

Liquid Democracy: An Analysis in Binary Aggregation and Diffusion*

Zoé Christoff* and Davide Grossi**

* Department of Philosophy, University of Bayreuth
zoe.christoff@gmail.com

** Department of Computer Science, University of Liverpool
d.grossi@liverpool.ac.uk

January 20, 2017

Abstract

The paper proposes an analysis of liquid democracy (or, delegable proxy voting) from the perspective of binary aggregation and of binary diffusion models. We show how liquid democracy on binary issues can be embedded into the framework of *binary aggregation with abstentions*, enabling the transfer of known results about the latter—such as impossibility theorems—to the former. This embedding also sheds light on the relation between delegation cycles in liquid democracy and the probability of collective abstentions, as well as the issue of individual rationality in a delegable proxy voting setting. We then show how liquid democracy on binary issues can be modeled and analyzed also as a specific process of dynamics of binary opinions on networks. These processes—called *Boolean DeGroot processes*—are a special case of the DeGroot stochastic model of opinion diffusion. We establish the convergence conditions of such processes and show they provide some novel insights on how the effects of delegation cycles and individual rationality could be mitigated within liquid democracy.

The study is a first attempt to provide theoretical foundations to the delegable proxy features of the liquid democracy voting system. Our analysis suggests recommendations on how the system may be modified to make it more resilient with respect to the handling of delegation cycles and of inconsistent majorities.

***Working paper:** The paper collects work presented at: *Dynamics in Logic IV*, TU Delft, November 2016; seminars at the Computer Science Departments of the University of Leicester and the University of Oxford, December 2016; *Dutch Social Choice Colloquium*, December 2016. The authors wish to thank the participants of the above workshops and seminars for many helpful suggestions. The authors wish also to thank Umberto Grandi for many insightful comments on an earlier version of this paper. Both authors acknowledge support for this research by EPSRC under grant EP/M015815/1.

Contents

1	Introduction	3
2	Binary Aggregation with Abstention	4
2.1	Opinions and Opinion Profiles	4
2.2	Aggregators	5
2.3	Agenda conditions	6
2.4	Properties of aggregators	7
2.5	Characterizing quota rules	8
2.6	Impossibility in Binary Aggregation with Abstentions	10
3	Liquid Democracy as Binary Aggregation	10
3.1	Binary Aggregation via Delegable Proxy	10
3.1.1	Proxy Opinions and Profiles	11
3.1.2	Delegation Graphs	11
3.1.3	Proxy Aggregators	12
3.2	Two Issues of Delegable Proxy	12
3.2.1	Cycles and Abstentions	12
3.2.2	Individual & Collective Rationality	14
3.3	Embedding in Binary Aggregation with Abstentions	15
3.3.1	One man—One vote	15
3.3.2	Impossibility	16
3.4	Section Summary	17
4	Liquid Democracy as Binary Opinion Diffusion	17
4.1	Binary aggregation and binary influence	18
4.1.1	DeGroot Processes and Opinion Diffusion	18
4.1.2	Boolean DeGroot processes	18
4.2	Convergence of BDPs	19
4.2.1	Preliminaries	20
4.2.2	Context: convergence in DeGroot processes	20
4.2.3	Two results	21
4.2.4	Liquid Democracy as a BDP	22
4.2.5	Cycles	23
4.3	Excursus: unanimity and 2-colorability	23
4.4	BDPs on logically interdependent issues	26
4.5	Section Summary	27
5	Fixpoint Logics for BDPs	28
5.1	Influence graphs as Kripke models	28
5.2	Modal μ -calculus	28
5.3	On the logic of convergence in BDPs	29
6	Conclusions	31
A	Binary aggregation (without abstention)	34
B	Relevant terminology from graph theory	35

1 Introduction

Liquid democracy [2] is a form of democratic decision-making considered to stand between direct and representative democracy. It has been used, advocated and popularized by local and even national parties (e.g., Demoex¹ in Sweden, and Piratenpartei² in Germany) to coordinate the behavior of party representatives in assemblies, as well as campaigns (e.g., Make Your Laws³ in the US). At its heart is voting via a delegable proxy, also called sometimes transitive proxy. For each issue submitted to vote, each agent can either cast its own vote, or it can delegate its vote to another agent—a proxy—and that agent can delegate in turn to yet another agent and so on. This differentiates liquid democracy from standard proxy voting [24, 27], where proxies cannot delegate their vote further. Finally, the agents that decided not to delegate their votes cast their ballots (e.g., under majority rule, or adaptations thereof), but their votes now carry a weight consisting of the number of all agents that, directly or indirectly, entrusted them with their vote.

Scientific context and contribution Analyses of standard (non-delegable) proxy voting from a social choice-theoretic perspective—specifically through the theory of spatial voting—have been put forth in [1] and [17]. Delegable proxy has not, to the best of our knowledge, been object of study so far, with the notable exception of [6] which focuses specifically on algorithmic aspects of a variant of liquid democracy (which the authors refer to as *viscous democracy*) with applications to recommender systems.

The objective of the paper is to provide a first analysis, via formal methods, of the liquid democracy voting system based on delegable proxy. This, we hope, should point to a number of future lines of research and stimulate further investigations into this and related systems.

Outline The paper starts in Section 2 by introducing some preliminaries on the theory of binary aggregation, which is the framework of reference for this study. It is then structured in two parts. This preliminary section presents also novel results on binary aggregation with abstentions. The first part (Section 3) studies voting in liquid democracy from the point of view of the delegation of voting power: we study delegable proxy aggregators using the machinery of binary and judgment aggregation. This allows us to shed novel light on some issues involved in the liquid democracy system, in particular: the issue of circular delegation, and the issue of individual irrationality when voting on logically interdependent issues. The second part (Sections 4 and 5) studies voting in liquid democracy as a very specific type of opinion diffusion on networks, whereby delegation is rather interpreted as the willingness to copy the vote of a trustee. We show that this perspective provides some interesting insights on how to address

¹demoex.se/en/

²www.piratenpartei.de

³www.makeyourlaws.org

the above mentioned issues of circular delegations and individual irrationality. Section 6 concludes the paper and outlines some on-going lines of research.

2 Binary Aggregation with Abstention

The formalism of choice for this paper is binary aggregation [14] with abstention.⁴ This preliminary section is devoted to its introduction.

2.1 Opinions and Opinion Profiles

A binary aggregation structure (*BA structure*) is a tuple $\mathcal{A} = \langle N, \mathbf{P}, \gamma \rangle$ where:

- $N = \{1, \dots, n\}$ is a non-empty finite set individuals s.t. $|N| = n \in \mathbb{N}$;
- $\mathbf{P} = \{p_1, \dots, p_m\}$ is a non-empty finite set of issues ($|\mathbf{P}| = m \in \mathbb{N}$), each represented by a propositional atom;
- $\gamma \in \mathcal{L}$ is an (integrity) constraint, where \mathcal{L} is the propositional language constructed by closing \mathbf{P} under a functionally complete set of Boolean connectives (e.g., $\{\neg, \wedge\}$).

An *opinion* function O is an assignment acceptance/rejection values (or, truth values) to the set of issues \mathbf{P} . Thus, $O(p) = \mathbf{0}$ (respectively, $O(p) = \mathbf{1}$) indicates that opinion O rejects (respectively, accepts) the issue p . Syntactically, the two opinions correspond to the truth of the literals p or $\neg p$. For $p \in \mathbf{P}$ we write $\pm p$ to denote one element from $\{p, \neg p\}$, and $\pm \mathbf{P}$ to denote $\bigcup_{p \in \mathbf{P}} \{p, \neg p\}$, which we will refer to as the *agenda* of \mathcal{A} . Allowing abstention in the framework of binary aggregation amounts to considering incomplete opinions: an *incomplete opinion* is a partial function from \mathbf{P} to $\{\mathbf{0}, \mathbf{1}\}$. We will study it as a function $O : \mathbf{P} \rightarrow \{\mathbf{0}, \mathbf{1}, *\}$ thereby explicitly denoting the undetermined value corresponding to abstention.

We say that the incomplete opinion of an agent i is *consistent* if the set of formulas $\{p \mid O_i(p) = \mathbf{1}\} \cup \{\neg p \mid O_i(p) = \mathbf{0}\} \cup \{\gamma\}$ can be extended to a model of γ (in other words, if the set is satisfiable). Intuitively, the consistency of an incomplete opinion means that the integrity constraint is consistent with i 's opinion on the issues she does not abstain about. We also say that an incomplete opinion is *closed* whenever the following is the case: *if* the set of propositional formulas $\{p \mid O_i(p) = \mathbf{1}\} \cup \{\neg p \mid O_i(p) = \mathbf{0}\} \cup \{\gamma\}$ logically implies p (respectively, $\neg p$), *then* $O_i(p) = 1$ (respectively, $O_i(p) = 0$). That is, individual opinions are closed under logical consequence or, in other words, agents cannot abstain on issues whose acceptance or rejection is dictated by their expressed opinions on other issues. The set of incomplete opinions is denoted \mathcal{O}^* and the set of consistent and closed incomplete opinions \mathcal{O}_c^* . As the latter are the opinions we are interested in, we will often refer to them simply as individual opinions.

⁴The standard framework of binary aggregation without abstention is sketched in the appendix for ease of reference.

An *opinion profile* $\mathbf{O} = (O_1, \dots, O_n)$ records the opinion, on the given set of issues, of every individual in N . Given a profile \mathbf{O} the i^{th} projection \mathbf{O} is denoted O_i (i.e., the opinion of agent i in profile \mathbf{O}). Let us introduce some more notation. We also denote by $\mathbf{O}(p) = \{i \in N \mid O_i(p) = \mathbf{1}\}$ the set of agents accepting issue p in profile \mathbf{O} , by $\mathbf{O}(\neg p) = \{i \in N \mid O_i(p) = \mathbf{0}\}$ and by $\mathbf{O}(\pm p) = \mathbf{O}(p) \cup \mathbf{O}(\neg p)$ the set of non-abstaining agents. We write $\mathbf{O} =_{-i} \mathbf{O}'$ to denote that the two profiles \mathbf{O} and \mathbf{O}' are identical, except for possibly the opinion of voter i .

2.2 Aggregators

Given a BA structure \mathcal{A} , an *aggregator* (for \mathcal{A}) is a function $F : (\mathcal{O}_c^*)^N \rightarrow \mathcal{O}^*$, mapping every profile of individual opinions to one collective (possibly incomplete) opinion.⁵ $F(\mathbf{O})(p)$ denotes the outcome of the aggregation on issue p . The benchmark aggregator is the *issue-by-issue strict majority rule* (**maj**), which accepts an issue if and only if the majority of the non-abstaining voters accept that issue:

$$\text{maj}(\mathbf{O})(p) = \begin{cases} \mathbf{1} & \text{if } |\mathbf{O}(p)| > |\mathbf{O}(\neg p)| \\ \mathbf{0} & \text{if } |\mathbf{O}(\neg p)| > |\mathbf{O}(p)| \\ * & \text{otherwise} \end{cases} \quad (1)$$

We will refer to this rule simply as ‘majority’.

Majority can be thought of as a quota rule. In general, quota rules in binary aggregation with abstention are of the form: accept when the proportion of *non-abstaining* individuals accepting is above the acceptance-quota, reject when the proportion of *non-abstaining* individuals is above the rejection-quota, and abstain otherwise.⁶

Definition 1 (Quota rules). *Let \mathcal{A} be an aggregation structure. A quota rule (for \mathcal{A}) is defined as follows, for any issue $p \in \mathbf{P}$, and any opinion profile $\mathbf{O} \in \mathcal{O}^*$:⁷*

$$F(\mathbf{O})(p) = \begin{cases} \mathbf{1} & \text{if } |\mathbf{O}(p)| \geq \lceil q_1(p) \cdot |\mathbf{O}(\pm p)| \rceil \\ \mathbf{0} & \text{if } |\mathbf{O}(\neg p)| \geq \lceil q_0(p) \cdot |\mathbf{O}(\pm p)| \rceil \\ * & \text{otherwise} \end{cases} \quad (2)$$

where for $x \in \{\mathbf{0}, \mathbf{1}\}$, q_x is a function $q_x : \mathbf{P} \rightarrow (0, 1] \subset \mathbb{Q}$ assigning a positive rational number smaller or equal to 1 to each issue, and such that, for each

⁵It is therefore worth stressing that, in this paper, we study aggregators that are resolute (that is, output exactly one value), even though they allow for collective abstention.

⁶There are several ways to think of quota rules in the presence of abstentions. Instead of a quota being a proportion of non-abstaining agents, one could for instance define rules with absolute quotas instead: accept when at least n agents accept, independently of how many agents do not abstain. In practice, voting rules with abstention are often a combination of those two ideas: accept an issue if a big enough proportion of the population does not abstain, and if a big enough proportion of those accept it.

⁷The definition uses the ceiling function $\lceil x \rceil$ denoting the smallest integer larger than x .

$p \in \mathbf{P}$:

$$q_x(p) > 1 - q_{(1-x)}(p), \quad (3)$$

A quota rule is called: uniform if, for all $p_i, p_j \in \mathbf{P}$, $q_x(p_i) = q_x(p_j)$; it is called symmetric if, for all $p \in \mathbf{P}$, $q_1(p) = q_0(p)$.

Notice that the definition excludes trivial quota.⁸ It should also be clear that, by (3) the above defines an aggregator of type $(\mathcal{O}_c^*)^N \rightarrow \mathcal{O}^*$ as desired.⁹ Notice also that if the rule is symmetric, then (3) forces $q_x > \frac{1}{2}$.

Example 1. The majority rule (1) is a uniform and symmetric quota rule where q_1 and q_0 are set to meet the equation $\lceil q_1(p) \cdot |\mathbf{O}(\pm p)| \rceil = \lceil q_0(p) \cdot |\mathbf{O}(\pm p)| \rceil = \left\lceil \frac{|\mathbf{O}(\pm p)|+1}{2} \right\rceil$, for any issue p and profile \mathbf{O} . This is achieved by setting the quota as $\frac{1}{2} < q_1, q_0 \leq \frac{1}{2} + \frac{1}{|N|} = \frac{|N|+1}{2|N|}$. More precisely one should therefore consider maj as a class of quota rules yielding the same collective opinions.

Example 2. The uniform and symmetric unanimity rule is defined by setting $q_1 = q_0 = 1$. A natural uniform but asymmetric variant of unanimity can be obtained by setting $q_1 = 1$ and $q_0 = \frac{1}{|N|}$.

Let us finally note an important difference between quota rules in binary aggregation with abstentions vs. without abstentions. In a framework without abstentions quota rules are normally defined by a unique acceptance quota q^1 , the rejection quota being uniquely determined as $q^0 = 1 - q^1$. As a consequence, the majority rule, when $|N|$ is odd, is the only unbiased quota rule in the standard framework. This is no longer the case when abstentions are considered. A novel characterization of the majority rule will be given in Section 2.5.

2.3 Agenda conditions

Definition 2 (simple/evenly negatable agenda). An agenda $\pm\mathbf{P}$ is said to be simple if there exists no set $X \subseteq \pm\mathbf{P}$ such that: $|X| \geq 3$, and X is minimally γ -inconsistent, that is:

- X is inconsistent with γ
- For all $Y \subset X$, Y is consistent with γ (or, γ -consistent).

An agenda is said to be evenly negatable if there exists a minimal γ -inconsistent set $X \subseteq \pm\mathbf{P}$ such that for a set $Y \subseteq X$ of even size, $X \setminus Y \cup \{\neg p \mid p \in Y\}$ is γ -consistent. It is said to be path-connected if there exists $p_1, \dots, p_n \in \pm\mathbf{P}$ such that $p_1 \models^c p_2, \dots, p_{n-1} \models^x p_n$ where $p_i \models^c p_{i+1}$ (conditional entailment) denotes that there exists $X \subseteq \pm\mathbf{P}$, which is γ -consistent with both p_i and $\neg p_{i+1}$, and such that $\{p\} \cup X \cup \{\gamma\}$ logically implies p_{i+1} .

⁸Those are quotas with value 0 (always met) or > 1 (never met). Restricting to non-trivial quota is not essential but simplifies our exposition.

⁹What needs to be avoided here is that both the acceptance and rejection quota are set so low as to make the rule output both the acceptance and the rejection of a given issue

We refer the reader to [18, Ch. 2] for a detailed exposition of the above conditions. We provide just a simple illustrative example.

Example 3. Let $\mathbf{P} = \{p, q, r\}$ and let $\gamma = (p \wedge q) \rightarrow r$. $\pm\mathbf{P}$ is not simple. The set $\{p, q, \neg r\} \subseteq \pm\mathbf{P}$ is inconsistent with γ , but none of its subsets is.

2.4 Properties of aggregators

We start by recalling some well-known properties of aggregators from the judgment aggregation literature, adapted to the setting with abstention:

Definition 3. Let \mathcal{A} be an aggregation structure. An aggregator $F : (\mathcal{O}_c^*)^N \rightarrow \mathcal{O}^*$ is said to be:

unanimous iff for all $p \in \mathbf{P}$, for all profiles \mathbf{O} and all $x \in \{0, 1, *\}$: if for all $i \in N$, $O_i(p) = x$, then $F(\mathbf{O})(p) = x$. I.e., if everybody agrees on a value, that value is the collective value.

anonymous iff for any bijection $\mu : N \rightarrow N$, $F(\mathbf{O}) = F(\mathbf{O}^\mu)$, where $\mathbf{O}^\mu = \langle O_{\mu(1)}, \dots, O_{\mu(n)} \rangle$. I.e., permuting opinions among individuals does not affect the output of the aggregator.

p-dictatorial iff there exists $i \in N$ (the p -dictator) s.t. for any profile \mathbf{O} , and all $x \in \{0, 1\}$, $O_i(p) = x$ iff $F(\mathbf{O})(p) = x$. I.e., there exists an agent whose definite opinion determines the group's definite opinion on p . If F is p -dictatorial, with the same dictator on all issues $p \in \mathbf{P}$, then it is called **dictatorial**.

p-oligarchic iff there exists $C \subseteq N$ (the p -oligarchs) s.t. $C \neq \emptyset$ and for any profile \mathbf{O} , and any value $x \in \{0, 1\}$, $F(\mathbf{O})(p) = x$ iff $O_i(p) = x$ for all $i \in C$. I.e., there exists a group of agents whose definite opinions always determine the group's definite opinion on p . If F is p -oligarchic, with the same oligarchs on all issues $p \in \mathbf{P}$, then it is called **oligarchic**.

monotonic iff, for all $p \in \mathbf{P}$ and all $i \in N$: for any profiles \mathbf{O}, \mathbf{O}' , if $\mathbf{O} =_{-i} \mathbf{O}'$: (i) if $O_i(p) \neq \mathbf{1}$ and $O'_i(p) \in \{1, *\}$, then: if $F(\mathbf{O})(p) = \mathbf{1}$, then $F(\mathbf{O}')(p) = \mathbf{1}$; and (ii) if $O_i(p) \neq \mathbf{0}$ and $O'_i(p) \in \{0, *\}$, then: if $F(\mathbf{O})(p) = \mathbf{0}$, then $F(\mathbf{O}')(p) = \mathbf{0}$. I.e., increasing support for a definite collective opinion does not change that collective opinion.

independent iff, for all $p \in \mathbf{P}$, for any profiles \mathbf{O}, \mathbf{O}' : if for all $i \in N$, $O_i(p) = O'_i(p)$, then $F(\mathbf{O})(p) = F(\mathbf{O}')(p)$. I.e., the collective opinion on each issue is determined only by the individual opinions on that issue.

neutral iff, for all $p, q \in \mathbf{P}$, for any profile \mathbf{O} : if for all $i \in N$, $O_i(p) = O_i(q)$, then $F(\mathbf{O})(p) = F(\mathbf{O})(q)$. I.e., all issues are aggregated in the same manner.

systematic iff it is neutral and independent. I.e., the collective opinion on issue p depends only on the individual opinions on this issue.

responsive iff for all $p \in \mathbf{P}$, there exist profiles \mathbf{O}, \mathbf{O}' such that $F(\mathbf{O})(p) = \mathbf{1}$ and $F(\mathbf{O}')(p) = \mathbf{0}$. I.e., the rule allows for an issue to be accepted for some profile, and rejected for some other.

unbiased iff for all $p \in \mathbf{P}$, for any profiles \mathbf{O}, \mathbf{O}' : if for all $i \in N$, $O_i(p) = \mathbf{1}$ iff $O'_i(p) = \mathbf{0}$ (we say that \mathbf{O}' is the “reversed” profile of \mathbf{O}), then $F(\mathbf{O})(p) = \mathbf{1}$ iff $F(\mathbf{O}')(p) = \mathbf{0}$. I.e., reversing all and only individual opinions on an issue p (from acceptance to rejection and from rejection to acceptance) results in reversing the collective opinion on p .

rational iff for any profile \mathbf{O} , $F(\mathbf{O})$ is consistent and closed. I.e., the aggregator preserves the constraints on individual opinions.

Example 4. It is well-known that majority is not rational in general. The standard example is provided by the so-called discursive dilemma, represented by the BA structure $\langle \{1, 2, 3\}, \{p, q, r\}, r \leftrightarrow (p \wedge q) \rangle$. The profile consisting of $O_1 \models p \wedge q \wedge r$, $O_2 \models p \wedge \neg q \wedge \neg r$, $O_3 \models \neg p \wedge q \wedge \neg r$, returns an inconsistent majority opinion $\text{maj}(\mathbf{O}) \models p \wedge q \wedge \neg r$ (cf. [18]).

Finally, let us defined also the following property. The **undecisiveness** of an aggregator F on issue p for a given aggregation structure is defined as the number of profiles which result in collective abstention on p :

$$u(F)(p) = |\{\mathbf{O} \in \mathcal{O}_c^* \mid F(\mathbf{O})(p) = *\}|. \quad (4)$$

2.5 Characterizing quota rules

As a typical example, consider the aggregator **maj**: it is unanimous, anonymous, monotonic, systematic, responsive and unbiased, but, as mentioned above, it is not rational in general. However, it can be shown (cf. [18, 3.1.1]) that aggregation by majority is collectively rational under specific assumptions on the constraint:

Fact 1. Let \mathcal{A} be a BA structure with a simple agenda. Then **maj** is rational.

Proof. If the agenda $\pm\mathbf{P}$ is simple, then all minimally inconsistent sets have cardinality 2, that is, are of the form $\{\varphi, \neg\varphi\}$ such that $\varphi \models \neg\varphi$ for $\varphi, \psi \in \mathbf{P}$. W.l.o.g. assume $\varphi = p_i$ and $\psi = p_j$. Suppose towards a contradiction that there exists a profile \mathbf{O} such that $\text{maj}(\mathbf{O})$ is inconsistent, that is, $\text{maj}(\mathbf{O})(p_i) = \text{maj}(\mathbf{O})(p_j) = \mathbf{1}$, and $\varphi \models \neg\psi$. By the definition of **maj** (1) it follows that $|\mathbf{O}(p_i)| > |\mathbf{O}(\neg p_i)|$ and $|\mathbf{O}(p_j)| > |\mathbf{O}(\neg p_j)|$. Since $p_i \models \neg p_j$ by assumption, and since individual opinions are consistent and closed, $|\mathbf{O}(\neg p_j)| \geq |\mathbf{O}(p_i)|$ and $|\mathbf{O}(\neg p_i)| \geq |\mathbf{O}(p_j)|$. From the fact that $|\mathbf{O}(p_i)| > |\mathbf{O}(\neg p_i)|$ we can thus conclude that $|\mathbf{O}(\neg p_j)| > |\mathbf{O}(p_j)|$. Contradiction. \square

May’s theorem [23] famously shows that for preference aggregation, the majority rule is in fact the *only* aggregator satisfying a specific bundle of desirable properties. A corresponding characterization of the majority rule is given in

judgment aggregation *without abstention*: when the agenda is simple, the majority rule is the only aggregator which is rational, anonymous, monotonic and unbiased [18, Th. 3.2]. We give below a novel characterization theorem, which takes into account the possibility of abstention (both at the individual and at the collective level). To the best of our knowledge, this is the first result of this kind in the literature on judgment and binary aggregation with abstention.

We first prove the following lemma:

Lemma 1. *Let F be a uniform and symmetric quota rule for a given \mathcal{A} . The following holds: $\frac{1}{2} < q_1 = q_0 \leq \frac{|N|+1}{2|N|}$ if and only if $F = \arg \min_G u(G)(p)$, for all $p \in \mathbf{P}$.*

Proof. The claim is proven by the following series of equivalent statements. (a) A uniform and symmetric quota rule F has quota such that $\frac{1}{2} < q_1 = q_0 \leq \frac{|N|+1}{2|N|}$. (b) A uniform and symmetric quota rule F has quota such that $\lceil q_1(p)|\mathbf{O}(\pm p)| \rceil = \lceil q_0(p)|\mathbf{O}(\pm p)| \rceil = \left\lceil \frac{|\mathbf{O}(\pm p)|+1}{2} \right\rceil$ for any profile \mathbf{O} and issue p . (c) For any $\mathbf{O} \in \mathcal{O}_c^*$ and $p \in \mathbf{P}$, $u(\mathbf{O})(p) = *$ if and only if $\mathbf{O}(p) = \mathbf{O}(\neg p)$, that is, an even number of voters vote and the group is split in half. (d) $F = \arg \min_G u(G)(p)$, for all $p \in \mathbf{P}$. \square

That is, the quota rule(s) corresponding to the majority rule (Example 1) is precisely the rule that minimizes undecisiveness.

We can now state and prove the characterization result:

Theorem 1. *Let $F : (\mathcal{O}_c^*)^N \rightarrow \mathcal{O}^*$ be an aggregator for a given \mathcal{A} . The following holds:*

1. *F is a quota rule if and only if it is anonymous, independent, monotonic, and responsive;*
2. *F is a uniform quota rule if and only if it is a neutral quota rule;*
3. *F is a symmetric quota rule if and only if it is an unbiased quota rule;*
4. *F is the majority rule \mathbf{maj} if and only if it is a uniform symmetric quota rule which minimizes undecisiveness.*

Proof. Claim 1 Left-to-right: Easily checked. Right-to-left: Let F be an anonymous, independent, monotonic, and responsive aggregator. By *anonymity* and *independence*, for any $p \in \mathbf{P}$, and any $\mathbf{O} \in \mathcal{O}_c^*$, the only information determining the value of $F(\mathbf{O})(p)$ are the integers $|\mathbf{O}(p)|$ and $|\mathbf{O}(\neg p)|$.

By *responsiveness*, there exists a non-empty set of profiles $S^1 = \{\mathbf{O} \in \mathcal{O}^* | F(\mathbf{O})(p) = \mathbf{1}\}$. Pick \mathbf{O} to be any profile in S^1 with a minimal value of $\frac{|\mathbf{O}(p)|}{|\mathbf{O}(\pm p)|}$ and call this value q_1 . Now let \mathbf{O}' be any profile such that $\mathbf{O}' = \neg_i \mathbf{O}$ and $\frac{|\mathbf{O}'(p)|}{|\mathbf{O}'(\pm p)|} > q_1$. This implies that $O_i(p) = \mathbf{0}$ and $O'_i(p) = \mathbf{1}$. By *monotonicity*, it follows that $F(\mathbf{O}')(p) = \mathbf{1}$. By iterating this argument a finite number of times we conclude that whenever $\frac{|\mathbf{O}(p)|}{|\mathbf{O}(\pm p)|} \geq q_1$, we have that $F(\mathbf{O})(p) = \mathbf{1}$. Given

that q_1 was defined as a minimal value, we conclude also that if $F(\mathbf{O})(p) = 1$, then $\frac{\mathbf{O}(p)}{\mathbf{O}(p^\pm)} \geq q_1$. The argument for q_0 is identical.

Claims 2 & 3 follow straightforwardly from the definitions of uniform quota rule (Definition 1) and of neutrality (Definition 3) and, respectively, from the definitions of symmetric quota rules (Definition 1) and of unbiasedness (Definition 3).

Claim 4 Left-to-right. Recall that **maj** is defined by quota $\frac{1}{2} < q_1 = q_0 \leq \frac{1}{2} + \frac{1}{|N|}$ (Example 1). It is clear that **maj** is uniform and symmetric. The claim then follows by Lemma 1. Right-to-left. By Lemma 1 if an aggregator minimizes undecisiveness then its quota are set as $\frac{1}{2} < q_1 = q_0 \leq \frac{1}{2} + \frac{1}{|N|}$. These quota define **maj** (Example 1). \square

By the above theorem and Fact 1, it follows that, on simple agendas, majority is the only rational aggregator which is also responsive, anonymous, systematic and monotonic.

2.6 Impossibility in Binary Aggregation with Abstentions

The following is a well-know impossibility result concerning binary aggregation with abstentions:

Theorem 2 ([11, 10]). *Let \mathcal{A} be a BA structure whose agenda is path connected and evenly negatable. Then if an aggregator $F : (\mathcal{O}_c^*)^N \rightarrow \mathcal{O}^*$ is independent, unanimous and collectively rational, then it is oligarchic.*

We will use this result to illustrate how impossibility results from binary aggregation with abstentions apply to delegable proxy voting on binary issues.

3 Liquid Democracy as Binary Aggregation

In this section we provide an analysis of liquid democracy by embedding it in the theory of binary aggregation with abstentions presented in the previous section. To the best of our knowledge, this is the first attempt at providing an analysis of delegable proxy voting using social-choice theoretic tools, with the possible exception of [16].

In what follows we will often refer to delegable proxy voting/aggregation simply as proxy voting/aggregation.

3.1 Binary Aggregation via Delegable Proxy

In binary aggregation with proxy, agents either express an acceptance/rejection opinion or *delegate* such opinion to a different agent.

3.1.1 Proxy Opinions and Profiles

Let a BA structure \mathcal{A} be given and assume for now that $\gamma = \top$, that is, all issues are logically independent. An opinion $O : \mathbf{P} \rightarrow \{\mathbf{0}, \mathbf{1}\} \cup N$ is an assignment of either a truth value or another agent to each issue in \mathbf{P} , such that $O_i(p) \neq i$ (that is, self-delegation is not an expressible opinion). We will later also require proxy opinion to be individually rational, in a precise sense (Section 3.2.2). For simplicity we are assuming that abstention is not a feasible opinion in proxy voting, but that is an assumption that can be easily lifted in what follows.

We call functions of the above kind *proxy opinions* to distinguish them from standard (binary) opinions, and we denote by \mathcal{P} the set of all proxy opinions, \mathcal{P}_c the set of all consistent proxy opinions, \mathcal{P}^N being the set of all proxy profiles.

3.1.2 Delegation Graphs

Each profile \mathbf{O} of proxy opinions (*proxy profile* in short) induces a delegation graph $G^{\mathbf{O}} = \langle N, \{R_p\}_{p \in \mathbf{P}} \rangle$ where for $i, j \in N$:

$$iR_pj \iff \begin{cases} O_i(p) = j & \text{if } i \neq j \\ O_i(p) \in \{\mathbf{0}, \mathbf{1}\} & \text{otherwise} \end{cases} \quad (5)$$

The expression iR_pj stands for “ i delegates her vote to j on issue p ”. Each R_p is a so-called functional relation. It corresponds to the graph of an endomap on N . So we will sometimes refer to the endomap $r_p : N \rightarrow N$ of which R_p is the graph. Relations R_p have a very specific structure and can be thought of as a set of trees whose roots all belong to cycles (possibly loops).

The weight of an agent i w.r.t. p in a delegation graph $G^{\mathbf{O}}$ is given by its indegree with respect to R_p^* (i.e., the reflexive and transitive closure of R_p):¹⁰ $w_i^{\mathbf{O}}(p) = |\{j \in N \mid jR_p^*i\}|$. This definition of weight makes sure that each individual carries the same weight, independently of the structure of the delegation graph. Alternative definitions of weight are of course possible and we will come back to this issue later.¹¹

For all $p \in \mathbf{P}$, we also define the function $g_p : N \rightarrow \wp(N)$ such that $g_p(i) = \{j \in N \mid jR_p^*i \text{ AND } \nexists k : jR_pk\}$. The function associates to each agent i (for a given issue p), the (singleton consisting of the) last agent reachable from i via a path of delegation on issue p , when it exists (and \emptyset otherwise). Slightly abusing notation we will use $g_p(i)$ to denote an agent, that is, the *guru* of i over p when $g_p(i) \neq \emptyset$. If $g_p(i) = \{i\}$ we call i a *guru* for p . Notice that $g_p(i) = \{i\}$ iff $r_p(i) = i$, that is, i is a guru of p iff it is a fixpoint of the endomap r_p .

If the delegation graph $G^{\mathbf{O}}$ of a proxy profile \mathbf{O} is such that, for some R_p , there exists no $i \in N$ such that i is a guru of p , we say that graph $G^{\mathbf{O}}$ (and profile \mathbf{O}) is *void* on p . Intuitively, a void profile on p is a profile where no voter expresses an opinion on p , because every voter delegates her vote to somebody else.

¹⁰ We recall that the reflexive transitive closure R^* of a binary relation $R \subseteq N^2$ is the smallest reflexive and transitive relation that contains R .

¹¹See also footnote 15 below.

Given a BA structure \mathcal{A} , a proxy aggregation rule (or proxy aggregator) for \mathcal{A} is a function $\text{pv} : \mathcal{P}^N \rightarrow \mathcal{O}^*$ that maps every proxy profile to one collective incomplete opinion. As above, $\text{pv}(\mathbf{O})(p)$ denotes the outcome of the aggregation on issue p .

3.1.3 Proxy Aggregators

The most natural form of voting via delegable proxy is a proxy version of the majority rule we discussed in Section 2:¹²

$$\text{pv}_{\text{maj}}(\mathbf{O})(p) = \begin{cases} \mathbf{1} & \text{if } \sum_{i \in \mathbf{O}(p)} w_i^{\mathbf{O}}(p) > \sum_{i \in \mathbf{O}(\neg p)} w_i^{\mathbf{O}}(p) \\ \mathbf{0} & \text{if } \sum_{i \in \mathbf{O}(\neg p)} w_i^{\mathbf{O}}(p) > \sum_{i \in \mathbf{O}(p)} w_i^{\mathbf{O}}(p) \\ * & \text{otherwise} \end{cases} \quad (6)$$

That is, an issue is accepted by proxy majority in profile \mathbf{O} if the sum of the weights of the agents who accept p in \mathbf{O} exceeds the majority quota, it is rejected if the sum of the weights of the agents who reject p in \mathbf{O} exceeds the majority quota, and it is undecided otherwise. It should be clear that $\sum_{i \in \mathbf{O}(p)} w_i^{\mathbf{O}}(p) = |\{i \in N \mid O_{g_i}(p) = \mathbf{1}\}|$ (and similarly for $\neg p$), that is, the sum of the weights of the gurus accepting (rejecting) p is precisely the cardinality of the set of agents whose gurus accept (reject) p .

In general, it should be clear that for any quota rule $F : \mathcal{O}_c^* \rightarrow \mathcal{O}^*$ a proxy variant pv_F of F can be defined via an obvious adaptation of (6).

To fix intuitions further about proxy voting it is worth discussing another example of aggregator, *proxy dictatorship*. It is defined as follows, for a given $d \in N$ (the dictator) any proxy profile \mathbf{O} and issue p :

$$\text{pv}_d(\mathbf{O})(p) = \begin{cases} \mathbf{O}_{g_p(d)} & \text{if } g_p(d) \neq \emptyset \\ * & \text{otherwise} \end{cases} \quad (7)$$

That is, in a proxy dictatorship, the collective opinion is the opinion of the guru of the dictator, when it exists, and it is undefined otherwise.

3.2 Two Issues of Delegable Proxy

3.2.1 Cycles and Abstentions

It should be clear from the definition of proxy aggregators like pv_{maj} , that such aggregators rely on the existence of gurus in the underlying delegation graphs. If the delegation graph R_p on issue p contains no guru, then the aggregator has access to no information in terms of who accepts and who rejects issue p . To avoid bias in favor of acceptance or rejection, such situations should therefore result in an undecided collective opinion. That is for instance the case of pv_{maj} . However, such situations may well be considered problematic, and the natural question arises therefore of how likely they are, at least in principle.

¹²On the importance of majority decisions in the current implementation of liquid democracy by liquid feedback cf. [2, p.106].

Fact 2. *Let \mathbf{A} be a BA structure where $\gamma = \top$ (i.e., issues are independent). If each proxy profile is equally probable (impartial culture assumption), then the probability that, for each issue p , the delegation graph R_p has no gurus tends to $\frac{1}{e^2}$ as n tends to infinity.*

Proof. The claim amounts to computing the probability that a random proxy profile \mathbf{O} induces a delegation graph R_p that does not contain gurus (or equivalently, whose endomap $r_p : N \rightarrow N$ has no fixpoints) as n tends to infinity. Now, for each agent i , the number of possible opinions on a given issue p (that is, functions $O : \{p\} \rightarrow \{\mathbf{0}, \mathbf{1}\} \cup N$) is $|(N \setminus \{i\}) \cup \{\mathbf{0}, \mathbf{1}\}| = n + 1$ (recall i cannot express “ i ” as an opinion). The number of opinions in which i is delegating her vote is $n - 1$. So, the probability that a random opinion of i about p is an opinion delegating i ’s vote is $\frac{n-1}{n+1}$. Hence the probability that a random profile consists only of delegated votes (no gurus) is $(\frac{n-1}{n+1})^n$. The claimed value is then established through this series of equations:

$$\begin{aligned} \lim_{n \rightarrow \infty} \left(\frac{n-1}{n} \right)^n &= \lim_{n \rightarrow \infty} \left(\frac{n}{n+2} \right)^n \\ &= \lim_{n \rightarrow \infty} \left(\frac{1}{\frac{n+2}{n}} \right)^n \\ &= \lim_{n \rightarrow \infty} \left(\frac{1}{1 + \frac{2}{n}} \right)^n \\ &= \lim_{n \rightarrow \infty} \left(\frac{1}{(1 + \frac{2}{n})^n} \right) \\ &= \frac{1}{\lim_{n \rightarrow \infty} (1 + \frac{2}{n})^n} \\ &= \frac{1}{e^2} \end{aligned}$$

This completes the proof. □

Now contrast the above simple fact with the probability that all agents abstain on an issue when each voter either expresses a $\mathbf{1}$ or $\mathbf{0}$ opinion or abstains (that is, the binary aggregation with abstentions setting studied earlier). In that case the probability that everybody abstains tends to 0 as n tends to infinity.

Fact 2 should obviously not be taken as a realistic estimate of the effect of cycles on collective abstention, as the impartial culture assumption is a highly idealized assumption. Election data should ideally be used to assess whether delegation cycles ever lead large parts of the electorate to ‘lose their vote’, possibly together with refinements of the above argument that take into consideration realistic distributions on proxy profiles, and therefore realistic delegation structures. Nonetheless, Fact 2 does flag a potential problem of cyclical delegations as sources of abstention which has, to the best of our knowledge, never been discussed. The mainstream position on cyclical delegations [2, Section 2.4.1] is:¹³

¹³Cf. also [3].

“The by far most discussed issue is the so-called circular delegation problem. What happens if the transitive delegations lead to a cycle, e.g. Alice delegates to Bob, Bob delegates to Chris, and Chris delegates to Alice? Would this lead to an infinite voting weight? Do we need to take special measures to prohibit such a situation? In fact, this is a nonexistent problem: A cycle only exists as long as there is no activity in the cycle in which case the cycle has no effect. As already explained [...], as soon as somebody casts a vote, their (outgoing) delegation will be suspended. Therefore, the cycle naturally disappears before it is used. In our example: If Alice and Chris decide to vote, then Alice will no longer delegate to Bob, and Chris will no longer delegate to Alice [...]. If only Alice decides to vote, then only Alice’s delegation to Bob is suspended and Alice would use a voting weight of 3. In either case the cycle is automatically resolved and the total voting weight used is 3.”

We will discuss later (Section 4) a possible approach to mitigate this issue by suggesting a different interpretation of liquid democracy in terms of influence rather than delegation.

3.2.2 Individual & Collective Rationality

In our discussion so far we have glossed over the issue of logically interdependent issues and collective rationality. The reason is that under the delegative interpretation of liquid democracy developed in this section individual rationality itself appears to be a more debatable requirement than it normally is in classical aggregation.

A proxy opinion O_i is *individually rational* if the set of formulas

$$\{\gamma\} \cup \{p \in \mathbf{P} \mid O_{g_p(i)}(p) = \mathbf{1}\} \cup \{\neg p \in \mathbf{P} \mid O_{g_p(i)}(p) = \mathbf{0}\} \quad (8)$$

is satisfiable (consistency), and if whenever (8) entails $\pm p$, then $\pm p$ belongs to it (closedness).¹⁴ That is, the integrity constraint γ is consistent with i ’s opinion on the issues she does not delegate on, and the opinions of her gurus (if they exist), and those opinions, taken together, are closed under logical consequence (w.r.t. the available issues).

The constraint in (8) captures, one might say, an idealized way of how delegation works: voters are assumed to be able to check or monitor how their gurus are voting, and always modify their delegations if an inconsistency arises. The constraint remains, however, rather counterintuitive under a delegative interpretation of proxy voting. Aggregation via delegable proxy has at least the potential to represent individual opinions as irrational (inconsistent and/or not logically closed).

Like in the case of delegation cycles we will claim that the interpretation of liquid democracy in terms of influence to be developed in Section 4, rather than

¹⁴Cf. the definition of individual opinions in Section 2.

in terms of delegation, makes individual rationality at least as defensible as in the classical case.

3.3 Embedding in Binary Aggregation with Abstentions

3.3.1 One man—One vote

Aggregation in liquid democracy as conceived in [2] should satisfy the principle that the opinion of every voter, whether expressed directly or through proxy, should be given the same weight:


“[...] in fact every eligible voter has still exactly one vote [...] unrestricted transitive delegations are an integral part of Liquid Democracy. [...] Unrestricted transitive delegations are treating delegating voters and direct voters equally, which is most democratic and empowers those who could not organize themselves otherwise” [2, p.34-36]

In other words, this principle suggests that aggregation via delegable proxy should actually be ‘blind’ for the specific type of delegation graph. Making this more formal, we can think of the above principle as suggesting that the only relevant content of a proxy profile is its translation into a standard opinion profile (with abstentions) via a function $t : \mathcal{P} \rightarrow \mathcal{O}^*$ defined as follows: for any $i \in N$ and $p \in \mathbf{P}$: $t(O_i(p)) = O_{g_p(i)}$ if $g_p(i) \neq \emptyset$ (i.e., if i has a guru for p), and $t(O_i(p)) = *$ otherwise. Clearly, if we assume proxy profiles to be individually rational, the translation will map proxy opinions into individually rational (consistent and closed) incomplete opinions. By extension, we will denote by $t(\mathbf{O})$ the incomplete opinion profile resulting from translating the individual opinions of a proxy profile \mathbf{O} .

The above discussion suggests the definition of the following property of proxy aggregators: a proxy aggregator pv has the **one man—one vote property** (or is a one man—one vote aggregator) if and only if $pv = t \circ F$ for some aggregator $F : \mathcal{O}_c^* \rightarrow \mathcal{O}^*$ (assuming the individual rationality of proxy profiles).¹⁵

The class of one man—one vote aggregators can therefore be studied simply as the concatenation $t \circ F$ where F is an aggregator for binary voting with abstentions, as depicted in the following diagram:

¹⁵It should be clear that not every proxy aggregator satisfies this property. By means of example, consider an aggregator that uses the following notion of weight accrued by gurus in a delegation graph. The weight $w(i)$ of i is $\sum_{j \in R^*(i)} \frac{1}{\ell(i,j)}$ where $\ell(i,j)$ denotes the length of the delegation path linking j to i . This definition of weight is such that the contribution of voters decreases as their distance from the guru increases. Aggregators of this type are studied in [6].



which gives us a handle to study a large class of proxy voting rules

Example 5. *Proxy majority pv_{maj} (6) is clearly a one man—one vote rule aggregator. It is easy to check that, for any proxy profile \mathbf{O} : $\text{pv}_{\text{maj}}(\mathbf{O}) = \text{maj}(t(\mathbf{O}))$. The same holds for proxy dictatorship (7). It is easy to see that proxy dictatorship pv_d is such that for any proxy profile \mathbf{O} : $\text{pv}_d(\mathbf{O}) = d(t(\mathbf{O}))$, where d is the standard dictatorship (of $d \in N$).*

It follows that for every proxy aggregator $\text{pv}_F = t \circ F$ the axiomatic machinery developed for standard aggregators can be directly tapped into. Characterization results then extend effortlessly to proxy voting, again providing a strong rationale for the use of majority in proxy aggregation:

Fact 3 (Characterization of proxy majority). *A one man—one vote proxy aggregator $\text{pv} = t \circ F$ for a given \mathcal{A} is proxy majority pv_{maj} iff F is anonymous, independent, monotonic, responsive, neutral and minimizes undecisiveness.*


Proof. This follows from the definition of t and Theorem 1. \square

It follows that on simple agendas and assuming the individual rationality of proxy profiles, proxy majority is the only rational aggregator which is anonymous, independent, monotonic, responsive, neutral and minimizes undecisiveness.

3.3.2 Impossibility

Similarly, there are many ways in which pursue the opposite embedding, from standard aggregation into proxy voting. For example, we can define a function $s : \mathcal{O}_c \rightarrow \mathcal{P}_c$ from opinion profiles to individually rational proxy profiles as follows. For a given opinion profile \mathbf{O} , and issue $p \in \mathbf{P}$ consider the set $\{i \in N \mid O_i(p) = *\}$ of individuals that abstain in \mathbf{O} and take an enumeration $1, \dots, m$ of its elements, where $m = |\{i \in N \mid O_i(p) = *\}|$. The function is defined as follows, for any $i \in N$ and $p \in \mathbf{P}$: $s(O_i(p)) = O_i(p)$ if $O_i(p) \in \{\mathbf{0}, \mathbf{1}\}$, $s(O_i(p)) = i + 1 \pmod m$, otherwise.¹⁶ A translation of this type allows to think of standard aggregators $F : \mathcal{O}_c^* \rightarrow \mathcal{O}^*$ as the concatenation $s \circ \text{pv}$, for some proxy aggregator pv :

¹⁶Notice that since self-delegation (that is, $O_i(p) = i$) is not feasible in proxy opinions, this definition of s works for profiles where, on each issue, either nobody abstains or at least two individuals abstain. A dummy voter can be introduced for that purpose.



Fact 4. *Let A be such that its agenda is path connected and evenly negatable. For any proxy aggregator pv , if $s \circ pv$ is independent, unanimous and collectively rational, then it is oligarchic.*

Proof. It follows directly from the definition of s and Theorem 2. □

3.4 Section Summary

The section has provided a very simple model of delegable proxy voting within the framework of binary aggregation. This has allowed us to put liquid democracy in perspective with an established body of results in the social choice theory tradition, and highlight two of its problematic aspects, which have so far gone unnoticed: the effect of cycles on collective indecisiveness, and the issue of preservation of individual rationality under delegable proxies.

An independent, purpose-built axiomatic analysis for liquid democracy focused on its more characteristic features (like the one man—one vote property) is a natural line of research, which we do not pursue here.

4 Liquid Democracy as Binary Opinion Diffusion

Proxy voting can also be studied from a different perspective. Imagine a group where, for each issue p , each agent copies the $\mathbf{0}, \mathbf{1}$ opinion of a unique personal “guru”. Imagine that this group does so repeatedly until all agents (possibly) reach a stable opinion. These new stable opinions can then be aggregated as the ‘true’ opinions of the individuals in the group. The collective opinion of a group of agents who either express a $\mathbf{0}, \mathbf{1}$ opinion or delegate to another agent is (for one man—one vote proxy aggregators) the same as the output obtained from a vote where each individual has to express a $\mathbf{0}, \mathbf{1}$ opinion but gets there by copying the opinion of some unique “guru” (possibly themselves). In this perspective, a proxy voting aggregation can be assimilated to a (converging) process of opinion formation.

The above interpretation of liquid democracy is explicitly put forth in [2]:

“While one way to describe delegations is the transfer of voting weight to another person, you can alternatively think of delegations as automated copying of the ballot of a trustee. While at assemblies with voting by a show of hands it is naturally possible to copy the

vote of other people, in Liquid Democracy this becomes an intended principle” [2, p. 22].

The current section develops an analysis of this interpretation, and highlights some of its advantages over the delegation-based one studied earlier.

4.1 Binary aggregation and binary influence

The section develops a very simple model of binary influence based on the standard framework of binary aggregation (see Appendix B for a concise presentation). For simplicity, in this section we assume agents are therefore not allowed to abstain, although this is not a crucial assumption for the development of our analysis.

4.1.1 DeGroot Processes and Opinion Diffusion

In [9], DeGroot proposes a simple model of step-by-step opinion change under social influence. The model combines two types of matrices. Assuming a group of n agents, a first $n \times n$ matrix represents the weighted influence network (who influences whom and how much), and a second $n \times m$ matrix represents the probability assigned by each agent to each of the m different alternatives. Both the agents’ opinion and the influence weights are taken within $[0, 1]$ and are (row) stochastic (each row sums up to 1). Given an opinion and an influence matrix, the opinion of each agent in the next time step is obtained through linear averaging.

Here we focus on a specific class of opinion diffusion processes in which opinions are binary, and agents are influenced by exactly one influencer, possibly themselves, of which they copy the opinion. The model captures a class of processes which lies at the interface of two classes of diffusion models that have remained so far unrelated: the stochastic opinion diffusion model known as DeGroot’s [9], and the more recent propositional opinion diffusion model due to [15]. The diffusion processes underpinning liquid democracy—which we call here Boolean DeGroot processes (BDPs)—are the $\{0, 1\}$ special case of the DeGroot stochastic processes and, at the same time, the special case of propositional opinion diffusion processes where each agent has access to the opinion of exactly one neighbor (cf. Figure 1).

4.1.2 Boolean DeGroot processes

Here we focus on the Boolean special case of a DeGroot process showing its relevance for the analysis of liquid democracy. Opinions are defined over a BA structure, and hence are taken to be binary. Similarly, we take influence to be modeled by the binary case of an influence matrix. Influence is of an “all-or-nothing” type and each agent is therefore taken to be influenced by exactly one agent, possibly itself. The graph induced by such a binary influence matrix (called *influence graph*) is therefore a structure $G = \langle N, R \rangle$ where $R \subseteq N^2$ is a binary relation where iRj is taken here to denote that “ i is influenced by j ”.




Figure 1: BDPs lie in the intersection of DeGroot processes and propositional opinion diffusion processes.

Such relation is serial ($\forall i \in N, \exists j \in N : iRj$) and functional ($\forall i, j, k \in N$ if iRj and iRk then $j = k$). So each agent i has exactly one successor (the influencer), possibly itself, which we denote $R(i)$. It should be clear that influence graphs are the same sort of structures we studied earlier in Section 3 under the label 'delegation graph'.

An *influence profile* $\mathbf{G} = (G_1, \dots, G_m)$ records how each agent is influenced by each other agent, with respect to each issue $p \in \mathbf{P}$. Given a profile \mathbf{G} the i^{th} projection G_i denotes the influence graph for issue p_i , also written G_p .

So let us define the type of opinion dynamics driving BDPs:

Definition 4 (BDP). *Now fix an opinion profile \mathbf{O} and an influence profile \mathbf{G} . Consider the stream $\mathbf{O}^0, \mathbf{O}^1, \dots, \mathbf{O}^n, \dots$ of opinion profiles recursively defined as follows:*

- *Base:* $\mathbf{O}_0 := \mathbf{O}$
- *Step:* for all $i \in N, p \in \mathbf{P}, O_i^{n+1}(p) := O_{R_p(i)}^n(p)$.

where $G_p = \langle N, R_p \rangle$. We call processes defined by the above dynamics Boolean DeGroot processes (BDPs).

It should be clear that the above dynamics is the extreme case of linear averaging applied on binary opinions and binary influence.

As noted above, BDPs are also the special case of processes that have recently been proposed in the multi-agent systems literature as *propositional opinion diffusion* processes [15], i.e., cases where 1) the aggregation rule is the unanimity rule (an agent adopts an opinion if and only if all her influencers agree on it), and 2) each agent has exactly one influencer. We will come back to the link with propositional opinion diffusion in some more detail later in Section 4.3.

4.2 Convergence of BDPs

When do the opinions of a group of individuals influencing each other stabilize? Conditions have been given, in the literature, for the general paradigms

of which BDPs are limit cases. This section introduces the necessary graph-theoretic notions and briefly recalls those results before giving a characterization of convergence for BDPs.

4.2.1 Preliminaries

We start with some terminology. We say that the stream of opinion profiles $\mathbf{O}^0, \mathbf{O}^1, \dots, \mathbf{O}^n, \dots$ *converges* if there exists $n \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, if $m \geq n$, then $\mathbf{O}^m = \mathbf{O}^n$.

We will also say that a stream of opinion profiles converges *for issue* p if there exists $n \in \mathbb{N}$ such that, for all $m \in \mathbb{N}$, if $m \geq n$, then $\mathbf{O}^m(p) = \mathbf{O}^n(p)$. Given a stream of opinion profiles starting at \mathbf{O} we say that agent $i \in N$ stabilizes in that stream for issue p if there exists $n \in \mathbb{N}$ such that $O_i^n(p) = O_i^m(p)$ for any $m > n$. So a BDP on influence graph \mathbf{G} starting with the opinion profile \mathbf{O} is said to converge if the stream $\mathbf{O}^0, \mathbf{O}^1, \dots, \mathbf{O}^n, \dots$ generated according to Definition 4 where $\mathbf{O} = \mathbf{O}^0$ converges. Similarly, A BDP is said to converge for issue p if its stream converges for p , and an agent i in the BDP is said to stabilize for p if it stabilizes for p in the stream generated by the BDP.

Notice first of all that influence graphs have a special shape:¹⁷

Fact 5. *Let G be an influence graph and C be a connected component of G . Then C contains exactly one cycle, and the set of nodes in the cycle is closed.*

Proof. Assume that C does not contain any cycle. Since N is finite and since no path can repeat any node, any path in C is finite too. Let i be the last element of (one of) the longest path(s) in C . Then i does not have any successor, which contradicts seriality. So C contains at least one cycle. Let S be the set of nodes of a cycle in C . Assume that S is not closed: for some $i \in S$ and $j \notin S$, iRj . Since S is a cycle, there is also some $k \in S$, such that iRk , which contradicts functionality. Therefore, the nodes of any cycle in C forms a closed set. Now assume that C contains more than one cycle. Since the nodes of each cycle forms a closed set, there is no path connecting any node inside a cycle to any node in any other cycle, which contradicts connectedness. So C contains a unique cycle, whose nodes form a closed set. \square

Intuitively, influence graphs of BDPs then look like sets of confluent chains aiming together towards common cycles.

4.2.2 Context: convergence in DeGroot processes

For the general case of DeGroot processes, an influence structure guarantees that any distribution of opinions will converge if and only if “every set of nodes that is strongly connected and closed is aperiodic” [20, p.233]. In the propositional opinion diffusion setting, sufficient conditions for stabilization have been given by [15, Th. 2]: on influence structures containing cycles of size at most one

¹⁷Please consult Appendix B for the relevant terminology from graph theory.

(i.e, only self-loops), for agents using an aggregation function satisfying (ballot-)monotonicity and unanimity¹⁸, opinions will always converge in at most at most $k + 1$ steps, where k is the diameter of the graph.¹⁹ The results below show how BDPs are an interesting limit case of both DeGroot and propositional opinion diffusion processes.

4.2.3 Two results

It must be intuitively clear that non-convergence in a BDP is linked to the existence of cycles in the influence graphs. However, from the above observation (Fact 5), we know that nodes in a cycle cannot have any influencers outside this cycle, and hence that cycles (including self-loops) can only occur at the “tail” of the influence graph. Hence, if the opinions in the (unique) cycle do not converge, which can only happen in a cycle of length ≥ 2 , the opinions of the whole population in the same connected component will not converge. The above implies that for any influence graphs with a cycle of length ≥ 2 , there exists a distribution of opinions which loops. This brings us back to convergence result for general (not necessarily Boolean) DeGroot processes. Indeed, for functional and serial influence graphs, a closed connected component is aperiodic if and only if its cycle is of length 1.

Fact 6. *Let \mathbf{G} be an influence profile. Then the following are equivalent:*

1. *The BDP converges for any opinion profile \mathbf{O} on \mathbf{G} .*
2. *For all $p \in \mathbf{P}$, G_p contains no cycle of length ≥ 2 .*
3. *For all $p \in \mathbf{P}$, all closed connected components of G_p are aperiodic.*

Proof. $\boxed{2) \Rightarrow 1)}$ Let $p \in \mathbf{P}$ and assume that G_p contains no cycle of length ≥ 2 and has diameter k . Let C_p be a connected component of G_p . By Fact 5, C_p contains a unique cycle, which, by assumption, is of length 1. Hence, C_p is aperiodic. Let i be the node in the cycle. The opinion of i will spread to all nodes in C_p after at most k steps. Therefore, all BDPs on G will converge after at most l steps, where l is the maximum within the set of diameters of G_p for all $p \in \mathbf{P}$. $\boxed{1) \Rightarrow 3)}$ We proceed by contraposition. Assume that for some $p \in \mathbf{P}$, a connected component C_p of G_p contains a cycle of length $k \geq 2$. By 5, this cycle is unique, and therefore the greatest common divisor of the cycles lengths of C_p is k , so C_p is not aperiodic. Let S be the set of nodes in the cycle. Let \mathbf{O} be such that for some $i, j \in S$ with distance d from i to j , $O_i(p) \neq O_j(p)$. Then $\mathbf{O}_i(p)$ will not converge, but enter a loop of size k : for

¹⁸Notice that the rule underpinning BDP, that is the ‘guru-copying’ rule on serial and functional graphs, trivially satisfies those constraints.

¹⁹A second sufficient condition for convergence is given by [15]: when agents use the unanimity aggregation rule, on irreflexive graphs with only vertex-disjoint cycles, such that for each cycle there exists an agent who has at least two influencers, opinions converge after at most N steps. Note that no BDP satisfies this second condition.

all $x \in \mathbb{N}$, $O_i^{x \times k}(p) \neq O_i^{(x \times k) + d}(p)$. Hence, \mathbf{O} does not converge. $\boxed{3) \Rightarrow 2)}$
 Trivial. \square

It is worth noticing that one direction (namely from 3 to 1) of the above result is actually a corollary of both the convergence result for DeGroot processes stated at the beginning of this section (cf. [20]), and of a known convergence result for propositional opinion diffusion [15, Th. 2], also stated earlier.

The above gives a characterization of the class of influence profiles on which *all* opinion streams converge. But we can aim at a more general result, characterizing the class of pairs of opinion and influence profiles which lead to convergence:

Theorem 3. *Let \mathbf{G} be an influence profile and \mathbf{O} be an opinion profile. Then the following statements are equivalent:*

1. *The BDP converges for \mathbf{O} on \mathbf{G} .*
2. *For all $p \in \mathbf{P}$, there is no set of agents $S \subseteq N$ such that: S is a cycle in G_p and there are two agents $i, j \in S$ such that $O_i(p) \neq O_j(p)$.*

Proof. $\boxed{1) \Rightarrow 2)}$ We proceed by contraposition. Let $p \in \mathbf{P}$, $S \subseteq N$ be a cycle in G_p , $i, j \in S$, and $O_i(p) \neq O_j(p)$. Let k be the length of the cycle and d be the distance from i to j . Then $O_i(p)$ will enter a loop of size k : for all $x \in \mathbb{N}$, $O_i^{xk}(p) \neq O_i^{xk+d}(p)$. $\boxed{2) \Rightarrow 1)}$ Assume $S \subseteq N$ be such that S is a cycle in G_p , and for all $i, j \in S$, $O_i(p) = O_j(p)$. Then, for all $j \in S$, and all $x \in \mathbb{N}$, $O_j^x(p) = O_i(p)$ and for all $f \in N \setminus S$ with distance d from f to i , for all $x \in \mathbb{N}$, such that $x \geq d$, $O_f^x(p) = O_i(p)$. \square

This trivially implies that the class of opinion profiles which guarantees convergence for *any* influence profile, is the one where everybody agrees on everything already. Note that the only stable distributions of opinions are the ones where, in each connected component in G , all members have the same opinion, i.e., on BDPs, converging and reaching a consensus (within each connected component) are equivalent, unlike in the stochastic case. Moreover, for an influence profile where influence graphs have at most diameter d and the smallest cycle in components with diameter d is of length c , it is easy to see that if a consensus is reached, it will be reached in at most $d - c$ steps, which is at most $n - 1$.

Finally observe that Theorem 3 subsumes Fact 6. If G_p contains only cycles of length 1 (second statement in Fact 6) then, trivially, no two agents in a cycle can disagree (second statement in Theorem 3).

4.2.4 Liquid Democracy as a BDP

We have seen (Section 3) that each proxy profile \mathbf{O} induces what we called a delegation graph $G^{\mathbf{O}} = \langle N, R_p \rangle$ for each issue p . Delegation graphs are the same sort of structures we referred to in the current section as influence graphs. So

each proxy profile \mathbf{O} can be associated to a BDP by simply assigning random $\mathbf{0}$ or $\mathbf{1}$ opinions to each voter delegating her vote in \mathbf{O} . It is then easy to show that for each connected component C of $G^{\mathbf{O}}$, if C has a guru with opinion x , then that component stabilizes in the BDP on opinion x for each assignment of opinions to the delegating agents in \mathbf{O} . Vice versa, if C stabilizes on value x in the BDP for each assignment of opinions to the delegating agents in \mathbf{O} , then C has a guru whose opinion is x . This establishes a direct correspondence between voting with delegable proxy and Boolean deGroot processes. However, BDPs offer an interesting and novel angle on the issue of cyclical delegations, to which we turn now.

4.2.5 Cycles

As discussed earlier (Section 3.2.1), cycles are a much discussed issue in liquid democracy. Its proponents tend to dismiss delegation cycles as a non-issue: since the agents forming a cycle delegate their votes, none of them is casting a ballot and the cycles get resolved essentially by not counting the opinions of the agents involved in the cycle [2]. We stressed this solution as problematic in the ‘vote-delegation’ interpretation of liquid democracy as it has the potential to discard large numbers of opinions. The elimination of cycles not only hides to aggregation the opinions of the agents involved in cycles, but also the opinions of agents that may be linked to any of those agents by a delegation path. In other words information about entire connected components in the delegation graph may be lost.

We argue that the ‘vote-copying’ interpretation of the system—formalized through BDPs—offers novel insights into possible approaches to cycles in delegable proxy. Theorems 3 and 6 offer an alternative solution by showing that not all cycles are necessarily bad news for convergence: cycles in which all agents agree still support convergence of opinions, and therefore a feasible aggregation of opinions by proxy. This suggests that alternative proxy voting mechanisms could be designed based on opinion convergence behavior rather than on weighted voting.

4.3 Excursus: unanimity and 2-colorability

In the above, we have worked at the intersection of two models of opinion diffusion, the DeGroot model, and the propositional opinion diffusion model. However, there is more to say about how the two frameworks relate.

Let us take a brief detour towards a generalisation of BDPs corresponding to the case of propositional opinion diffusion with the unanimity rule, where agents can have several influencers and change their opinions only if all their influencers disagree with them. This means that we relax the functionality constraint on influence graphs. We will show how the two frameworks meet again: some non-stabilizing opinion cases under the unanimity rule correspond to a special class among the ‘semi-Boolean’ cases of DeGroot processes where opinions are still binary but influence does not need to be.

We define the dynamics of opinions under the unanimity rule in the obvious way:

Definition 5 (UP). *Fix an opinion profile \mathbf{O} and a (serial but non-necessarily functional) influence profile \mathbf{G} . Consider the stream $\mathbf{O}^0, \mathbf{O}^1, \dots, \mathbf{O}^n, \dots$ of opinion profiles recursively defined as follows:*

- *Base: $\mathbf{O}_0 := \mathbf{O}$*
- *Step: for all $i \in N$ and all $p \in \mathbf{P}$:*

$$\mathbf{O}_i^{n+1}(p) = \begin{cases} O_i^n(p) & \text{if for some } j, k \in R_p(i), O_j^n(p) \neq O_k^n(p) \\ O_j^n(p) & \text{otherwise, where } j \in R_p(i) \end{cases} \quad (9)$$

where $G_p = \langle N, R_p \rangle$. We call processes defined by the above dynamics Unanimity Processes (UPs).

We give a sufficient condition for non-convergence of UPs:

Lemma 2. *Let G be a (serial and non-necessarily functional) influence profile and \mathbf{O} be an opinion profile, such that, for some $p \in \mathbf{P}$, for all $i, j \in C$, where C is a connected component of G_p : if $i \in R_p(j)$, then $O_i(p) \neq O_j(p)$. Then \mathbf{O} does not converge in UP.*

Proof. Let G be a (serial and non-necessarily functional) influence profile, and \mathbf{O} be an opinion profile, such that, for some $p \in \mathbf{P}$, for all $i, j \in C$ with C a connected component of G_p : if $i \in R_p(j)$, then $O_i(p) \neq O_j(p)$. Then, by definition of UPs, for all $i \in C$, $O_i^1(p) \neq O_i(p)$, and by repeating the same argument, for all $n \in \mathbb{N}$, $O_i^{n+1}(p) \neq O_i^n(p)$. \square

Intuitively, the above condition for non-convergence corresponds to a situation of global maximal disagreement: *all* agents (of a connected component) disagree with *all* their influencers. Recall that a graph is properly k -colored if each node is assigned exactly one among k colors and no node has a successor of the same color, and consider the two possible opinions on issue p as colors. The above result can be reformulated in terms of proper 2 colorings, as follows: if for some $p \in \mathbf{P}$, \mathbf{O} properly colors G_p , then \mathbf{O} does not converge. In such a case, all agents will change their opinion on p at every step, entering an oscillation of size 2. So the maximal state of disagreement is the maximally unstable case of the dynamics. Note that this limit case of opinion distribution is yet another special case of DeGroot processes, another example within the intersection between the two frameworks of propositional opinion diffusion and DeGroot.

The possibility of such a distribution of opinions on p relies on the influence graph G_p being 2-colorable, which is again a requirement about the lengths of its cycles: it is 2-colorable if and only if it contains no cycle of odd length. However, non 2-colorability is not a sufficient condition for convergence of UPs in general: a simple cycle of three agents, for instance, is not 2-colorable but does not guarantee convergence either (as illustrated above with the convergence

conditions for BDPs). Nevertheless, there is a class of influence profiles for which being 2-colorable is a necessary condition of non-convergence of UPs, the *symmetric* ones:

Lemma 3. *Let \mathbf{G} be a symmetric (serial and non-necessarily functional) influence profile and \mathbf{O} be an opinion profile. The following statements are equivalent:*

1. \mathbf{O} converges in UP on \mathbf{G} ;
2. For all $p \in \mathbf{P}$, for all connected component C of G_p , there are $i, j \in C$, such that $i \in R_p(j)$, and $O_i(p) = O_j(p)$, where $G_p = \langle N, R_p \rangle$.

Proof. $\boxed{2) \Rightarrow 1)}$ Assume that for any $p \in \mathbf{P}$, for any connected component C of G_p , there exist $i, j \in C$, such that $R_p(j)$ and $O_i(p) = O_j(p)$. By definition of UP, this implies that $O_i(p)$ is stable, and that all agents with distance $\leq k$ will be stable after at most k steps. $\boxed{1) \Rightarrow 2)}$ This follows from Lemma 2. \square

This means that opinions on a given p will converge if and only if two agents influencing each other on p already agree on it. We can therefore, as we did for BDPs, characterize the class of influence profiles for which all (symmetric) opinion profiles converge in UPs:

Theorem 4. *Let \mathbf{G} be a symmetric (serial and non-necessarily functional) influence profile. The following statements are equivalent:*

1. All opinion profiles \mathbf{O} , converge in UPs on \mathbf{G} .
2. For all $p \in \mathbf{P}$, and all connected components of $C \subseteq G_p$, C is not 2-colorable (contains cycle(s) of odd length), where $G_p = \langle N, R_p \rangle$.

Proof. $\boxed{2) \Rightarrow 1)}$ Let $p \in \mathbf{P}$ and C be connected component of G_p with diameter k . Let C contain a cycle of length c , with c odd. Let \mathbf{O} be an arbitrary opinion profile. Since c is odd, there exist $i, j \in S$ such that $j \in R_p(i)$ and $O_i(p) = O_j(p)$. By definition of UP, this implies that $O_i(p)$ is stable, and that all agents with distance $\leq k$ will be stable after at most k steps. Hence, \mathbf{O} converges. $\boxed{1) \Rightarrow 2)}$ This follows from Lemma 3. \square

Note that, while the basic modal language cannot capture graph 2-colorability, it can capture non 2-colorability, and therefore capture the class of symmetric (serial and non-necessarily functional) influence profiles which guarantee convergence of UPs. We leave the detail out for space reason.

We have shown that, for UPs in general, convergence (in a connected component) is not guaranteed if it contains no odd cycles, and that symmetric UPs guarantee convergence as soon as they contain some odd cycle. However, containing an odd cycle is a very “easy” requirement for a real-life influence network to meet (it corresponds to a non-zero clustering coefficient). By contrast, recall that BDPs guarantee convergence (on *any* opinion profile) only when they contain only cycles of size 1, which is a rather implausible requirement to be satisfied on real influence networks.

4.4 BDPs on logically interdependent issues

So far we have assumed the aggregation to happen on a set of issues without constraint (or rather with $\gamma = \top$). In this subsection we study what happens in the presence of a constraint $\gamma \neq \top$. BDPs on aggregation structures with constraints may lead individuals to update with logically inconsistent opinions. But the diffusion perspective whereby agents copy the opinion of trustees rather than delegating their voting right better lends itself to an assumption of individual rationality.

The following processes are simple adaptations of BDPs where agents update their opinions only if the opinions of their influencers, on the respective issues, are consistent with the constraint.²⁰

Definition 6. Fix an opinion profile \mathbf{O} , an influence profile \mathbf{G} , and a constraint γ . Consider the stream $\mathbf{O}^0, \mathbf{O}^1, \dots, \mathbf{O}^n, \dots$ of opinion profiles recursively defined as follows:

- Base: $\mathbf{O}_0 := \mathbf{O}$
- Step: for all $i \in N, p \in \mathbf{P}$,

$$O_i^{n+1}(p) := \begin{cases} O_{R_p(i)}^n(p) & \text{if } \bigwedge_{p \in \mathbf{P}} O_{R_p(i)}^n(p) \wedge \gamma \text{ is consistent} \\ O_i^n(p) & \text{otherwise} \end{cases}$$

where $G_p = \langle N, R_p \rangle$. We call processes defined by the above dynamics individually rational BDPs.

Individually rational BDPs converge in some cases in which BDPs do not. There are cases in which there is disagreement in the cycles but the process still converges, because of the safeguard towards individual rationality built into the dynamics.

Example 6. Consider the following example. Let $N = \{1, 2\}$, $\mathbf{P} = \{p, q\}$ and $\gamma = \{p \leftrightarrow \neg q\}$. Let then $G = \langle N, \{R_i\}_{i \in \mathbf{P}} \rangle$ be as follows: $1R_q1, 2R_q2, 1R_p2$ and $2R_p1$. Finally let \mathbf{O} be such that $O_1(p) = O_2(q) = \mathbf{1}$, $O_2(p) = O_1(q) = \mathbf{0}$. Voters 1 and 2 form a non-unanimous cycle, but \mathbf{O} is a stable opinion profile.

The example shows that direction 1) \Rightarrow 2) of Theorem 3 does not hold for individually rational BDPs: some individually rational BDPs may stabilize even in the presence of disagreement within a cycle. Intuitively, the reason why this happens is that individually rational BDPs that stabilize even when disagreements occur within cycles do so because their cycles are not "synchronized". In the above example, given the constraint $p \leftrightarrow \neg q$, the only way to get stabilization starting from a situation respecting the constraint is to have a cycle of influence for q which goes 'in the opposite direction' from the one from p , all other cases would amount to violate the constraint.

²⁰Other update policies are of course possible. A recent systematic investigation of opinion diffusion on interconnected issues is [7].

Beyond this simple example, we want to find out what happens with more complex constraints and what are the conditions for individually rational BDPs to converge. Let us first show that direction 2) \Rightarrow 1) of Theorem 3 still holds, that is, individually rational BDPs without disagreement in their cycles always converge:

Theorem 5. *Let \mathbf{G} be an influence profile, \mathbf{O} be an opinion profile, and γ a constraint. Then the following holds: if for all $p \in \mathbf{P}$, for all $S \subseteq N$ such that S is a cycle in G_p , and all $i, j \in S$: $O_i(p) = O_j(p)$, then the individually rational BDP for \mathbf{O} , \mathbf{G} and γ converges in at most k steps, where $k \leq \max\{\text{diam}(G_p) \mid p \in P\}$.*

Proof. Assume that for all $p \in \mathbf{P}$, for all $S \subseteq N$ such that S is a cycle in G_p , for all $i, j \in S$: $O_i(p) = O_j(p)$. Consider an arbitrary $i \in N$. Let $k_i(p)$ be the distance from i to the closest agent in a cycle of G_p , and let k_i denote $\max\{k_i(p) \mid p \in P\}$. We show that for any $k_i \in \mathbb{N}$, $O_i^{k_i}$ is stable.

- If $k_i = 0$: i is its only influencer, therefore O_i^0 is stable.
- If $k_i = n + 1$: Assume that for all agents j such that $k_j = n$, $O_j^{k_j}$ is stable. This implies that all influencers of i are stable. We need to consider the following cases:
 1. If $\bigwedge_{p \in \mathbf{P}} O_{R_p(i)}^m(p) \wedge \gamma$ is not consistent, then it will never be: O_i^n is stable.
 2. If $\bigwedge_{p \in \mathbf{P}} O_{R_p(i)}^m(p) \wedge \gamma$ is consistent, then O_i^{n+1} is stable.

This completes the proof. □

4.5 Section Summary

In this section we studied a very simple class of opinion diffusion processes on networks (Boolean DeGroot processes, BDPs), which precisely capture the vote-copying behavior suggested by a standard interpretation of the liquid democracy system. Interestingly these processes lie at the interface of two so far unconnected network diffusion models: the well-known DeGroot processes—of which BDPs constitute the binary special case—and of propositional opinion diffusion processes—of which BDPs constitute the special case where the set of neighbors is a singleton. We established necessary and sufficient conditions for convergence, which can be captured in modal fixpoint logics as we will show in the next section. We argued that these results provide a novel angle on the issue of delegation cycles in liquid democracy.

There are a number of further questions concerning, especially, individually rational BDPs that we leave for future investigations: What are the necessary conditions for their stabilization? What opinions are reachable? And, in particular, when is a consensus reached? Finally, one could consider other types of influence policies than the one used in individually rational BDPs. For instance, agents may be allowed to ‘pass through’ an inconsistent state at some point, in

which case one can wonder under which conditions the process can still converge to a consistent state. Indeterministic policies would also make sense, where an agent confronted with inconsistent opinions from her influencers keeps one of the closest consistent opinions set, rather than not being influenced at all (cf. [7]).

5 Fixpoint Logics for BDPs

In this section we show how a well-established logic for formal verification can be readily used to specify and reason about properties of BDPs, and in particular their convergence. The logic is the so-called μ -calculus. This points to a so-far unexplored interface between fixpoint logics and models of opinion dynamics—like the DeGroot model and propositional opinion diffusion. The section moves some first steps in that direction along the lines of another recent work [4], where the μ -calculus, and extensions thereof, have been applied to the study of dynamical systems.

5.1 Influence graphs as Kripke models

We treat influence graphs as Kripke (multi-relational) models [22, 8].

Definition 7. *We call an influence model a tuple $\mathcal{M} = \langle N, \mathbf{G}, \mathbf{O} \rangle$ where $\mathbf{G} = (G_{p_1}, \dots, G_{p_m})$ is an influence profile, and $\mathbf{O} : \mathbf{P} \rightarrow 2^N$ is an opinion profile over \mathbf{P} , that is, a valuation function.*

One can therefore easily interpret a modal language over influence models, where modalities are interpreted on the accessibility relations in \mathbf{G} . That is, to each graph G_p we associate modalities $[p]$ and $\langle p \rangle$. We will give the details below, but let us immediately note that the class of (possibly infinite) influence graphs would then be characterized by the following properties, for any $p \in \mathbf{P}$:

$$[p]\varphi \rightarrow \langle p \rangle \varphi \quad (\text{seriality}) \quad (10)$$

$$\langle p \rangle \varphi \rightarrow [p]\varphi \quad (\text{functionality}) \quad (11)$$

More precisely, for any influence profile $\mathbf{G} = (G_{p_1}, \dots, G_{p_m})$, formula $[p_i]\varphi \rightarrow \langle p_i \rangle \varphi$ (respectively, $\langle p_i \rangle \varphi \rightarrow [p_i]\varphi$) is valid in such graph—that is, true in any pointed influence model built on such graph—if and only if each G_{p_i} consists of a serial (respectively, functional) relation.²¹ Put otherwise, on serial and functional graphs the modal box and diamond are equivalent.

5.2 Modal μ -calculus

$$\mathcal{L}^\mu : \varphi ::= p \mid \perp \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle p \rangle \varphi \mid \mu p. \varphi(p)$$

The language of the μ -calculus expands the basic modal language with at least fixpoint operator μ . Here is the BNF of the language:

²¹These are known results from modal correspondence theory (cf. [5]).

where p ranges over \mathbf{P} and $\varphi(p)$ indicates that p occurs free in φ (i.e., it is not bounded by fixpoint operators) and under an even number of negations.²² In general, the notation $\varphi(\psi)$ stands for ψ occurs in φ . The usual definitions for Boolean and modal operators apply. Intuitively, $\mu p.\varphi(p)$ denotes the smallest formula p such that $p \leftrightarrow \varphi(p)$. The greatest fixpoint operator ν can be defined from μ as follows: $\nu p.\varphi(p) := \neg\mu p.\neg\varphi(\neg p)$.

We interpret \mathcal{L}^μ on influence models as follows:

Definition 8. *Let $\varphi \in \mathcal{L}^\mu$. The satisfaction of φ by a pointed influence model (\mathcal{M}, i) is inductively defined as follows:*

$$\begin{aligned} \mathcal{M}, i &\not\models \perp \\ \mathcal{M}, i &\models p \iff i \in \mathbf{O}(p), \text{ for } p \in \mathbf{P} \\ \mathcal{M}, i &\models \neg\varphi \iff i \notin \|\varphi\|_{\mathcal{M}} \\ \mathcal{M}, i &\models \varphi_1 \wedge \varphi_2 \iff i \in \|\varphi_1\|_{\mathcal{M}} \cap \|\varphi_2\|_{\mathcal{M}} \\ \mathcal{M}, i &\models \langle p \rangle \varphi \iff i \in \{j \mid \exists k : jG_p k \ \& \ k \in \|\varphi\|_{\mathcal{M}}\} \\ \mathcal{M}, i &\models \mu p.\varphi(p) \iff i \in \bigcap \{X \in 2^N \mid \|\varphi\|_{\mathcal{M}[p:=X]} \subseteq X\} \end{aligned}$$

where $\|\varphi\|_{\mathcal{M}[p:=X]}$ denotes the truth-set of φ once $\mathbf{O}(p)$ is set to be X . As usual, we say that: φ is valid in a model \mathcal{M} iff it is satisfied in all points of \mathcal{M} , i.e., $\mathcal{M} \models \varphi$; φ is valid in a class of models iff it is valid in all the models in the class.

We list some relevant known results about \mathbf{K}^μ . The logic has a sound and (weakly) complete axiom system [28]. The satisfiability problem of \mathbf{K}^μ is decidable [26]. The complexity of the model-checking problem for \mathbf{K}^μ is known to be in $\text{NP} \cap \text{co-NP}$ [19]. It is known that the model-checking problem for a formula of size m and alternation depth d on a system of size n can be solved by the natural fixpoint-approximation algorithm with (time) complexity of $O((m \cdot n)^{d+1})$ [12], where the alternation depth of a formula of \mathcal{L}^μ is the maximum number of μ/ν alternations in a chain of nested fixpoint subformulas.²³ Finally, the μ -calculus is known to be invariant for bisimulation (cf. [5]). It is actually known to correspond to the bisimulation-invariant fragment of monadic second-order logic [21].

5.3 On the logic of convergence in BDPs

Each stream of opinion profiles $\mathbf{O}^0, \mathbf{O}^1, \dots, \mathbf{O}^n, \dots$ corresponds to a stream of influence models $\mathcal{M}^0, \mathcal{M}^1, \dots, \mathcal{M}^n, \dots$

From the point of view of an influence model $\mathcal{M} = \langle N, \mathbf{G}, \mathbf{O} \rangle$ the BDP dynamics of Definition 4 can therefore be recast in terms of updates of the valuation function \mathbf{O} as follows:

²²This syntactic restriction guarantees that every formula $\varphi(p)$ defines a set transformation which preserves \subseteq , which in turn guarantees the existence of least and greatest fixpoints by the Knaster-Tarski fixpoint theorem (cf. [25]).

²³The reader is referred to, e.g. [13], for the precise definition.

- Base: $\mathbf{O}^0 := \mathbf{O}$
- Step: $\mathbf{O}^{n+1}(p) := \|[p]p\|_{\mathcal{M}^n}$.

That is, the interpretation of p at step $n + 1$ is the interpretation of $[p]p$ at step n . Equivalently, the interpretation of $\neg p$ at step $n + 1$ is the interpretation of $[p]\neg p$ at step n .

Lemma 4. *Let $\mathcal{M} = \langle N, \mathbf{G}, \mathbf{O} \rangle$ be an influence model. The two following statements are equivalent:*

1. $i \in N$ is stable for p ;
2. The pointed model (\mathcal{M}, i) satisfies:²⁴

$$\text{stb}(p) := \nu x. \pm p \wedge [p]x \quad (12)$$

Proof. First of all observe that, by the semantics of the μ -calculus, formula $\text{stb}(p)$ denotes the largest fixpoint of function $\pm p \wedge [p](\cdot)$, that is, formula $[p^*]\pm p$ where $[p^*]$ is the modal box interpreted over the reflexive and transitive closure of G_p . $\boxed{1 \Rightarrow 2}$ Assume that i is stable for p and suppose towards a contradiction that $\mathcal{M}, i \not\models \text{stb}(p)$. By what said above, it follows that there exists a j such that $\mathbf{O}_i(p) \neq \mathbf{O}_j(p)$ which is connected by a finite G_p path to i . By the functionality of influence models and the dynamics of Definition 4 then at some stage n in the stream of opinion profiles it should hold that $\mathbf{O}_i^n(p) = \mathbf{O}_j(p)$, against the assumption that i be stable for p . $\boxed{2 \Rightarrow 1}$ Assume $\mathcal{M}, i \models \text{stb}(p)$. By what said above, this implies that there exists no j such that $\mathbf{O}_i(p) \neq \mathbf{O}_j(p)$ which is connected by a finite G_p path to i . It follows that in the stream generated by the BDP dynamics i cannot change its opinion, and hence it is stable. \square

Theorem 6. *Let $\mathcal{M} = \langle N, \mathbf{G}, \mathbf{O} \rangle$ be an influence model. The two following statements are equivalent:*

1. $i \in N$ stabilizes for issue $p \in \mathbf{P}$;
2. The pointed model (\mathcal{M}, i) satisfies:

$$\mu x. \text{stb}(p) \vee [p]x \quad (13)$$

Proof. First of all observe that, by the semantics of the μ -calculus $\mu x. \text{stb}(p) \vee [p]x$ denotes the smallest fixpoint of equation $x \leftrightarrow \text{stb}(p) \vee [p]x$. By the Knaster-Tarski theorem and the fact that influence models are finite, we can compute such fixpoint as $\bigcup_{0 \leq n < \omega} \|\text{stb}(p)^n\|$ where $\|\text{stb}(p)^0\| = \|\text{stb}(p) \vee [p]\perp\|$ (notice that $[p]\perp \leftrightarrow \perp$ on influence models) and $\|\text{stb}(p)^{n+1}\| = \|\text{stb}(p) \vee [p]\text{stb}(p)^n\|$. So, by Lemma 4 i belongs to $\|\mu x. \text{stb}(p) \vee [p]x\|$ either i is stable for issue p or has

²⁴Notice that $\pm p$ is used as a variable ranging over $\{p, \neg p\}$. Technically the above formula is to be read as a scheme for $\nu x. p \wedge [p]x$ and $\nu x. \neg p \wedge [p]x$.

access in a finite number of steps to an agent who is stable for p . $\boxed{1) \Rightarrow 2)}$
 Assume that i stabilizes for issue $p \in \mathbf{P}$. So there exists a stage n in the stream of profiles generated through Definition 4 at which $\mathbf{O}_i^n(p) = \mathbf{O}_i^m(p)$ for all $m > n$. By Lemma 4, $\langle N, \mathbf{G}, \mathbf{O}^n \rangle, i \models \text{stb}(p)$. It follows that i is connected through a finite G_p -path to an agent j such that $\mathcal{M}, j \models \text{stb}(p)$. By what established above we thus have that $\mathcal{M}, i \models \mu x. \text{stb}(p) \vee [p] x$. $\boxed{2) \Rightarrow 1)}$ Assume $\mathcal{M}, i \models \mu x. \text{stb}(p) \vee [p] x$. It follows that i is connected through a finite G_p -path to an agent j such that $\mathcal{M}, j \models \text{stb}(p)$. By Lemma 4 j is therefore stable and therefore i will stabilize for p . \square

So the formula that expresses the stabilization of the agents' opinions on one issue is $\mu x. (\nu y. \pm p \wedge [p] y) \vee [p] x$. Informally, the theorem states that in a BDP an agent reaches a stable opinion if and only if it has an indirect influencer (linked by an influence path) whose all direct and indirect influencer have the same opinion. Notice that such formula has alternation depth 0. So an off-the-shelf model-checking algorithm for the μ -calculus can check stabilization in time $O(m \cdot n)$ with n being the size of the model and m the size of the formula.

Now confront this with the earlier Theorem 6. Since the convergence of the BDP is equivalent to the stabilization of all agents on all issues p (either on p or $\neg p$), we have the following corollary:

Corollary 1. *The BDP for an opinion profile \mathbf{O} based on influence graph \mathbf{G} converges if and only if*

$$\langle N, \mathbf{G}, \mathbf{O} \rangle, i \models U \left(\bigwedge_{p \in \mathbf{P}} \mu x. \text{stb}(p) \vee [p] x \right) \quad (14)$$

for any agent $i \in N$, where U denotes the universal modality (cf. [5]).

So the above formula characterizes the property of convergence for a BDP. Since the process of voting in a liquid democracy system can be modeled by a BDP, the formula also characterizes precisely when voting by delegable proxy results in a $\mathbf{1}$ or $\mathbf{0}$ opinion on a given issue.

6 Conclusions

The paper has moved the first steps towards the development of theoretical foundations for the voting system of liquid democracy based on delegable proxy.

We have pursued two lines of research linked to two interpretations commonly associated to the proxy character of liquid democracy: the delegation of voting right to trustees, vs. the copying of the votes of influencers. The first interpretation has led us to develop a simple model of liquid democracy based on the theory of binary and judgment aggregation. This has allowed us to study liquid democracy as a form of binary aggregation with abstentions. The second interpretation has led us to study liquid democracy through extremely simple

models of opinion diffusion corresponding to the Boolean special case of the stochastic processes of opinion diffusion known as DeGroot processes. We have argued that studying aggregation in liquid democracy through this lens offers important advantages with respect to the handling of delegation cycles and the preservation of individual rationality. Through this second perspective we have also shown how off-the-shelf logical techniques can be used to analyze properties (such as convergence) of the diffusion process underpinning liquid democracy.

References

- [1] Dan Alger. Voting by proxy. *Public Choice*, 126(1-2):1–26, jan 2006.
- [2] J. Behrens, A. Kistner, A. Nitsche, and B. Swierczek. Principles of Liquid Feedback. 2014.
- [3] Jan Behrens. Cyclical delegations: Myths or disaster. *Liquid Democracy Journal*, 3, 2015.
- [4] Johan van Benthem. Oscillations, Logic, and Dynamical Systems. In Sujata Ghosh and Jakub Szymanik, editors, *The Facts Matter. Essays on Logic and Cognition in Honour of Rineke Verbrugge*, pages 9–22. College Publications, 2015.
- [5] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge University Press (CUP), 2001.
- [6] Paolo Boldi, Francesco Bonchi, Carlos Castillo, and Sebastiano Vigna. Viscous democracy for social networks. *Communications of the ACM*, 54(6):129, jun 2011.
- [7] Sirin Botan. Propositional opinion diffusion with constraints. *ILLC Master of Logic Thesis*, 2016.
- [8] Zoé Christoff and Jens Ulrik Hansen. A logic for diffusion in social networks. *Journal of Applied Logic*, 13(1):48–77, mar 2015.
- [9] Morris H. Degroot. Reaching a Consensus. *Journal of the American Statistical Association*, 69(345):118–121, mar 1974.
- [10] Franz Dietrich and Christian List. Judgment aggregation without full rationality. *Social Choice and Welfare*, 31(1):15–39, jul 2007.
- [11] Elad Dokow and Ron Holzman. Aggregation of binary evaluations with abstentions. *Journal of Economic Theory*, 145(2):544–561, mar 2010.
- [12] E. Allen Emerson. Model Checking and the Mu-calculus. In *Descriptive Complexity and Finite Models, Proceedings of a DIMACS*, 1996.

- [13] E. Allen Emerson, Charanjit S. Jutla, and A. Prasad Sistla. On model checking for the mu-calculus and its fragments. *Theoretical Computer Science*, 258(1-2):491–522, may 2001.
- [14] Umberto Grandi and Ulle Endriss. Lifting integrity constraints in binary aggregation. *Artif. Intell.*, 199:45–66, 2013.
- [15] Umberto Grandi, Emiliano Lorini, and Laurent Perrussel. Propositional Opinion Diffusion. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, pages 989–997, Richland, SC, 2015. International Foundation for Autonomous Agents and Multiagent Systems.
- [16] James Green-Armytage. Direct democracy by delegable proxy. 2005.
- [17] James Green-Armytage. Direct voting and proxy voting. *Const Polit Econ*, 26(2):190–220, dec 2014.
- [18] Davide Grossi and Gabriella Pigozzi. Judgment Aggregation: A Primer. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 8(2):1–151, mar 2014.
- [19] Erich Grädel and Martin Otto. On logics with two variables. *Theoretical Computer Science*, 224(1-2):73–113, aug 1999.
- [20] Matthew O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, NJ, USA, 2008.
- [21] David Janin and Igor Walukiewicz. On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. In *CONCUR '96: Concurrency Theory*, pages 263–277. Springer, 1996.
- [22] Fenrong Liu, Jeremy Seligman, and Patrick Girard. Logical dynamics of belief change in the community. *Synthese*, 191(11):2403–2431, 2014.
- [23] Kenneth O. May. A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decision. *Econometrica*, 20(4):680, oct 1952.
- [24] James C. Miller. A program for direct and proxy voting in the legislative process. *Public Choice*, 7-7(1):107–113, sep 1969.
- [25] Colin Stirling. Modal Mu-Calculus. In *Texts in Computer Science*, pages 103–132. Springer, 2001.
- [26] Robert S. Streett and E. Allen Emerson. The propositional mu-calculus is elementary. In *Automata Languages and Programming*, pages 465–472. Springer, 1984.
- [27] Gordon Tullock. Computerizing politics. *Mathematical and Computer Modelling*, 16(8-9):59–65, aug 1992.

[28] Igor Walukiewicz. Completeness of Kozen's Axiomatisation of the Propositional mu-Calculus. *Information and Computation*, 157(1-2):142–182, feb 2000.

A Binary aggregation (without abstention)

The formalism of choice for this paper is binary aggregation [14]. A binary aggregation structure (*BA structure*) is a tuple $\mathcal{A} = \langle N, \mathbf{P}, \gamma \rangle$ where:

- $N = \{1, \dots, n\}$ is a finite set individuals s.t. $|N| = n \in \mathbb{N}$;
- $\mathbf{P} = \{p_1, \dots, p_m\}$ is a finite set of issues ($|\mathbf{P}| = m \in \mathbb{N}$), each represented by a propositional atom;
- $\gamma \in \mathcal{L}$ is an (integrity) constraint, where \mathcal{L} is the propositional language constructed by closing \mathbf{P} under a functionally complete set of Boolean connectives (e.g., $\{\neg, \wedge\}$)

An *opinion* $O : \mathbf{P} \rightarrow \{\mathbf{0}, \mathbf{1}\}$ is an assignment of truth values to the set of issues \mathbf{P} , and the set of all opinions is denoted by \mathcal{O} . The opinion of an agent i is said to be “consistent” whenever $O_i \models \gamma$, that is, i 's opinion satisfies the integrity constraint. The set of all consistent opinions is denoted $\mathcal{O}_c = \{O \in \mathcal{O} \mid O \models \gamma\}$. Thus, $O(p) = \mathbf{0}$ (respectively, $O(p) = \mathbf{1}$) indicates that opinion O rejects (respectively, accepts) the issue p . Syntactically, the two opinions correspond to the truth of the literals p or $\neg p$. For $p \in \mathbf{P}$ we write $\pm p$ to denote one element from $\{p, \neg p\}$. An *opinion profile* $\mathbf{O} = (O_1, \dots, O_n)$ records the opinion, on the given set of issues, of every individual in N . Given a profile \mathbf{O} the i^{th} projection \mathbf{O} is denoted O_i (i.e., the opinion of agent i in profile \mathbf{O}). We also denote by $\mathbf{O}(p) = \{i \in N \mid O_i(p) = \mathbf{1}\}$ the set of agents accepting issue p in profile \mathbf{O} and by $\mathbf{O}(p^-) = \{i \in N \mid O_i(p) = \mathbf{0}\}$.

Given a BA structure \mathcal{A} , an aggregation rule (or *aggregator*) for \mathcal{A} is a function $F : (\mathcal{O}_c)^N \rightarrow \mathcal{O}$, mapping every profile of consistent opinions to one collective opinion in \mathcal{O} . $F(\mathbf{O})(p)$ denotes the outcome of the aggregation on issue p . A benchmark aggregator is *issue-by-issue strict majority rule* (**maj**), which accepts an issue if and only if the majority of the population accepts it:

$$\text{maj}(\mathbf{O})(p) = \mathbf{1} \iff |\mathbf{O}(p)| \geq \frac{|N| + 1}{2}. \quad (15)$$

It is well-known that aggregation by majority does not preserve consistency. The standard example is provided by the discursive dilemma, represented by the BA structure $\langle \{1, 2, 3\}, \{p, q, r\}, r \leftrightarrow p \wedge q \rangle$. The profile consisting of $O_1 \models p \wedge q \wedge r$, $O_2 \models p \wedge \neg q \wedge \neg r$, $O_3 \models \neg p \wedge q \wedge \neg r$, returns an inconsistent majority opinion $\text{maj}(\mathbf{O}) \models p \wedge q \wedge \neg r$.

B Relevant terminology from graph theory

Let $G = \langle N, R \rangle$ be a graph and R^* be the transitive and symmetric closure of R . A *path* is a sequence of nodes $\langle i_1, \dots, i_k \rangle$, such that, for all $l \in \{1, \dots, k\}$, $i_l R i_{l+1}$. The *distance* between two nodes i, j is the length of the shortest path $\langle i, \dots, j \rangle$ between them. The *diameter* of a graph is the maximal distance between any two nodes related by a path. A *cycle* is a path of length k such that $i_1 = i_k$. A set of nodes $S \subseteq N$ is said to be:

a cycle in G if all elements in S are in one cycle of length $|S|$,

connected if for any $i, j \in S$: $i R^* j$,

strongly connected if for any $i, j \in S$: there is a path $\langle i, \dots, j \rangle$,

closed if for any $i \in S, j \notin S$, it is not the case that $i R j$,

a connected component if for any $i, j \in N$: $i R^* j$ if and only if $i, j \in S$,

aperiodic if the greatest common divisor of the lengths of its cycles is 1.