Tracking Truth by Weighting Proxies in Liquid Democracy

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

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

ABSTRACT

We study wisdom-of-the-crowd effects in liquid democracy on networks where agents are allowed to apportion parts of their voting weight to different proxies. We show that in this setting—unlike in the standard one where voting weight is delegated in full to only one proxy—it becomes possible to construct delegation structures that optimize the truth-tracking ability of the group. Focusing on group accuracy we contrast this centralized solution with the setting in which agents are free to choose their weighted delegations by greedily trying to maximize their own individual accuracy. While equilibria with weighted delegations may be as bad as with standard delegations, they are never worse and may sometimes be better. To gain further insights into this model we experimentally study quantal response delegation strategies on random networks. We observe that weighted delegations can lead, under specific conditions, to higher group accuracy than simple majority voting.

KEYWORDS

Liquid Democracy; Weighted Delegations; Epistemic Social Choice

ACM Reference Format:

Yuzhe Zhang and Davide Grossi. 2022. Tracking Truth by Weighting Proxies in Liquid Democracy. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, IFAAMAS, 9 pages.

1 INTRODUCTION

Liquid democracy [5] is a form of proxy voting [2, 10, 20, 31, 37] where proxy votes are also 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 retain their votes cast their ballots, which now carry the weight given by the number of delegations they accrued. The system has been implemented in decision-support tools like LiquidFeedback [3] and has been object of much research (see [34] for an 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. This functionality is available in some implementations of liquid democracy (e.g. on the platform Congressus of the French Pirate Party). We first 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 four 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 3). *Second*, we extend the strategic model of liquid democracy developed in [4, 38] to the setting involving weighted delegations. 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 4). 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 probabilistic one (Theorem 7). *Fourth*, we provide experimental evidence, via simulations, of high truth-tracking performance of weighted delegations even in decentralized settings, if agents are boundedly rational according to the quantal response model [29].

Related Work. Three main lines of research in liquid democracy may be broadly identified. First, papers have pointed to potential weakenesses of voting with transitive delegations and suggested alternative schemes for liquid democracy, specifically focusing on delegation methods. Problems the literature has focused on include: delegation cycles and the failure of individual rationality in multi-issue voting [7, 9]; poor accuracy of group decisions as compared to those achievable via direct voting in non-strategic settings [8, 26], as well as strategic ones [4]; issues related to power [38]; and impossibility results concerning proxy selection [22]. In response to these issues research has focused on the development of better behaved delegation schemes, e.g.: delegations with some level of centralized control [26]; delegations with preferences over trustees [7] or over gurus [17]; multiple delegations [19]; complex delegations like delegations to a majority of trustees [11]; dampened delegations [6]; breadth-first delegations [28]. Second, papers have focused on computational aspects of some of the themes mentioned above, like the computation of equilibria in delegation games [17] or the resolution of cycles [15]. Finally, implementations of liquid democracy for real-world applications have been studied [27, 34].

Our paper is a contribution to the first line of research mentioned above and is most directly related to [8, 26], which studied the truth-tracking properties of liquid democracy as opposed to direct voting. In particular [26] showed that no 'local' probabilistic procedure for proxy-selection can guarantee that liquid democracy is, at the same time, never less accurate (in large enough graphs) and sometimes strictly more accurate than direct simple majority voting. These negative results have been further strengthened along

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online.

several lines in [8]. Like these papers we study delegations from a probabilistic perspective, but we use probabilistic delegations to leverage existing generalizations of the Condorcet jury theorem ([32, 36]) 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 performance. 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 [8, 26] the idea of weighted delegations has been proposed and studied in other papers on liquid democracy and proxy voting (see [1, 13, 19]). Our paper further develops the theory of weighted delegation schemes in a truth-tracking context.

Outline. Section 2 introduces our model and Section 3 describes a centralized mechanism to use weighted delegations in order to achieve optimal group accuracy. Section 4 studies weighted delegations from a game-theoretic perspective and focuses on group accuracy in equilibrium. Section 5 studies an alternative model of weighted delegations based on the transfer of voting power. Section 6 presents our experimental results and Section 7 concludes.

2 WEIGHTING PROXIES

2.1 Binary Truth-Tracking with Delegations

Our model is based on the binary voting setting for truth-tracking [12, 16, 21]. The setting has already been applied to the study of liquid democracy by [4, 8, 26].

A finite set of agents $N = \{1, 2, \ldots, n\}$ have 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. Both options are assumed to have equal prior probability. The agents' ability to make the right choice is represented by the agent's *accuracy* $q_i \in [0.5, 1]$, for $i \in N$, corresponding to the conditional probability of i accepting (respectively, rejecting) the issue if it is indeed best to accept (respectively, reject) it. Each agent is endowed with one vote and the result of such an election is determined by (weighted) majority. Agents may accrue voting power through delegations. When agent i delegates to agent j we write $d_i = j$. Then $\mathbf{d} = (d_1, d_2, \ldots, d_n)$ is called a *delegation profile* (or simply a *profile*) and is a vector describing each agent's delegation. Equivalently, delegation profiles can be thought of as maps $\mathbf{d} : N \longrightarrow N$, where $\mathbf{d}(i) = d_i$. When $d_i = i$, agent i votes on her own behalf. We call such an agent a *guru*. Any agent who is not a guru, is called a *delegator*. For profile $\mathbf{d}$, and $C \subseteq N$, $Gu(C, \mathbf{d})$ denotes the set of all gurus from C in $\mathbf{d}$, i.e., $Gu(C, \mathbf{d}) = \{i \in C \mid d_i = i\}$. We write $Gu(\mathbf{d})$ instead of $Gu(N, \mathbf{d})$ to denote the set of all gurus. A profile $\mathbf{d}$ can be represented by a directed graph (the *delegation graph* of $\mathbf{d}$). An edge $\longrightarrow$ from agent i to j exists whenever $d_i = j$. A path in $\mathbf{d}$ from i to j, i.e., $i \longrightarrow i_1 \longrightarrow \cdots \longrightarrow i_k \longrightarrow j$, is called a *delegation chain*. When a delegation chain exists from i to a guru j we denote i's guru by $d_i^* = \mathbf{d}^*(i) = j$. A *delegation cycle* is a chain where the first and last agents coincide. A cycle of length one ($i \longrightarrow i$) is called a *loop*.

The weight accrued by an agent via delegations in $\mathbf{d}$ is:

$$w_{\mathbf{d}}(i) = \begin{cases} |\{j \in N \mid \mathbf{d}^*(j) = i\}| & \text{if } i \in Gu(\mathbf{d}) \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

As gurus are the only ones voting in $\mathbf{d}$ they accrue the power transferred by delegators through delegation chains. Observe that because of delegation cycles $\sum_{i \in N} w_{\mathbf{d}}(i)$ may be smaller than n. That is, voting weight is lost by delegation cycles. This is in line with the intuition that agents not linked to a guru fail to relay their votes to the mechanism (see also [9]).

Finally, we assume delegations to be constrained by a network represented by an undirected graph $R = \langle N, E \rangle$. In this paper we assume R to be connected: any pair of agents $i, j \in N$ are linked by a path in R. For $i \in N$, $R(i)$ denotes the *neighborhood* of i, i.e., $R(i) = \{i\} \cup \{j \in N \mid (i, j) \in E\}$. Agents are able to delegate only to agents in their neighborhoods. We write $R'(i)$ to denote $R(i) \backslash \{i\}$.

2.2 Weighted Delegations

We generalize the above setting by allowing agents to apportion parts of their voting power to different proxies: i's delegation amounts now to a stochastic vector $\mathbf{D}_i = (D_{i1}, \ldots, D_{in}) \in \mathbb{R}^n_{\geq 0}$ with $\sum_{j \in N} D_{ij} = 1$. We call such delegations *weighted delegations.* A profile of weighted delegations (*weighted profile*) is an $n \times n$-dimensional stochastic matrix $\mathbf{D} = (\mathbf{D}_1, \ldots, \mathbf{D}_n)$, and $\mathbb{D}$ is the collection of all such profiles. A standard delegation profile $\mathbf{d}$ corresponds then to a degenerate stochastic matrix where each row contains only one 1 entry. We will be referring to standard delegations also as *pure delegations*. We use notation $(\mathbf{D}'_i, \mathbf{D}_{-i})$ to refer to a profile obtained from $\mathbf{D}$ by replacing $\mathbf{D}_i$ with $\mathbf{D}'_i$.

A weighted profile $\mathbf{D}$ defines a weighted directed graph $G(\mathbf{D}) = \left(N, \xrightarrow{x}\right)$ where for any pair of $i, j \in N$, a directed edge $i \xrightarrow{D_{ij}} j$ from i to j with weight D_{ij} exists whenever $D_{ij} > 0$. Weighted delegation chains, cycles and loops can then be defined on these graphs as we did for graphs of pure delegations.

2.3 Agents' Weights after Delegations

Several interpretations of a voter's weight become possible under weighted delegations. Here we deal with two interpretations. The first one is based on a probabilistic interpretation of the weights, and it will be the one we use to develop our framework. Later, in Section 5, we are going to interpret weights also as direct transfers of shares of voting weight, and compare the two approaches.

Each weighted profile $\mathbf{D}$ can be thought of as describing a probability distribution over pure profiles where the probability of a pure profile $\mathbf{d}$ is $\Pr(\mathbf{d}) = \Pi_{i \in N} D_{i\mathbf{d}(i)}$. The *weight transfer* of i in $\mathbf{D}$, is the vector $\boldsymbol{t}^i_{\mathbf{D}} = (t_{\mathbf{D}}(i, 1), \ldots, t_{\mathbf{D}}(i, n))$ describing how i's weight is distributed in expectation among all guru agents, where:

$$\boldsymbol{t}_{\mathbf{D}}(i, j) = \sum_{\mathbf{d} \in s(\mathbf{D})} \mathbb{1}_{\mathbf{d}^*(i) = j} \Pr(\mathbf{d}) \tag{2}$$

where $j \in N$, $s(\mathbf{D})$ denotes the support of (the probability distribution over pure profiles induced by) $\mathbf{D}$, and $\mathbb{1}_{\mathbf{d}^*(i) = j}$ is the indicator that j is the guru of i in $\mathbf{d}$, i.e., $\mathbb{1}_{\mathbf{d}^*(i) = j} = 1$ if $\mathbf{d}^*(i) = j$, otherwise 0. We will refer to this as the *expected weight approach* (cf. the notion of local delegation mechanisms in Kahng et al. [26]).

Under Equation (2) the weight that agent i accrues consists simply of the sum of the weights she receives from all agents:

$$w_{\mathbf{D}}(i) = \sum_{j \in N} t_{\mathbf{D}}(j, i). \tag{3}$$

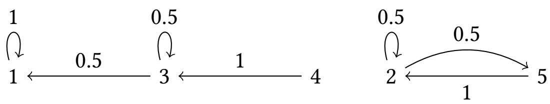

Figure 1: Delegation graph of Example 1

Equation (3) defines, for each $\mathbf{D}$, vector $\mathbf{w_D} = (w_\mathbf{D}(1), \ldots, w_\mathbf{D}(n))$, which we call the *weight distribution* of $\mathbf{D}$, assigning a weight to each agent. Then we denote by $Gu(\mathbf{D}) = \{i \in N \mid w_\mathbf{D}(i) > 0\}$ the set of gurus in $\mathbf{D}$, that is, the set of agents with positive weight in the weight distribution of $\mathbf{D}$. It is worth observing that $\sum_{i \in N} w_\mathbf{D}(i)$ can be less than n, because agents may end up having no guru in pure profiles where they delegate into delegation cycles. In such cases, the agent loses the weight corresponding to the probability attached to such pure profiles.

Example 1. *Consider a set of agents $N = \{1, 2, 3, 4, 5\}$, with weighted profile $\mathbf{D}$, such that $D_{11} = 1$ (i.e., maintaining her full voting weight), $\mathbf{D}_2 = (\ldots, D_{22} = 0.5, \ldots, D_{25} = 0.5)$, $\mathbf{D}_3 = (D_{31} = 0.5, \ldots, D_{33} = 0.5, \ldots)$, $D_{43} = 1$, and $D_{52} = 1$. The delegation graph is as in Figure 1. By the expected weight approach, for the component consisting of agents 1, 3 and 4, agent 4 always fully delegates to 3, while agent 3 is expected to keep half of the delegation from 4 and half of her own weight, then $w_\mathbf{D}(4) = 0$ and $w_\mathbf{D}(3) = 1 \times 0.5 + 1 \times 0.5 = 1$. The remaining weight in this component is delegated to agent 1, i.e., $w_\mathbf{D}(1) = 2$. Then for the component consisting of 2 and 5, the two agents form a delegation cycle with probability 0.5, while in the other possible standard profile (with probability 0.5), agent 2 is the only guru. Therefore $w_\mathbf{D}(2) = 1$ and $w_\mathbf{D}(5) = 0$. So $\mathbf{w_D} = (2, 1, 1, 0, 0)$. Observe that $\sum_{i \in N} w_\mathbf{D}(i) = 4 < n = 5$, since agents 2 and 5 are caught in a cycle with probability 0.5, thereby losing one unit of weight.*

2.4 Group Accuracy with Weighted Delegations

Each agent i with positive weight in the weight distribution $\mathbf{w_D}$ votes with accuracy q_i and weight $w_\mathbf{D}(i)$. We are interested in the accuracy of the group decision when these votes are aggregated by weighted majority. This accuracy is nothing but the probability that a coalition of gurus $C \subseteq Gu(\mathbf{D})$ with majority weight, i.e., such that $\sum_{i \in C} w_\mathbf{D}(i) > \sum_{i \in Gu(\mathbf{D}) \setminus C} w_\mathbf{D}(i)$, contains agents that all vote correctly, while all agents in $Gu(\mathbf{D}) \setminus C$ vote incorrectly. That is:

$$q_\mathbf{D} = \sum_{C \in \mathcal{W}(\mathbf{D})} \prod_{i \in C} q_i \prod_{i \in Gu(\mathbf{D}) \setminus C} (1 - q_i), \tag{4}$$

where $\mathcal{W}(\mathbf{D})$ is the set of *winning coalitions*, i.e., $\mathcal{W}(\mathbf{D}) = \{C \subseteq N \mid \sum_{i \in C} w_\mathbf{D}(i) > \sum_{i \in Gu(\mathbf{D}) \setminus C} w_\mathbf{D}(i)\}$. In case of ties, if $\sum_{i \in C} w_\mathbf{D}(i) = \sum_{i \in Gu(\mathbf{D}) \setminus C} w_\mathbf{D}(i)$, one of the two coalitions at random is added to $\mathcal{W}(\mathbf{D})$. When dealing with pure delegation profiles $\mathbf{d}$ we will also write $q_\mathbf{d}$ for the group accuracy determined by $\mathbf{d}$.

Example 2 (Example 1 cnt'd). *We use Example 1 with accuracy $\boldsymbol{q} = (0.9, 0.9, 0.6, 0.6, 0.6)$. Consider the expected weight approach, i.e., $\mathbf{w_D} = (2, 1, 1, 0, 0)$ ($Gu(\mathbf{D}) = \{1, 2, 3\}$). Then all winning coalitions are $\mathcal{W}(\mathbf{D}) = \{\{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$ (assume agent $\{2, 3\}$ is randomly chosen between the tied coalitions $\{1\}$ and $\{2, 3\}$). Then take an instance of the winning coalition $\{1, 2\}$, such that the accuracy of this winning coalition is computed by assuming agent 1 and 2 vote correctly, i.e., with probabilities $q_1 = 0.9$ and $q_2 = 0.9$, while agent 3 votes incorrectly with probability $1 - q_3 = 0.4$. That is $q_1 q_2 (1 - q_3) = 0.324$. The group accuracy is the sum of all accuracies of these winning coalitions, i.e., $q_\mathbf{D} = 0.324 + 0.054 + 0.054 + 0.486 = 0.918$.*

3 CENTRALIZED WEIGHTED DELEGATIONS

3.1 Optimal Truth-Tracking

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 with a specific choice of weights:

Theorem 1 (Grofman et al. [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 distribution, for any winning coalition $N_1 \subseteq N$ (with $N - N_1 = N_2$), it holds that $\Pi_{i \in N_1} q_i \Pi_{i \in N_2} (1 - q_i) > \Pi_{i \in N_1} (1 - q_i) \Pi_{i \in N_2} q_i$, because taking the logarithm of both sides of the equivalent formula $\Pi_{i \in N_1} \frac{q_i}{1-q_i} > \Pi_{i \in N_2} \frac{q_i}{1-q_i}$, we get $\sum_{i \in N_1} \log\left(\frac{q_i}{1-q_i}\right) > \sum_{i \in N_2} \log\left(\frac{q_i}{1-q_i}\right)$, as desired. That is, a coalition with probability to make a correct decision (N_1) higher than its complement (N_2) always has more weight. □

We can leverage Theorem 1 to solve the optimal delegation problem: given a set of agents with different individual accuracies, what is the weighted delegation graph that maximizes group accuracy? We develop an answer to this question in two steps. First, to fix intuitions, we provide a solution for complete networks (Algorithm 1), and then move to connected networks (Algorithm 2).

3.2 Centralized Delegations in Complete Nets

In complete networks all agents can delegate to all other agents. We propose an algorithm that uses one-hop weighted delegations to reallocate weight from the less accurate to the more accurate voters in the group. We define the optimal weight of each $i \in N$ by:

$$w_i^\star = n \cdot \frac{\log \frac{q_i}{1-q_i}}{\sum_{j \in N} \log \frac{q_j}{1-q_j}}. \tag{5}$$

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 $w_i^\star > 1$ apportion their full weight to themselves, and have each agent j with $w_j^\star < 1$ apportion share $w_i^\star - 1$ of the excess weight $1 - w_j^\star$ to each agent i, normalized by the total excess weight of the i agents. Notice that if all agents are equally accurate ($N = N_3$), Algorithm 1 returns the trivial profile. We use Example 2 to illustrate Algorithm 1.

Example 3 (Example 2 cnt'd). *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 $w_i^\star$ is computed by normalizing the above vector by the entire weight, which is in total 5 votes, and for $i \in N$, $w_i^\star$ are $(1.958, 1.958, 0.3613, 0.3613, 0.3613)$. Hence we can observe that $w_1^\star$ and $w_2^\star$ 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 $w_i^\star = 0.3613$ to agents*

Algorithm 1 Optimal delegations in complete networks

Input: $\mathbf{w}^\star$
Initialize: $N_1 = \{i \in N \mid w_i^\star < 1\}$, $N_2 = \{i \in N \mid w_i^\star > 1\}$, $N_3 = \{i \in N \mid w_i^\star = 1\}$, $w = \sum_{i \in N_2}(w_i^\star - 1)$.
Delegate:
- For $i \in N_3$: $D_{ii} = 1$.
- For $i \in N_2$: $D_{ii} = 1$.
- For $i \in N_1$: for all $j \in N_2$, $D_{ii} = w_i^\star$, $D_{ij} = (1 - w_i^\star)\frac{w_j^\star - 1}{w}$

Return: $\mathbf{D}$

1 and 2 equally since $w_1^\star = w_2^\star$. 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 $i \in N$, $w_{\mathbf{D}}(i) = w_i^\star$.*

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

Proof. Observe first that agents in N_1 have optimal weight below 1, and this requires them to delegate part of their weight to agents in N_2 whose optimal weight is above 1. Then for all agents in N_3, their optimal weight is exactly 1. Therefore they just need to be single gurus to reach their optimal weight, i.e., for all $i \in N_3$, by the algorithm, $D_{ii} = 1$ and $w_{\mathbf{D}}(i) = 1 = w_i^\star$. We then consider N_1 and N_2 in turn. Note that no weight is lost in the returned weighted profile since the only cycles are loops. First consider all agents in N_1. By Algorithm 1, for all $i \in N_1$, $D_{ii} = w_i^\star$, and for all $j \in N \setminus \{i\}$, $D_{ji} = 0$, thus $w_{\mathbf{D}}(i) = D_{ii} = w_i^\star$. For the surplus weight of i, i.e., $1 - w_i^\star$, she delegates a proportion of it to each agent in N_2. The proportion is decided by $\frac{w_j^\star - 1}{w}$ for all $j \in N_2$, that is agent j is expected to receive weight amount of $(1 - w_i^\star)\frac{w_j^\star - 1}{w}$ from all $i \in N_1$. Then any agent $j \in N_2$, in total, is expected to receive $\sum_{i \in N_1}(1 - w_i^\star)\frac{w_j^\star - 1}{w}$. Moreover, $D_{jj} = 1$, which indicates j is expected to delegate all her weight to herself, i.e., amount of 1. Notice that for all agents in N_1 and N_2, $\sum_{k \in N_1 \cup N_2} w_k^\star = |N_1| + |N_2|$, and hence $\sum_{k \in N_1}(1 - w_k^\star) = |N_1| - \sum_{k \in N_1} w_k^\star = \sum_{k \in N_1 \cup N_2} w_k^\star - |N_2| - \sum_{k \in N_1} w_k^\star = \sum_{k \in N_2} w_k^\star - |N_2| = \sum_{k \in N_2}(w_k^\star - 1)$. Therefore j collects $\sum_{k \in N_2}(w_k^\star - 1)\frac{w_j^\star - 1}{w} + 1 = w_j^\star$. □

The next example shows how optimal accuracy via weighted delegations may be higher than that achievable via pure delegations.

Example 4 (Example 3 cnt'd). *Let us continue with Example 3. 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 3.*

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

Remark 1. *It is worth discussing Algorithm 1 in the context of the GreedyCap algorithm of [26]. GreedyCap is a local probabilistic delegation algorithm, with a centralized element: a cap on the maximal number of delegations, which avoids the creation of too much correlation among voters and thus preserving wisdom-of-the-crowd effects. Algorithm 1 implements a fully centralized approach to group accuracy by assuming delegations to be centrally determined.*

Algorithm 2 Optimal delegations in general networks

Input: $\mathbf{w}^*$, $R = \langle N, E \rangle$
Initialize: $A = 0^{n \times n}$, $\forall i \in N_1$, $j \in N$, $D_{i,j} = 0$, $N_1 = \{i \in N \mid w_i^\star < 1\}$, $N_2 = \{i \in N \mid w_i^\star > 1\}$.
Determine Paths: For all $i, j \in N$ ($i \neq j$), select an arbitrary acyclic path L_{ij} in R.
Label Required Transfer of Weight: For all $(k, k') \in L_{ij}$:
- For $i \in N_1$, $j \in N_2$: $e_{k,k'}^{i,j} = (1 - w_i^\star)\frac{w_j^\star - 1}{\sum_{\ell \in N_2}(w_\ell^\star - 1)}$.
- For $i \in N$ and $j \in R'(i)$: $\mathbf{A}_{i,j} = \sum_{k \in N_1, k' \in N_2} e_{i,j}^{k,k'} - \sum_{k \in N_2, k' \in N_1} e_{j,i}^{k,k'}$, if $\sum_{k \in N_1, k' \in N_2} e_{i,j}^{k,k'} - \sum_{k \in N_2, k' \in N_1} e_{j,i}^{k,k'} > 0$.

Remove Cycles: For $c \in C(\mathbf{A})$, and $(k, k') \in c$: $\mathbf{A}_{k,k'} = \mathbf{A}_{k,k'} - \min_{(\ell,\ell') \in c}(\mathbf{A}_{\ell,\ell'})$.
Decide Weights: For $i \in N$ and $j \in R'(i)$:
- If $\mathbf{A}_{i,i'} > 0$: $D_{i,i'} = \mathbf{A}_{i,i'} / (\sum_{\ell \in R'(i), \mathbf{A}_{\ell,i} > 0} \mathbf{A}_{\ell,i} + 1)$;
- $D_{i,i} = (\sum_{\ell \in R'(i), \mathbf{A}_{\ell,i} > 0} \mathbf{A}_{\ell,i} + 1 - \sum_{\ell \in R'(i), \mathbf{A}_{i,\ell} > 0} \mathbf{A}_{i,\ell}) / (\sum_{\ell \in R'(i), \mathbf{A}_{\ell,i} > 0} \mathbf{A}_{\ell,i} + 1)$.

Return: $\mathbf{D}$

3.3 Centralized Delegations in Connected Nets

We extend now Algorithm 1 to the case of connected networks.

We fix some notation before introducing the algorithm. Given a non-negative matrix $\mathbf{A} \in \mathbb{R}_{\geq 0}^{n \times n}$, $G(\mathbf{A}) = \langle N, E(\mathbf{A}) \rangle$, denotes the directed graph of $\mathbf{A}$ as above. Let $C(\mathbf{A})$ denote all cycles in $G(\mathbf{A})$.

Algorithm 2 generalizes the idea of Algorithm 1 to connected networks as follows. Similar to Algorithm 1, each agent, expected to delegate (i.e., in N_1), transfers part of her excess weight to an agent, who expects to receive delegations (i.e., in N_2). Then the **Determine Paths** component first decides an acyclic path between each such pair of agents, and **Label Required Transfer of Weight** component labels the expected transfer weight amount on each edge on the path. Hence for each node, some incoming and outgoing edges are labeled with weight (one edge might have several labels), and we aggregate the net expected transfer amount between the node and each neighbor. Note that cycles may exist to introduce noise, therefore we break every cycle by subtracting the minimum amount among all edge labels in the cycle, by the component **Remove Cycles**. Finally by the **Decide Weights** component, each agent decides her delegation strategy by computing the proportion of each expected outgoing weight in total incoming weight (including her initial weight). An illustration of the algorithm follows.

Example 5. *Consider a network with 4 agents $N = \{1, 2, 3, 4\}$, with accuracies $\boldsymbol{q} = (0.5, 0.9, 0.6, 0.9)$. Let R be as in Figure 2 (top). By Equation* (5) *the optimal weight distribution is* $\mathbf{w}^\star = (0, 1.831, 0.338, 1.831)$, *and therefore $N_1 = \{1, 3\}$ and $N_2 = \{2, 4\}$.*

We first determine the weight transfer paths, and assume that $L_{1,2} = \{(1, 2)\}$, $L_{1,4} = \{(1, 2), (2, 3), (3, 4)\}$, $L_{3,4} = \{(3, 4)\}$, and

$L_{3,2} = \{(3,4),(4,2)\}$. *Then we compute weight transfers. For instance, agent* 1 *should transfer amount* $(1 - w^\star(1)) \frac{w^\star(2)-1}{\sum_{i \in N_2}(w^\star(i)-1)} = 0.5$ *to agent* 2. *Therefore for each path in* $L_{1,2}$, *i.e.,* $(1,2)$, $e^{1,2}_{1,2} = 0.5$. *Similarly we label all required transfer weights and aggregate amounts on each edge, and obtain the labeled graph in Figure 2 (middle). Observe that there is a cycle formed by agents* 2, 3 *and* 4. *Such cycle is broken by subtracting amount* 0.169 *from each edge in the cycle and obtain the acyclic graph in Figure 2 (bottom).*

Finally we compute the weighted profile. Let us take agent 4 *as an instance. Since* 4 *is expected to transfer* 0.162 *to agent* 2, *while simultaneously receive* 0.662 *from* 3, 0.162 *takes* 9.75% *among her received amount* 0.662 *plus her initial weight amount of* 1. *We thus obtain* $D_{42} = 9.75\%$ *and* $D_{44} = 1 - D_{42} = 90.25\%$. *Similarly we have* $D_{12} = 100\%$, $D_{22} = 100\%$, $D_{34} = 66.2\%$, *and* $D_{33} = 33.8\%$.

We can prove that Algorithm 2 achieves optimal accuracy:

Theorem 3. *Algorithm 2 outputs an element of* $\arg\max_{\mathbf{D} \in \mathbb{D}} q_{\mathbf{D}}$.

Proof. First recall that $\mathbf{w}^\star$ is the optimal weight distribution given by Theorem 1. Then notice that in the algorithm, we transfer the amount of $(1 - w^\star_i) \frac{w^\star_j - 1}{\sum_{\ell \in N_2}(w^\star_j - 1)}$ from each $i \in N_1$ to each $j \in N_2$, by the **Label Required Transfer of Weight** component, since we can find a path from any agent in N_1 to any agent in N_2 in the connected network R. These transfers are not cyclical because of **Remove Cycles** which takes care of removing such cycles.

So at this point the algorithm has constructed an acyclic graph encoding the required transfer $\mathbf{A}_{ij}$ of expected weight between any pair of agents i and j. This amount needs to be normalized for each agent by the **Decide Weights** routine. For any agent $i \in N$, if she is required to transfer positive weight $A_{i,i'}$ to a neighbor $i' \in R(i)$, the weighted strategy $\mathbf{D}_{ii'}$ should be the proportion of $\mathbf{A}_{ii'}$ in her total required incoming weight, plus her original endowed weight, i.e., $\sum_{\ell \in R'(i), \mathbf{A}_{\ell,i} > 0} \mathbf{A}_{\ell,i} + 1$. The obtained weighted profile $\mathbf{D}$ thus ensures that for any pair of agents $i \in N_1$ and $j \in N_2$, $\boldsymbol{t}_{\mathbf{D}}(i,j) = (1 - w^\star_i) \frac{w^\star_j - 1}{\sum_{\ell \in N_2}(w^\star_\ell - 1)}$. Therefore by Equation (3), $\boldsymbol{w}_{\mathbf{D}} = \boldsymbol{w}^\star$ as desired. □

We observe that the connectedness of R, which we assume throughout the paper, is necessary for Theorem 3 to go through (an example to this effect is provided in the appendix). Observe also that if R is complete, the path L_{ij} between any pair of agents

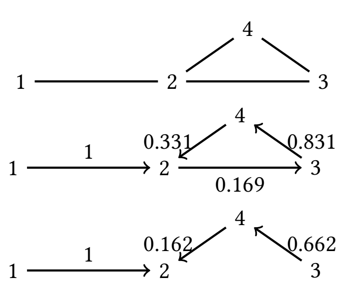

Figure 2: Network underlying Example 5 (top) and depiction of intermediate steps of Algorithm 5 (middle and bottom)

$i \in N_1$ and $j \in N_2$ can be selected to be the one-hop edge (i,j). Algorithm 2 reduces then to Algorithm 1.

4 DECENTRALIZED DELEGATIONS

Algorithms 1 and 2 provide us with tractable (with time complexity polynomial in N) centralized mechanisms to achieve weighted delegations that are truth-tracking optimal. We move now to define a setting in which agents decide their delegations autonomously, assuming they greedily aim at maximizing their own individual accuracy. We are interested in determining—analytically in this section, and empirically later in Section 6—the effects of decentralized delegations on the truth-tracking performance of the group.

4.1 Weighted Delegation Games

Agents' utilities. The utility that an agent i obtains if j acts as her guru is given by $u_i : N \to \mathbb{R}$. We assume here that such utility is given simply by the accuracy that i inherits through delegation, that is: $u_i(\mathbf{d}) = q_{\mathbf{d}^*(i)}$ when $\mathbf{d}^*(i)$ exists, and $u_i(\mathbf{d}) = 0$ otherwise. In other words, given a pure profile $\mathbf{d}$, i's utility is the accuracy of her guru under $\mathbf{d}$ [4]. In weighted profiles each agent may transfer weight to several gurus so the above setting can be extended by assigning to i a utility equal to the average of the accuracies of i's gurus, weighted by the weights that i transfers to those gurus.

Formally, given a weighted profile $\mathbf{D}$ and its associated weight transfer profile $\boldsymbol{t}_{\mathbf{D}}(i)$ for agent i (recall Equation (2)), i's utility is given by:

$$U_i(\mathbf{D}) = \sum_{j \in N} q_j \boldsymbol{t}_{\mathbf{D}}(i,j). \tag{6}$$

Observe that vector $\boldsymbol{t}_{\mathbf{D}}(i)$ can be interpreted as a probability distribution over i's gurus when none of i's weight is lost due to cycles. Equation (6) then gives us the expected individual accuracy of i in D or, in other words, the expectation $\mathbb{E}(u_i)$ over u_i given $\mathbf{D}$.

Delegation Games. Equipped with the notion of utility we move to define delegation games as structures $G = \langle N, R, S, U \rangle$, where $N = \{1, 2, \ldots, n\}$ is the set of agents, R is an undirected connected graph, $S = \{s_i\}_{i \in N}$ is the strategy space of each agent $i \in N$, where $s_i = \mathbf{D}_i$ such that $\sum_{j \in R(i)} D_{ij} = 1$, and U is the function defined by Equation (6). Observe that, since U_i equals the expectation over u_i given the distribution over pure profiles induced by a weighted 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 [33]) we therefore know that such games always have Nash Equilibria (NE), and we call such equilibria U-NE. We will also write $\mathcal{E}(G)$ to denote the set of all U-NE of G. It has already been shown that the pure delegation variant of these games (where $\mathbf{D}_i$ is a degenerate probability vector) also always admits a NE [4, Th. 1].

An important feature of U-NE is that they contain only weighted delegation cycles of length 1 (loops):

Lemma 1. *In a U-NE the only weighted delegation cycles are loops.*

Proof sketch. In any cycle (excluding loops), any agent obtains utility through all caught-in-cycle agents' delegations out of the cycle, while some weight is still lost because of the cycle. Then the agent, who can obtain maximal utility by her delegations out of the cycle, can always be strictly better off by breaking the cycle. □

4.2 Group Accuracy in Equilibrium

So what is the truth-tracking quality of equilibria in weighted delegation games? We answer this question by first comparing group accuracy $q_{\mathbf{D}}$ when $\mathbf{D}$ is a U-NE, w.r.t. when $\mathbf{D}$ is a pure delegation equilibrium. We then establish a bound on how bad group accuracy may be in equilibrium using price of anarchy.

Weighted vs. pure equilibria. Weighted delegations make it possible to achieve higher group accuracy in equilibrium by balancing weight among maximally accurate agents. As a result NE in delegation games with weighted profiles can be shown to be never worse that NE with pure delegations and to be better in some cases. Given a delegation game on N, let us denote with $N^* = \{i \in N \mid \forall j \in N, q_i \geq q_j\}$ the set of maximally accurate agents in N.

Theorem 4. *Any profile $\mathbf{D}^*$ of a game G, such that for all $i \in N^*$ $\mathbf{w}_{\mathbf{D}^*}(i) = \frac{n}{|N^*|}$, is a U-NE and $\mathbf{D}^* \in \arg\max_{\mathbf{D} \in \mathcal{E}(G)} q_{\mathbf{D}}$.*

Proof. First we show that if a weighted profile $\mathbf{D}$ is a U-NE, for all $i \in N \setminus N^*$, $w_{\mathbf{D}}(i) = 0$, that is the entire weight is concentrated in N^*. Notice that no cycle exists in $\mathbf{D}$ by Lemma 1. Assume the above condition does not hold. Since the network is connected, each agent has access to any agents in N^*. There must exists some agent $j \in N \setminus N^*$, such that $\mathbf{D}_{jj} > 0$, and she can change her strategy to have gurus only in N^* to be better off. Therefore for any NE $\mathbf{D}$, $\sum_{i \in N^*} w_{\mathbf{D}}(i) = n$. Then by Theorem 1, the NE $\mathbf{D}$, such that for all $i \in N^*$, $w_{\mathbf{D}}(i) = n/|N^*|$, optimizes group accuracy for N^*. Furthermore, by the Condorcet jury theorem, which states that larger group of agents (with homogeneous accuracy higher than 0.5) enhances group accuracy (see [36, Th. 1]), no NE $\mathbf{D}$ in which $Gu(\mathbf{D})$ is a strict subset of N^* has higher group accuracy. □

Then, the following example shows that there exist delegation games in which $\mathbf{D}^*$ has strictly better group accuracy than any equilibrium in pure delegation strategies.

Example 6. *Consider a delegation game where there are 7 agents, 5 of which have maximal accuracy $q^* = 0.9$. Any pure delegation 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., $0.00081 = {q^*}^2(1-q^*)^3 < {q^*}^3(1-q^*)^2 = 0.00729$. So the resulting group accuracy, 0.98496, is strictly worse than that of $\mathbf{D}^*$, $q_{\mathbf{D}^*} = 0.99144$ in Theorem 4.*

Price of Anarchy. We define the price of anarchy of a game G as:

$$\text{PoA}(G) = \frac{\max_{\mathbf{D} \in \mathbb{D}} q_{\mathbf{D}}}{\min_{\mathbf{D} \in \mathcal{E}(G)} q_{\mathbf{D}}}. \quad (7)$$

When restricting to the price of anarchy in games with pure delegations (and therefore one pure delegation NE) we refer to PoA^{pure}. So Equation (7) gives us a measure of how much group accuracy is 'lost' in equilibrium, in the worst case, with respect to what would be achievable via Algorithm 2.

Theorem 5. *When $|N| \to \infty$, $\text{PoA} \to \frac{1}{q^*}$, 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 $\mathbf{D}'$ be the U-NE profile for which $q_{\mathbf{D}'}$ is minimal. $\mathbf{D}'$ is the case when all agents delegate to the same guru, which has accuracy q^*. Then, since each $q_i \in (0.5, 1]$, by the law of large numbers as $|N| \to \infty$, $q_{\mathbf{D}} \to 1$ and by construction $q_{\mathbf{D}'} = q^*$. □

The same argument can be applied to the setting with pure delegations, obtaining the same asymptotic value for PoA^{pure}.

In the non-asymptotic case, since weighted delegations enable optimal group accuracy (Theorem 2) while pure delegations do not (Example 4), the PoA in delegation games with weighted delegations is trivially higher than that in the case of pure delegations:

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

5 WEIGHTS AS SHARES OF POWER

We have so-far developed our theory based on the expected weight approach of Equations (2) and (3). This is not the only way in which agents' weights can be interpreted under weighted delegations. In this section we briefly highlight another interpretation, which we call *limit weight approach*, and relate it to the previous one.

5.1 The Limit Weight Approach

Each weighted profile $\mathbf{D}$ describes the direct transfer of voting weight between any two agents: D_{ij} is the share of i's power transfered to j. The indirect transfer of weight, via transitive delegations, is described therefore by the powers of $\mathbf{D}$. For example, $\mathbf{D}^2_{ij} = \sum_{k \in N} D_{ik} D_{kj}$ is the share of power transfered in two steps from i to j. In this view the weight transfer of agent i consists of the transfer of i's weight in the limit, described by the vector

$$\mathring{\boldsymbol{t}}_{\mathbf{D}}(i) = \lim_{k \to \infty} \mathbf{1}_i \mathbf{D}^k \quad (8)$$

when such limit exists,[1] and where $\mathbf{1}_i$ is the n-dimensional vector where all elements are 0s except for the i-th one which is 1. This approach describes to whom i's original weight of 1 'flows' in $\mathbf{D}$, and an agent's weight is then (cf. Equations (2) and (3)):

$$\mathring{w}_{\mathbf{D}}(i) = \sum_{j \in N} \mathring{\boldsymbol{t}}_{\mathbf{D}}(j)_i, \quad (9)$$

defining the weight distribution $\mathring{\mathbf{w}}_{\mathbf{D}} = (\mathring{w}_{\mathbf{D}}(1), \ldots, \mathring{w}_{\mathbf{D}}(n))$. We refer to this as the *limit weight approach.*

The two approaches are compared in the following two examples:

Example 7. *As in Example 1, consider $N = \{1, 2, 3, 4, 5\}$, with weighted profile $\mathbf{D}$, such that $D_{11} = 1$ (i.e., maintaining her full voting weight), $\mathbf{D}_2 = (\ldots, D_{22} = 0.5, \ldots, D_{25} = 0.5)$, $\mathbf{D}_3 = (D_{31} = 0.5, \ldots, D_{33} = 0.5, \ldots)$, $D_{43} = 1$, and $D_{52} = 1$ (recall Figure 1). By the limit weight approach, for the component consisting of 1, 3 and 4, since agent 1 is the only absorbing node in the component (in the delegation graph), all weight in that component 'flows' to her. Hence $\mathring{w}_{\mathbf{D}}(1) = 3$, while $\mathring{w}_{\mathbf{D}}(3) = \mathring{w}_{\mathbf{D}}(4) = 0$. Then for the component consisting of 2 and 5, we observe that it is strongly connected and aperiodic because of the loop at 2. Then in the limit, agent 2 always keeps half of the stabilized weight and delegates half to agent 5, while agent 5 always delegates the amount back to 2. Therefore $\mathring{w}_{\mathbf{D}}(2) = 2 \times \frac{2}{3} = 4/3$*

[1] It is worth remarking that the inexistence of $\lim_{k \to \infty} \mathbf{D}^k$ does not necessarily imply the inexistence of $\lim_{k \to \infty} \mathbf{1}_i \mathbf{D}^k$.

and $\mathring{w}_{\mathbf{D}}(5) = 2/3$. *So* $\mathring{\mathbf{w}}_{\mathbf{D}} = (3, 4/3, 0, 0, 2/3)$. *Unlike in the expected weight case (Example 1) there is therefore no weight loss in* $\mathring{\mathbf{w}}_{\mathbf{D}}$.

Example 8 (Example 1 cnt'd). *We use Example 1 with accuracy* $\boldsymbol{q} = (0.9, 0.9, 0.6, 0.6, 0.6)$ *and, unlike in Example 2 where we applied the expected weight approach, we apply the limit weight approach. The weight distribution is* $\mathring{w}_{\mathbf{D}} = (3, 4/3, 0, 0, 2/3)$ *and the group accuracy is then* $q_{\mathbf{D}} = 0.9$ *since agent* 1 *is the dictator (agent* 1 *alone is a winning coalition), and is therefore lower than that in the expected weight case.*

Remark 2. *The limit weight approach is related to the so-called influence matrices studied in the literature on the DeGroot model [14, 25] and on power in organizations [18, 23, 24]. Both these strands of literature define limit influence notions as done in Equation* (8).

5.2 Expected vs. Limit Weight

Sufficient conditions for equivalence. The two notions of weight coincide under specific conditions on the delegation graph:

Theorem 6. *Let* $\mathbf{D}$ *be a weighted delegation profile. If all cycles in* $\mathbf{D}$ *are loops of weight* 1*, then for all* $i \in N$: $\boldsymbol{t}_{\mathbf{D}}(i) = \mathring{\boldsymbol{t}}_{\mathbf{D}}(i)$.

Proof. First of all note that since all cycles contained in the delegation graph are loops, each closed strongly connected component of the delegation graph is aperiodic.[2] We can then apply known theorems about the convergence of stochastic matrices (e.g., [35, Theorem 12], [25, Theorem 8.1]) to conclude that for all $i \in N$, $\lim_{k\to\infty} \mathbf{1}^i \mathbf{D}^k$ exists. Then for any agent $i_1 \in N$, part of her weight is accrued by some agents delegating their full weight to themselves (loops of weight 1). Let one such agent be $i_\ell \in N$. Assume a path from i_1 to i_ℓ is $i_1 \to i_2 \to \cdots \to i_\ell$. Then the amount of weight flowing from i_1 to i_ℓ through this path is $\prod_{1\leq j<\ell-1} D_{j j+1}$, and under the expected weight approach, $\boldsymbol{t}_{\mathbf{D}}(i_1, i_\ell)$ is the sum of weights flowing through all paths from i_1 to i_ℓ, which is identical to the i_ℓ's element of $\lim_{k\to\infty} \mathbf{1}^{i_1} \mathbf{D}^k$. We conclude that $\boldsymbol{t}_{\mathbf{D}}(i_1) = \mathring{\boldsymbol{t}}_{\mathbf{D}}(i_1)$. □

That is, if delegations are such that all gurus delegate their full weights to themselves, and no other cycle exists, then Equations (3) and (9) define the same values. A direct corollary of the result is that expected and limit weight coincide in pure delegation profiles.

Equilibria in expected and limit weight. We can then use Theorem 6 to establish that NE in expected weight are also equilibria in limit weight. Recall that we refer to the former as U-NE. We will refer then to equilibria under the limit weight approach as $\mathring{U}$-NE, where the utility function is defined as (cf. Equation (6)):

$$\mathring{U}_i(\mathbf{D}) = \sum_{j\in N} q_j \mathring{\boldsymbol{t}}_{\mathbf{D}}(i)_j. \tag{10}$$

Theorem 7. *If a weighted profile* $\mathbf{D}$ *is a* U*-NE of a delegation game* G*, it is also a* $\mathring{U}$*-NE of* G *where* U *is replaced by* $\mathring{U}$.

Proof sketch. The only cycles are loops in $\mathbf{D}$ (Lemma 1). If all loops are of weight 1, $\mathbf{D}$ is a $\mathring{U}$-NE (Theorem 6). If not all loops are with weight 1, we can still show that no agent has a $\mathring{U}$-better response because all gurus are maximally accurate under $\mathring{w}$. □

[2] We recall that a component of a directed graph is a subgraph, which is: strongly connected if there exists a path from every node to every node; it is aperiodic if there are no two cycles in the graph whose length is divided by an integer larger than 1; it is closed if there exists no edge from the component to a node outside it.

The other direction of the theorem does not hold, however:

Example 9. *Consider* $N = \{1, 2\}$ *with* $\boldsymbol{q} = (0.9, 0.9)$ *and they are connected in the underlying network. Profile* $D_{11} = 0.5$, $D_{12} = 0.5$, *and* $D_{21} = 1$ *is a* $\mathring{U}$*-NE. Since the delegation graph is aperiodic, under the limit weight approach, each agent still obtains maximal utility of* 0.9*. However, the profile contains a cycle of weight* 0.5 *and cannot therefore be a* U*-NE by Lemma 1.*

6 EXPERIMENTS

To gain further insights into the effects of decentralized weighted delegations on group accuracy, under both the expected and limit weight approaches, we proceed with a set of experiments.

6.1 Experimental Setting

Agents are constrained in their delegations by a random network which will be treated as a parameter. Agents try to maximize their own utility as per Equation (6). However, they are assumed to be boundedly rational and achieve this maximization only imperfectly. To this aim we model agents' strategic behavior with the so-called quantal response model [29], which has already been applied successfully to other strategic contexts in social choice (e.g., see [30]).

Logit quantal response. The quantal response model assumes that agents choose their strategies with noise. The probability (belief distribution) of choosing a pure delegation is positively related to the utility of that delegation, and agents respond to the others' strategies assuming that all agents have the same belief distribution, until an equilibrium is reached.

More precisely, we assume a special case of the quantal response model, known as *logit quantal response* (LQR). Based on a given weighted profile $\mathbf{D}$, an agent $i \in N$ responds in that profile by changing her individual weighted strategy to $\mathbf{D}' = (\mathbf{D}'_i, \mathbf{D}_{-i})$, such that for any neighbor $j \in R(i)$,

$$D'_{ij} = \frac{e^{\lambda U_i(D_{ij}=1,\mathbf{D}_{-i})}}{\sum_{k\in R(i)} e^{\lambda U_i(D_{ik}=1,\mathbf{D}_{-i})}}, \tag{11}$$

where λ is a parameter indicating the error level agents are subject to. So $\lambda = 0$ corresponds to a uniformly random choice, while as $\lambda \to \infty$ the choices approach optimality. Note that we replace $U_i(\mathbf{D})$ in Equation (11) by $\mathring{U}_i(\mathbf{D})$ for the limit weight approach.

We implement an iterated LQR model starting with the trivial profile $\mathbf{D}^0$ where for all $i \in N$, $D^0_{ii} = 1$. Agents then apply LQR iteratively according to a fixed round robin sequence. By way of illustration, let agent 1 be the first agent responding to $\mathbf{D}^0$. Her response would be $\mathbf{D}^1 = (\mathbf{D}^1_1, \mathbf{D}^0_{-1})$, such that for all $j \in R(1)$, $D^1_{1j} = \frac{e^{\lambda q_j}}{\sum_{k\in R(1)} e^{\lambda q_k}}$, since for all $k \in R(1)$, $U_1(D_{1k} = 1, \mathbf{D}^0_{-1}) = q_k$. Then agent 2 responds to $\mathbf{D}^1$ by LQR, and so and on until no agent changes her strategy any more, reaching a so-called LQR equilibrium. By [29, Th. 2], we know that as $\lambda \to \infty$ this LQR equilibrium converges to one of the Nash equilibria of the delegation game.

Parameters. In the simulations we manipulate the following parameters. We consider the two weight models: the expected weight model underpinning utility U_i of Equation (6), and the limit weight model underpinning utility $\mathring{U}_i$ of Equation (10). We then consider different error levels in the LQR model by varying

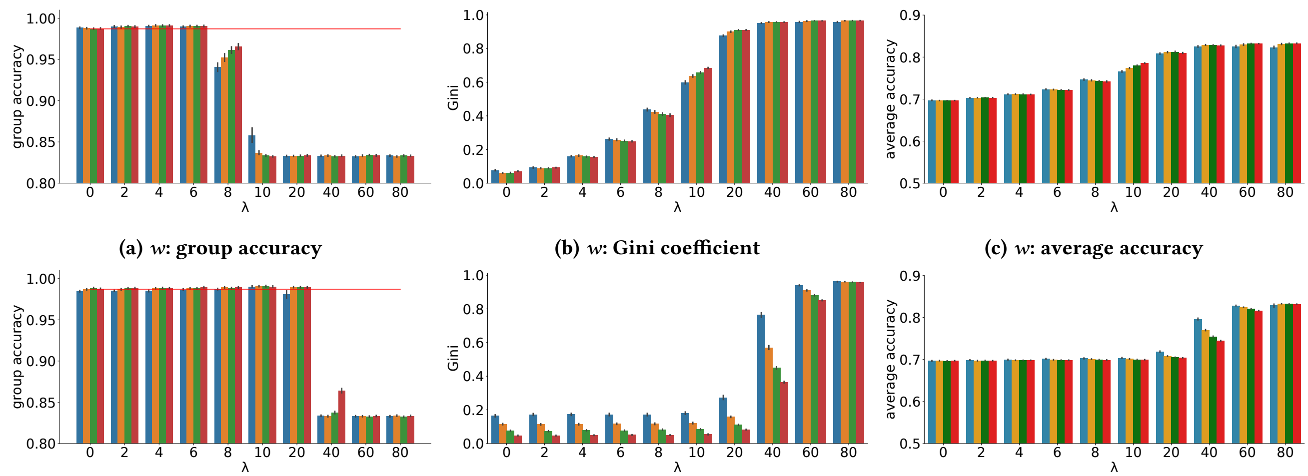

(a) w**: group accuracy** **(b)** w**: Gini coefficient** **(c)** w**: average accuracy**

(d) $\mathring{w}$**: group accuracy** **(e)** $\mathring{w}$**: Gini coefficient** **(f)** $\mathring{w}$**: average accuracy**

Figure 3: Estimated group accuracy, average accuracy, and Gini coefficient for iterated LQR in expected weight games w (top) and limit weight games $\mathring{w}$ (bottom) on random networks. The red line on the leftmost plots denotes the estimated accuracy of simple majority. Underlying network density p: ■ : 0.8, ■ : 0.6, ■ : 0.4, ■ : 0.2

$\lambda \in \{0, 2, 4, 6, 8, 10, 20, 40, 60, 80\}$. Finally, we consider different levels of density in the underlying network varying the probability of any two nodes being connected with $p \in \{0.2, 0.4, 0.6, 0.8\}$.

Criteria. We study the effects of the above parameters on three properties of weighted delegation profiles: group accuracy, average accuracy, and Gini coefficient. Group accuracy is defined in Equation (4). The average accuracy is simply the weighted mean of all gurus' accuracies, based on a given weight distribution. Finally, the Gini coefficient measures the equality of the weight distribution: the higher the index is, the more unequal the distribution is.

Setup. We set $n = 30$. Agents' accuracies are independently drawn from the same Gaussian distribution ($\mu = 0.7$, $\sigma = 0.075$) and values are forced within the $[0.5, 1]$ range. For each parameter configuration we perform 50 runs to obtain our data. As group accuracy involves exponential-time computations, we estimate it via a Monte Carlo approximation sampling $2^{n-1}/100$, i.e., 5368709 times, random coalitions for each computation. The experiments have been run on a university cpu cluster with 1GB memory.

6.2 Findings

We highlight two findings. *First* (leftmost plots), group accuracy q_{D} remains high for noisy response strategies and decreases sharply as λ grows above a certain threshold (8 in the expected and 20 in limit weight setting). This is due to the fact that as λ grows, LQR approximates a pure delegation response, which is known to lead to lower group accuracy [4, 38]. Importantly, however, if λ takes specific values (in $\{2, 4, 6\}$ for w, or 10 for $\mathring{w}$) the experiments show that weighted delegations achieve group accuracy that outperforms the accuracy of a simple majority (i.e., the standard one-man-one-vote wisdom of the crowd). The difference between the two has been tested as significant (threshold set to 0.05) via t-test between 50 samples of the approximated group accuracy for each parameter combination (all p's and $\lambda \in \{2, 4, 6\}$, resp. $\{10\}$, under w, resp. $\mathring{w}$) and 50 samples of the approximated group accuracy for simple majority. We consider this finding particularly interesting in the context of current literature (e.g., [8, 26]) as it points to the possibility of decentralized delegation schemes that support, rather than hinder, wisdom-of-the-crowd effects in liquid democracy.

Second, the trends on group accuracy are matched by the Gini coefficient (middle plots), which grows from near 0 values (equal distribution of power) to near 1 values (full inequality) as λ grows. Intuitively, as quantal responses tend towards pure delegations, the group is able to identify the most accurate agents, who then accrue all weight. Also, identifying the most accurate agents becomes easier as the network grows denser. In line with these trends, the average accuracy (rightmost plots) grows from $\mu = 0.7$ to the accuracy of the most accurate agents (~ 0.83) as λ grows.

7 CONCLUSIONS

We studied a variant of liquid democracy with weighted proxies. Interpreting weights as probabilities we showed that centralized delegations enable optimal group accuracy, and that decentralized delegations may enable better equilibria than in the pure delegation case. We complemented these findings with experimental observations showing how weighted delegations may boost group accuracy also in decentralized settings with boundedly rational agents.

The work presented relies on a connectedness assumption on the underlying network, which we aim at lifting in future work. Although we provided some results about the limit weight approach to weighted delegations, much more has to be understood about that setting. It would also be interesting to investigate experimentally the effects of different network classes on group accuracy, expanding the scope of our simulations. Finally, analytical results about quantal response and group accuracy (e.g., PoA w.r.t. quantal response equilibria) appear worth pursuing.

REFERENCES

[1] Ben Abramowitz and Nicholas Mattei. 2019. 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*. 3–10.

[2] Dan Alger. 2006. Voting by proxy. *Public Choice* 126, 1-2 (jan 2006), 1–26. https://doi.org/10.1007/s11127-006-3059-1

[3] Jan Behrens, Axel Kistner, Andreas Nitsche, and Björn Swierczek. 2014. *Principles of Liquid Feedback.* Interaktieve Demokratie.

[4] Daan Bloembergen, Davide Grossi, and Martin Lackner. 2019. On rational delegations in liquid democracy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1796–1803.

[5] Christian Blum and Christina Isabel Zuber. 2016. Liquid democracy: Potentials, problems, and perspectives. *Journal of Political Philosophy* 24, 2 (2016), 162–182.

[6] Paolo Boldi, Francesco Bonchi, Carlos Castillo, and Sebastiano Vigna. 2011. Viscous democracy for social networks. *Commun. ACM* 54, 6 (2011), 129–137.

[7] Markus Brill and Nimrod Talmon. 2018. Pairwise Liquid Democracy.. In *Proceedings of the International Joint Conference on Artificial Intelligence.* 137–143.

[8] Ioannis Caragiannis and Evi Micha. 2019. 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.* 116–122.

[9] Zoé Christoff and Davide Grossi. 2017. 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)*, Vol. 251. EPTCS, 134–150.

[10] Gal Cohensius, Shie Mannor, Reshef Meir, Eli Meirom, and Ariel Orda. 2017. Proxy Voting for Better Outcomes. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems.* International Foundation for Autonomous Agents and Multiagent Systems, 858–866.

[11] Rachael Colley, Umberto Grandi, and Arianna Novaro. 2020. Smart Voting. In *Proceeding of the the 29th International Joint Conference on Artificial Intelligence (IJCAI), 2020.*

[12] Marquis de Condorcet, M.J.A.N. de C. 1785. *Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix.* Imprimerie Royale, Paris.

[13] Jonas Degrave. 2014. Resolving multi-proxy transitive vote delegation. *CoRR* abs/1412.4039 (2014). arXiv:1412.4039 http://arxiv.org/abs/1412.4039

[14] Morris H. Degroot. 1974. Reaching a Consensus. *J. Amer. Statist. Assoc.* 69, 345 (mar 1974), 118–121. https://doi.org/10.1080/01621459.1974.10480137

[15] Palash Dey, Arnab Maiti, and Amatya Sharma. 2021. On Parameterized Complexity of Liquid Democracy. In *Algorithms and Discrete Applied Mathematics - 7th International Conference, CALDAM 2021, Rupnagar, India, February 11-13, 2021, Proceedings (Lecture Notes in Computer Science, Vol. 12601)*, Apurva Mudgal and C. R. Subramanian (Eds.). Springer, 83–94.

[16] Edith Elkind and Arkadii Slinko. 2016. Rationalizations of Voting Rules. In *Handbook of Computational Social Choice.* Cambridge University Press, Chapter 8, 169–196.

[17] Bruno Escoffier, Hugo Gilbert, and Adèle Pass-Lanneau. 2020. Iterative Delegations in Liquid Democracy with Restricted Preferences. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI'20).* 1926–1933.

[18] J. French. 1956. A Formal Theory of Social Power. *Psychological Review* 61 (1956), 181–194.

[19] Paul Gölz, Anson Kahng, Simon Mackenzie, and Ariel D Procaccia. 2018. The Fluid Mechanics of Liquid Democracy. In *Proceedings of the 14th Conference on Web and Internet Economics (WINE'18).* 188–202.

[20] James Green-Armytage. 2015. Direct voting and proxy voting. *Constitutional Political Economy* 26, 2 (2015), 190–220.

[21] Bernard Grofman, Guillermo Owen, and Scott L. Feld. 1983. Thirteen theorems in search of the truth. *Theory and Decision* 15, 3 (1983), 261–278.

[22] Jacqueline Harding. 2021. Proxy selection in transitive proxy voting. *Social Choice and Welfare* (2021). https://doi.org/10.1007/s00355-021-01345-8

[23] X. Hu and L. Shapley. 2003. On Authority Distributions in Organizations: Controls. *Games and Economic Behavior* 45 (2003), 153–170.

[24] X. Hu and L. Shapley. 2003. On Authority Distributions in Organizations: Equilibrium. *Games and Economic Behavior* 45 (2003), 132–152.

[25] M. O. Jackson. 2008. *Social and Economic Networks.* Princeton University Press.

[26] Anson Kahng, Simon Mackenzie, and Ariel Procaccia. 2018. Liquid Democracy: An Algorithmic Perspective. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI'18).*

[27] Christoph Kling, Jérome 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 *Proceedings of the International Conference on Weblogs and Social Media.*

[28] Grammateia Kotsialou and Luke Riley. 2020. Incentivising Participation in Liquid Democracy with Breadth-First Delegation. In *Proceedings of the International Foundation for Autonomous Agents and Multiagent Systems (AAMAS '20)* (Auckland, New Zealand). IFAAMAS, Richland, SC, 638—644.

[29] R. McKelvey and T. Palfrey. 1995. Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior* 10 (1995), 6–38.

[30] R. Meir. 2018. . Morgan & Claypool.

[31] James C Miller. 1969. A program for direct and proxy voting in the legislative process. *Public choice* 7, 1 (1969), 107–113.

[32] Shmuel Nitzan and Jacob Paroush. 1982. Optimal Decision Rules in Uncertain Dichotomous Choice Situations. *International Economic Review* 23, 2 (1982), 289–297.

[33] Martin J. Osborne and Ariel Rubinstein. 1994. *A Course in Game Theory.* MIT Press.

[34] Alois Paulin. 2020. An Overview of Ten Years of Liquid Democracy Research. In *Proceedings of the 21st International Conference on Digital Government Research.*

[35] Anton V. Proskurnikov and Roberto Tempo. 2017. A Tutorial on Modeling and Analysis of Dynamic Social Networks. Part I. *Annual Reviews in Control* 43 (2017), 65–79.

[36] Lloyd Shapley and Bernard Grofman. 1984. Optimizing group judgmental accuracy in the presence of interdependencies. *Public Choice* 43, 3 (1984), 329–343.

[37] Gordon Tullock. 1992. Computerizing politics. *Mathematical and Computer Modelling* 16, 8-9 (aug 1992), 59–65.

[38] Yuzhe Zhang and Davide Grossi. 2021. Power in Liquid Democracy. In *Proceedings of the AAAI conference on Artificial Intelligence.*