

On Rational Delegations in Liquid Democracy

Daan Bloembergen
Centrum Wiskunde & Informatica
Amsterdam, The Netherlands

Davide Grossi
Bernoulli Institute
University of Groningen
Groningen, The Netherlands

Martin Lackner
TU Wien
Vienna, Austria

Abstract

Liquid democracy is a proxy voting method where proxies are delegable. We propose and study a game-theoretic model of liquid democracy to address the following question: when is it rational for a voter to delegate her vote? We study the existence of pure-strategy Nash equilibria in this model, and how group accuracy is affected by them. We complement these theoretical results by means of agent-based simulations to study the effects of delegations on group's accuracy on variously structured social networks.

Introduction

Liquid democracy (Blum and Zuber 2016) is an influential proposal in recent debates on democratic reforms in both Europe and the US. Several grassroots campaigns, as well as local parties, experimented with this novel type of decision making procedure. Examples include the German Piratenpartei¹ and the EU Horizon 2020 project WeGov-Now (Boella et al. 2018), which have incorporated the LiquidFeedback² platform in their decision making, as well as grass-roots organizations such as the Democracy Earth Foundation³. Liquid democracy is a form of proxy voting (Miller 1969; Tullock 1992; Alger 2006; Green-Armytage 2015; Cohensius et al. 2017) where, in contrast to classical proxy voting, proxies are delegable (or transitive, or transferable). Suppose we are voting on a binary issue, then each voter can either cast her vote directly, or she can delegate her vote to a proxy, who can again either vote directly or, in turn, delegate to yet another proxy, and so forth. Ultimately, the voters that decided not to delegate cast their ballots, which now carry the weight given by the number of voters who entrusted them as proxy, directly or indirectly.

Contribution The starting point of our analysis is an often cited feature of liquid democracy: transitive delegations reduce the level of duplicated effort required by direct voting, by freeing voters from the need to invest effort in order to vote accurately. The focus of the paper is the decision-making problem that voters, who are interested in casting an

accurate vote, face between voting directly, and thereby incurring a cost in terms of effort invested to learn about the issue at hand, or delegating to another voter in their network, thereby avoiding costs. We define a game-theoretic model, called *delegation game*, to represent this type of interaction. We establish pure strategy Nash equilibrium existence results for classes of delegation games, and study the quality of equilibria in terms of the average accuracy they enable for the population of voters, both analytically and through simulations. Proofs of the two main results (Theorems 1 and 2) are presented in full, while we provide proofs of the simpler secondary results only as supplementary material.⁴

By means of simulations we also study the effects of different network structures on delegation games in terms of: performance against direct voting, average accuracy and the probability of a correct majority vote, the number and quality of voters acting as ultimate proxies (so-called gurus) and, finally, the presence of delegation cycles. To the best of our knowledge, this is the first paper providing a comprehensive study of liquid democracy from a game-theoretic angle.

Related Work Although the idea of delegable proxy was already sketched by Dodgson (1884), only a few very recent papers have studied aspects of liquid democracy in the (computational) social choice theory (Brandt et al. 2016) literature. Kling et al. (2015) provide an analysis of election data from the main platform implementing a liquid democracy voting system (Liquid Feedback) for the German Piratenpartei. They focus on network theoretic properties emerging from the structure of delegations—with particular attention to the number of highly influential gurus or ‘super-voters’. Inspired by their experimental analysis, Gözl et al. (2018) propose and analyze a variant of the liquid democracy scheme able to restrict reliance on super-voters. Skowron et al. (2017) study an aspect of the Liquid Feedback platform concerning the order in which proposals are ranked and by which they are brought to the attention of the community. Boldi et al. (2011) investigate applications of variants of the liquid democracy voting method (called viscous democracy) to recommender systems. Brill (2018) presents some research directions in the context of liquid democracy. A general, more philosophical discussion of liquid democracy is provided by Blum and Zuber (2016).

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://www.piratenpartei.de/>

²<https://liquidfeedback.org/>

³<https://www.democracy.earth/>

⁴<https://arxiv.org/abs/1802.08020>

More directly related to our investigations is the work by Christoff and Grossi (2017) and, especially, by Kahng, Mackenzie, and Procaccia (2018). The first paper studies liquid democracy as an aggregator—a function mapping profiles of binary opinions to a collective opinion—in the judgment aggregation and binary voting tradition (Grossi and Pigozzi 2014; Endriss 2016). The focus of that paper are the unintended effects that transferable proxies may have due to delegation cycles, and due to the failure of rationality constraints normally satisfied by direct voting.

The second paper addresses one of the most cited selling arguments for liquid democracy: delegable proxies guarantee that better informed agents can exercise more weight on group decisions, thereby increasing their quality. Specifically, Kahng, Mackenzie, and Procaccia (2018) study the level of accuracy that can be guaranteed by liquid democracy (based on vote delegation with weighted majority) vs. direct voting by majority. Their key result consists in showing that no ‘local’ procedure to select proxies can guarantee that liquid democracy is, at the same time, never less accurate (in large enough graphs) and sometimes strictly more accurate than direct voting. In contrast to their work, we assume that agents incur costs (effort) when voting directly, and on that basis we develop a game-theoretic model. Also, we assume agents aim at tracking their own type rather than an external ground truth, although we do assume such a restriction in our simulations to better highlight how the two models are related and to obtain insights applicable to both.

Preliminaries

Types, Individual Accuracy and Proximity

We are concerned with a finite set of agents (or voters, or players) $N = \{1, \dots, n\}$ having to decide whether $x = 1$ or $x = 0$. For each agent one of these two outcomes is better, but the agent is not necessarily aware which one. We refer to this hidden optimal outcome as the *type of agent i* and denote it by $\tau_i \in \{0, 1\}$. Agents want to communicate their type truthfully to the voting mechanism, but they know it only imperfectly. This is captured by the *accuracy q_i* of an agent i : q_i determines the likelihood that, if i votes directly, she votes according to her type τ_i . We assume that an agent’s accuracy is always ≥ 0.5 , i.e., at least as good as a coin toss.

We distinguish two settings depending on whether the agents’ types are deterministic or probabilistic. A *deterministic type profile* $T = \langle \tau_1, \dots, \tau_n \rangle$ simply collects each agent’s type. In *probabilistic type profiles* types are independent random variables drawn according to a distribution \mathbb{P} . Given a probabilistic type profile, the likelihood that any two agents $i, j \in N$ are of the same type is called the *proximity $p_{i,j}$* where $p_{i,j} = \mathbb{P}(\tau(i) = \tau(j)) = \mathbb{P}(\tau(i) = 1) \cdot \mathbb{P}(\tau(j) = 1) + (1 - \mathbb{P}(\tau(i) = 1)) \cdot (1 - \mathbb{P}(\tau(j) = 1))$. In the probabilistic setting we assume agents know such value although, importantly, they do not know \mathbb{P} . In a deterministic type profile, we have $p_{i,j} = 1$ if $\tau_i = \tau_j$ and $p_{i,j} = 0$ otherwise. Following standard equilibrium theory, our theoretical results assume agents act as if they have access to the accuracy of each agent. More realistically, in our simulations we assume agents have access to such information only with respect to

neighbors on an underlying interaction structure.

Interaction Structure and Delegations

Agents are nodes in a network (directed graph) represented by a relation $R \subseteq N^2$. For $i \in N$ the neighborhood of i in $\langle N, R \rangle$ is denoted $R(i)$, i.e., the agents that are directly connected to i . Agents have the choice of either voting themselves, thereby relying solely on their own accuracy, or delegating to an agent in their neighborhood. A *delegation profile* is a vector $\mathbf{d} = \langle d_1, \dots, d_n \rangle \in N^n$. Given a delegation profile \mathbf{d} we denote by d_i the *proxy* selected by i in \mathbf{d} . Clearly a delegation profile can be viewed as a functional graph on N or, equivalently, as a map in $\mathbf{d} : N \rightarrow N$ where $\mathbf{d}(i) = d_i$. When the iterated application of \mathbf{d} from i reaches a fixed point we denote such fixed point as d_i^* and call it i ’s *guru* (in \mathbf{d}). In the following, we write N^* to denote the set of voters whose delegation does not lay on a path ending on a cycle, i.e., the set of voters i for which d_i^* exists. We write $\mathbf{d}' = (\mathbf{d}_{-i}, j)$ as a short form for $\mathbf{d}' = \langle d_1, \dots, d_{i-1}, j, d_{i+1}, \dots, d_n \rangle$.

As agents may only be able to observe and interact with their direct network neighbors, structural properties of the interaction network may play a role in the model dynamics. In our simulations we will focus on undirected graphs (that is, R will be assumed to be symmetric, as social ties are normally mutual) consisting of one single connected component (that is, N^2 is included in the reflexive transitive closure of R). Under these assumptions, we consider four typical network structures that are well represented in the literature on social network analysis (cf. Jackson 2008): 1) the *random network*, in which each pair of nodes has a given probability of being connected (Erdős and Rényi 1959); 2) the *regular network*, in which all nodes have the same degree; 3) the *small world network*, which features a small average path length and high clustering (Watts and Strogatz 1998); and 4) the *scale free network*, which exhibits a power law degree distribution (Barabási and Albert 1999).⁵

A Model of Rational Delegations

Individual Accuracy under Delegable Proxy

Each agent i has to choose between two options: either to vote herself with accuracy q_i or to delegate, thereby inheriting the accuracy of another voter (unless i is involved in a delegation cycle). These choices are recorded in the delegation profile \mathbf{d} and can be used to compute the individual accuracy for each agent $i \in N^*$ as follows:

$$q_i^*(\mathbf{d}) = \begin{cases} q d_i^* \cdot p_{i,d_i^*} + (1 - q d_i^*) \cdot (1 - p_{i,d_i^*}) & \text{if } i \in N^* \\ 0.5 & \text{if } i \notin N^* \end{cases} \quad (1)$$

In (1) i ’s accuracy equals the likelihood that i ’s guru has the same type and votes accurately plus the likelihood that i ’s guru has the opposite type and fails to vote accurately. Note that if i votes directly, i.e., $d_i = i$, then $q_i^*(\mathbf{d}) = q_i$.

⁵Although random and regular graphs are not generally applicable to real-world settings, they serve as a useful baseline to illustrate the effects of network structure on delegations.

Observe that if i 's delegation leads to a cycle ($i \notin N^*$), i 's accuracy is set to 0.5. The rationale for this assumption is the following. If an agent delegates into a cycle, even though she knows her own accuracy and she actively engages with the voting mechanism by expressing a delegation, she fails to pass information about her type to the mechanism. No information is therefore available to decide about her type.

It may be worth observing that, by (1), in a deterministic type profile we have that $p_{i,j} \in \{0, 1\}$ and therefore i 's accuracy reduces to: $q_{d_i^*}$ if $i \in N^*$ and $\tau(i) = \tau(d_i^*)$; $1 - q_{d^*(i)}$ if $i \in N^*$ and $\tau(i) \neq \tau(d_i^*)$; and 0.5 if $i \notin N^*$.

Before introducing our game theoretic analysis, we make the following observation. Agents have at their disposal an intuitive strategy to improve their accuracy: simply delegate to a more accurate neighbor. We say that a delegation profile \mathbf{d} is *positive* if for all $j \in N$ either $d_j = j$ or $q_j^*(\mathbf{d}) > q_j$. Furthermore, we say that a delegation from i to a neighbor j is *locally positive* if $q_j \cdot p_{i,j} + (1 - q_j) \cdot (1 - p_{i,j}) > q_i$.

Proposition 1. *Let \mathbf{d} be a positive delegation profile. Further, let $s, t \in N$, $d_s = s$, and $\mathbf{d}' = (\mathbf{d}_{-s}, t)$, i.e., agent s votes directly in \mathbf{d} and delegates to t in \mathbf{d}' . If the delegation from s to t is locally positive, then \mathbf{d}' is positive.*

However, locally positive delegations do not necessarily correspond to optimal delegations. This can be easily seen in an example where agent i is not a neighbor of a very competent agent j , but would have to delegate via an intermediate agent k (who delegates to j). If this intermediate agent k has a lower accuracy than i , then the delegation from i to k would not be locally positive, even though it is an optimal choice. So utility-maximization may require information which is inherently non-local (accuracy of 'far' agents).

Delegation Games

We assume that each agent i has to invest an effort e_i to manifest her accuracy q_i . If she delegates, she does not have to spend effort. Agents aim therefore at maximizing the trade-off between the accuracy they can achieve (either by voting directly or through proxy) and the effort they spend. Under this assumption, the binary decision set-up with delegable proxy we outlined above can be used to define a natural game—called *delegation game*— $G = \langle N, \mathbb{P}, R, \Sigma_i, u_i \rangle$, with $i \in N$, where N is the set of agents, \mathbb{P} is the (possibly degenerate) distribution from which the types of the agents in N are drawn, R the underlying network as defined above, $\Sigma_i \in N$ is the set of strategies of agent i (voting, or choosing a specific proxy), and

$$u_i(\mathbf{d}) = \begin{cases} q_i^*(\mathbf{d}) & \text{if } d_i \neq i \\ q_i - e_i & \text{if } d_i = i \end{cases} \quad (2)$$

is agent i 's utility function. The utility i extracts from a delegation profile equals the accuracy she inherits through proxy or, if she votes, her accuracy minus the effort spent.⁶ In delegation games we assume that $q_i - e_i \geq 0.5$ for all $i \in N$. This is because if $q_i - e_i < 0.5$, then i would prefer a random effortless choice over taking a decision with effort.

⁶No utility is accrued for gaining voting power in our model.

	vote	delegate (to 1)
vote	$q_1 - e_1, q_2 - e_2$	$q_1 - e_1, q_1$
delegate (to 2)	$q_2, q_2 - e_2$	0.5, 0.5

Table 1: A two-player delegation game. The row player is agent 1 and the column player is agent 2.

A few comments about the setup of (2) are in order. First of all, as stated earlier, we assume agents to be truthful. They do not aim at maximizing the chance their type wins the vote, but rather to relay their type to the mechanism as accurately as possible.⁷ Secondly, notice that the utility an agent extracts from a delegation profile may equal the accuracy of a random coin toss when the agent's delegation ends up into a delegation cycle (cf. (1)). If this happens the agent fails to relay information about her type, even though she acted in order to do so. This justifies the fact that 0.5 is also the lowest payoff attainable in a delegation game.

The following classes of delegation games will be used in the paper: games with *deterministic profiles*, i.e., where \mathbb{P} is degenerate and all players are assigned a crisp type from $\{0, 1\}$; *homogeneous* games, where all players have the same (deterministic) type;⁸ and *effortless voting* games, where for each $i \in N$ we have $e_i = 0$.

As an example, a homogeneous game in matrix form is given in Table 1, where $N = \{1, 2\}$, $R = N^2$ and the distribution yields the deterministic type profile $T = \langle 1, 1 \rangle$. Interestingly, if we assume that $q_i - e_i > 0.5$ with $i \in \{1, 2\}$, and that⁹ $q_{-i} > q_i - e_i$ (i.e., the opponent's accuracy is higher than the player's individual accuracy minus her effort), then the game shares the ordinal preference structure of the class of anti-coordination games: players need to avoid coordination on the same strategy (one should vote and the other delegate), with two coordination outcomes (both players voting or both delegating) of which the second (the delegation cycle) is worst for both players. Notice that, were the underlying network not complete (i.e., $R \subset N^2$), the matrix would be shrunk by removing the rows and columns corresponding to the delegation options no longer available.

The introduction of effort has significant consequences on the delegation behavior of voters, and we will study it in depth in the coming sections. It is worth noting immediately that the assumptions of Proposition 1 no longer apply, since agents may prefer to make delegations that are not locally positive due to the decreased utility of voting directly.

Existence of Equilibria in Delegation Games

In this section we study the existence of pure strategy Nash Equilibria (NE) in two classes of delegation games. NE

⁷Notice however that our modeling of agents' utility remains applicable in this form even if agents are not truthful but the underlying voting rule makes truthfulness a dominant strategy—such as majority in the binary voting setting used here.

⁸This is the type of interaction studied, albeit not game-theoretically, by Kahng, Mackenzie, and Procaccia (2018) and normally assumed by jury theorems (Grofman, Owen, and Feld 1983).

⁹We use here the usual notation $-i$ to denote i 's opponent.

describe how ideally rational voters would resolve the effort/accuracy trade-off. Of course, such voters have common knowledge of the delegation game structure—including, therefore, common knowledge of the accuracies of ‘distant’ agents in the underlying network. Our simulations will later lift some of such epistemic assumptions built into NE.

Deterministic Types In the following we provide a NE existence result for games with deterministic type profiles.

Theorem 1. *Delegation games with deterministic type profiles always have a (pure strategy) NE.*

Proof. First of all, observe that since the profile is deterministic, for each pair of agents i and j , $p_{i,j} \in \{0, 1\}$. The proof is by construction. First, we partition the set of agents N into $N_1 = \{i \in N \mid \tau(i) = 1\}$ and $N_0 = \{i \in N \mid \tau(i) = 0\}$. We consider these two sets separately; without loss of generality let us consider N_1 . Further we consider the network $R_1 = \{(i, j) \in N_1 \times N_1 : (i, j) \in R\}$. Since (N_1, R_1) can be seen as a directed graph, we can partition it into Strongly Connected Components (SCCs). If we shrink each SCC into a single vertex, we obtain the condensation of this graph; note that such a graph is a directed acyclic graph (DAG). We construct a delegation profile \mathbf{d} by traversing this DAG bottom up, i.e., starting with leaf SCCs.

Let $S \subseteq N_1$ be a set of agents corresponding to a leaf SCC in the condensation DAG. We choose an agent i in S that has maximum $q_i - e_i$. Everyone in S (including i) delegates to i . Now let $S \subseteq N_1$ be a set of agents corresponding to an inner node SCC in the condensation DAG and assume that we have already defined the delegation for all SCCs that can be reached from S . As before, we identify an agent $i \in S$ with maximum $q_i - e_i$. Further, let $T \subseteq N_1 \setminus S$ be the set of all voters j that can be reached from S in (N_1, R_1) , and for which $q_j^* > q_i - e_i$. We distinguish two cases. (i) If $T \neq \emptyset$, then we choose an agent $k \in T$ with $q_k^* = \max_{j \in T} q_j^*$ and all agents in S directly or indirectly delegate to k . (ii) If $T = \emptyset$, all agents in S delegate to i . This concludes our construction (as for N_0 the analogous construction applies); let \mathbf{d} be the corresponding delegation profile.

It remains to verify that this is indeed a NE: Let i be some agent in an SCC S , and, without loss of generality, let $i \in N_1$. Observe that since we have a deterministic profile, if agent i changes her delegation to j , then i 's utility changes to $q_j^*(\mathbf{d})$ if $i \in N_1$ and $1 - q_j^*(\mathbf{d})$ if $i \in N_0$. First, note that for all agents $k \in N$, $q_k^*(\mathbf{d}) \geq q_k - e_k \geq 0.5$. Hence, we can immediately exclude that for agent i delegating to an agent in $j \in N_0$ is (strictly) beneficial, as it would yield an accuracy of at most $1 - q_j^* \leq 0.5$. Towards a contradiction assume there is a beneficial deviation to an agent $j \in N_1$, i.e., there is an agent $j \in R(i) \cap N_1$ with $q_j^*(\mathbf{d}) > q_i^*(\mathbf{d})$. Let us now consider the three cases: (1) $d_i = i$, (2) $d_i^* \in S$ but $d_i \neq i$, and (3) $d_i^* \notin S$. In case (1), everyone in S delegates to i . Hence, if $j \in S$, a cycle would occur yielding a utility of 0.5, which is not sufficient for a beneficial deviation. If a delegation to $j \notin S$ is possible but was not chosen, then by construction $q_j^* \leq q_i - e_i$ and hence this deviation is not beneficial. We conclude that in case (1) such an agent j cannot exist. In case (2), everyone in S delegates to d_i^* .

Hence, if $j \in S$, then $d_j^* = d_i^*$, a contradiction. If $j \notin S$, the same reasoning as before applies and hence also here we obtain a contradiction. In case (3), by construction, $d_i^* \notin S$ had been chosen to maximise accuracy, hence $j \in S$. Since for all $k \in S$, $d_k^* = d_i^*$, only a deviation to i itself can be beneficial, i.e., $j = i$. However, since d_i^* was chosen because $q_{d^*(i)}^* > q_i - e_i$, no beneficial deviation is possible. \square

It follows that also homogeneous games always have NE.

Effortless Voting Effortless voting ($e_i = 0$ for all $i \in N$) is applicable whenever effort is spent in advance of the decision and further accuracy improvements are not possible.

Theorem 2. *Delegation games with effortless voting always have a (pure strategy) NE.*

Proof. We prove this statement by showing that the following procedure obtains a NE: We start with a strategy profile in which all players vote directly, i.e., player i 's strategy is i . Then, we iteratively allow players to choose their individual best response strategy to the current strategy profile. Players act sequentially in arbitrary order. If there are no more players that can improve their utility by changing their strategy, we have found a NE. We prove convergence of this procedure by showing that a best response that increases the player's utility never decreases the utility of other players.

We proceed by induction. Assume that all previous best responses have not reduced any players' utility (IH). Assume player i now chooses a best response that increases her utility. Let \mathbf{d} be the delegation profile; further, let $d_i^* = s$. By assumption, i 's utility started with $q_i - e_i = q_i$ and has not decreased since, i.e., $u_i(\mathbf{d}) \geq q_i$. Since i 's best response strictly increases i 's utility, it cannot be a delegation to herself. So let a delegation to $j \neq i$ be i 's best response and consider profile $\mathbf{d}' = (\mathbf{d}_{-i}, j)$. Further, let $d_j^* = t$, i.e., i now delegates to j and by transitivity to t , i.e., $d_i'^* = d_j'^* = t$. Let k be some player other than i . We define the delegation path of k as the sequence $(\mathbf{d}(k), \mathbf{d}(\mathbf{d}(k)), \mathbf{d}(\mathbf{d}(\mathbf{d}(k))), \dots)$. If k 's delegation path does not contain i , then k 's utility remains unchanged, i.e., $u_k(\mathbf{d}') \geq u_k(\mathbf{d})$. If k 's delegation path contains i , then k now delegates by transitivity to t , i.e., we have $d_k^* = s$ and $d_k'^* = t$. By Equation (2), we have

$$u_k(\mathbf{d}) = q_s \cdot p_{k,s} + (1 - q_s) \cdot (1 - p_{k,s}) \quad \text{and} \quad (3)$$

$$u_k(\mathbf{d}') = q_t \cdot p_{k,t} + (1 - q_t) \cdot (1 - p_{k,t}). \quad (4)$$

We have to show that k 's utility does not decrease, i.e., $u_k(\mathbf{d}') \geq u_k(\mathbf{d})$, under the assumption that i chooses a best response, i.e., $u_i(\mathbf{d}') > u_i(\mathbf{d})$, with:

$$u_i(\mathbf{d}) = q_s \cdot p_{i,s} + (1 - q_s) \cdot (1 - p_{i,s}) \quad \text{and} \quad (5)$$

$$u_i(\mathbf{d}') = q_t \cdot p_{i,t} + (1 - q_t) \cdot (1 - p_{i,t}). \quad (6)$$

In the following we will often use the fact that, for $a, b \in [0, 1]$, if $ab + (1 - a)(1 - b) \geq 0.5$, then either $a, b \in [0, 0.5]$ or $a, b \in [0.5, 1]$. By IH, since accuracies are always at least 0.5, it holds that $u_i(\mathbf{d}) \geq q_i \geq 0.5$ and by Equation (5) we have $q_s \cdot p_{i,s} + (1 - q_s) \cdot (1 - p_{i,s}) \geq 0.5$ and hence $p_{i,s} \geq 0.5$. Analogously, Equation (3) implies that $p_{k,s} \geq 0.5$.

Furthermore, we use the fact that

$$p_{k,i} = p_{k,s}p_{s,i} + (1 - p_{k,s})(1 - p_{s,i}) + (-2(2x_k - 1) \cdot (2x_i - 1) \cdot (x_s - 1)x_s) \quad (7)$$

where $x_j = \mathbb{P}(\tau(j) = 1)$ for $j \in \{k, i, s\}$. Observe that, by the definition of utility in (2), the assumptions made on \mathbf{d} and \mathbf{d}' , and the fact that for $a, b \in [0, 1]$, if $ab + (1 - a)(1 - b) \geq 0.5$, then either $a, b \in [0, 0.5]$ or $a, b \in [0.5, 1]$. So we have that either $x_j \geq 0.5$ for $j \in \{k, i, s\}$, or $x_j \leq 0.5$ for $j \in \{k, i, s\}$. We work on the first case. The other case is symmetric. Let also $\gamma_{k,s,i} = -2(2x_k - 1) \cdot (2x_i - 1) \cdot (x_s - 1)x_s$. From the above it follows that $0.5 \geq \gamma_{k,i,s} \geq 0$. Furthermore, given that $p_{i,s} = p_{s,i} \geq 0.5$, we can also conclude that $p_{k,i} \geq 0.5$. Now by substituting

$$p_{k,s} = p_{k,i}p_{i,s} + (1 - p_{k,i})(1 - p_{i,s}) + \underbrace{(-2(2x_k - 1) \cdot (2x_s - 1) \cdot (x_i - 1)x_i)}_{\gamma_{k,i,s}}$$

in Equation (3), we obtain

$$u_k(\mathbf{d}) = (2p_{k,i} - 1) \underbrace{(2q_s p_{i,s} - q_s - p_{i,s} + 1)}_{u_i(\mathbf{d})} + 1 - p_{k,i} + \gamma_{k,i,s}(2q_s - 1). \quad (8)$$

Similarly, using the appropriate instantiation of (7) for x_j with $j \in \{k, i, t\}$, by substituting $p_{k,i} \cdot p_{i,t} + (1 - p_{k,i})(1 - p_{i,t}) + \gamma_{k,i,t}$ for $p_{k,t}$ in Equation (4) we obtain

$$u_k(\mathbf{d}') = (2p_{k,i} - 1) \cdot \underbrace{(2q_t p_{i,t} - q_t - p_{i,t} + 1)}_{u_i(\mathbf{d}')} + 1 - p_{k,i} + \gamma_{k,i,t}(2q_t - 1). \quad (9)$$

Now observe that, since $p_{k,i} \geq 0.5$ we have that $(2p_{k,i} - 1) \geq 0$. It remains to compare $\gamma_{k,i,s}(2q_s - 1)$ with $\gamma_{k,i,t}(2q_t - 1)$, showing the latter is greater than the former. Observe that both expressions have a positive sign. We use the fact that $ab + (1 - a)(1 - b) < cd + (1 - c)(1 - d)$ implies $cd > ab$ under the assumption that $a, b, c, d \in [0.5, 1]$. On the basis of this, and given that $u_i(\mathbf{d}') > u_i(\mathbf{d})$, we obtain that $q_s \cdot p_{i,s} < q_t \cdot p_{i,t}$ and therefore that $q_s \cdot x_s < q_t \cdot x_t$, from which we can conclude that

$$\underbrace{(-2(2x_k - 1) \cdot (2x_s - 1) \cdot (x_i - 1)x_i)}_{\gamma_{k,i,s}} \cdot (2q_s - 1) < \underbrace{(-2(2x_k - 1) \cdot (2x_t - 1) \cdot (x_i - 1)x_i)}_{\gamma_{k,i,t}} \cdot (2q_t - 1).$$

It follows that the assumption $u_i(\mathbf{d}') > u_i(\mathbf{d})$ (player i chose a best response that increased her utility) together with Equations (8) and (9) implies that $u_k(\mathbf{d}') > u_k(\mathbf{d})$ (and a fortiori that $u_k(\mathbf{d}') \geq u_k(\mathbf{d})$). We have therefore shown that if some player chooses a best response, the utility of other players does not decrease. This completes the proof. \square

Discussion The existence of NE in general delegation games remains an interesting open problem. It should be noted that the proof strategies of both Theorems 1 and 2 do not work in the general case. Without a clear dichotomy of

type it is not possible to assign delegations for all agents in an SCC (as we do in the proof of Theorem 1). And the key property upon which the proof of Theorem 2 hinges (that a best response of an agent does not decrease the utility of other agents) fails in the general case due to the presence of non-zero effort. Finally, it should also be observed that Theorem 2 (as well as Proposition 1) essentially depend on the assumption that types are *independent* random variables. If this is not the case (e.g., because voters' preferences are correlated), delegation chains can become undesirable.

Example 1. Consider the following example with agents 1, 2 and 3. The probability distribution \mathbb{P} is defined as $\mathbb{P}(\tau(1) = 1 \wedge \tau(2) = 1 \wedge \tau(3) = 0) = 0.45$, $\mathbb{P}(\tau(1) = 0 \wedge \tau(2) = 1 \wedge \tau(3) = 1) = 0.45$, and $\mathbb{P}(\tau(1) = 1 \wedge \tau(2) = 1 \wedge \tau(3) = 1) = 0.1$. Consequently, $p_{1,2} = 0.55$, $p_{2,3} = 0.55$, and $p_{1,3} = 0.1$. Let us assume that the agents' accuracies are $q_1 = 0.5001$, $q_2 = 0.51$, and $q_3 = 0.61$. A delegation from agent 1 to 2 is locally positive as $q_2 \cdot p_{1,2} + (1 - q_2) \cdot (1 - p_{1,2}) = 0.501 > q_1$. Furthermore, a delegation from 2 to 3 is locally positive as $q_3 \cdot p_{2,3} + (1 - q_3) \cdot (1 - p_{2,3}) = 0.511 > q_2$. However, the resulting delegation from 1 to 3 is not positive since $q_3 \cdot p_{1,3} + (1 - q_3) \cdot (1 - p_{1,3}) = 0.412$.

Quality of Equilibria in Delegation Games

In delegation games players are motivated to maximize the tradeoff between the accuracy they acquire and the effort they spend for it. A natural measure for the quality of a delegation profile is, therefore, how accurate or informed a random voter becomes as a result of the delegations in the profile, that is, the average accuracy (i.e., $\bar{q}^*(\mathbf{d}) = \frac{1}{n} \sum_{i \in N} q_i^*(\mathbf{d})$) players enjoy in that profile. One can also consider the utilitarian social welfare $\text{SW}(\mathbf{d}) = \sum_{i \in N} u_i(\mathbf{d})$ of a delegation profile \mathbf{d} . This relates to average accuracy as follows: $\bar{q}^*(\mathbf{d}) = \frac{\text{SW}(\mathbf{d}) + \sum_{i|d(i)=i} e_i}{n}$.

It can immediately be noticed that equilibria do not necessarily maximize average accuracy. On the contrary, in the following example NE yields an average accuracy of close to 0.5, whereas an average accuracy of almost 1 is achievable.

Example 2. Consider an n -player delegation game where all players have the same type and $(i, j) \in R$ for all j and all $i > 1$, i.e., player 1 is a sink in R and cannot delegate to anyone, but all others can delegate to everyone. Further, we have $e_1 = 0$ and $e_i = 0.5 - \epsilon$ for $i \geq 2$. The respective accuracies are $q_1 = 0.5 + 2\epsilon$ and $q_i = 1$. If player $i \geq 2$ does not delegate, her utility is $0.5 + \epsilon$. Hence, it is always more desirable to delegate to player 1 (which yields a utility of $0.5 + 2\epsilon$ for i). Consider now the profiles in which all players delegate to player 1 (either directly or transitively). Player 1 can only vote directly (with utility $0.5 + 2\epsilon$). All such profiles are NE with average accuracy $0.5 + 2\epsilon$. If, however, some player $j \geq 2$ chose to vote herself, all players (except 1) would delegate to j thereby obtaining an average accuracy of $1 - \frac{0.5 - 2\epsilon}{n}$, which converges to 1 for $n \rightarrow \infty$. This is not a NE, as j could increase her utility by delegating to 1.

The findings of the example can be made more explicit by considering a variant of the price of anarchy for delegation

games, based on the above notion of average accuracy. That is, for a given delegation game G , the price of anarchy (PoA) of G is given by $PoA(G) = \frac{\max_{\mathbf{d} \in N^n} \bar{q}^*(\mathbf{d})}{\min_{\mathbf{d} \in NE(G)} \bar{q}^*(\mathbf{d})}$, where $NE(G)$ denotes the set of pure-strategy NE of G .

Fact 1. *PoA is bounded below by 1 and above by 2.*

An informative performance metrics for liquid democracy is the difference between the group accuracy after delegations versus the group accuracy achieved by direct voting. This measure, called *gain*, was introduced and studied by Kahng, Mackenzie, and Procaccia (2018). Here we adapt it to our setting as follows: $G(G) = (\min_{\mathbf{d} \in NE(G)} \bar{q}^*(\mathbf{d})) - \bar{q}$ where $\bar{q} = \bar{q}^*(\langle 1, \dots, n \rangle)$. That is, the gain in the delegation game G is the average accuracy of the worst NE minus the average accuracy of the profile in which no voter delegates. It turns out that the full performance range is possible:

Fact 2. *G is bounded below by -0.5 and above by 0.5 .*

The above bounds for PoA and gain provide only a very partial picture of the performance of liquid democracy when modeled as a delegation game. The next section provides a more fine-grained perspective on the effects of delegations.

Simulations

We simulate the delegation game described above in a variety of settings. We restrict ourselves to homogeneous games. This allows us to relate our results to those of Kahng, Mackenzie, and Procaccia (2018). Our experiments serve to both verify and extend the theoretical results of the previous section. In particular we simulate the best response dynamics employed in the proof of Theorem 2 and show that these dynamics converge even in the setting with effort, which we could not establish analytically. In addition, we investigate the dynamics of a one-shot game scenario, in which all agents need to select their proxy simultaneously at once.

Setup We generate graphs of size $N = 250$ of each of the four topologies *random*, *regular*, *small world*, and *scale free*, for different average degrees, while ensuring that the graph is connected. Agents’ accuracy and effort are initialized randomly with $q_i \in [0.5, 1]$ and $q_i - e_i \geq 0.5$. We average results over 2500 simulations for each setting (25 randomly generated graphs \times 100 initializations). Agents correctly observe their own accuracy and effort, and the accuracy of their neighbors. The game is homogeneous, so proximities are 1.

Each agent i selects from her neighborhood $R(i)$ (which includes i herself) the agent j that maximizes her expected utility following Equation 2. We compare two scenarios. The *iterated best response* scenario follows the procedure of the proof of Theorem 2, in which agents sequentially update their proxy to best-respond to the current profile \mathbf{d} using knowledge of their neighbors’ accuracy $q_i^*(\mathbf{d})$. In the *one-shot game* scenario all agents choose their proxy only once, do so simultaneously, and based only on their neighbors’ accuracy. The latter setup captures more closely the epistemic limitations that agents face in liquid democracy.

Iterated Best Response Dynamics

These experiments complement our existence theorems. They offer insights into the effects of delegations on average

Degree	4	8	12	16	20	24
BR updates (effortless)	298.1 (18.2)	261.7 (11.1)	254.2 (6.9)	251.6 (4.5)	250.6 (3.3)	250.0 (2.6)
Full passes (effortless)	3.6 (0.5)	3.0 (0.1)	2.9 (0.2)	2.8 (0.4)	2.7 (0.5)	2.5 (0.5)
BR updates (with effort)	294.7 (18.4)	259.4 (10.6)	252.9 (6.6)	250.9 (4.8)	250.2 (4.9)	249.9 (4.3)
Full passes (with effort)	3.6 (0.5)	3.0 (0.3)	2.8 (0.6)	2.6 (0.8)	2.4 (0.9)	2.4 (1.0)

Table 2: Total number of best response updates by individual agents and corresponding full passes over the network required for convergence. Reported are the mean (std.dev.) over all network types. *Note:* not all agents update their delegation at each full pass, but any single update requires an additional pass to check whether the best response still holds.

Degree	4	8	12	16	20	24
$\max_j q_j$	0.8908	0.8908	0.8904	0.8909	0.8904	0.8910
$\bar{q}^*(\mathbf{d})$	0.8906	0.8903	0.8897	0.8899	0.8890	0.8893

Table 3: Comparing the maximum accuracy across all agents and the mean accuracy under delegation \mathbf{d} for different network degrees, averaged across network types. The differences are statistically significant (paired t-test, $p = 0.05$).

voter’s accuracy in equilibrium, and on the effects of different network structures on how such equilibria are achieved.

We initialize $q_i \sim \mathcal{N}(0.75, 0.05)$ and first investigate the case in which $e_i = 0$ for all i (effortless voting). Across all combinations of network types and average degrees ranging from 4 to 24, we find that the best response dynamics converges, as predicted by Theorem 2, and does so optimally with $d_i^* = \arg \max_j q_j$ for all i . We observe minimal differences between network types, but see a clear inverse relation between average degree and the number of iterations required to converge (Table 2, top). Intuitively, more densely connected networks facilitate agents in identifying their optimal proxies.

We accumulate results across all network types and compare the effortless setting to the case in which effort is taken into account. When we include effort $e_i \sim \mathcal{N}(0.025, 0.01)$, we still observe convergence in all cases and, interestingly, the number of iterations required does not change significantly (Table 2, bottom). Although the process no longer results in an optimal equilibrium, each case still yields a single guru j with $q_j \approx \max_k q_k$ (less than 1% error) for all $k \in N$. In this scenario, the inclusion of effort means that a best response update of agent i no longer guarantees non-decreasing accuracy and utility for all other agents, which was a key property in the proof of Theorem 2. This effect becomes stronger as the average network degree increases, and as a result higher degree networks allow a greater discrepancy between the maximal average accuracy achievable and the average accuracy obtained at stabilization (Table 3).

In lower degree graphs (e.g. degree 4) we further observe differences in convergence speed between the four different network types which coincide with differences between

the average path lengths in those graphs: a shorter average distance between nodes yields a lower number of best response updates. This is intuitive, as larger distances between nodes mean longer delegation chains, but we have not yet conducted statistical tests to verify this hypothesis.

One-Shot Delegation Games

Here we study one-shot interactions in a delegation game: all agents select their proxy (possibly themselves) simultaneously among their neighbors; no further response is possible. This contrasts the previous scenario in which agents could iteratively improve their choice based on the choices of others. While Kahng, Mackenzie, and Procaccia (2018) study a probabilistic model, we instead assume that agents deterministically select as proxy the agent $j \in R(i)$ that maximizes their utility, as above. We compare \bar{q} and \bar{q}^* (the average network accuracy without and with delegation, respectively), as well as the probability of a correct majority vote under both direct democracy P_D and liquid democracy P_L where gurus carry as weight the number of agents for whom they act as proxy. The difference $P_L - P_D$ is again similar to the notion of *gain* (Kahng, Mackenzie, and Procaccia 2018). In line with Condorcet’s jury theorem (see for instance Grofman, Owen, and Feld 1983) $P_D \rightarrow 1$ as $N \rightarrow \infty$, and indeed for $N = 250$ we obtain $P_D \approx 1$.

First we again look at the effortless setting. Figure 1 (top) shows both metrics for the four different network types and for different average degrees. We observe that while \bar{q}^* (**d**) increases as the network degree increases (and in fact is always higher than \bar{q} without delegation), the probability of a correct majority outcome, P_L , simultaneously decreases. This confirms the analysis of Kahng, Mackenzie, and Procaccia (2018). We also observe that the number of gurus decreases exponentially as the degree increases (Figure 2, left). Simply put, giving all voting weight to a small group of gurus increases the chance of an incorrect majority vote, assuming that gurus have a less than perfect accuracy.

When we include effort (Figure 1, bottom), thereby moving away from the model of Kahng, Mackenzie, and Procaccia (2018), we observe a drastic decrease in average network accuracy combined with a lower probability of a correct majority outcome under liquid democracy, with both decreasing as network degree increases. The main reason is the existence of delegation cycles in this case. This contrasts the best response setting above where agents could iteratively reconsider their choice of proxy and thus avoid cycles. Now, even with relatively low effort (mean 0.025), up to *half* of all agents are stuck in a cycle (and thereby fail to pass information about their type) when the degree increases. This confirms results on the probability of delegation cycles from Christoff and Grossi (2017) and stresses the importance of cycle resolution in concrete implementations of liquid democracy.

Figure 1 further highlights differences between the four network types. Scale free networks yield a lower probability of a correct majority outcome across all degrees, as well as a larger number of gurus with a lower average accuracy and longer delegation chains (Figure 2, right). Intuitively, this indicates that one-shot interactions in scale free networks are




Figure 1: **Top:** Without effort. **Bottom:** With effort $e_i \sim \mathcal{N}(0.025, 0.01)$. **Left:** mean accuracy under liquid democracy, \bar{q}^* (**d**). The solid (dashed) line shows the mean (std. dev.) of the initial accuracy q ; the dotted line shows $\max_i q_i$. **Right:** probability of a correct majority vote under **d**.




Figure 2: Percentage of guru nodes under **d** (left) and mean distance between (non-guru) nodes and their gurus (right).

more likely to end up in a local optimum. In contrast, small world networks have short average distances and thus agents are more likely to be close to their optimal guru. Finally, our experiments highlight a key feature of liquid democracy: the trade-off between a reduction in (total) effort against a loss in average voting accuracy.

Conclusions and Future Work

The paper introduced delegation games as a first game-theoretic model of liquid democracy. Both our theoretical and experimental results showed that voting effort is a key ingredient for understanding how delegations form and what their effects are. Our empirical findings provided further insights into the influence of interaction networks on the quality of collective decisions in liquid democracy.

The paper opens up several directions of research. A gen-

eral NE existence theorem is the main open question. Our model can then be generalized in many directions, e.g.: by making agents' utility dependent on voting outcomes; by dropping the independence assumption on agents' types; or by assuming the voting mechanism has better than 0.5 accuracy in identifying the types of agents involved in cycles.

Acknowledgements

We are indebted to the anonymous reviewers of IJ-CAI/ECAI'18 and AAAI'19 for many helpful comments on earlier versions of this paper. We are also grateful to the participants of the LAMSADE seminar at Paris Dauphine University, and the THEMA seminar at University Cergy-Pontoise where this work was presented, for many helpful comments and suggestions. Daan Bloembergen has received funding in the framework of the joint programming initiative ERA-Net Smart Energy Systems' focus initiative Smart Grids Plus, with support from the European Union's Horizon 2020 research and innovation programme under grant agreement No 646039. Davide Grossi was partially supported by EPSRC under grant EP/M015815/1. Martin Lackner was supported by the European Research Council (ERC) under grant number 639945 (ACCORD) and by the Austrian Science Foundation FWF, grant P25518 and Y698.

References

Alger, D. 2006. Voting by proxy. *Public Choice* 126(1-2):1–26.

Barabási, A.-L., and Albert, R. 1999. Emergence of scaling in random networks. *science* 286(5439):509–512.

Blum, C., and Zuber, C. I. 2016. Liquid democracy: Potentials, problems, and perspectives. *Journal of Political Philosophy* 24(2):162–182.

Boella, G.; Francis, L.; Grassi, E.; Kistner, A.; Nitsche, A.; Noskov, A.; Sanasi, L.; Savoca, A.; Schifanella, C.; and Tsampoulatidis, I. 2018. WeGovNow: A map based platform to engage the local civic society. In *WWW '18 Companion: The 2018 Web Conference Companion*, 1215–1219. Lyon, France: International World Wide Web Conferences Steering Committee.

Boldi, P.; Bonchi, F.; Castillo, C.; and Vigna, S. 2011. Viscous democracy for social networks. *Communications of the ACM* 54(6):129–137.

Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

Brill, M. 2018. Interactive democracy. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 1183–1187. International Foundation for Autonomous Agents and Multiagent Systems.

Christoff, Z., and Grossi, D. 2017. Binary voting with delegable proxy: An analysis of liquid democracy. In *Proceedings of TARK'17*, volume 251, 134–150. EPTCS.

Cohensius, G.; Mannor, S.; Meir, R.; Meirom, E.; and Orda, A. 2017. Proxy voting for better outcomes. In *Proceedings*

of the 16th Conference on Autonomous Agents and Multi-Agent Systems, 858–866. International Foundation for Autonomous Agents and Multiagent Systems.

Dodgson, C. L. 1884. *The Principles of Parliamentary Representation*. Harrison and Sons.

Endriss, U. 2016. Judgment aggregation. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press.

Erdős, P., and Rényi, A. 1959. On random graphs, I. *Publicationes Mathematicae (Debrecen)* 6:290–297.

Gözl, P.; Kahng, A.; Mackenzie, S.; and Procaccia, A. D. 2018. The fluid mechanics of liquid democracy. In *WADE'18*, *arXiv:1808.01906*.

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

Grofman, B.; Owen, G.; and Feld, S. 1983. Thirteen theorems in search of truth. *Theory and Decision* 15:261–278.

Grossi, D., and Pigozzi, G. 2014. *Judgment Aggregation: A Primer*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.

Jackson, M. O. 2008. *Social and Economic Networks*. Princeton University Press.

Kahng, A.; Mackenzie, S.; and Procaccia, A. 2018. Liquid democracy: An algorithmic perspective. In *Proc. 32nd AAAI Conference on Artificial Intelligence (AAAI'18)*.

Kling, C. C.; Kunegis, J.; Hartmann, H.; Strohmaier, M.; and Staab, S. 2015. Voting behaviour and power in online democracy: A study of liquidfeedback in germany's pirate party. In *Ninth International AAAI Conference on Web and Social Media*.

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

Skowron, P.; Lackner, M.; Brill, M.; Peters, D.; and Elkind, E. 2017. Proportional rankings. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, 409–415.

Tullock, G. 1992. Computerizing politics. *Mathematical and Computer Modelling* 16(8-9):59–65.

Watts, D. J., and Strogatz, S. H. 1998. Collective dynamics of 'small-world' networks. *Nature* 393(6684):440.