A Contribution to the Critique of Liquid Democracy

Ioannis Caragiannis[1] and **Evi Micha**[2]
[1]University of Patras, Greece
[2]University of Toronto, Canada
caragian@ceid.upatras.gr, emicha@cs.toronto.edu

Abstract

Liquid democracy, which combines features of direct and representative democracy has been proposed as a modern practice for collective decision making. Its advocates support that by allowing voters to delegate their vote to more informed voters can result in better decisions. In an attempt to evaluate the validity of such claims, we study liquid democracy as a means to discover an underlying ground truth. We revisit a recent model by Kahng *et al.* [2018] and conclude with three negative results, criticizing an important assumption of their modeling, as well as liquid democracy more generally. In particular, we first identify cases where natural local mechanisms are much worse than either direct voting or the other extreme of full delegation to a common dictator. We then show that delegating to less informed voters may considerably increase the chance of discovering the ground truth. Finally, we show that deciding delegations that maximize the probability to find the ground truth is a computationally hard problem.

1 Introduction

Liquid democracy is a recent trend that aims to modernize the way we act as citizens. Together with other interesting proposals that lie under the umbrella of interactive democracy (e.g., see [Brill, 2018] and the references therein), it is the response to recent disappointing findings by political scientists that democracy on the Internet era uses tools from previous centuries [Mancini, 2015]. The traditional standard of electing representatives to vote on behalf of the citizens on any possible issue seems so obsolete today, while the seemingly ideal alternative of direct voting is rather impractical [Green-Armytage, 2015]. Existing systems like LiquidFeedback [Behrens *et al.*, 2014], used by the Pirate Parties in Germany and elsewhere for internal decision making, support the concept of liquid democracy by combining traditional approaches to voting in novel ways.

The main idea of liquid democracy is to allow citizens to be involved actively in everyday decision making within a society, borrowing the main feature of direct democracy. Following the most fundamental principle of democracy, every citizen has the right to vote for every given issue. However, there might be issues for which a citizen does not feel comfortable to vote; here, liquid democracy exploits the main advantage of *representative democracy*. A citizen may delegate her vote to another citizen who is believed to be more informed about the given issue. A citizen may collect many delegations and can either vote on their behalf or transfer all these delegations together with her right to vote to another citizen, and so on. A vote has a weight indicating the total number of voters it represents. Even though the setting does not constrain the kind of voting rules that can be used to decide the election outcome, *weighted majority* is the usual practice.

Liquid democracy aspires to strengthen a key advantage of *direct democracy*. In a vote for a given issue, it involves as voters all those members of the society that are best-informed for the particular issue. In addition, it provides the mechanism that will allow the vote of these members to receive support by less-informed or indifferent citizens. Ideally, the final vote will be only among well-informed individuals, and their opinions will be weighted by the will of the citizens who preferred to delegate their vote instead of actively vote.

The benefits seem to be higher when there is a *ground truth* (i.e., a *correct* answer) and voting aims to discover it. In this case, the best-informed members of the society will most probably cast the correct answer to the issue at hand (or something close to it) as their vote. Then, via liquid democracy, the society will delegate the quest for the correct decision to the right group of citizens. It sounds like a safe bet for optimal decision making. Is it really so?

To study how effective liquid democracy is in discovering the truth, we consider a model that was first proposed and studied by Kahng *et al.* [2018]. We consider elections with two alternatives: one correct alternative T and an incorrect one F. Every voter has an associated quantity p_i indicating how well-informed the voter is. The voters are nodes in a *social network*, modeled with a directed graph. A directed edge indicates a voter (the origin) who follows another one (the destination of the edge). Each voter can decide to vote directly or delegate this right to another voter. If she chooses to vote, she casts a vote for T with probability p_i or a vote for F with probability $1 - p_i$. If she decides to delegate, she transfers her right to vote together with all delegations she has received by other voters to one of the voters she follows. The outcome of the election is decided using a weighted majority.

Kahng *et al.* [2018] introduced the concept of a *delegation mechanism* to describe the collective behavior of all voters by defining an action for every voter. For a voter, an action can be either to cast her vote directly or to delegate this right to one of the voters she follows (who is also defined by the delegation mechanism). A crucial (and reasonable, at first glance) assumption made by Kahng *et al.* [2018] is that delegations are from less to more informed voters. Under (a refinement of) this assumption, we focus on *local* delegation mechanisms (in Section 3), where the action of each voter depends on her neighborhood only. We show that such mechanisms can be very inefficient; our proof involves profiles where it is impossible for a local mechanism to decide whether almost all nodes should simultaneously delegate to a common dictator or they should use direct voting instead.

Next, somewhat surprisingly, we show that delegating only to better-informed voters may harm the probability of discovering the ground truth significantly (see Section 4). This does not only question the particular modeling assumption of Kahng *et al.* [2018] but, more importantly, one of the central arguments the advocates of liquid democracy support. We remark that our proofs in Sections 3 and 4 do not use transitive delegations and hence apply to paradigms that are simpler than liquid democracy such as *proxy voting* [Green-Armytage, 2015].

Finally, in Section 5, we consider non-local mechanisms that decide delegations by exploiting the global structure of the social network. We show that the *optimal delegation problem* (ODP) of maximizing the probability to find the ground truth by coordinating delegations is not only computationally hard but also hard to approximate with a substantial additive constant. The criticism by this complexity result should be clear: if it is hard to discover the ground truth in a coordinated way, why should the voters be expected to find it by acting independently?

A discussion of related work follows. Preliminary definitions are given in Section 2. Our technical results are presented in Sections 3, 4, and 5. We conclude in Section 6.

1.1 Related Work

The first mentions of combining direct and representative democracy for electing members of a house of representatives date back to the 19th century (e.g., in [Dodgson, 1884]). However, proposals to use delegative voting in practical ways are more recent [Tullock, 1967; Miller, 1969] and coincide with a vision for progress in computers. Behrens [2017] discusses the origins of liquid democracy in some detail. Undoubtedly, it was the dramatic growth of the Internet, the broad spread of smartphones and portable devices, and the popularity of social networking that has made liquid democracy look so attractive today.

Early studies on liquid democracy, as we know it today and have defined above, include the papers by Ford [2002], Green-Armytage [2005], Yamakawa *et al.* [2007], Allen [2008], and Boldi *et al.* [2011]. Very recently, Green-Armytage [2015], Christoff and Grossi [2017], Zhang and Zhou [2017], Brill and Talmon [2018], and Bloembergen *et al.* [2019] consider models to address foundational aspects from logical, game-theoretic, and social choice-theoretic viewpoints. Among papers studying abstract models of liquid democracy, the paper by Kahng *et al.* [2018] is the closest to ours and is discussed in more detail later.

In spite of the attention it enjoys recently by academics and political parties, liquid democracy has received some criticism [Blum and Zuber, 2016]. Among other complaints, situations where some voters attract so many delegations that essentially act as dictators have been observed in practice [Kling *et al.*, 2015]. Gölz *et al.* [2018] study this problem from a computational point of view and propose fractional delegations as a solution. On the other hand, we remark that dictatorships can be very effective in discovering the ground truth. So, our criticism of liquid democracy is of a much different flavor.

Viewing votes as noisy estimates of ground truth and then studying how voting can be used to discover this ground truth is a key approach in computational social choice [Brandt *et al.*, 2016]. With origins from Marquis de Condorcet [1785], the founding father of social choice theory, and further refined by Young [1988], it has received much attention recently, with the papers by [Conitzer and Sandholm, 2005] and [Caragiannis *et al.*, 2016] being indicative examples. See Elkind and Slinko [2016] for an introduction to the approach and a survey of related results.

2 Preliminaries

We consider elections with two alternatives: a correct alternative T (the ground truth) and an incorrect one F. There is a set $\mathcal{N}$ of n voters (or *agents*), who are members of a social network. The social network is represented by a directed graph $G(\mathcal{N}, E)$ having the agents as nodes. A directed edge $(i, j) \in E$ from agent i to another agent j implies that agent i "follows" agent j in the social network. We call the set of agents that agent i follows the *neighborhood* of agent i.

An agent can select between the following two actions. She can either cast her vote or delegate her right to vote to another agent in her neighborhood. Delegations are transitive. An agent who has received delegations from some of her followers can delegate her right to vote together with all the delegations she has received to one of the agents she follows. We assume the existence of a mechanism that prevents delegations from becoming cyclic (i.e., with an agent i, who delegates to agent j, who in turn delegates to agent k, who delegates back to agent i). In this way, delegations partition the nodes of the social network into a forest of trees, each consisting of edges that are directed toward the root. The root of each tree corresponds to an agent who casts a vote on behalf of all agents in her tree, i.e., all agents who directly or indirectly delegated their voting right to her. Whether the vote is for T or for F is decided by the root agent.

As delegations move voting power from agent to agent, votes have different weights. A vote has a weight that is equal to the total number of nodes in the tree rooted at the agent who casts the vote. The election outcome is computed using weighted majority on all votes. Votes are decided according to the agents' competency levels. Every agent i has an associated competency level p_i, indicating the probability that she will cast a vote for T when she decides to act so. We denote by $\mathbf{p}$ the vector of competency levels.

Following Kahng *et al.* [2018], we use the term *delegation mechanism* to refer to the collective actions of all agents. The following series of events take place:

1. On input the social network G and the vector $\mathbf{p}$ of competency levels, the delegation mechanism M decides (possibly, in a randomized way) the delegation forest.
2. The root agents cast their votes according to their competency levels.
3. The outcome of the election is computed using weighted majority of the votes.

We denote by $P_M(G, \mathbf{p})$ the probability that the outcome of the above process is alternative T.

We are particularly interested in delegation mechanisms that are *local*. Local mechanisms decide the actions of each agent independently. The decision for agent i depends only on her neighborhood and her followers, on their competency levels, as well as on the competency level of agent i. We denote by DV (standing for direct voting) the local delegation mechanism which lets all agents cast their votes directly.

3 Ineffectiveness of Local Mechanisms

Kahng *et al.* [2018], in their model, assume that voters can only delegate their votes to neighbors that have strictly higher competency level. In particular, they assume that there is a constant parameter α so that delegations from an agent i to another agent j in i's neighborhood only if $p_j > p_i + \alpha$.

We have two objections to this assumption. We describe the first one here and relegate the second one to Section 4. Consider an agent i with competency level $p_i = 10\%$, who follows another agent j with $p_j = 70\%$. Would agent i ever delegate to agent j in practice? A competency level of 10% for agent i can be interpreted as a 90% confidence that the correct alternative is F. Then, according to agent i's view, agent j is highly misinformed. Fixing the assumption of Kahng *et al.* [2018] is easy. We refine the sufficient condition for a delegation from agent i to agent j to be $p_j > p_i + \alpha$ if $p_i \geq 1/2$ and $p_j < p_i - \alpha$ if $p_i \leq 1/2$. We use the term α-delegations to refer to delegations obeying this rule. Essentially, under α-delegations, a voter can choose to delegate to someone who confirms her intuition and looks more expert.

Assuming that agents can only α-delegate, can local mechanisms recover the ground truth with as high probability as possible? We answer this question negatively by considering a particular network G consisting of a *guru* agent, agents that can delegate to the guru only (the *followers*), and isolated agents (the *partisans*). Two extreme mechanisms in this network are (1) the full delegation mechanisms (FD), in which any follower delegates to the guru and (2) the direct voting mechanism DV (with no delegations).

We will show that any local mechanism M is considerably worse than one of these two extreme mechanisms for some profile $\mathbf{p}$ of competence levels. Formally, we will prove the following statement.

Theorem 1. *Let M be a local delegation mechanism, $\alpha \geq 0$ and $\delta > 0$. There exists a profile $G, \mathbf{p}$ such that $\max\{P_{DV}(G, \mathbf{p}), P_{FD}(G, \mathbf{p})\} - P_M(G, \mathbf{p}) \geq 1/2 - \alpha - \delta$.*

Theorem 1 complements an impossibility result of Kahng *et al.* [2018], stating that no delegation mechanism outperforms DV. Specifically for small values of α, Theorem 1 implies that any mechanism can behave almost as poorly as the trivial mechanism that ignores the agents and selects one of the two alternatives equiprobably.

Recall that local mechanisms decide whether an agent delegates her vote or not, based on the local network structure and possibly on additional information about the competence levels of the agents in the local structure. Furthermore, locality requires that delegations are decided independently for each agent. We use the same competence level for all followers so that the profile is fully symmetric from every follower's point of view. This restricts the space of mechanisms we have to consider to those that use the same probability d to decide independently whether a follower will delegate her vote to the guru. This probability does not depend on the partisans.

Proof of Theorem 1. The number of agents is odd. The competence levels are either $p \in [1/2, 1)$ and $q > p + \alpha$ or $p \in (0, 1/2]$ and $q < p - \alpha$ for the followers and the guru, resepectively. There are $k = \left\lceil \max\{\frac{4}{\alpha}, \frac{24}{p^2}\} \ln \frac{4}{\delta} \right\rceil$ followers.

For every follower i, denote by X_i the Bernoulli random variable indicating whether she delegates her vote to the guru ($X_i = 1$) or not ($X_i = 0$). Let $X = \sum_{i=1}^{k} X_i$ be the total number of delegations. Let Y_i be the Bernoulli random variable indicating whether follower i will cast T as her vote if she is asked to do so by mechanism M (this is independent of whether agent i will eventually delegate her vote to the guru or not). Let $Y = \sum_{i=1}^{k} Y_i$ and $U = \sum_{i=1}^{k} (1 - X_i) Y_i$. Intuitively, Y is the total number of T-opinions by the followers while U is the total number of T-votes by the followers who do not delegate their vote to the guru. Also, let Z denote the event that the guru casts a vote for alternative T.

We will use the well-known Hoeffding inequality for bounding random variables like X, Y, and U.

Lemma 1 (Hoeffding 1963)**.** *Let S be the number of successes in N trials of a Bernoulli random variable, which takes value 1 with probability r (and 0 with probability $1-r$). Then, for every $\epsilon > 0$:*

$$\Pr[S \geq (r + \epsilon)N] \leq \exp(-2\epsilon^2 N) \tag{1}$$

$$\Pr[S \leq (r - \epsilon)N] \leq \exp(-2\epsilon^2 N) \tag{2}$$

We distinguish between two cases for mechanism M:

Case I: mechanism M uses $d < 1/2$

Consider the profile $\mathbf{p}$ with $\frac{n+1}{2} - k - 1$ of the partisans voting for T (with certainty) and the rest $\frac{n-1}{2}$ voting for F. Then, mechanism FD would return alternative T with probability q, i.e., the probability that the guru votes for 1.

We will show instead that M returns alternative T as the winner with probability at most δ. Observe that alternative T wins (below, we use the auxiliary symbol W to denote this event) if and only if all followers that do not delegate, as well as the guru, vote for T. We have

$$P_M(G, \mathbf{p}) = \Pr[W | X < 3k/4] \cdot \Pr[X < 3k/4]$$

$$
\begin{aligned}
&+\Pr[W|X \geq 3k/4]\cdot \Pr[X \geq 3k/4] \\
&\leq \Pr[W|X < 3k/4] + \Pr[X \geq 3k/4] \\
&= \Pr[U = k - X|X < 3k/4]\cdot \Pr[Z|X < 3k/4] \\
&\quad + \Pr[X \geq 3k/4] \\
&\leq \Pr[U = k - X|X < 3k/4] + \Pr[X \geq 3k/4]
\end{aligned} \tag{3}
$$

To bound $\Pr[X \geq 3k/4]$ we use inequality (1) from Lemma 1 with $S = X$, $N = k$, $\epsilon = 3/4 - d$, and $r = d$. Thus,

$$
\begin{aligned}
\Pr[X \geq 3k/4] &\leq \exp(-2(3/4-d)^2 k) \\
&\leq \exp(-k/8) \leq \delta/4,
\end{aligned} \tag{4}
$$

where the last two inequalities follow by the definition of k and d. Now, observe that $\Pr[U = k - X|X < 3k/4]$ is equal to the probability that the random variable Y_i is equal to 1 for every follower i that does not delegate, given that there are at least $k/4$ such followers. By the definition of p and q, we always have $0 < p < 1 - \alpha$. Thus,

$$
\begin{aligned}
\Pr[U = k - X|X < 3k/4] &< p^{k/4} < (1-\alpha)^{k/4} \\
&\leq \exp(-\alpha k/4) \leq \delta/4.
\end{aligned} \tag{5}
$$

Again, the last inequality follows by the definition of k. By inequalities (3), (4), and (5), we have that

$$P_{\text{FD}}(G, \mathbf{p}) - P_M(G, \mathbf{p}) \geq q - \delta/2. \tag{6}$$

Case II: mechanism M uses $d \geq 1/2$

Let $\ell = \left\lceil pk - \sqrt{2k \ln \frac{4}{\delta}} \right\rceil$; the definition of k (which is strictly higher than $\frac{2}{p^2} \ln \frac{4}{\delta}$) implies that $\ell > 0$. Consider the profile in which $\frac{n+1}{2} - \ell$ of the partisans vote for T (with certainty) and the rest vote for F.

We will first show that mechanism DV returns T as the winner with probability at least $1-\delta/4$. A sufficient condition for this is when at least ℓ followers cast a vote of T. Hence,

$$P_{\text{DV}}(G, \mathbf{p}) \geq \Pr[Y > \ell] = 1 - \Pr[Y \leq \ell]. \tag{7}$$

By applying inequality (2) from Lemma 1 for $S = Y$, $N = k$, $r = p$, and $\epsilon = \sqrt{\frac{\ln \frac{4}{\delta}}{2k}}$ (these yield $(r - \epsilon)N \geq \ell$), we obtain that $\Pr[Y \leq \ell] \leq \Pr[Y \leq (r-\epsilon)N] \leq \delta/4$; then, (7) implies that $P_{\text{DV}}(G) \geq 1 - \delta/4$.

In contrast, we now show that the probability that mechanism M yields T as the winner is at most $q + \delta/4$ and, hence,

$$P_{\text{DV}}(G, \mathbf{p}) - P_M(G, \mathbf{p}) \geq 1 - q - \delta/2. \tag{8}$$

Indeed, using again W to denote the event that T wins,

$$
\begin{aligned}
P_M(G, \mathbf{p}) &= \Pr[W|Z]\cdot\Pr[Z] + \Pr[W|\overline{Z}]\cdot\Pr[\overline{Z}] \\
&\leq \Pr[Z] + \Pr[W|\overline{Z}].
\end{aligned}
$$

By definition, $\Pr[Z] = q$; we complete the proof of (8) by showing that $\Pr[W|\overline{Z}] \leq \delta/4$. When the guru casts a vote for F, alternative T is the winner if and only if at least ℓ of the followers that do not delegate cast a vote for T, i.e., $\Pr[W|\overline{Z}] = \Pr[U \geq \ell]$. Observe that U is a sum of independent Bernoulli random variables, each of them having probability $(1-d)p$ of taking value 1. We have

$$
\begin{aligned}
\Pr[W|\overline{Z}] = \Pr[U \geq \ell] &\leq \Pr[U \geq \frac{1}{2}pk + \sqrt{\frac{k \ln \frac{4}{\delta}}{2}}] \\
&\leq \Pr[U \geq (1-d)pk + \sqrt{\frac{k \ln \frac{4}{\delta}}{2}}] \leq \delta/4.
\end{aligned}
$$

The first inequality follows by the definition of k (which is at least $\frac{24}{p^2} \ln \frac{4}{\delta}$). The second inequality is due to the fact that $d \geq 1/2$. The last inequality follows by applying inequality (1) from Lemma 1 with $S = U$, $N = k$, $r = (1-d)p$, and $\epsilon = \sqrt{\frac{\ln \frac{4}{\delta}}{2k}}$.

The proof of Theorem 1 now follows by inequalities (6) and (8) after setting specifically $p = 1/2$ and either $q = 1/2 + \alpha + \delta/2$ or $q = 1/2 - \alpha - \delta/2$. □

4 The Curse of α-Delegations

Our second objection to α-delegations is way more important. In our next example, we present a profile where delegating to less-informed agents is highly beneficial.

Example 1. Consider a set $\mathcal{N}$ of n agents (with odd n), connected through a social network $G(\mathcal{N}, E)$ as follows: $(n-1)/2$ agents are isolated while the remaining ones form a star. The competency levels $\mathbf{p}$ are $p_c = 1-2\epsilon$ for the center and $p_\ell = 1 - \epsilon$ for each leaf agent, where $\epsilon > 0$ is a negligibly small constant. The isolated nodes have competency level $p_i = 0$.

First consider the scenario in which only delegations to more informed agents are allowed. Under this scenario, no leaf agent will delegate to the center of the star and the delegation mechanism DV will be used. Majority of the ground truth will be achieved only if all agents of the star vote for T. The probability that this happens is $P_{\text{DV}}(G, \mathbf{p}) = p_c \cdot p_\ell^{(n-1)/2} \leq (1-\epsilon)^{(n+1)/2}$, which approaches 0 as n grows.

Now, let us consider the delegation mechanism M which makes all leaf agents delegate to the center of the star. The center has weight $(n+1)/2$ now and the probability that the ground truth is the outcome of weighted majority is $P_M(G, \mathbf{p}) = 1 - 2\epsilon$. □

Let us now attempt to develop some intuition on how to detect beneficial delegations by considering a very small example.

Example 2. Consider the five-agent social network and the competency levels shown in Figure 1. All possible combinations of actions (and corresponding delegation mechanisms) are as follows:

- Direct voting. The probability that T wins is the probability that at least two the agents 2-5 vote for T. Tedius calculations yield $P_{\text{DV}} = 79.5\%$.
- M_1: agent 4 delegates to 2. Then, T wins when agent 2 casts a vote for T or when agent 2 casts a vote for

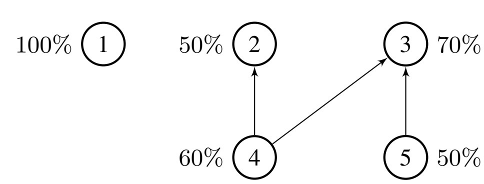

Figure 1: The social network and competency levels for Example 2.

F and both agents 3 and 5 cast votes for T. We get $P_{M_1} = 67.5\%$.

- M_2: agent 4 delegates to 3. Now, T wins when agent 3 casts a vote for T or when agent 3 casts a vote for F and both agents 2 and 5 cast votes for T. We get $P_{M_2} = 77.5\%$.
- M_3: agent 5 delegates to 3. T wins when agent 3 casts a vote for T or when agent 3 casts a vote for F and both agents 2 and 4 cast votes for T. We get $P_{M_3} = 79\%$.
- M_4: agents 4 and 5 delegate to 3. Now, T wins when agent 3 casts a vote for T. Hence, $P_{M_4} = 70\%$.
- M_5: agent 4 delegates to 2 and agent 5 delegates to 3. In this case, T wins when either agent 2 or agent 3 (or both) cast a vote for T. Hence, $P_{M_5} = 85\%$.

Mechanism M_5, which does not use only α-delegations is by far the best. In contrast, most of the other mechanisms use α-delegations. □

Is there some recipe that can be used to identify quickly the best possible mechanism in Example 2? Unfortunately, it does not seem so. To make things worse, in the next section, we show that, in general, this is not due to the different competency levels and difficulties in the computation of winning probabilities, but it has mostly to do with the structure of the social network.

5 The Complexity of Optimal Delegations

We now relax the restrictions posed by the model assumptions and locality. We assume that each agent can delegate to any of the agents she follows. Given information about the social network and the competency levels, can we coordinate (using a hypothetical non-local mechanism) delegations that maximize the probability of returning T as the winner? We call this optimization problem the *optimal delegation problem* (*ODP*) and also refer to the maximized probability as the optimal ODP value. We have the next negative statement supporting further the claim that liquid democracy may not be effective in discovering the ground truth.

Theorem 2. *Approximating the optimal value of ODP within an additive term of* $1/16$ *is NP-hard.*

We will use a polynomial-time reduction from 2P2N-3SAT, the special case of 3SAT consisting of 3-CNF formulas in which every boolean variable appears four times: twice as positive and twice as negative literal. 2P2N-3SAT is NP-hard [Yoshinaka, 2005]. Let ϕ be an instance of 2P2N-3SAT with n boolean variables $x_1, x_2, ..., x_n$ and m clauses $C_1, C_2, ..., C_m$. Since each clause has exactly three literals, it holds that $n = 3m/4$. Given ϕ, our reduction constructs a social network G_ϕ together with the agents' competence levels $\mathbf{p}$ with the following properties:

- If ϕ is satisfiable, there are delegations in G_ϕ that lead to selecting alternative T with probability $3/4$.
- If ϕ is not satisfiable, any delegations lead to selecting alternative T with probability at most $11/16$.

The graph G_ϕ is as follows: For every variable x_i, G_ϕ has a *variable gadget* subgraph as the one shown in Figure 2. The

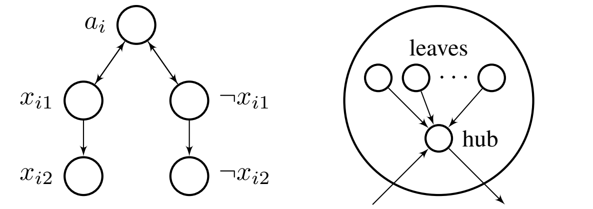

Figure 2: The variable gadget corresponding to variable x_i (left) and a supernode (right).

nodes at the left (respectively, right) segment corresponds to the two occurrences of the literal x_i (resp., of the literal $\neg x_i$). All variable gadgets are stacked by connecting the nodes x_{i2} and $\neg x_{i2}$ of the variable gadget for x_i to the node a_{i+1} of the variable gadget for x_{i+1}.

The graph G_ϕ contains several *supernodes*. A supernode is a star subgraph consisting of a *hub* and *leaves* connecting to it (see Figure 2). All edges connecting a supernode with the rest of G_ϕ are incoming to or outgoing from the hub.

For each clause c_j of ϕ, graph G_ϕ has a *clause gadget* C_j, consisting of a supernode with $10m$ nodes. The graph G_ϕ has four additional supernodes:

- H_1 of size $70m^2 - m$, connected to supernode C_1,
- L_1 of size $20m^2$,
- L_2 of size $20m^2$, connected to node a_1, and
- H_2 of size $80m^2 - 11m/4$. The nodes x_{n2} and $\neg x_{n2}$ are connected to H_2.

For every literal, we label arbitrarily its two appearances in ϕ with 1 and 2, respectively. For a literal x_i (resp., $\neg x_i$) that appears in a clause c_j with label $\ell \in \{1, 2\}$, we connect supernode C_j to node $x_{i\ell}$ (resp., $\neg x_{i\ell}$) and node $x_{i\ell}$ (resp., $\neg x_{i\ell}$) to supernode C_{j+1} if $i \leq n-1$ or to supernode L_1 if $i = n$. G_ϕ has an additional isolated node Q. The agent corresponding to node Q votes for T with certainty. All other nodes have a competency level of $1/2$. An example of G_ϕ is given in Figure 3.

The total number of nodes in G_ϕ is $1 + 200m^2$. Observe that, besides node Q, the graph G_ϕ has two nodes (the hubs of supernodes L_1 and H_2) with out-degree equal to 0. Furthermore, from any other node of G_ϕ besides Q, there is a path to either of these two out-degree-0 nodes. Hence, the two corresponding agents can amass all delegations from all agents besides Q. If, in addition, each of these agents gets exactly $100m^2$ delegations, the probability that alternative T is selected as the winner of weighted majority is exactly $3/4$ (i.e., the probability that at least one of the centers of L_1 and H_2 votes for alternative T; recall that agent Q certainly votes for T). The next lemma shows that this probability is the maximum possible, and for every other set of delegations, the probability that alternative T is selected as the winner of weighted majority is at most $11/16$. The proof is long and technical; it is omitted due to lack of space. It uses detailed arguments about how an even number (the total weight from all agents besides Q) can be expressed as a sum of integers (each corresponding to the weight of an agent who casts a vote).

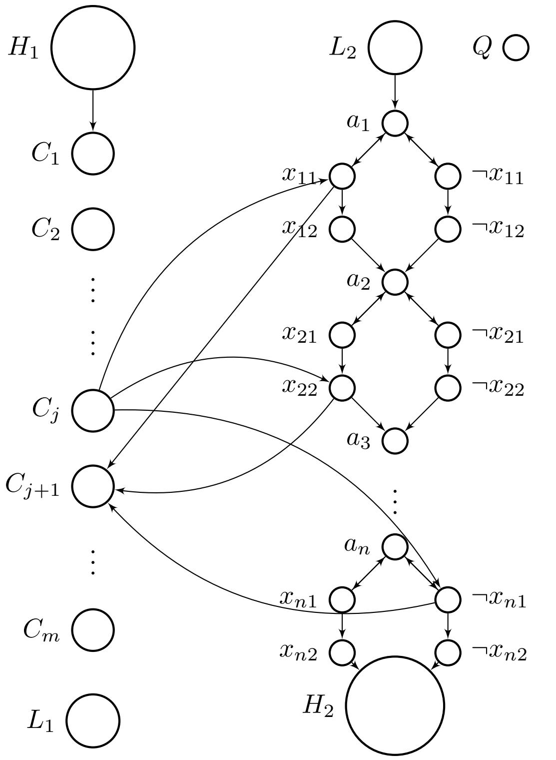

Figure 3: An example of G_ϕ. The edges from supernode C_j to the variable gadgets and then back to supernode C_{j+1} correspond to clause $c_j = x_1 \vee x_2 \vee \neg x_n$. The labels of the corresponding literals are 1, 2, and 1, respectively.

Lemma 2. *If each of the hubs of supernodes L_1 and H_2 amass weight equal to $100m^2$, then the probability that T is the winner is $3/4$. For any other set of delegations, this probability is at most $11/16$.*

Now, the proof of Theorem 2 follows by Lemma 3.

Lemma 3. *The hubs of supernodes L_1 and H_2 can amass $100m^2$ delegations each if and only if ϕ is satisfiable.*

Proof. If ϕ is satisfiable, we will use a truth assignment to show that the hubs of supernodes L_1 and H_2 can amass $100m^2$ delegations each, as follows: First, at all supernodes, the leaves delegate their votes to their hub. Then, the hub of H_1 delegates its weight of $70m^2 - m$ to the hub of supernode C_1.

For every clause c_j, we arbitrarily select one of the true literals of the clause, and identify the corresponding node to which the hub of C_j is connected as the *active* variable node of c_j (we select arbitrarily in case more than one literals of c_j are true). For $j = 1, ..., m-1$, the hub of supernode C_j delegates its weight to the active variable node of c_j, which in turn delegates its weight to the hub of supernode C_{j+1} if $j \neq m$ or of supernode L_1 if $j = m$. In this way, the hub of L_1 amasses weight equal to $100m^2$, where a weight of $70m^2 - m$ comes from the hub of H_1, $10m^2$ comes from the hubs of C_j for $j \in [m]$, m comes from the active variable nodes, and $20m^2 - 1$ come from its leaves.

The hub of H_2 amasses the remaining weight of $100m^2$ as follows: The hub of L_2 delegates its weight to node a_1. For $i = 1, ..., n$, node a_i delegates to x_{i1} if x_i is false and to $\neg x_{i1}$ otherwise. Every non-active variable node x_{i1} (resp., $\neg x_{i1}$) delegates its votes to x_{i2} (resp., $\neg x_{i2}$) if x_i is false (resp., true), otherwise it delegates to a_i. Every non-active variable node x_{i2} or $\neg x_{i2}$ delegates its votes to a_{i+1} if $i \neq n$ and to the hub of H_2 if $i = n$. In this way, the hub of H_2 amasses weight equal to $100m^2$, where a weight of $20m^2$ comes from the hub of L_1, $11m/4$ comes from the nodes in the variable gadgets besides the active variable nodes, and $80m^2 - 11m/4 - 1$ comes from its leaves.

Conversely, if the hubs of L_1 and H_2 amass $100m^2$ delegations each, we will show that ϕ is satisfiable. Clearly, in this case, the leaves of any supernode delegate their votes to their hub. Observe that if the hub of H_2 amassed the weight of H_1, then its weight would exceed $100m^2$. So, the hub of L_1 amasses the weight of H_1. Then, the hub of L_1 cannot amass the weight of L_2 as well, since its weight would then exceed $100m^2$. So, the hub of H_2 amasses the weight of L_2. Hence, the hub of H_2 acquires weight of $100m^2 - 11m/4$ from itself, its leaves and L_2. Thus, it cannot get any weight from clause supernodes as each of them has weight equal to $10m > 11m/4$. So, the remaing weight of $11m/4$ that the hub of H_2 gets comes from nodes in variable gadgets. As a result, the hub of L_1 amasses the weight of itself, its leaves, of supernode H_1, all clause supernodes, and exactly m nodes in variable gadgets.

Hence, the weight of L_2 propagates to the hub of H_2 through a "delegation path" that crosses each variable gadget corresponding to variable x_i either through the left segment with node a_i delegating to x_{i1}, x_{i1} to x_{i2}, and x_{i2} to a_{i+1} (or to the hub of H_2 if $i = n$) or through the right segment with node a_i delegating to $\neg x_{i1}$, $\neg x_{i1}$ to $\neg x_{i2}$, and $\neg x_{i2}$ to a_{i+1} (or, similarly, to the hub of H_2). Then, the hub of every clause supernode delegates to a variable node that does not belong to the segment of the variable gadget that is used by the L_2-to-H_2 delegation path. Hence, by setting the variable x_i to true if the L_2-to-H_2 crosses the variable gadget of x_i from the right and to false if its crosses it from the left, we get an assignment that satisfies all clauses of ϕ. □

6 Conclusion

Our results indicate that liquid democracy can be very ineffective in discovering the ground truth. Of course, this conclusion might be due to modeling assumptions here and in previous work, e.g., in [Kahng *et al.*, 2018]. However, modeling of similar flavor has been used extensively in social choice theory to show that even simple voting rules discover ground truths very effectively (e.g., see Caragiannis *et al.* 2016). We interpret our results as an indication that the drawback is in the combination of weighted majority with spontaneous delegations in the social network without any central control. In this way, we question the implementations of today's liquid democracy systems.

References

[Allen, 2008] Michael Allen. A medium of assent and its fit with society. *The ITP News*, 4:12–13, 2008.

[Behrens *et al.*, 2014] Jan Behrens, Axel Kistner, Andreas Nitsche, and Björn Swierczek. *The Principles of LiquidFeedback*. 2014.

[Behrens, 2017] Jan Behrens. The origins of liquid democracy. *The Liquid Democracy Journal*, 5, 2017.

[Bloembergen *et al.*, 2019] Daan Bloembergen, Davide Grossi, and Martin Lackner. On rational delegations in liquid democracy. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019.

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

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

[Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[Brill and Talmon, 2018] Markus Brill and Nimrod Talmon. Pairwise liquid democracy. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 137–143, 2018.

[Brill, 2018] Markus Brill. Interactive democracy. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1183–1187, 2018.

[Caragiannis *et al.*, 2016] Ioannis Caragiannis, Ariel D. Procaccia, and Nisarg Shah. When do noisy votes reveal the truth? *ACM Transactions on Economics and Computation*, 4(3):15:1–15:30, 2016.

[Christoff and Grossi, 2017] Zoe 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)*, pages 134–150, 2017.

[Conitzer and Sandholm, 2005] Vincent Conitzer and Tuomas Sandholm. Common voting rules as maximum likelihood estimators. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 145–152, 2005.

[de Condorcet, 1785] Marquis de Condorcet. Essai sur l'application de l'analyse à la probabilité de décisions rendues à la pluralité de voix. Imprimerie Royal, 1785. Facsimile published in 1972 by Chelsea Publishing Company, New York.

[Dodgson, 1884] Charles L. Dodgson. *Principles of Parliamentary Representation*. Harrison and Sons, London, 1884.

[Elkind and Slinko, 2016] Edith Elkind and Arkadii Slinko. Rationalizations of voting rules. In *Handbook of Computational Social Choice*, pages 169–196. 2016.

[Ford, 2002] Bryan Ford. Delegative democracy. Unpublished manuscript, 2002.

[Gölz *et al.*, 2018] Paul Gölz, 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)*, pages 188–202, 2018.

[Green-Armytage, 2005] James Green-Armytage. Direct democracy by delegable proxy. Unpublished manuscript, 2005.

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

[Hoeffding, 1963] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

[Kahng *et al.*, 2018] Anson Kahng, Simon Mackenzie, and Ariel D. Procaccia. Liquid democracy: An algorithmic perspective. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1095–1102, 2018.

[Kling *et al.*, 2015] Christoph Carl 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 9th AAAI International Conference on Web and Social Media (ICWSM)*, pages 208–217, 2015.

[Mancini, 2015] Pia Mancini. Why it is time to redesign our political system. *European View*, 14(1):69–75, 2015.

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

[Tullock, 1967] Gordon Tullock. *Toward a Mathematics of Politics*. University of Michigan Press, 1967.

[Yamakawa *et al.*, 2007] Hiroshi Yamakawa, Motohiro Yoshida, and Michiko Tsuchiya. Toward delegated democracy: Vote by yourself, or trust your network. *International Journal of Human and Social Sciences*, 1(2), 2007.

[Yoshinaka, 2005] Ryo Yoshinaka. Higher-order matching in the linear lambda calculus in the absence of constants is np-complete. In *Proceedings of the 16th International Conference on Term Rewriting and Applications (RTA)*, pages 235–249, 2005.

[Young, 1988] H. Peyton Young. Condorcet's theory of voting. *The American Political Science Review*, 82(4):1231–1244, 1988.

[Zhang and Zhou, 2017] Bingsheng Zhang and Hongsheng Zhou. Statement voting and liquid democracy. In *Proceedings of the 36th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 359–361, 2017.