++}$.

Lemma 1. *NE in GD games are acyclic.*

Proof. Note that this proof is for general underlying networks, and it henceforth holds for complete networks. We prove by contradiction. Assume that in a weighted strategy NE \mathbf{D} , there is a cycle in the corresponding delegation graph, i.e., there is a sequence of agents $(i_1, i_2, \dots, i_\ell)$ (let N^ℓ denote the set of these agents) such that $D_{i_j, i_{j+1}} > 0$ for all $j \in [\ell - 1]$ and $D_{i_\ell, i_1} > 0$. Observe that since \mathbf{D} is a weighted strategy NE, in the delegation graph of \mathbf{D} , for any agent $i \in N$, any edge from the agent is on a path to the highest-accuracy agent, who is reachable from i in the underlying network. Otherwise some agent in the network must be able to change her strategy to delegate to the reachable highest-accuracy agent, since all agents have identical utility function. Assume that the highest accuracy the set of agents N^ℓ can reach is q^* , that is in the delegation graph, there is a path from any agent in N^ℓ to an agent with accuracy of q^* . Then we consider two possible cases: (1) An agent $i \in N^\ell$ has accuracy of q^* . Then by adjusting her weighted strategy to $(\mathbf{D}_{-i}, (\dots, D'_{ii} = 1, \dots))$, i

can strictly improve her utility since she obtains 0 in the profiles with the delegation cycle by probability of $\prod_{a \in [\ell]} D_{i_a i_{a+1}} D_{i_\ell i_1}$. Contradiction. (2) All agents in N^ℓ has lower accuracy than q^* . Then in N^ℓ , we can always find an agent, for instance i_1 without loss of generality, such that $\exists k \in R(i_1) \setminus N^\ell$, the edge (i_1, k) is on a path in the delegation graph to an agent with accuracy of q^* . By adjusting i_1 's weighted strategy to $(\mathbf{D}_{-i_1}, (\dots, D'_{i_1 k} = 1, \dots))$, i_1 can also strictly improve her utility. This goes against the assumption that \mathbf{D} is a NE. Contradiction. \square

5 Social Welfare in Greedy Delegation Games

We measure social welfare in GD games by group accuracy and study this measure of welfare in equilibrium.

5.1 Group Accuracy in Equilibrium

Weighted delegations make it possible to achieve wisdom of the crowd effects in equilibrium by balancing weight among maximally accurate agents. As a result NE in GD games with weighted profiles can be shown to be never worse than NE with pure delegations and to be better in some cases.

Theorem 3. *Let a GD game be given and let $N^* \subseteq N$ be the set of m agents with maximal accuracy. Then any weighted profile \mathbf{D} , such that for any $i \in N^*$, $w_{\mathbf{D}}(i) = n/m$, is the NE with maximal accuracy.*

Proof. Assume that the set of agents N^* has identical highest accuracy of q^* . First notice that any weighted profile \mathbf{D} , such that $\sum_{i \in N^*} w_{\mathbf{D}}(i) = n$, is a mixed NE, since any agent in N then has a set of gurus in N^* and acquires utility of q^* . Therefore any weighted profile \mathbf{D}^* , such that for any $i \in N^*$, $w_{\mathbf{D}^*}(i) = n/m$, is a mixed NE.

Then in \mathbf{D}^* , we consider any bipartition $(N_1, N_2) \in \mathcal{B}(N^*)$ (where $\mathcal{B}(N^*)$ denotes the set of all bipartitions in N^*), where N_1 is the majority, i.e., $\sum_{a \in N_1} w_{\mathbf{D}^*}(a) > \sum_{a \in N_2} w_{\mathbf{D}^*}(a)$. Therefore the accuracy of this bipartition is $\prod_{a \in N_1} q_a \prod_{a \in N_2} (1 - q_a)$. Since $\forall i, j \in N^*$, $q_i = q_j = q^*$ and $w_{\mathbf{D}^*}(i) = w_{\mathbf{D}^*}(j)$, we have $|N_1| > |N_2|$, which leads to

$$q^{*|N_1|}(1 - q^{*|N_2|}) = \prod_{a \in N_1} q_a \prod_{a \in N_2} (1 - q_a) > \prod_{a \in N_2} q_a \prod_{a \in N_1} (1 - q_a) = q^{*|N_2|}(1 - q^{*|N_1|}).$$

Therefore if N^* is the set of gurus, \mathbf{D}^* maximizes the majority accuracy. By the Condorcet jury theorem, which states that larger group of agents (with homogeneous accuracy higher than 0.5) enhances group accuracy (see [31, Th. 1]), no NE \mathbf{D} in which $Gu(\mathbf{D})$ is a strict subset of N^* has higher group accuracy. \square

Then the following example shows that there exist GD games in which \mathbf{D}^* has strictly better group accuracy than any NE with pure delegations.

Example 5. *Consider a GD game, where there are 7 agents, 5 agents with maximal accuracy (q^*). In this example, any pure delegations NE with maximal group accuracy, would involve a pair of maximally accurate agents who get each a weight of 2. These two agents form a winning coalition, but they have a lower group accuracy than the remaining three gurus, i.e., $q^{*2}(1 - q^*)^3 < q^{*3}(1 - q^*)^2$. So the resulting group accuracy is strictly worse than that of \mathbf{D}^* .*

5.2 Price of Anarchy

We define the price of anarchy of GD games as follows:

$$\text{PoA} = \frac{\max_{\mathbf{D} \in \mathbb{D}} q_{\mathbf{D}}}{\min_{\mathbf{D} \in \text{NE}} q_{\mathbf{D}}} \quad (6)$$

where NE denotes the set of equilibria. When restricting to the price of anarchy in games with pure delegations (and therefore one pure delegations NE) we refer to PoA^{pure} .

Equation (6) gives us a measure of how much group accuracy is 'lost' in equilibrium, in the worst case, with respect to what would be centrally achievable via Algorithm 1.

Theorem 4. *When $|N| \rightarrow \infty$, $\text{PoA} \rightarrow \frac{1}{q^*}$, where q^* is the accuracy of a maximally accurate agent in N .*

Proof. Let \mathbf{D} be the weighted profile for which $q_{\mathbf{D}}$ is maximal and let $\dot{\mathbf{D}}$ be the equilibrium profile for which $q_{\dot{\mathbf{D}}}$ is minimal. This is the case when all agents delegate to the same guru, which has accuracy q^* since it is in equilibrium by assumption. Then, by [21, Th. 13], and the Condorcet jury theorem, as $|N| \rightarrow \infty$, $q_{\mathbf{D}} \rightarrow 1$ and by construction $q_{\dot{\mathbf{D}}} = q^*$. \square

It is worth noticing that the same argument can be applied to the setting with pure delegations, obtaining the same asymptotic value for the price of anarchy.

Intuitively in asymptotic cases, the PoA coincides whether weighted strategy is allowed or not. Then we consider cases, in which $|N|$ is not asymptotic. That is the agent size is given. The following result shows that weighted strategy outperforms pure strategy in aspect of PoA. Observe that, in the non-asymptotic case, since weighted delegations enable optimal group accuracy (Theorem 2) while pure delegations do not (Example 3), the price of anarchy in GD games with weighted delegations is trivially higher than in the case of pure delegations.

Corollary 1. $\text{PoA} \geq \text{PoA}^{pure}$

6 Discussion: Weights as Shares of Votes

We interpreted the weight assigned to proxies in a probabilistic way. However, such weights could also be interpreted as an actual allocation of shares of voting power to proxies (cf. the notion of 'splittable' delegations discussed in [19]). The two views lead to different insights, and in this section we briefly develop this second perspective and relate it to the probabilistic interpretation.

6.1 Shares of Votes

We move now to interpret weights as shares of voting rights apportioned to proxies. Before any delegation happens each agent is endowed with 1 vote and may decide to transfer shares of this 1 vote to other agents. So in that initial situation, if i delegates to j then \mathbf{D}_{ij} is the share of i 's vote that is transferred to agent j . But in general, i may also receive delegations from other agents, then D_{ij} would correspond to the share of i 's current voting power that is transferred to j .

Under this interpretation weighted profiles describe the extent of direct transfers of voting power between nodes. Formally, by the assumption that each agent is endowed with 1 vote, we denote the weight of all agents by a vector $w \in \mathbb{R}^n$ and the initial weight is $w^{(0)} = (1, \dots, 1)$. Agents then delegate according to \mathbf{D} and obtain the delegated weight

$$w^{(1)} = w^{(0)}\mathbf{D}$$

where the weight of agent i is $w_i^{(1)} = \sum_{j \in N} w_j^{(0)} D_{ji}$. This captures the one-step transfer of the initial voting power. But in liquid democracy these transfers are not only direct: by transferring power to agent k that in turn transfers power to agent i , an agent j indirectly transfers shares of her power to i . So for each positive integer t we can compute the t -steps transfers as $w^{(t+1)} = w^{(t)}\mathbf{D}$. We say then that the weighted profile \mathbf{D} converges if for some positive integer k , $w^{(k+1)} = w^{(k)}$.

The above process describes a Markov chain and we can use established results to obtain sufficient conditions for the convergence of weighted profiles (cf. [15] and [30, 29]) such that,

if a weighted profile \mathbf{D} is irreducible⁴ and aperiodic, then we can obtain a stationary weight apportionment in finite steps.⁵

If \mathbf{D} converges, and we can therefore obtain the ‘stationary’ transfers of voting power within the group, we can also attach a utility for each agent i given \mathbf{D} , as follows. Consider $\chi_i^{(0)} \in \mathbb{R}^n$, where the i -th element is 1. That is $\chi_i^{(0)}$ denotes the initial weight apportionment of agent i , where the totality of i ’s vote is allocated to i herself. Then for any positive integer t , $\chi_i^{(t+1)} = \chi_i^{(t)}\mathbf{D}$ denotes the delegation of i ’s weight at round t . If the Markov chain of \mathbf{D} satisfies the above condition, we have that $\exists k \in \mathbb{Z}_{++}$, $\chi_i^{(k+1)} = \chi_i^{(k)} = \chi_i^{(s)}$, where $\chi_i^{(s)}$ denotes the stationary weight distribution of agent i . In this case, we say that \mathbf{D} converges on i . Then in $\chi_i^{(s)}$, each positive element is a guru of i with the corresponding weight. Let then $\chi_i^{(s)}(j)$ denote the share of i ’s initial vote that is ultimately apportioned to j , namely the j -th element in $\chi_i^{(s)}$, for any $j \in N$.

Remark 3. *It is worth noticing that this notion of agents’ weight trivially coincides with the probabilistic one developed in the earlier sections when delegations are only one-hop, that is, when delegation chains are of length at most one. In particular this makes Algorithm 1 relevant for both models.*

Now that we have a way to compute the weight of agents we can define their utility as a mapping $\hat{U} : \mathbb{D} \rightarrow (N \rightarrow \mathbb{R})$, such that

$$\hat{U}_i(\mathbf{D}) = \begin{cases} \sum_{j \in N} u_i(j)\chi_i^{(s)}(j) & \text{if } \mathbf{D} \text{ converges on } i \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

Observe that we treat the failure to converge, and therefore the lack of a stationary weight apportionment for an agent, as a situation in which the agent gets 0 utility as she does not get representation by any guru. Observe furthermore that \mathbf{D} might also converge on some $i \in N$, even though it does not satisfy the above condition. For example in the delegation graph of \mathbf{D} , i is not linked to a periodic strongly connected component.

We illustrate this alternative way of interpreting weighted delegations by discussing the game that would result from Example 4 using the utility function \hat{U} of equation (7). As follows we call \hat{U} -NE the NE under the utility of \hat{U} .

Example 6. *Consider the same 2-player game of Example 4 but now using equation (7) to determine the utility function. We first consider the strategy of agent 1, given the weighted strategy of agent 2 is $D_2 = (y_{21}, y_{22})$. If $y_{21} \in (0, 1]$, agent 1 would vote on her own behalf, otherwise \mathbf{D} is not aperiodic. On the other hand, if $y_{21} = 0$, i.e., agent 2 is a self-voter, agent 1 would choose to delegate full weight to agent 2 since $x_{12} > x_{11}$. We can obtain symmetric best response strategy for agent 2 given $D_1 = (y_{11}, y_{12})$. Then both agents’ optimal strategies are drawn in Fig. 2. The \hat{U} -NE are $\begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix}$ and $\begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix}$.*

The example shows that the two notions of utility in equations (4) and (7) lead to different types of games, with different equilibria. As we show below, the two notions of utility are independent generalizations of the notion of utility under pure delegations.

6.2 U - and \hat{U} -Equilibria in GD Games

We finally establish a relation between the two types of equilibria in the class of GD games. Note that the following result holds for general underlying networks.

Theorem 5. *For any GD game, if a weighted profile \mathbf{D} is a U -NE it is also a \hat{U} -NE.*

⁴Recall that a Markov chain is irreducible if each state is accessible by each other state.

⁵Some periodic Markov chains may converge in infinite steps. We focus on convergence in finite steps as this has a natural interpretation in our setting as an iterated process in which shares of voting power are delegated until a stationary weight distribution (a fixpoint) is reached.




Figure 2: Equilibria for the game of Example 6

Proof. For any agent $i \in N$, assume the highest-accuracy agent i can access in the underlying network is $i^* \in N$ (there might be multiple highest-accuracy agents and similar argument can be applied), with accuracy of q^* . Since \mathbf{D} is a mixed NE and by Lemma 1, we have that for any i and j , if $D_{ij} > 0$, any path containing edge (i, j) terminates at i^* . Therefore we have that \mathbf{D} converges and i 's full weight is delegated to i^* in $\lim_{k \rightarrow \infty} w_i^{(k)}$. That is, i cannot further improve her utility, which implies \mathbf{D} is a \hat{U} -NE. \square

Note that the above theorem does not hold in the other direction. Two-player games offer a simple example. Infinitely many \hat{U} -NE exist where the most accurate agent is a sink in \mathbf{D} . However, the only U -NE in this case is the profile where the maximally accurate agent is a sink and the other agent delegates to her with probability 1.

7 Conclusions and Future Work

The paper has studied, from a truth-tracking perspective, a variant of liquid democracy where proxies are weighted by a probability of delegation or, alternatively, by a share of the delegator's voting weight. We showed that this type of weighted delegations enable, through central coordination, optimal group accuracy and, in a decentralized setting, equilibria with better group accuracy.

The work presented relies on one strong assumption, namely that agents can delegate to every other agent, so that the interaction network is complete. It will be a priority to lift this assumption in future work. In particular, it has been shown [9] that, in the general case in which delegations are constrained by an underlying network, it is intractable to approximate the optimal group accuracy $q_{\mathbf{d}}$ that could be obtained by centrally determining a standard delegation graph \mathbf{d} . Adapting our Algorithm 1 to such general setting would provide an interesting angle from which to approach this optimal delegation problem.

Then, while we studied weighted delegations in a majority voting context from a wisdom of the crowd perspective, it would be interesting to test them in settings involving different types of voting mechanisms. An example are random dictatorships of the sort used, for instance, in proof-of-stake blockchain protocols where stakes can be delegated to proxies in the network (e.g., Tezos [3]).

Acknowledgments The authors would like to thank the anonymous reviewers of COM-SOC'21 for several helpful comments on an earlier draft of this paper. Davide Grossi wishes to acknowledge partial support for this research by the Hybrid Intelligence Center, a 10-year programme funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research, <https://hybrid-intelligence-centre.nl>.

References

- [1] Ben Abramowitz and Nicholas Mattei. Flexible representative democracy: An introduction with binary issues. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 3–10, 2019.
- [2] Dan Alger. Voting by proxy. *Public Choice*, 126(1-2):1–26, jan 2006.
- [3] Victor Allombert, Mathias Bourgoïn, and Julien Tesson. Introduction to the tezos blockchain. In *Proceedings of the 17th International Conference on High Performance Computing & Simulation, HPCS 2019, Dublin, Ireland, July 15-19, 2019*, pages 1–10, 2019.
- [4] Jan Behrens, Axel Kistner, Andreas Nitsche, and Björn Swierczek. *Principles of Liquid Feedback*. Interaktieve Democratie, 2014.
- [5] Daan Bloembergen, Davide Grossi, and Martin Lackner. On rational delegations in liquid democracy. In *Proceedings of the AAI Conference on Artificial Intelligence*, volume 33, pages 1796–1803, 2019.
- [6] Christian Blum and Christina Isabel Zuber. Liquid democracy: Potentials, problems, and perspectives. *Journal of Political Philosophy*, 24(2):162–182, 2016.
- [7] Paolo Boldi, Francesco Bonchi, Carlos Castillo, and Sebastiano Vigna. Viscous democracy for social networks. *Communications of the ACM*, 54(6):129–137, 2011.
- [8] Markus Brill and Nimrod Talmon. Pairwise liquid democracy. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 137–143, 2018.
- [9] Ioannis Caragiannis and Evi Micha. A contribution to the critique of liquid democracy. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 116–122, 2019.
- [10] Zoé Christoff and Davide Grossi. Binary voting with delegable proxy: An analysis of liquid democracy. In *Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK’17)*, volume 251, pages 134–150. EPTCS, 2017.
- [11] Gal Cohensius, Shie Mannor, Reshef Meir, Eli Meir, and Ariel Orda. Proxy voting for better outcomes. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 858–866. International Foundation for Autonomous Agents and Multiagent Systems, 2017.
- [12] Rachael Colley, Umberto Grandi, and Arianna Novaro. Smart voting. In *Proceeding of the the 29th International Joint Conference on Artificial Intelligence (IJCAI), 2020*, 2020.
- [13] Marquis de Condorcet, M.J.A.N. de C. *Essai sur l’Application de l’Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. Imprimerie Royale, Paris, 1785.
- [14] Jonas Degraeve. Resolving multi-proxy transitive vote delegation. *CoRR*, abs/1412.4039, 2014.
- [15] Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- [16] Palash Dey, Arnab Maiti, and Amatya Sharma. On parameterized complexity of liquid democracy. In Apurva Mudgal and C. R. Subramanian, editors, *Algorithms and Discrete Applied Mathematics - 7th International Conference, CALDAM 2021, Rupnagar, India, February 11-13, 2021, Proceedings*, volume 12601 of *Lecture Notes in Computer Science*, pages 83–94. Springer, 2021.

- [17] Edith Elkind and Arkadii Slinko. Rationalizations of voting rules. In *Handbook of Computational Social Choice*, chapter 8, pages 169–196. Cambridge University Press, 2016.
- [18] Bruno Escoffier, Hugo Gilbert, and Adèle Pass-Lanneau. Iterative delegations in liquid democracy with restricted preferences. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI’20)*, pages 1926–1933, 2020.
- [19] Paul Gözl, Anson Kahng, Simon Mackenzie, and Ariel D Procaccia. The fluid mechanics of liquid democracy. In *Proceedings of the 14th Conference on Web and Internet Economics (WINE’18)*, pages 188–202, 2018.
- [20] James Green-Armytage. Direct voting and proxy voting. *Constitutional Political Economy*, 26(2):190–220, 2015.
- [21] Bernard Grofman, Guillermo Owen, and Scott L. Feld. Thirteen theorems in search of the truth. *Theory and Decision*, 15(3):261–278, 1983.
- [22] Anson Kahng, Simon Mackenzie, and Ariel D. Procaccia. Liquid democracy: An algorithmic perspective. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI’18)*, 2018.
- [23] Christoph Kling, Jérôme Kunegis, Heinrich Hartmann, Markus Strohmaier, and Steffen Staab. Voting behaviour and power in online democracy: A study of LiquidFeedback in Germany’s Pirate Party. In *Proceedings of the International Conference on Weblogs and Social Media*, 2015.
- [24] Grammateia Kotsialou and Luke Riley. Incentivising participation in liquid democracy with breadth-first delegation. In *Proceedings of the International Foundation for Autonomous Agents and Multiagent Systems (AAMAS ’20)*, pages 638–644, Richland, SC, 2020. IFAAMAS.
- [25] James C Miller. A program for direct and proxy voting in the legislative process. *Public choice*, 7(1):107–113, 1969.
- [26] Shmuel Nitzan and Jacob Paroush. Optimal decision rules in uncertain dichotomous choice situations. *International Economic Review*, 23(2):289–297, 1982.
- [27] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [28] Alois Paulin. An overview of ten years of liquid democracy research. In *Proceedings of the 21st International Conference on Digital Government Research*, 2020.
- [29] Anton V. Proskurnikov and Roberto Tempo. A tutorial on modeling and analysis of dynamic social networks. Part I. *Annual Reviews in Control*, 43:65–79, 2017.
- [30] Felix R. Gantmacher. *The Theory of Matrices*. AMS Chealsea Publishing, 1959.
- [31] Lloyd Shapley and Bernard Grofman. Optimizing group judgmental accuracy in the presence of interdependencies. *Public Choice*, 43(3):329–343, 1984.
- [32] Gordon Tullock. Computerizing politics. *Mathematical and Computer Modelling*, 16(8-9):59–65, aug 1992.
- [33] Yuzhe Zhang and Davide Grossi. Power in liquid democracy. In *Proceedings of the AAAI conference on Artificial Intelligence*, 2021.

Yuzhe Zhang
University of Groningen
Groningen, the Netherlands
Email: yoezy.zhang@rug.nl

Davide Grossi
University of Groningen
Groningen, the Netherlands
University of Amsterdam
Amsterdam, the Netherlands
Email: d.grossi@rug.nl