# Pairwise Liquid Democracy

**Markus Brill**[1]  and  **Nimrod Talmon**[2]

[1] TU Berlin, Berlin, Germany
[2] Ben-Gurion University, Be'er Sheva, Israel
brill@tu-berlin.de, talmonn@bgu.ac.il

## Abstract

In a liquid democracy, voters can either vote directly or delegate their vote to another voter of their choice. We consider ordinal elections, and study a model of liquid democracy in which voters specify partial orders and use several delegates to refine them. This flexibility, however, comes at a price, as individual rationality (in the form of transitive preferences) can no longer be guaranteed. We discuss ways to detect and overcome such complications. Based on the framework of distance rationalization, we introduce novel variants of voting rules that are tailored to the liquid democracy context.

## 1 Introduction

Liquid democracy is based on the paradigm of *delegative voting* and can be seen as a middle ground between direct democracy and representative democracy (see, e.g., [Blum and Zuber, 2016; Behrens *et al.*, 2014]). Under this paradigm, each voter can choose to either vote directly, or to select another voter acting as her delegate. Delegation is transitive, in the sense that the delegate can choose to delegate the vote further, resulting in delegation paths along which voting weight is accumulated. The voting weight of a voter who decides to actually vote is then given by the number of voters who—directly or indirectly—delegated their vote to this voter.

Liquid democracy has been popularized by the German Pirate Party, which was one of the first organizations to employ the *LiquidFeedback* software [Behrens *et al.*, 2014] for decision making within the party. Similar tools have been used in other political parties, such as the Spanish party *Partido de Internet*, the local Swedish party *Demoex*, and the Argentinian "Net Party" (*Partido de la Red*).

Liquid democracy is appealing due to the flexibility it offers to voters. In this paper, we aim to increase this flexibility further by introducing the concept of *pairwise liquid democracy*. Specifically, we consider *ordinal* elections, where each voter needs to rank-order a set of alternatives. This setting, which is standard in social choice theory, is in fact used in *LiquidFeedback* [Behrens *et al.*, 2014, Chapter 14.12]. We propose that voters in an ordinal election can make use of delegations in a fine-grained manner, by specifying a partial

order and then using *several delegates* to refine different parts of that partial order. We illustrate the idea with an example.

**Example 1.** Consider the alternatives $a$, $b$, $c$, $d$, and assume that a voter $v$ feels strongly about $a$ being the best alternative, but is not sure how the remaining three alternatives compare. Moreover, voter $v$ trusts another voter $v'$ to know whether $b$ is better than $c$ (maybe $b$ and $c$ differ mainly in one aspect, on which $v'$ happens to be an expert) and voter $v$ moreover thinks that the comparison between $c$ and $d$ could be best performed by a third voter $v''$ (maybe $v''$ has a lot of experience with both $c$ and $d$). For the sake of the example, say that voter $v'$ decides that $b$ is better than $c$ and voter $v''$ decides that $c$ is better than $d$. Then, the resulting ranking of $v$, after taking the delegations into account, would be $a \succ b \succ c \succ d$.[1]

Notice the flexibility voter $v$ used in her delegations, and how $v'$ and $v''$ helped to refine her initial partial order. This flexibility is useful, e.g., for voters who (a) really care about the outcome of an election, but (b) do not feel well-informed or competent enough to compare certain pairs of alternatives. This flexibility, as might be suspected, comes at a price: Refining a partial order by delegating certain pairwise comparisons to different delegates might result in intransitive preference orders (see Figure 1 for a concrete example).

In this paper, we study such combinatorial complications, and discuss ways to overcome them. Specifically, we consider the problem of deciding whether a given delegation graph is consistent (i.e., whether, after taking into account delegations, all preference orders are transitive). For inconsistent delegation graphs, we then consider several approaches to cope with intransitive preference orders. Finally, we identify a rich family of voting rules which operate directly on such delegation graphs. These rules generalize standard voting rules and could be considered more appropriate for the pairwise liquid democracy setting.

An important appeal of liquid democracy is the flexibility it offers to participants. We argue that extending this flexibility further, specifically by allowing pairwise delegations,

---

[1] Note that, in this example, voter $v$ did not explicitly delegate the comparison between $b$ and $d$. In a sense, voter $v$ got lucky that the decisions by her delegates $v'$ and $v''$ resulted in a consistent ranking. Later we will concentrate on situations in which voters are required to delegate all pairwise comparisons that are not decided by themselves; we will see that this might result in certain inconsistencies.

holds a significant potential for advancing the liquid democracy paradigm.

RELATED WORK. While some of the ideas behind liquid democracy can be traced back to early works of [Dodgson, 1884] and [Miller, 1969], the idea has been mainly developed since the early 2000s (see, e.g., [Ford, 2002]). For an historical overview of the development of the idea, we refer to the surveys by [Ford, 2014] and [Behrens, 2017].

In recent years, liquid democracy has gained increasing attention from the scientific community. [Boldi *et al.*, 2011] studied a variant of liquid democracy called *viscous democracy*, which uses a discount factor for dampening the impact of long delegation paths. Attempts to formally establish the virtues of liquid democracy have been undertaken by [Green-Armytage, 2015] and [Cohensius *et al.*, 2017], who considered spatial voting models and showed that liquid democracy outperforms direct democracy under certain conditions. On the other hand, [Kahng *et al.*, 2018] have established that liquid democracy with so-called *local delegation mechanisms* cannot truly outperform direct democracy with respect to the ability of recovering a ground truth.

Perhaps the most related work to our paper is that of [Christoff and Grossi, 2017], who consider delegative voting in the binary aggregation framework. In this setting, there are several binary issues and voters can delegate their vote on individual issues to different voters, potentially resulting in voter preferences that are not "individually rational" because they violate an integrity constraint. Our model of ordinal voting can be seen as a special case, where issues are pairwise comparisons between alternatives and the integrity constraint requires the transitivity of preferences. Focussing on the important special case of ordinal voting allows us to study specific problems, and to come up with specific solutions, that might not manifest in the more general framework.

## 2 Preliminaries

Let $N = \{1, \ldots, n\}$ be a finite set of $n$ voters and $A = \{a, b, c, \ldots\}$ a finite set of $|A| = m$ alternatives. The preferences of voter $i \in N$ are modeled as an asymmetric[2] binary relation $R_i \subset A \times A$. The interpretation of $(a, b) \in R_i$, which is denoted $a \succ_i b$ (or $a \succ b$ if $i$ is obvious from the context), is that voter $i$ strictly prefers alternative $a$ over alternative $b$. A preference relation $R_i$ is *complete* if $a \succ_i b$ or $b \succ_i a$ for every pair of distinct $a, b \in A$, and it is *transitive* if $a \succ_i b$ and $b \succ_i c$ implies $a \succ_i c$ for all distinct $a, b, c \in A$. For a preference relation $R_i$ and a subset $A' \subseteq A$ of alternatives, we let $R_i|_{A'}$ denote the restriction of $R_i$ to $A' \times A'$.

A *ranking* (or *linear order*) over $A$ is a transitive and complete preference relation and is denoted $a_1 \succ a_2 \succ \ldots \succ a_m$, with the understanding that $a_j \succ a'_j$ if and only if $j < j'$. A *weak ranking* over $A$ is a transitive preference relation that can be described by an ordered partition $(A_1, A_2, \ldots, A_k)$ and the assumption that $a \succ b$ if and only if there exits $j$ and $j'$ such that $a \in A_j$, $b \in A_{j'}$, and $j < j'$. For example, $\{a, b\} \succ c \succ \{d, e\}$ denotes the weak ranking with
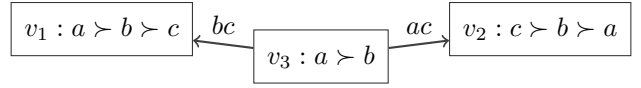


Figure 1: Intransitive preferences due to pairwise delegations.

$A_1 = \{a, b\}$, $A_2 = \{c\}$, and $A_3 = \{d, e\}$. Each set $A_i$ in the ordered partition is called an *indifference class*.

A preference relation $R_i$ *contains a preference cycle* if there exists a $k \geq 3$ and alternatives $a_1, a_2, \ldots, a_k$ such that $a_1 \succ_i a_2$, $a_2 \succ_i a_3$, ..., $a_k \succ_i a_1$. A complete preference relation is transitive iff it does not contain a preference cycle.

A *preference profile* is a list $R = (R_1, \ldots, R_n)$ containing a preference relation $R_i$ for every voter $i \in N$. A preference profile $R$ is *complete* (resp., *transitive*) if every preference relation in it is complete (respectively, transitive). A preference profile $R' = (R'_1, \ldots, R'_n)$ is an *extension* of a preference profile $R = (R_1, \ldots, R_n)$ if $R_i \subseteq R'_i$ for all $i \in N$.

## 3 Pairwise Delegations

We introduce the model of *pairwise liquid democracy*, which strictly generalizes the standard liquid democracy paradigm.

Let $\mathcal{P}$ denote the set of all unordered pairs of distinct alternatives. For a pair $\{a, b\} \in \mathcal{P}$, we usually write $ab$; in particular, $ab$ and $ba$ refer to the same unordered pair. Every voter $i \in N$ divides the set $\mathcal{P}$ into two sets: the set of *internal pairs* $\mathcal{P}_{\text{int}}(i)$ and the set of *external pairs* $\mathcal{P}_{\text{ext}}(i)$. For every internal pair $ab$, the voter specifies her (strict) preferences between the two alternatives in question ($a \succ_i b$ or $b \succ_i a$).[3] For every external pair $cd$, the voter designates another voter $j \in N \setminus \{i\}$ as the *pairwise delegate* for that pair.[4]

PAIRWISE DELEGATION GRAPHS. The collection of all internal and external pairs can be represented by a directed graph with $n$ vertices, where each vertex $v_i$ corresponds to a voter $i \in N$. Every node $v_i \in N$ is labeled with the preferences over the internal pairs $\mathcal{P}_{\text{int}}(i)$ of voter $i$. A directed arc from $v_i$ to $v_j$ has a label $ab \in \mathcal{P}$ and corresponds to the external pair $ab \in \mathcal{P}_{\text{ext}}(i)$ that voter $i$ delegates to voter $j$. Notice that there could be parallel edges (if $i$ delegates several pairs to $j$). We refer to such a graph as a *pairwise delegation graph*.

Pairwise delegation graphs can sometimes lead to intransitive preference relations, as the next example demonstrates.

**Example 2.** Consider the pairwise delegation graph $G$ depicted in Figure 1 with $n = 3$ voters and alternatives $a, b, c$.

---

[2]Asymmetry requires that $(a, b) \in R_i$ implies $(b, a) \notin R_i$ for all $a, b \in A$. It also implies *irreflexivity*, i.e., $(a, a) \notin R_i$ for all $a \in A$.

[3]Internally specified preferences are assumed to be transitive.

[4]We usually assume that each voter specifies every pair in $\mathcal{P}$ either as an internal pair or as an external pair. Sometimes, however, we consider instances (such as the one constructed in the proof of Theorem 2) without specifying explicitly whether certain pairs are internal or external. In those cases, such "unspecified" pairs can be arbitrarily set to either internal or external pairs without impacting the construction. A simple, general way of reducing an instance with unspecified pairs to one without such pairs is to add some additional, dummy voters: For each voter $i$ containing at least one unspecified pair, add a voter $i'$, initially with the same internal and external pairs as voter $i$, and add each unspecified pair of $i$ as an external pair for both $i$ and $i'$, who delegate those pairs to each other. The modified instance preserves (weak) consistency (see Definitions 1 and 2).

Voter 1 internally declares the ranking $a \succ_1 b \succ_1 c$ and voter 2 internally declares the ranking $c \succ_2 b \succ_2 a$. Voter 3 internally declares $a \succ_3 b$, delegates $bc$ to voter 1, and delegates $ac$ to voter 2. As a result, the preference relation of voter 3 contains a preference cycle: $a \succ_3 b$, $b \succ_3 c$, $c \succ_3 a$.

The following notation will be useful for the remainder of the paper. For a voter $i \in N$ and an external pair $ab \in \mathcal{P}_{\text{ext}}(i)$, let $del_{ab}(i)$ denote the *ultimate delegate* of $i$ with respect to $ab$; it can be found by following the path in $G$ starting in $v_i$ and following the $ab$-labeled arcs. There are two possibilities: either the path leads to a voter $j \in N \setminus \{i\}$ with $ab \in \mathcal{P}_{\text{int}}(j)$ or the path leads to a cycle (which we call an $ab$-cycle). In the former case, we let $del_{ab}(i) = j$ (note that $j$ is uniquely determined by $G$); in the latter case, we write $del_{ab}(i) = \emptyset$ and say that $i$ has no ultimate delegate with respect to $ab$.

The arc-set of a pairwise delegation graph $G$ is the union of $\binom{m}{2}$ arc-sets, each corresponding to a pair in $\mathcal{P}$. For a pair $ab \in \mathcal{P}$, the $ab$-*graph* is the subgraph of $G$ containing only those arcs labeled $ab$. $G$ *does not contain pairwise delegation cycles* if, for each $ab \in \mathcal{P}$, $G$ does not contain an $ab$-cycle.

FROM GRAPHS TO PREFERENCE PROFILES. For a pairwise delegation graph $G$, $R(G)$ is the preference profile resulting from resolving all delegations. That is, for each voter $i \in N$ and each external pair $ab$ of $i$, we find the ultimate delegate $del_{ab}(i)$ (if it exists) and add the corresponding ordered pair (either $(a, b)$ or $(b, a)$) to $R_i$. If $del_{ab}(i) = \emptyset$ (implying that there is an $ab$-cycle), however, then neither $(a, b)$ nor $(b, a)$ is added to $R_i$. Indeed, whenever $G$ contains pairwise delegation cycles, then some ultimate delegates are not defined and, consequently, $R(G)$ is not complete.

A preference profile $R$ *respects* a pairwise delegation graph $G$ if for all voters $i \in N$ and all $ab \in \mathcal{P}_{\text{ext}}(i)$, $R_i|_{ab} = R_j|_{ab}$, where $j$ is the voter that $i$ delegates $ab$ to. (If $del_{ab}(i) = \emptyset$, then $R_i|_{ab} = \emptyset$.) By definition, $R(G)$ respects $G$. Sometimes we are interested in complete and transitive preference profiles respecting $G$. The set of all such profiles is denoted by $\hat{R}(G)$. Figure 1 illustrates that $\hat{R}(G)$ may be empty. A necessary condition for a profile to be contained in $\hat{R}(G)$ is to extend $R(G)$. If $G$ does not contain pairwise delegation cycles and $\hat{R}(G) \neq \emptyset$, then $\hat{R}(G) = \{R(G)\}$.

## 4 Detecting Intransitivities

Example 2 demonstrates that pairwise delegations can potentially yield intransitive preferences for individual voters, even when the internal pairs satisfy transitivity. Whether such intransitive preferences arise depends on the pairwise delegation graph in question, and it might not be straightforward to check whether a given graph yields intransitivities or not. Next, we analyze the detection of intransitivities from a computational complexity perspective. A first step in checking for inconsistencies in voters' preferences is to propagate pairwise comparisons along delegation arcs and check whether this results in violations of transitivity. If this is not the case, we call the pairwise delegation graph weakly consistent.

**Definition 1** (weak consistency). A pairwise delegation graph $G$ is *weakly consistent* if no preference relation in $R(G)$ contains a preference cycle.
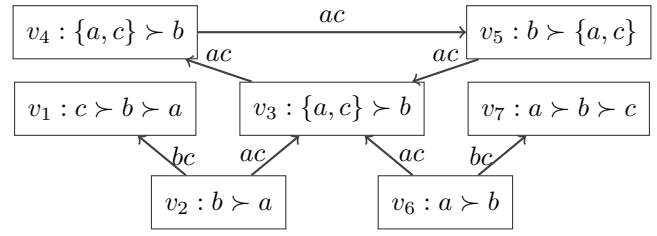


Figure 2: A weakly consistent pairwise delegation graph that is not consistent. There are two ways to "fix" the $ac$-cycle: either set $a \succ c$ for all voters involved in (or delegating to) the cycle, resulting in intransitive preferences for voter 2; or set $c \succ a$ for all these voters, resulting in intransitive preferences for voter 6.

For example, the pairwise delegation graph in Figure 1 violates weak consistency.

Weak consistency can be checked efficiently: It is sufficient to observe that, for each voter $i \in N$ and for each pair $ab \in \mathcal{P}_{\text{ext}}(i)$, the ultimate delegate $del_{ab}(i)$ can be computed efficiently, e.g., using a depth first search.

**Observation 1.** *It can be checked in polynomial time whether a pairwise delegation graph is weakly consistent.*

Even if pairwise delegations do not result in preference cycles, it could be the case that $R(G)$ cannot be extended to a complete and transitive preference profile respecting all delegations. In other words, the set $\hat{R}(G)$ may be empty even if $G$ is weakly consistent. This is illustrated in Figure 2. The following stronger notion of consistency forbids such "hidden" inconsistencies.

**Definition 2** (consistency). A pairwise delegation graph $G$ is *consistent* if $\hat{R}(G) \neq \emptyset$; i.e., if there exists a complete and transitive $G$-respecting extension of $R(G)$.

Recall that each incompleteness of $R(G)$ (i.e., each preference relation in $R(G)$ for which some pairwise comparisons are unknown) is caused by a pairwise delegation cycle. Thus, a pairwise delegation graph is consistent if and only if, for every pairwise delegation cycle of $G$, the preferences over the respective pair can be "fixed" in a way that does not result in intransitive individual preferences. In particular, profiles in $\hat{R}(G)$ have the property that all voters that are contained in an $ab$-cycle (and also those contained in an $ab$-path leading to that cycle) have identical preferences over $ab$ (either $a \succ b$ or $b \succ a$). In the absence of pairwise delegation cycles, $R(G)$ is complete (as ultimate delegates are always defined); consequently, the two consistency notions coincide.

**Proposition 1.** *A pairwise delegation graph not containing pairwise delegation cycles is weakly consistent if and only if it is consistent.*

It turns out that, in contrast to checking *weak* consistency, checking consistency is computationally intractable.

**Theorem 2.** *Deciding whether a given pairwise delegation graph is consistent is NP-complete.*

*Proof.* Membership in NP follows directly from Proposition 1, by guessing the decisions made on the cycles. Next, we provide a reduction from the NP-hard problem 3SAT, in
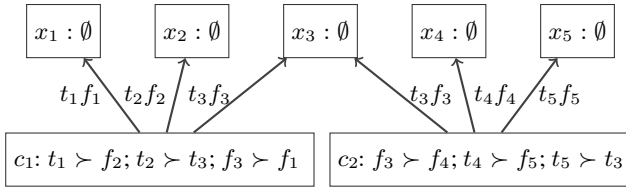
Figure 3: Example for the reduction described in the proof of Theorem 2, corresponding to $\phi = C_1 \wedge C_2$, where $C_1 = x_1 \vee x_2 \vee \bar{x}_3$ and $C_2 = \bar{x}_3 \vee x_4 \vee x_5$.

which we have a Boolean 3-CNF formula $\phi$ with $m$ clauses $C_j$, $j \in [m]$, over $n$ variables $x_i$, $i \in [n]$, and shall decide whether $\phi$ is satisfiable. Given an instance $\phi$ of 3SAT, we create a pairwise delegation graph $G$ as follows.

For each variable $x_i$, $i \in [n]$, we create a voter $x_i$ and two alternatives, $t_i$ and $f_i$. For each clause $C_j$, $j \in [m]$, we create a voter $c_j$. If $C_j = (x_{j1} \vee x_{j2} \vee x_{j3})$, then voter $c_j$ delegates the pair $t_{j1}f_{j1}$ to voter $x_{j1}$, the pair $t_{j2}f_{j2}$ to voter $x_{j2}$, and the pair $t_{j3}f_{j3}$ to voter $x_{j3}$. The idea is that the decision of voter $x_i$ for the pair $t_if_i$ shall correspond to the sign of the literal of $x_i$: $t_i \succ f_i$ corresponds to a positive literal, while $f_i \succ t_i$ corresponds to a negative literal. The internal pairs of voter $c_j$ correspond to the signs of its literals. To explain how, we consider several cases, depending on its form; for notational simplicity, assume that $C_j$ contains the variables $x_1$, $x_2$, and $x_3$, and consider the form of its negation, $\neg C_j$:

- If $\neg C_j = (x_1 \wedge x_2 \wedge x_3)$, then $c_j$ internally decides $f_1 \succ t_2$, $f_2 \succ t_3$, and $f_3 \succ t_1$.

- If $\neg C_j = (x_1 \wedge x_2 \wedge \neg x_3)$, then $c_j$ internally decides $f_1 \succ t_2$, $f_2 \succ f_3$, and $t_3 \succ t_1$.

- If $\neg C_j = (x_1 \wedge \neg x_2 \wedge \neg x_3)$, then $c_j$ internally decides $f_1 \succ f_2$, $a_2 \succ f_3$, and $t_3 \succ t_1$.

- If $\neg C_j = (\neg x_1 \wedge \neg x_2 \wedge \neg x_3)$, then $c_j$ internally decides $t_1 \succ f_2$, $t_2 \succ f_3$, and $t_3 \succ f_1$.

The idea is that, as satisfying $\neg C_j$ would cause $\phi$ to be unsatisfied, the internal decisions of $c_j$, combined with those decisions of the voters corresponding to the literals of $C_j$ which would satisfy $\neg C_j$ would cause a preference cycle in $c_j$'s preference relation. This finishes the reduction. An example is given in Figure 3. (Pairs not explicitly specified as either internal or external pair can be set arbitrarily without introducing preference cycles; see also Footnote 4.)

Correctness follows as any $G$-respecting complete extension of $R(G)$ shall decide, for each $x_i$, whether $t_i \succ f_i$ or $f_i \succ t_i$; these decisions would propagate to the $c_j$'s, causing a preference cycle in each $c_j$ for which $\neg C_j$ is satisfied. $\square$

**Remark 1.** There is an elegant SAT-formulation of the consistency problem. This is also of practical importance, as it allows the use of efficient SAT solvers. For each vertex $v$ and each pair $ab$, define a binary variable $v_{a,b}$, with the intended meaning that $v_{a,b} = 0$ if $a \succ_v b$ in a complete extension of $R(G)$, and $v_{a,b} = 1$ otherwise. For any delegation arc $(v, u)$ labeled $ab$, we add an arc-clause: $(u_{a,b} \to v_{a,b}) \wedge (\overline{u_{a,b}} \to \overline{v_{a,b}})$. For each vertex $v$ and each triplet of alternatives $a, b, c$ we add the transitivity constraints: $(v_{a,b} \wedge v_{b,c}) \to v_{a,c}$ and

$(\overline{v_{a,b}} \wedge \overline{v_{b,c}}) \to \overline{v_{a,c}}$. The corresponding formula is satisfiable if and only if $G$ is consistent; in this case, the variables encode a profile in $\hat{R}(G)$.

## 5 Coping with Intransitivities

It is unfortunate that allowing pairwise delegations might result in intransitive preferences. In this section we suggest several ways of coping with this problem. In particular, we explore approaches to prevent intransitive preferences by restricting allowed delegations, ways to circumvent intransitive preferences by ignoring delegations, and the possibility of consolidating intransitive preferences into weak rankings.[5]

### 5.1 Restricting Allowed Delegations

Intransitive preferences are the result of the flexibility of pairwise liquid democracy. It is therefore natural to consider restrictions to our general model, specifically tuning this flexibility to avoid such individually irrational outcomes.

Taken to the extreme, we might say that if a voter delegates some pair to another voter, then she shall delegate all $\binom{m}{2}$ pairs to that voter; indeed, this is a very cumbersome way to describe the standard (non-pairwise) delegation on which liquid democracy is based: Each voter can either specify her own ranking, or delegate her complete vote to another voter. A more flexible approach that still prevents intransitivities consists in letting each voter specify a weak ranking together with a list of delegates, one delegate for each indifference class of the weak ranking. The weak ranking of the voter would then be completed into a linear order by ranking the alternatives in each indifference class according to the preferences of the corresponding delegate. Assuming no pairwise delegation cycles, this approach is guaranteed to result in a complete and transitive preference profile. Note, however, that this approach restricts the freedom of voters to some extent: coming up with a weak ranking already requires a lot of internal decisions, and a voter might not be willing/competent to make those decisions.

### 5.2 Modifying Existing Delegations

Another approach to prevent intransitive preferences consists in modifying the delegation graph. Specifically, intransitivities can be circumvented by removing (ignoring) some pairwise delegations. E.g., recall the inconsistent pairwise delegation graph of Figure 1 and observe that removing any one of the two external pairs yields a consistent graph.

A natural objective is to make a pairwise delegation graph consistent by removing as few delegations as possible, as it would hopefully correspond to only a negligible change to the graph. This motivates the following computational problem, which we refer to as the *consistency modification problem*:

---

[5]This list is by no means exhaustive. For instance, [Christoff and Grossi, 2017] suggest to avoid inconsistencies by employing an opinion diffusion approach [DeGroot, 1974]. Under this approach, a voter incorporates the decisions of delegates only if it does not introduce inconsistencies. Though we do not consider this idea further in this paper, we note that opinion diffusion can be applied on the level of pairwise comparisons of alternatives [Brill *et al.*, 2016].

*Given a pairwise delegation graph $G$ and an integer $b \geq 0$, can at most $b$ arcs be removed from $G$ in order to make $G$ consistent?*

The consistency modification problem can be defined for both weak consistency and consistency as the desired goal. The budget $b$ plays an important role, as with a high enough budget any graph can be made consistent (e.g., the empty graph is always consistent). The problem is intractable.

**Theorem 3.** *The consistency modification problem is NP-complete, both for weak consistency and for consistency.*

*Proof sketch.* NP-hardness follows from a straightforward reduction from *Feedback Arc Set on Tournaments*. Given a tournament $T = (W, E)$, we construct a pairwise delegation graph $G$ with voter set $N = v^* \cup \{v_e : e \in E\}$. The set of alternatives is defined by the set $W$ of vertices of $T$. Voter $v^*$ has only external pairs: for every directed edge $e = (x, y)$ of the tournament $T$, voter $v^*$ delegates the pair $xy$ to voter $v_e$, who has only internal pairs, and in particular prefers $x$ to $y$.

Now we can make $G$ consistent (or weakly consistent) by removing $b$ external pairs if and only if we can make $T$ acyclic by removing $b$ directed edges from $E$. □

**Remark 2.** The consistency modification problem can be efficiently solved using MaxSAT solvers, by constructing the formula described in Remark 1, setting the weight of the arc-clauses to 1, and setting the weight of other clauses to $\infty$. Indeed, a maximum weight MaxSAT solution corresponds to a solution of the consistency modification problem, with unsatisfied arc-clauses corresponding to deleted arcs.

### 5.3 Consolidating into Weak Rankings

A further way to cope with intransitive voter preferences is to first "consolidate" those intransitivities into weak rankings, and then apply a voting rule which can cope with weak rankings. One way to accomplish this is by applying a *tournament solution* (e.g., the *top cycle* or the *Copeland set*), which takes an asymmetric and complete[6] binary relation on a set of alternatives as input and outputs a non-empty subset of alternatives [Laslier, 1997; Brandt *et al.*, 2016]. These can be used iteratively to construct a weak rankings as follows: First, the tournament solution is applied to the set of all alternatives and the resulting subset of alternatives then defines the top-most indifference class of the weak ranking. These alternatives are then removed from consideration and the tournament solution is applied to the set of remaining alternatives to determine the next indifference class, and so on. This approach has been analyzed by [Bouyssou, 2004] and (for the special case of two indifference classes) [Yang, 2017].

## 6 Voting Rules for Delegation Graphs

The need to prevent, consolidate, or circumvent intransitivities can be rendered moot by employing a voting rule that accepts intransitive voter preferences. In particular, this is the

case for all pairwise voting rules (a voting rule is *pairwise*[7] if it only depends on anonymized comparisons between pairs of alternatives). The motivation for this approach is rather straightforward: The premise of pairwise liquid democracy is that domain expertise in the form of high-accuracy pairwise comparisons should be exploited; and pairwise voting rules take *all* pairwise comparisons into account.

An important example of a pairwise rule is *Kemeny's rule* [Kemeny, 1959], which enjoys a strong axiomatic foundation [Young and Levenglick, 1978] and has been established as the maximum likelihood estimator for a natural noise model [Young, 1995]. Under Kemeny's rule, for every ranking $r$ we compute the minimum number of *pairwise swaps* that are necessary to transform the preference profile into one where *all* preference relations of voters coincide with $r$, and we pick the top-ranked alternative from the ranking $r$ for which this number is the smallest.

However, pairwise voting rules—like all standard voting rules—have a drawback when applied to the pairwise liquid democracy setting: they ignore the structure of the pairwise delegation graph. Motivated by this observation, we initiate the study of so-called *liquid voting rules*, which take pairwise delegation graphs as input. Since pairwise delegation graphs can be viewed as generalizations of preference profiles (a preference profile corresponds to a pairwise delegation graph without external arcs), liquid voting rules generalize standard voting rules. In the following, we propose a generic way to "liquidize" voting rules. Our approach relies on the concept of distance rationalization.

### 6.1 Distance Rationalization

*Distance rationalization* (DR) is a very general framework in social choice theory [Elkind and Slinko, 2016]. Intuitively, a voting rule is distance rationalizable if there is a "distance function" and a "consensus class" such that winning alternatives coincide with consensus alternatives of the closest consensus profile, measuring closeness by the distance function.

Formally, a *distance function* is a metric on preference relations[8] and a *consensus class* $\mathcal{C}$ is a pair $(\mathcal{R}, \mathcal{W})$ where $\mathcal{R}$ is a set of preference profiles and $\mathcal{W}$ is a function mapping every $R \in \mathcal{R}$ to a winning alternative. The intuition is that the winner $\mathcal{W}(R)$ of a consensus profile $R \in \mathcal{R}$ is "obvious." A voting rule is $(d, \mathcal{C})$-*distance rationalizable* for distance $d$ and consensus class $\mathcal{C} = (\mathcal{R}, \mathcal{W})$ if, for every profile $R$, the rule selects the winner $\mathcal{W}(R')$ of the consensus profile $R' \in \mathcal{R}$ that is closest to $R$ according to $d$. If several closest consensus profiles exist, then all corresponding consensus winners will be tied winners. For instance, Kemeny's rule is $(d, \mathcal{C})$-distance rationalizable for $d$ being the *swap distance* (aka *bubble sort distance*) and $\mathcal{C} = (\mathcal{E}, \mathcal{W})$ being the *unanimous consensus class*, consisting of all complete and transitive preference profiles in which all preference relations are identical.

---

[6]Completeness can be relaxed [Brandt *et al.*, 2018]; therefore, this approach is feasible also for incomplete voter preferences resulting from pairwise delegation cycles.

[7]Pairwise voting rules are also known as *C2 functions* [Fishburn, 1977] or *weighted tournament solutions* [Fischer *et al.*, 2016].

[8]That is, we consider what [Elkind and Slinko, 2016] call *votewise* distances. A distance function $d$ on the set of preference relations naturally extends to a distance function on the set of preference profiles by letting $d(R, R') = \sum_{i \in N} d(R_i, R_i')$.

## 6.2 Liquid Distance Rationalization

We generalize the DR framework by incorporating operations on pairwise delegation graphs. The basic idea is as follows. While the standard DR framework considers only the space of preference profiles and defines distance functions and consensus classes within that space, the liquid DR framework also considers the space $\mathcal{G}$ of all pairwise delegation graphs. These two spaces are related via the mappings $R(\cdot)$ and $\hat{R}(\cdot)$ defined in Section 3. Specifically, the distance function is defined on $\mathcal{G}$, whereas the consensus classes remain in the space of preference profiles. Applying function $\hat{R}(\cdot)$ allows us to transfer the notion of consensus to the graph space.

**Definition 3.** Let $\mathcal{C} = (\mathcal{R}, \mathcal{W})$ a consensus class. A pairwise delegation graph $G$ is a *consensus graph* if $\hat{R}(G) \cap \mathcal{R} \neq \emptyset$.

That is, a consensus graph is a pairwise delegation graph $G$ whose corresponding preference profile $R(G)$ can be extended to a consensus preference profile. The function $\mathcal{W}$ can be extended to consensus graphs by letting $\mathcal{W}(G) = \mathcal{W}(R)$, where $R \in \hat{R}(G) \cap \mathcal{R}$. (If $|\hat{R}(G) \cap \mathcal{R}| > 1$, then all corresponding consensus winners are contained in $\mathcal{W}(G)$.)

We are now ready to define liquid distance rationalizability.

**Definition 4.** Let $d$ be a distance function on $\mathcal{G}$ and $\mathcal{C} = (\mathcal{R}, \mathcal{W})$ a consensus class. A liquid voting rule is $(d, \mathcal{C})$-*distance rationalizable* if, for any pairwise delegation graph $G$, the rule selects the winner(s) $\mathcal{W}(G')$ of the consensus graph $G'$ that is closest to $G$ according to $d$. (If several closest consensus graphs exist, then all corresponding consensus winners tie as winners.)

That is, a liquid distance rationalizable voting rule takes as input a pairwise delegation graph $G$ and finds a closest graph $G'$ such that $\hat{R}(G')$ contains a consensus profile; the rule then outputs the consensus winner of this profile.[9]

A natural distance function on $\mathcal{G}$ counts the number of swaps of *internal pairs* that are needed to transform one graph to another. The appeal of this distance function, denoted $d_{\text{int}}$, has to do with delegations: If some voters delegate a pair $ab$ to a voter $i$, and we swap voter $i$'s internal pair $ab$, then this swap will be propagated—by pairwise delegations—to all voters having $i$ as their ultimate $ab$-delegate. Thus, a single pairwise swap of an internal pair might lead to many swaps of that pair in other voters. It might therefore be possible to "reach" a consensus graph $G'$ by making only few internal swaps.

## 6.3 Liquid Voting Rules

An interesting example of a distance rationalizable liquid voting rule can be obtained by combining the distance function $d_{\text{int}}$ with the unanimous consensus class. We call the resulting liquid voting rule *LiquidKemeny*; it can be seen as an adaptation of Kemeny's rule to the pairwise liquid democracy setting. LiquidKemeny takes a pairwise delegation graph and, for every ranking $r$, computes the minimum number of *pairwise swaps of internal pairs* such that the resulting graph can

---

[9]Since the output of a liquid distance rationalizable voting rule may depend on the structure of the pairwise delegation graph, such rules generally violate the *one man–one vote property* as discussed by [Christoff and Grossi, 2017].
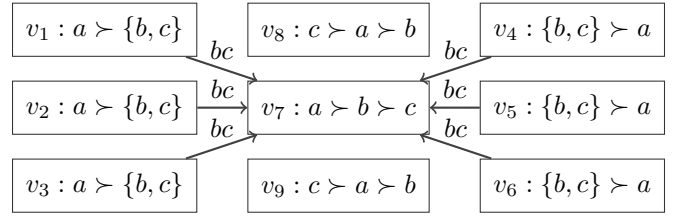


Figure 4: A pairwise delegation graph $G$ with 9 voters. LiquidKemeny selects alternative $c$, because $G$ can be turned into a consensus graph $G'$ via 8 pairwise swaps of internal pairs. The corresponding consensus profile satisfies $c \succ_i a \succ_i b$ for all $i \in N$. No other consensus graph can be reached with at most 8 pairwise internal swaps. (For comparison: Kemeny's rule, applied to the profile $R(G)$, selects winner $a$. The optimal ranking is $a \succ b \succ c$ and 10 pairwise swaps are necessary to transform $R(G)$ into a unanimous profile.)

be extended to a profile in which all preference relations are identical to $r$. The rule then selects the top-ranked alternative of the ranking $r$ for which this number is smallest. That is, in this liquid variant of Kemeny's rule, we do *not* count all disagreements between a ranking and preference relations of voters, but rather only those disagreements between a ranking and the *internal pairs* of voters. Figure 4 illustrates this rule.

There are other rules that can be distance rationalized with the swap distance [Elkind and Slinko, 2016]. For each such rule, we can immediately define their liquid version by generalizing the domain to the space of pairwise delegation graphs and replacing the swap distance with $d_{\text{int}}$. For instance, this approach yields liquid versions of Dodgson's rule and Borda's rule. Whenever the delegation graph is empty, the liquid version of a rule coincides with the original version.

**Remark 3.** The distance function $d_{\text{int}}$ is based on the implicit assumption that the "cost" of swapping an internal pair of a voter does not depend on the number of other voters to which this swap will be propagated through delegations. This, however, violates the intuition that swapping pairs for voters with many delegations should be *more costly* than doing so for a voter who attracts only few delegations. Indeed, the appeal of pairwise liquid democracy is closely connected to the promise of effective utilization of expert knowledge, thus it can be argued that the preferences of voters attracting many delegations should be given increased attention. This argument can be accounted for by augmenting $d_{\text{int}}$ with a *cost function* $c_G : N \times \mathcal{P} \to \mathbb{R}_+$ that, with respect to a pairwise delegation graph $G$, assigns a cost for swapping each internal pair of each voter. Intuitively, a high value of $c_G(i, ab)$ means that the opinion of voter $i$ on the pairwise comparison between $a$ and $b$ is valued highly. In fact, the rules we have considered can be recovered by defining the cost function appropriately: Let $\#d(i, ab)$ denote the number of voters that have voter $i$ as their ultimate $ab$-delegate, and observe that setting $c_G(i, ab) = 1$ corresponds to LiquidKemeny, while setting $c_G(i, ab) = \#d(i, ab)$ corresponds to the standard version of Kemeny's rule. A middle ground could be based on cost functions that take the length of delegation paths (or other structural properties of $G$) into account. An interesting way to do that has been proposed by [Boldi *et al.*, 2011].

# 7 Conclusion and Outlook

We proposed a generalized model of liquid democracy where voters can delegate pairwise comparisons, and explored approaches to deal with the possibility of intransitive preferences that might arise as a result of the increased flexibility that the model offers. We described a generalization of distance rationalization and suggested a framework for constructing "liquid voting rules" working directly on delegation graphs. Next we mention some directions for future research.

The most immediate direction is to study further liquid voting rules and their properties, e.g., by considering metrics on delegation graphs (including *graph edit distances* [Gao *et al.*, 2010]), or by "liquidizing" other classes of voting rules (e.g., *generalized scoring rules* [Xia, 2013]).

Further, one could consider other generalizations of liquid democracy. A particularly interesting approach is "statement voting" [Zhang and Zhou, 2017], where voters can specify *delegation rules* (e.g., "I delegate to $j$ unless $j$ delegates further"), resulting in an even greater flexibility.

## Acknowledgments

## References

[Behrens *et al.*, 2014] J. Behrens, A. Kistner, A. Nitsche, and B. Swierczek. *The Principles of LiquidFeedback*. 2014.

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

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

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

[Bouyssou, 2004] D. Bouyssou. Monotonicity of 'ranking by choosing': A progress report. *Social Choice and Welfare*, 23(2):249–273, 2004.

[Brandt *et al.*, 2016] F. Brandt, M. Brill, and P. Harrenstein. Tournament solutions. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 3. Cambridge University Press, 2016.

[Brandt *et al.*, 2018] F. Brandt, M. Brill, and P. Harrenstein. Extending tournament solutions. *Social Choice and Welfare*, 2018. Forthcoming.

[Brill *et al.*, 2016] M. Brill, E. Elkind, U. Endriss, and U. Grandi. Pairwise diffusion of preference rankings in social networks. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 130–136, 2016.

[Christoff and Grossi, 2017] Z. Christoff and D. Grossi. Binary voting with delegable proxy: An analysis of liquid democracy. In *Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 134–150, 2017.

[Cohensius *et al.*, 2017] G. Cohensius, S. Mannor, R. Meir, E. Meirom, and A. Orda. Proxy voting for better outcomes. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 858–866, 2017.

[DeGroot, 1974] M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

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

[Elkind and Slinko, 2016] E. Elkind and A. Slinko. Rationalizations of voting rules. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 8. Cambridge University Press, 2016.

[Fischer *et al.*, 2016] F. Fischer, O. Hudry, and R. Niedermeier. Weighted tournament solutions. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 4. Cambridge University Press, 2016.

[Fishburn, 1977] P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.

[Ford, 2002] B. Ford. Delegative democracy. Unpublished manuscript. Available at `http://www.brynosaurus.com/deleg/deleg.pdf`, 2002.

[Ford, 2014] B. Ford. Delegative democracy revisited. Blog post, 2014.

[Gao *et al.*, 2010] X. Gao, B. Xiao, D. Tao, and X. Li. A survey of graph edit distance. *Pattern Analysis and applications*, 13(1):113–129, 2010.

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

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

[Kemeny, 1959] J. G. Kemeny. Mathematics without numbers. *Daedalus*, 88:577–591, 1959.

[Laslier, 1997] J.-F. Laslier. *Tournament Solutions and Majority Voting*. Springer-Verlag, 1997.

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

[Xia, 2013] L. Xia. Generalized scoring rules: a framework that reconciles Borda and Condorcet. *ACM SIGecom Exchanges*, 12(1):42–48, 2013.

[Yang, 2017] Y. Yang. Approval voting with intransitive preferences. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1769–1771, 2017.

[Young and Levenglick, 1978] H. P. Young and A. Levenglick. A consistent extension of Condorcet's election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.

[Young, 1995] H. P. Young. Optimal voting rules. *Journal of Economic Perspectives*, 9(1):51–64, 1995.

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