Cycles in Liquid Democracy: A Game-Theoretic Justification

Markus Brill, Rachael Colley, Anne-Marie George, Grzegorz Lisowski, Georgios Papasotiropoulos, and Ulrike Schmidt-Kraepelin

Abstract

Liquid democracy is a voting scheme in which individuals either vote directly or delegate their vote to others. A common critique is that *delegation cycles* can occur, seemingly resulting in unused voting power. Yet, practitioners argue that delegation cycles are not only unproblematic but are even intentionally formed by participants. This divergent view stems from differing interpretations of delegations: in practice, delegations serve as *backup options* that can be overridden at any time by direct voting, whereas the literature often treats voting and delegating as mutually exclusive. Bringing theory closer to reality, we introduce a probabilistic model that captures strategic behavior under uncertainty. Within this model, we study the existence and structure of Nash equilibria, revealing that delegation cycles naturally emerge. We further examine the quality of equilibria via a Price-of-Anarchy approach. To complement these findings, we perform computational experiments using best-response dynamics.

1 Introduction

Liquid democracy is a flexible voting system that allows voters to vote directly or *delegate* their vote to another participant, who can vote on their behalf. Delegations are *transitive*, meaning that delegated votes can be delegated further, creating *delegation chains*. The voter at the end of a chain casts ballots on behalf of everyone in the chain. This system combines the advantages of *direct democracy* and *representative democracy* by giving voters the freedom to choose their own mode of participation [8].

Over the past decade, academic interest in liquid democracy has grown rapidly [26]. However, practitioners have noted that some parts of this literature overlook or misinterpret key aspects of liquid democracy as it is implemented in practice (see 3 and 4, Sec. 2.4.1). While liquid democracy is often applied over an extended period of time to an ongoing stream of decisions, much of the literature models it as a one-time event. Specifically, this divergence has led to differing views on the issue of *delegation cycles*. In the literature *delegating* and *voting* are frequently modeled as mutually exclusive options. Consequently, *delegation cycles* (i.e., a situation where a voter i delegates to j, and this delegation eventually returns to i through a chain of delegations) are viewed as problematic because none of the voters in the cycle eventually cast a vote, leaving their collective voting weight unused. In contrast, Behrens et al. [5] propose an alternative interpretation (informed by their experience with the *LiquidFeedback* platform): voters specify *default delegations* that remain fixed across multiple decisions. These default delegations serve as a fallback whenever voters do not cast a vote; if a voter casts a vote, their default delegation is ignored. Under this interpretation, like-minded voters who trust one another may intentionally form delegation cycles with their default delegations. These cycles rarely result in lost votes for actual decisions, since all the voting weight in the cycle is used as long as at least one of the involved voters casts a vote.

To bridge theory and practice, we introduce the *default delegation model*. Voters declare delegations for a future election, which will be used only if they do not participate themselves. To model uncertainty about who will participate in this election, each voter has a fixed probability of participating in the election. A voter's *utility* thus depends on both (i) other voters' delegation choices and (ii) who actually turns out to vote. We assume that participants live in a (one-dimensional) metric space and prefer to be represented by those who are close to them. Moreover, voters can have different *tolerance levels* towards being represented by far-away voters. While the model is not explicitly temporal, analyzing strategic behavior in this probabilistic setting can shed light on long-term delegation dynamics. In a world with a stream of similar elections, where each voter participates in only a fraction of them and participation is not highly correlated, equilibrium states in our model can serve as proxies for long-term behavior across the election stream.

We provide theoretical evidence supporting the practical observation that delegation cycles naturally arise among rational users of liquid democracy platforms. More concretely, within the default delegation model, we provide a game-theoretic analysis of how voters can strategically set their default delegations to ensure that their voting power ends up with casting voters whose preferences closely align with their own. We study the (i) *existence*, (ii) *structure*, and (iii) *quality* of (pure) Nash equilibria and demonstrate that, under mild assumptions, delegation cycles are necessarily being formed. Our theoretical findings are complemented by computational experiments, providing additional insights into best-response dynamics in our model.

1.1 Related Work

Recent research in (computational) social choice and beyond shows a growing interest in liquid democracy, with various models and methodologies emerging. We focus on studies that are related to ours in terms of these two aspects.

The Role of Cyclic Delegations. A key branch of the literature addresses delegation cycles, typically viewing them as undesirable and proposing solutions to eliminate them [9, 11, 14, 15, 19, 21, 22, 24, 27, 28]. Importantly, all of these works distinguish between delegating and casting voters and focus on axiomatic and algorithmic aspects, rather than strategic behavior. Notably, the work by Markakis and Papasotiropoulos [24], like ours, was directly motivated by the study of Behrens et al. [5]. They study a temporal model with delegation updates over discrete time-steps. By contrast, in our paper, the temporal aspect is captured through the probabilistic model for ballot casting. Other approaches prevent cycles by design: Abramowitz and Mattei [1] disallow transitive delegations; Kahng et al. [20] only allow delegations to voters with higher competence; Caragiannis and Micha [10] assume the existence of a mechanism preventing cyclic delegations.

One-Dimensional Spatial Models. To model voters' preferences over potential delegates, we assume that voters are positioned in a one-dimensional metric space. This is a common modeling choice (often representing ideological alignment) that has been employed in delegative voting settings by Yamakawa et al. [29], Green-Armytage [18], Cohensius et al. [12], Escoffier et al. [17], and Anshelevich et al. [2].

Strategic Delegation Behavior. A prominent line of research in liquid democracy employs a game-theoretic perspective. Indicatively, this includes works that analyze Nash equilibria of delegation games [6, 7, 16, 17], provide worst-case guarantees [25], or study voting power [30, 13]. However, the games these works consider are significantly different from the one we analyze. None of them model the strategic choice of delegates under probabilistic voter participation. For instance, prior utility models are based on the effort of voting [7] or are analyzed in purely deterministic settings [16, 17].

1.2 Our Contribution

A central contribution of our work is conceptual: we introduce the novel *default delegation model* for liquid democracy, which captures and explains strategic delegation behavior under uncertain participation — a feature inherent in real-world liquid democracy systems.

We analyze the following aspects of the model:

Existence of Nash Equilibria. Our extensive computational experiments suggest that Nash equilibria are prevalent across a broad range of synthetic instances. On the negative side, we identify instances where a Nash equilibrium does not exist, even in simple settings with only three voters, or where all voters have identical tolerance levels. On the positive side, we establish the existence of Nash equilibria in several special cases or slight variants of our original model.

Structure of Nash Equilibria. We prove that, under mild assumptions, strategic voters form delegation cycles in equilibrium. More precisely, every non-trivial component of an equilibrium delegation graph contains exactly one cycle. Furthermore, we show that relaxing any of the assumptions invalidates the result, and we provide additional insights into the structure of delegations at equilibrium. In more general settings, computational experiments reveal that the vast majority of components contain cycles. The width of these cycles appears to be proportional to voters' tolerance levels and inversely proportional to the number of voters.

Quality of Nash Equilibria. We evaluate the quality of Nash equilibria primarily in terms of their social welfare, i.e., the total utility they achieve, and we measure the *Price of Anarchy (PoA)*, i.e., the ratio between the best possible social welfare and the social welfare of equilibria. While we prove that the PoA is generally unbounded, we also provide strong positive results: for non-degenerate instances, the difference between the two quantities is bounded, and as voting probabilities increase or tolerance levels decrease, the welfare in equilibrium approaches the optimal social welfare. Moreover, notably, our experiments show that Nash equilibria often achieve close-to-optimal social welfare.

Omitted proofs and further details can be found in the full version of this paper.

2 The Default Delegation Model

We consider a finite set V of *voters* using a liquid democracy platform. Each voter nominates a *default delegate* for a future election, in which any voter may choose to vote or abstain. If a voter $i \in V$ does *not* vote, their voting power is passed to their default delegate, continuing transitively until a voter who casts a vote is reached — this voter is called i's *ultimate delegate*. If no one in the delegation chain votes, i has no ultimate delegate and their voting power is lost.

Default Delegations. For each voter i, we let $d(i) \in V$ denote their *default delegate*. Self-nominations ($d(i) = i$) are allowed and interpreted as abstentions from nominating a default delegate. Each *delegation profile* $\boldsymbol{d} = (d(i))_{i \in V}$ naturally corresponds to a (directed) *delegation graph* $G_{\boldsymbol{d}} = (V, \{(i, d(i)) \mid i \in V\})$ whose edges correspond to default delegations. Thus, each vertex of $G_{\boldsymbol{d}}$ has out-degree exactly 1 and self-nominations correspond to self-loops.

Ultimate Delegates. Assume that the set of voters casting a vote is known, and let $X \subseteq V$ denote this set. The default delegations of voters in X are irrelevant. Therefore, to resolve delegations for this election — that is, to determine which voters in $V \setminus X$ are ultimately represented by which casting voters in X under $\boldsymbol{d}$ — it suffices to consider the subgraph of $G_{\boldsymbol{d}}$ that only contains delegations from non-casting voters. For each non-casting voter $i \in V \setminus X$, we can identify their ultimate delegate by following the (unique) directed walk in this graph starting from i. If this walk leads to a casting voter $j \in X$, then j is the ultimate delegate of i. If the walk leads to a cycle or a self-loop,[1] then i has no ultimate delegate. Each casting voter has a voting weight in the examined election equal to the number of voters they are the ultimate delegate for, themselves included.

Probabilistic Participation. A crucial ingredient of our model is the assumption that voters do not know which other voters are casting a vote in the future election. Rather, when choosing a default

[1]We use the term "cycle" exclusively for closed paths involving *at least two vertices*, excluding self-loops from this definition.

delegate, they need to consider different possibilities of where their vote "ends up". To capture this uncertainty, we use a simple probabilistic model. Namely, we assume that each voter $i \in V$ casts a vote with a fixed probability $p_i \in [0, 1]$. We let $\boldsymbol{p} = (p_i)_{i \in V}$ denote the profile of vote-casting probabilities. A voter i with $p_i \in \{0, 1\}$ is called *deterministic*. This modeling choice reflects the idea that, in practice, many elections occur, and voters are not aware of or engaged in every single one of them. The probability p_i then can be thought of as the fraction of elections in which voter i typically participates. This probabilistic approach provides a simple way to model varying voting behavior without requiring complex assumptions about individual awareness or decision-making for each election.

For a given delegation profile $\boldsymbol{d}$ and a voter $i \in V$, we can now derive the probability distribution over i's ultimate delegates. Let $\pi(\boldsymbol{d}, i)$ denote the longest *simple path* in $G_{\boldsymbol{d}}$ starting at i. Formally, $\pi(\boldsymbol{d}, i)$ is the unique sequence $(y_0, y_1, \cdots, y_k)$ of distinct voters starting with $y_0 = i$ and satisfying the following:

(i) $d(y_{\ell-1}) = y_\ell$ and $y_\ell \notin \{y_0, y_1, y_2, \cdots, y_{\ell-1}\}$ for $\ell \leqslant k$,

(ii) $d(y_k) \in \{y_0, y_1, y_2, \cdots, y_k\}$.

Voter i's ultimate delegate is the first casting voter along the path $\pi(\boldsymbol{d}, i)$. Therefore, for $\ell \in \{0, 1, \cdots, k\}$, the probability that voter y_ℓ is the ultimate delegate of i is given by

$$p_{y_\ell} \cdot \prod_{r=0}^{\ell-1} (1 - p_{y_r}). \tag{1}$$

The ultimate delegate of i is undefined with probability $\prod_{r=0}^{k} (1 - p_{y_r})$, which we interpret as i's ballot being lost.

Distance and Tolerance. To evaluate and compare different delegation options, we assume that each voter's utility depends on the alignment between their preferences and those of their ultimate delegate. Alignment is defined in terms of the Euclidean distance between voters along a one-dimensional ideological space, represented as the interval $[0, 1]$. Each voter $i \in V$ is associated with a fixed position $x_i \in [0, 1]$, reflecting their ideological stance, which remains constant throughout all elections. Let $\boldsymbol{x} = (x_i)_{i \in V}$ denote *positions*. The utility of a voter decreases as the distance between their position and that of their ultimate delegate increases, capturing the notion that voters prefer representatives who are ideologically closer to themselves. Furthermore, each voter $i \in V$ is associated with a tolerance parameter $\tau_i \geqslant 0$. This parameter represents the maximum distance the voter is willing to accept between their own position and that of their ultimate delegate, while still deriving positive utility from delegating. Let $\text{dist}(i, j) = |x_i - x_j|$ denote the distance between voters i and j, and let $\boldsymbol{\tau} = (\tau_i)_{i \in V}$. The *acceptability set* of voter i is given by: $\mathcal{A}_i(\boldsymbol{x}, \boldsymbol{\tau}) = \{j \in V \mid \text{dist}(i, j) \leqslant \tau_i\}$.

Instances. Given the voters' positions $\boldsymbol{x}$, their voting probabilities $\boldsymbol{p}$, and their tolerance parameters $\boldsymbol{\tau}$, we define an instance as the triple $\mathcal{I} = \langle \boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\tau} \rangle$. To avoid ties, we will always assume that our instances are in *general position*, meaning that no two voters share the same position and that no voter $j \neq i$ is at distance exactly τ_i from voter i.[2] Further, we will sometimes assume that j belongs to $\mathcal{A}_i(\boldsymbol{x}, \boldsymbol{\tau})$ if and only if i belongs to $\mathcal{A}_j(\boldsymbol{x}, \boldsymbol{\tau})$. If this holds for all pairs of voters, then we say that the instance satisfies *mutual acceptance*. A special case of mutual acceptance instances are those where all voters have identical tolerance parameters ($\tau_i = \tau_j$ for all $i, j \in V$); we will refer to such instances as *symmetric*. We often assume that voters are ordered by increasing x_i, and then specify $\boldsymbol{p}$ and $\boldsymbol{\tau}$ as vectors corresponding to that order.

Voter Utility. The utility of a voter from a realization of the random election is defined as their tolerance minus the distance to their ultimate delegate, or 0 if the ultimate delegate is undefined. In

[2] If an instance is not in general position, then slightly perturbing entries of $\boldsymbol{x}$ will bring the instance into general position.

other words, voters rank potential ultimate delegates by proximity and prefer abstaining over delegating to a voter outside their acceptability set. Formally, given positions $\boldsymbol{x}$, tolerances $\boldsymbol{\tau}$, and the set X of casting voters, the utility of voter i under a delegation profile $\boldsymbol{d}$ is $\tau_i - \text{dist}(i, j)$, where $j \in X$ is the ultimate delegate of i, or 0 if no ultimate delegate is defined. Due to probabilistic participation, the ultimate delegate of the voter i is a random variable (distributed according to Expression (1)). To account for this randomness, we define the *expected utility* of voter i as the weighted sum of their utilities over all possible ultimate delegates. Specifically, the expected utility of voter i (henceforth, simply *utility*) in an instance $\mathcal{I}$ can be expressed as

$$u_i(\boldsymbol{d}, \mathcal{I}) = \sum_{\ell=0}^{k} (\tau_i - \text{dist}(i, y_\ell)) \cdot p_{y_\ell} \cdot \prod_{r=0}^{\ell-1} (1 - p_{y_r}), \tag{2}$$

where $(y_0, y_1, \cdots, y_k)$ are the voters along the path $\pi(\boldsymbol{d}, i)$. When the instance $\mathcal{I}$ is clear from the context, we will refer to $u_i(\boldsymbol{d}, \mathcal{I})$ simply as $u_i(\boldsymbol{d})$.

The *social welfare* of a delegation profile $\boldsymbol{d}$ in an instance $\mathcal{I}$ is the sum over voter utilities, $SW(\boldsymbol{d}, \mathcal{I}) = \sum_{i \in V} u_i(\boldsymbol{d}, \mathcal{I})$. The profile maximizing the social welfare among all possible delegation profiles for $\mathcal{I}$ will be called *optimal* and denoted by $\boldsymbol{d}_{SW}(\mathcal{I})$, or simply $\boldsymbol{d}_{SW}$.

Example 1. *Consider an instance with six voters, $V = \{A, B, C, D, E, F\}$. The positions $\boldsymbol{x}$ and voting probabilities $\mathbf{p}$ are depicted in the figure below, together with delegation graph $G_{\boldsymbol{d}}$ of the delegation profile $\boldsymbol{d}$ with $d(A) = B$, $d(B) = C$, $d(C) = A$, $d(D) = E$, $d(E) = D$, $d(F) = F$.*

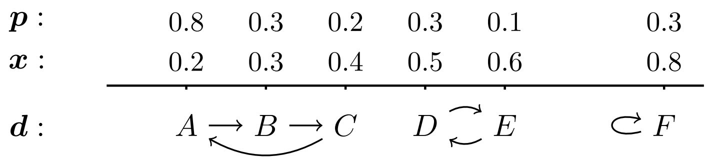

Suppose that for a given election, the set of casting voters is $\{A, F\}$. This situation happens with probability $p_A \cdot p_F \cdot \Pi_{i \in \{B,C,D,E\}}(1 - p_i) = 0.8 \cdot 0.3 \cdot 0.7 \cdot 0.8 \cdot 0.7 \cdot 0.9 \approx 0.085$. In this scenario, A has a voting weight of 3 and F has a voting weight of 1. The voting weights of D and E cannot be allocated to a casting voter and are lost.

Assume that $\tau_i = 0.25$ for all $i \in V$. Then the acceptability set of voter D is $\mathcal{A}_D(\boldsymbol{x}, \boldsymbol{\tau}) = \{B, C, D, E\}$. The expected utility of voters A and F w.r.t. the delegation profile $\boldsymbol{d}$ is $u_A(\boldsymbol{d}) = 0.8 \cdot 0.25 + 0.2 \cdot 0.3 \cdot (0.25 - 0.1) + 0.2 \cdot 0.7 \cdot 0.2 \cdot (0.25 - 0.2) \approx 0.21$ and $u_F(\boldsymbol{d}) = 0.3 \cdot 0.25 = 0.075$.

3 Existence of Nash Equilibria

We start our game-theoretic analysis of the default delegation model. Using the utility model described by Expression (2), we define the concept of best responses and profitable deviations in a standard way. We denote by $\boldsymbol{d}_{-i}$ the profile $\boldsymbol{d}$ not including the choice of i, and by $(\boldsymbol{d}_{-i}, d'(i))$ the delegation profile in which all voters except i delegate according to $\boldsymbol{d}$, whereas i delegates to $d'(i)$.

Definition 1. For a voter $i \in V$, $d'(i)$ is a *best response* to delegation profile $\boldsymbol{d}$ if and only if it maximizes $u_i(\boldsymbol{d}_{-i}, \cdot)$. We say that $d'(i)$ is a *profitable deviation* from $\boldsymbol{d}$ for voter i if $u_i(\boldsymbol{d}_{-i}, d'(i)) > u_i(\boldsymbol{d})$.

Building upon the concept of profitable deviations, we are ready to define (pure) Nash equilbria.

Definition 2. A delegation profile $\boldsymbol{d}$ is a *Nash equilibrium (NE)* if no voter $i \in V$ has a profitable deviation from $\boldsymbol{d}$.

We illustrate Nash equilibria in the default delegation model with the help of our initial example, which highlights that equilibria are not necessarily unique and that different equilibria may have different graph-theoretic structures.

Example 1 Continued. *The delegation profile* $\boldsymbol{d}$ *from Example 1 is a Nash equilibrium, since each voter chooses a best response, as shown in Table 1. Note that equilibrium delegations need not go to the closest voter.*

	A	B	C	D	E	F
A	0.200	**0.210**	0.202	0.195	0.194	0.179
B	0.159	0.075	**0.163**	0.083	0.081	0.023
C	**0.089**	0.086	0.050	0.089	0.086	0.014
D	0.052	0.086	0.075	0.075	**0.085**	0.064
E	-0.084	-0.043	-0.055	**0.066**	0.025	0.039
F	-0.134	-0.102	-0.111	0.067	0.069	**0.075**

Table 1: Expected utility for deviations from $\boldsymbol{d}$ in Example 1. The entry in cell (i, j) corresponds to $u_i(\boldsymbol{d}_{-i}, j)$ and the entries corresponding to best responses are indicated in bold.

Interestingly, $\boldsymbol{d}$ *is not the only NE of this instance. It can be verified that* $\boldsymbol{d}'$, *the delegation graph* $G_{\boldsymbol{d}'}$ *of which follows, is also a NE.*

$$\begin{tikzcd} \boldsymbol{d}': & A \arrow[rr, bend right] & B \arrow[l] & C \arrow[l] & D \arrow[l] & E \arrow[r, bend left] & F \arrow[l, bend left] \end{tikzcd}$$

We highlight that two equilibria might significantly differ. Indicatively, $G_{\boldsymbol{d}'}$ *has two (weakly) connected components, in contrast to* $G_{\boldsymbol{d}}$. *Moreover,* $\boldsymbol{d}$ *and* $\boldsymbol{d}'$ *differ in terms of social welfare, casting voters' weights and expected number of votes that are lost.*

Experimental Analysis. To get a first impression on whether Nash equilibria exist in general, we carried out computational experiments using a *best-response dynamic*. That is, the process starts with some delegation profile (e.g., a random one) and then iterates over the voters, updating their delegation whenever there is a profitable deviation. The process stops when no voter can make a profitable deviation, which results in a Nash equilibrium by definition. Interestingly, running our best-response dynamic on 20,000 different (mutual acceptance) instances for various values of n, $\boldsymbol{x}$, $\boldsymbol{p}$, and $\boldsymbol{\tau}$, and starting profiles, has always led to the identification of a Nash equilibrium.

Non-Existence of Equilibria. In contrast to what we observed in our computational experiments sketched above, Nash equilibria do not always exist in the default delegation model. To showcase this, we provide the following example and later strengthen the result in Theorem 2.

Example 2. *The instance* $\mathcal{I}$ *with* $V = \{A, B, C, D\}$, $\boldsymbol{x} = (0, 0.05, 0.1, 0.5)$, $\boldsymbol{p} = (0.4, 0.05, 0.2, 0.4)$, $\boldsymbol{\tau} = (1, 0, 0.2, 0)$ *does not admit a Nash equilibrium. For the contrary, assume* $\boldsymbol{d}$ *is a Nash equilibrium in* $\mathcal{I}$. *Then in* d *the voters* B *and* D *must delegate to themselves as* $\tau_B = \tau_D = 0$, *i.e., delegating to any other voter lowers their utility compared to self-delegation. Further, since* $\tau_C = 0.2$, *dist*$(C, D) = 0.4$ *and* $d(D) = D$, *voter* C *does not delegate to* D. *Both* A *and* C *do not delegate to themselves. This is because* $\tau_A -$ *dist*$(A, B) > 0$, $\tau_C -$ *dist*$(C, B) > 0$, *and* $d(B) = B$, *implying that choosing to delegate to* B *provides better utility than self-delegation. In the table appearing in Figure 1, we show the utilities of* A *and* C *in all profiles that were not ruled out by the previous reasoning. It is routine to check that none of the possible delegation profiles* $\boldsymbol{d}$ *is a Nash equilibrium.*

Given this impossibility result, we focus on subclasses of the default delegation model or slight variations towards obtaining positive results on the existence of Nash equilibria.

	$d(A) = B$	$d(A) = C$	$d(A) = D$
$d(C) = A$	.428 / .075	.508 / .072	.52 / .033
$d(C) = B$	.428 / .044	.53 / .044	.52 / .044

Figure 1: The normal form representation of the game induced by the instance $\mathcal{I}$, described in Example 2. The rows correspond to the possible choices of voter C and the columns to the choices of voter A, for specifying $\boldsymbol{d}$.

3.1 Special Cases

We discuss three special cases of the default delegation model and draw a complete picture of whether these restrictions suffice to guarantee the existence of equilibria. More precisely, we study (i) *deterministic* instances, i.e., those instances where $p_i \in \{0, 1\}$ for all $i \in V$, (ii) *mutual acceptance* instances, i.e., those instances where $j \in \mathcal{A}_i(\boldsymbol{x}, \boldsymbol{\tau})$ if and only if $i \in \mathcal{A}_j(\boldsymbol{x}, \boldsymbol{\tau})$ for all $i, j \in V$, and (iii) instances with *few voters*, i.e., those instances where $|V|$ is upper bounded by a constant. On the positive side, such restrictions can guarantee the existence of a Nash equilibrium.

Theorem 1. *For each of the following restrictions, any instance $\mathcal{I}$ is guaranteed to contain a Nash equilibrium:*

(i) $\mathcal{I}$ is deterministic,

(ii) $\mathcal{I}$ has two voters, i.e., $|V| = 2$, or

(iii) $\mathcal{I}$ satisfies mutual acceptance and $|V| \leqslant 3$.

In the case of deterministic instances, the profile where every non-casting voter (i.e., with $p_i = 0$) delegates to their closest casting voter (i.e., with $p_j = 1$) in their acceptability set is a Nash equilibrium. For two-voter instances, we show that the expected utility of a voter is not influenced by the delegation choice of the other. For mutual acceptance instances with three voters, we propose a greedy algorithm for finding an equilibrium. While these restrictions are strong, we complement Theorem 1 by showing that relaxing them even slightly invalidates the result.

Theorem 2. *There exists an instance $\mathcal{I}$ for which no Nash equilibrium exists if:*

(i) $\mathcal{I}$ has three voters, i.e., $|V| = 3$, or

(ii) $\mathcal{I}$ satisfies mutual acceptance and $|V| = 4$.

We remark that statement (ii) of Theorem 2 holds even for *symmetric* instances, i.e., with voters of equal tolerance.

3.2 Variants of the Model

In response to the negative results of Example 2 and Theorem 2, we discuss two variants of our model that guarantee the existence of Nash equilibria.

Leftists and Rightists. In the default delegation model, a voter accepts representation by voters positioned both to their left and to their right, only dependent on their distance. We introduce a variant of the model, where each voter selects a direction and accepts only representation by voters in that direction, in which case, the utility is still determined by the distance, in line with Expression (2). Depending on the direction selected, we refer to a voter as *leftist* or *rightist*.

In Example 2, voter A can (trivially) be considered a rightist and voter C can be considered a leftist as $\mathcal{A}_C(\mathcal{I})$ only includes voters to their left. Thus, on a profile with both leftists and rightists, a Nash equilibrium is not guaranteed to exist. In contrast, Theorem 3 shows that any instance with only leftists or only rightists contains a Nash equilibrium.

For the sake of concreteness, we define an example of a utility function that induces leftist voters. Namely, replace $(\tau_i - \text{dist}(i, y_\ell))$ in Expression (2) (intuitively, the utility of voter i for being represented by y_ℓ in a specific election) by:

$$\begin{cases} \tau_i - \text{dist}(i, y_\ell), & \text{if } y_\ell \text{ is left of } i, \text{ i.e., } x_{y_\ell} < x_i, \\ -\text{dist}(i, y_\ell), & \text{if } y_\ell \text{ is right of } i, \text{ i.e., } x_{y_\ell} > x_i. \end{cases}$$

We remark that Theorem 3 holds for any utility model that assigns negative utility to representation on the one side and utility equal to $\tau_i - \text{dist}(i, y_\ell)$ to the other.

Theorem 3. *Every instance in which the voters are all leftists or all rightists admits a Nash equilibrium.*

The proof of Theorem 3 constructs a Nash equilibrium by starting from a profile where everyone delegates to themselves, and then finding best responses for all voters sequentially in order of their position.

Proxy Voting. We now move to another variant of the model in which Nash equilibria are guaranteed to exist. In the *proxy voting* setting, we restrict the number of voters on any path $\pi(\boldsymbol{d}, i)$ that leads to a casting voter. Specifically, no such path is allowed to contain more than two voters (including voter i themselves). Hence, we effectively restrict the strategy space of the voters based on the actions of the other voters. This restriction is reminiscent[3] of the well-established framework of proxy voting [12, 2], a variant of liquid democracy in which voters are divided into delegating and casting voters and delegation chains may contain at most 2 voters.

In Example 2, delegation chains of three voters arose. By forbidding such chains, we effectively eliminate the issue that leads to the non-existence of equilibria. In the proxy voting setting, we guarantee the existence of Nash equilibria, leading to a dichotomy in the maximum allowable delegation chain length to ensure the existence of Nash equilibria.

Theorem 4. *In the proxy voting setting, every instance admits a Nash equilibrium.*

4 Structure of Nash Equilibria

We now focus on the structural properties of equilibria. In particular, we are interested in the existence of cycles in delegation graphs corresponding to Nash equilibria. Our first result establishes that delegation cycles are the rule, rather than the exception. This aims to provide a game-theoretical justification for the behavior of voters observed in practice.

Theorem 5. *Consider a mutual acceptance instance $\mathcal{I}$ without deterministic voters. Then, for every Nash Equilibrium $\boldsymbol{d}$ of $\mathcal{I}$, it holds that every weakly connected component of $G_{\boldsymbol{d}}$ with more than a single vertex has exactly one cycle.*

[3]The settings are not identical as we allow for length 2 cycles.

Proof Sketch. For contradiction, assume there is an acyclic weakly connected component W of $G_{\boldsymbol{d}}$ with at least two voters. In this case, W would form a tree with a "sink" voter i such that $d(i) = i$. By analyzing the incentives of voter i, we derive that $j \notin \mathcal{A}_i(\boldsymbol{x}, \boldsymbol{\tau})$ for any j such that $d(j) = i$. However, by the mutual acceptance assumption, it must also hold that $i \notin \mathcal{A}_j(\boldsymbol{x}, \boldsymbol{\tau})$, which contradicts $\boldsymbol{d}$ being an equilibrium. Uniqueness follows from the fact that each vertex in the component has out-degree 1. □

When the assumptions of Theorem 5 do not hold, cycles do not necessarily exist in every equilibrium.

Observation 6. *Cycles are not guaranteed to exist in $G_{\boldsymbol{d}}$, where $\boldsymbol{d}$ is a Nash equilibrium of an instance that is not of mutual acceptance or where deterministic voters exist.*

Yet, in mutual acceptance instances, at least one equilibrium with a cyclic structure is guaranteed to exist, even with some deterministic voters. The proof is similar to that of Theorem 5 and includes the observation that deterministic voters may be indifferent towards delegation options.

Theorem 7. *Consider a mutual acceptance instance $\mathcal{I}$ admitting a NE. Then, there exists a NE $\boldsymbol{d}$ of $\mathcal{I}$ in which every weakly connected component of $G_{\boldsymbol{d}}$ with more than a single vertex has exactly one cycle.*

Without mutual acceptance, the existence of equilibria exhibiting cycles is not guaranteed. For instance, consider an instance with two non-deterministic voters such that A accepts B, but B does not accept A. Then, there is a unique equilibrium in which A delegates to B and B self-loops.

Returning to the case where the assumptions of Theorem 5 hold, we now aim to further analyze the structure of equilibria by turning our attention to delegations "entering" a cycle. Specifically, for a weakly connected component W, let $\mathcal{C}(W)$ denote the set of voters forming the cycle within that component, and let $\mathcal{L}(W)$ and $\mathcal{R}(W)$ denote the sets of voters of W positioned to the left and right of the cycle, respectively. Formally, $\mathcal{L}(W) = \{i \in W : x_i < x_j \text{ for all } j \in \mathcal{C}(W)\}$ and $\mathcal{R}(W) = \{i \in W : x_i > x_j \text{ for all } j \in \mathcal{C}(W)\}$.

Theorem 8. *Consider a mutual acceptance instance $\mathcal{I}$ without deterministic voters and a Nash equilibrium $\boldsymbol{d}$ of $\mathcal{I}$. Consider a weakly connected component W of $G_{\boldsymbol{d}}$ that consists of more than a single vertex, and let $\mathcal{C}(W)$ denote the cycle in W. There is at most one vertex $v_L \in \mathcal{L}(W)$ with $d(v_L) \in \mathcal{C}(W)$ and at most one vertex $v_R \in \mathcal{R}(W)$ with $d(v_R) \in \mathcal{C}(W)$. Moreover, in $G_{\boldsymbol{d}}$, $\mathcal{L}(W)$ and $\mathcal{R}(W)$ form in-trees rooted at v_L and v_R, respectively.*

Thus, the cycle $\mathcal{C}(W)$ has a unique "entry point" v_L for voters in $\mathcal{L}(W)$, and all voters in $\mathcal{L}(W)$ have delegation paths to v_L (analogously for v_R and $\mathcal{R}(W)$). It might be tempting to conjecture that these entry points v_L and v_R delegate to the leftmost and rightmost voters in $\mathcal{C}(W)$, respectively, or that all voters in $\mathcal{L}(W)$ (respectively, $\mathcal{R}(W)$) form a simple delegation path. However, in the extended version of this paper we show that this is not generally the case. Therefore, a significant strengthening of the structural description offered by Theorem 8 is unlikely.

Experimental Analysis. Our theoretical results do not specify how large delegation cycles are, or how often they occur in instances not satisfying the assumptions of Theorem 5. To shed light on these questions, we conducted computational simulations. In particular, we examined the *size* (i.e., number of vertices) and *width* (i.e., maximum distance between two vertices) of cycles and weakly connected components and we observe that, as tolerance levels decrease, cycle size and width, as well as component width, decline gradually. Moreover, as n increases, the average cycle and component width decreases, with voters in the same component — especially cycles — having closely aligned positions. The proportion of voters with self-loops remains stable at around 5%. The average cycle size stays around 4.9 across instances and grows as n increases. Notably, nearly all weakly connected components with more than one vertex contain a cycle, indicating that the pattern identified theoretically for mutual acceptance instances (see Theorem 5) also appears in general, randomly generated instances.

5 Quality of Nash Equilibria

We next focus on evaluating the quality of equilibria. We use a Price-of-Anarchy approach to compare the social welfare of Nash equilibria to the optimal social welfare [23]. Before that, we compare the structure of social-welfare-maximizing delegation graphs to Nash equilibria, observing an interesting contrast.

Observation 9. *There exist mutual acceptance instances without deterministic voters in which social welfare maximising delegations do not induce cycles; e.g., in the symmetric instance with $V = \{A, B, C\}$, $\boldsymbol{x} = (0.12, 0.5, 0.88)$, $\boldsymbol{p} = (0.1, 0.9, 0.1)$, and $\tau_i = 0.4$ for all $i \in V$, social welfare is maximized if A and C delegate to B, who self-loops.*

Recall that $\boldsymbol{d}_{\mathtt{SW}}(\mathcal{I})$ is a profile maximizing social welfare and let $\boldsymbol{d}_{\mathtt{NE}}(\mathcal{I})$ be a Nash equilibrium of $\mathcal{I}$ achieving the *lowest* social welfare among all NE. We define the *Price of Anarchy (PoA)* of an instance $\mathcal{I}$ in the standard way,

$$PoA(\mathcal{I}) = \frac{SW(\boldsymbol{d}_{\mathtt{SW}}(\mathcal{I}))}{SW(\boldsymbol{d}_{\mathtt{NE}}(\mathcal{I}))},$$

and show that this ratio can be arbitrarily large.

Theorem 10. *The Price of Anarchy of default delegation instances is unbounded.*

Proof sketch. We prove the statement by describing a family of instances parametrized by ε and n. In the limit for $\varepsilon \to 0$ and $n \to \infty$, the worst Nash equilibrium $\boldsymbol{d}$ satisfies $SW(\boldsymbol{d}) \to 0$. Further, we show there is a parameterized profile $\boldsymbol{d}'$ such that the social-welfare-maximizing delegation profile $\boldsymbol{d}_{\mathtt{SW}}$ satisfies $SW(\boldsymbol{d}_{\mathtt{SW}}) \geqslant SW(\boldsymbol{d}') \to \frac{e-1}{e}\lambda$, where $\lambda > 0$ corresponds to some fixed value and e is Euler's number.

Fix a value $\lambda > 0$ as well as values $\varepsilon \in (0, 1)$ and $n \in \mathbb{N}$ such that $\lambda > \varepsilon/n$. Define $\mathcal{I}_{\varepsilon,n}$ as the instance with voters $\{1, 2, \ldots, n+1\}$ and $\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\tau}$ as specified in Figure 2. Note that voters are equidistant: $\text{dist}(i, i+1) = \varepsilon/n$, for all $i \in [n]$.

The profile $\boldsymbol{d}$ with $d(1) = 2$ and $d(i) = i$ for all $i \geqslant 2$ is a Nash equilibrium. Since $u_i(\boldsymbol{d}) = 0$ for all $i \geqslant 2$, we get

$$SW(\boldsymbol{d}) = u_1(\boldsymbol{d}) = \lambda\varepsilon + (1-\varepsilon)(\lambda - \frac{\varepsilon}{n})\frac{1}{n} \xrightarrow[\varepsilon \to 0]{n \to \infty} 0.$$

For profile $\boldsymbol{d}'$ with $d'(n+1) = n+1$ and $d'(i) = i+1$ for $i \leqslant n$, we get $u_i(\boldsymbol{d}') \xrightarrow{n\to\infty} 0$ for all $i \geqslant 2$ and

$\tau_1 = \lambda$	$\tau_2 = 0$	$\tau_3 = 0$	$\cdots$	$\tau_n = 0$	$\tau_{n+1} = 0$
$p_1 = \varepsilon$	$p_2 = 1/n$	$p_3 = 1/n$	$\cdots$	$p_n = 1/n$	$p_{n+1} = 1$
$x_1 = 0$	$x_2 = \varepsilon/n$	$x_3 = 2\varepsilon/n$	$\cdots$	$x_n = (n-1)\varepsilon/n$	$x_{n+1} = \varepsilon$

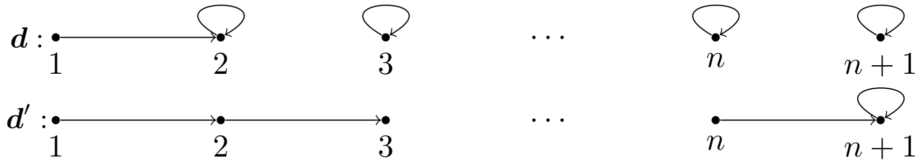

Figure 2: Illustration of the instance $\mathcal{I}_{\varepsilon,n}$ from the proof of Theorem 10.

$$u_1(\boldsymbol{d}') = \lambda\varepsilon + (1-\varepsilon)(\lambda - \frac{\varepsilon}{n})\frac{1}{n} + (1-\varepsilon)(\lambda - \frac{2\varepsilon}{n})(1-\frac{1}{n})\frac{1}{n} + \cdots$$
$$+ (1-\varepsilon)(\lambda - \frac{(n)\varepsilon}{n})(1-\frac{1}{n})^{n-1}\frac{1}{n} \xrightarrow[\varepsilon\to 0]{} \lambda\frac{1}{n}\sum_{i=0}^{n-1}(1-\frac{1}{n})^i.$$

Computing the limit of this expression for $n \to \infty$, we get

$$SW(\boldsymbol{d}') = \sum_{i=1}^{n+1} u_i(\boldsymbol{d}') \xrightarrow[\varepsilon\to 0]{n\to\infty} \frac{e-1}{e}\lambda. \qquad \square$$

At first glance, Theorem 10 is a negative result concerning the quality of NE. However, the constructed instances have certain characteristics, such as a voter with a very low voting probability and all but one voter having acceptability sets limited to themselves, while voter 1 has $\mathcal{A}_i(\boldsymbol{x}, \boldsymbol{\tau}) = V$. Moreover, the social welfare of the two delegation profiles in the proof of Theorem 10 exhibit a relatively small absolute difference ($\frac{e}{e-1}\lambda \approx 0.632\lambda = 0.632\sum_{i\in V}\tau_i$). This suggests that measuring the difference rather than the ratio may yield better conclusions about the quality of equilibria.

We define the *additive Price of Anarchy* of an instance $\mathcal{I}$ as $PoA^+(\mathcal{I}) = SW(\boldsymbol{d}_{\mathtt{SW}}(\mathcal{I})) - SW(\boldsymbol{d}_{\mathtt{NE}}(\mathcal{I}))$, and proceed with the following positive results on both the multiplicative and the additive Price of Anarchy.[4]

Theorem 11. *For every instance $\mathcal{I}$, $PoA(\mathcal{I}) \leqslant 1/p_{\min}$ and $PoA^+(\mathcal{I}) \leqslant (1-p_{\min})\sum_{i\in V}\tau_i$, where $p_{\min} = \min_{i\in V}\{p_i\}$.*

Theorem 11 asserts that higher voting probabilities correlate with better Nash equilibria in terms of social welfare. Furthermore, and perhaps surprisingly, the smaller the tolerance levels, the better the additive PoA bound.

In the extended version of this paper, we also assess the expected number of votes cast in equilibria, demonstrating that, unlike in other liquid democracy frameworks where cycles are criticized for resulting in ballot loss, in our setting, they effectively help mitigate lost voting power. Moreover, results on the structure of optimal delegation profiles, which complement Observation 9, can be found in the extended version of the paper, highlighting further differences in their structure compared to Nash equilibria and the profiles minimizing vote loss.

Experimental Analysis. To complement our worst-case bounds, we examined how the social welfare achieved by Nash equilibria compares to the optimal social welfare in randomly generated instances. Since identifying $\boldsymbol{d}_{SW}$ is computationally infeasible when n is large, we approximate it by the sum of each voter's expected utility under their optimal delegation profile, denoted by $\mathtt{ODP}(\mathcal{I})$. Formally, $\mathtt{ODP}(\mathcal{I}) = \sum_{i\in V} u_i(\mathbf{d}^{\mathbf{i}*})$, where $\mathbf{d}^{\mathbf{i}*}$ is a delegation profile maximizing the utility of voter i. In the extended version of the paper, we show that $G_{\mathbf{d}^{\mathbf{i}*}}$ contains a path that starts in i and passes through all vertices in $\mathcal{A}_i(\boldsymbol{x}, \boldsymbol{\tau})$ in increasing order of distance to i. This value serves as an upper bound, $SW(\boldsymbol{d}_{SW}) \leqslant \mathtt{ODP}(\mathcal{I})$.

We generated 100 instances for each $n \in \{20, 50, 100, 200\}$, with values for $\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\tau}$ chosen uniformly at random. For $n = 50$, we furthermore tested 5 tolerance vectors $\boldsymbol{\tau}$, scaling each by 0.75 and 0.5 to assess the effect of different tolerance levels. For each instance, we computed $\mathtt{ODP}(\mathcal{I})$ as an upper bound on the social welfare and a Nash equilibrium $\boldsymbol{d}_{BR}$ via best-response dynamics. Table 2 shows the average ratios $SW(\boldsymbol{d}_{BR})/ODP(\mathcal{I})$. As a baseline, we also include the social welfare achieved by the delegation profile $\boldsymbol{d}_{\mathtt{dir}}$ ("direct voting") in which every voter self-loops.

[4]While PoA^+ is not explicitly normalized by utility, the bound in Theorem 11 implicitly is, since $\sum_{i\in[n]}\tau_i$ is a trivial upper bound for social welfare.

	Number of Voters				$\tau_i \in [0, \tau_{\max}]$ with $\tau_{\max} =$		
	20	50	100	200	1	0.75	0.5
$\boldsymbol{d}_{BR}$	97.6%	98.8%	99.4%	99.7%	98.9%	98.9%	98.6%
$\boldsymbol{d}_{\texttt{dir}}$	53.6%	49.6%	50.4%	50.5%	51.2%	51.6%	52.3%

Table 2: The average social welfare achieved by $\boldsymbol{d}_{BR}$ and $\boldsymbol{d}_{\texttt{dir}}$ in our experiments, as a percentage of $\texttt{ODP}(\mathcal{I})$.

As expected, the Nash equilibrium profiles outperform the direct voting profiles, which consistently reach only around 50% of $\texttt{ODP}(\mathcal{I})$. The average social welfare achieved by $\boldsymbol{d}_{BR}$ is remarkably high ($\geqslant$98% of $\texttt{ODP}(\mathcal{I})$) and gets closer to $\texttt{ODP}(\mathcal{I})$ as n increases. Given that $\texttt{ODP}(\mathcal{I})$ upper bounds the optimal social welfare, we conclude that the Nash equilibria in our model have an almost optimal social welfare.

6 Conclusion

In this paper, we introduced the default delegation model and used it to provide a novel game-theoretic perspective on strategic delegation decisions in liquid democracy. We revealed how delegation cycles naturally emerge among rational participants, offering a justification for their existence.

Our model leads to several avenues for future research. One immediate direction is to explore the computational complexity of finding Nash equilibria or delegation profiles maximizing social welfare. In our experiments, we use best-response dynamics to find equilibria; however, these algorithms are not guaranteed to converge. It would also be interesting to define voters' alignment based on more general metric spaces; e.g., a two-dimensional Euclidean space. Preliminary experiments reveal that the structure of Nash equilibria becomes more complicated in that setting. Considering alternative utility functions — such as normalized ones, non-linear functions of distance, or those incorporating voting costs — could yield further insights. Ballot casting probability could also be part of the strategy space of a voter. Finally, while we primarily focused on cycles, *long delegation paths* remain an important and underexplored aspect of liquid democracy, as they may be associated with eroding trust in ultimate delegates. Potentially, this could be explored with a model similar to ours.

Acknowledgments

We thank Andreas Nitsche for helpful discussions. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 101002854). Georgios Papasotiropoulos was supported by the European Union (ERC, PRO-DEMOCRATIC, 101076570). Views and opinions expressed are however those of the author only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. Rachael Colley has been supported by the ANR JCJC project SCONE (ANR 18-CE23-0009-01) and the grant EP/X013618/1 from the Engineering and Physical Sciences Research Council.

European Research Council
Established by the European Commission

References

[1] Ben Abramowitz and Nicholas Mattei. Flexible representative democracy. *Social Choice and Welfare*, 64(1):263–308, 2025.

[2] Elliot Anshelevich, Zack Fitzsimmons, Rohit Vaish, and Lirong Xia. Representative proxy voting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5086–5093, 2021.

[3] Jan Behrens. Circular delegations–myth or disaster? *The Liquid Democracy Journal*, (3), 2015.

[4] Jan Behrens, Axel Kistner, Andreas Nitsche, and Björn Swierczek. *The principles of LiquidFeedback*. Interaktive Demokratie, 2014.

[5] Jan Behrens, Axel Kistner, Andreas Nitsche, and Björn Swierczek. The temporal dimension in the analysis of liquid democracy delegation graphs. *The Liquid Democracy Journal*, (7), 2022.

[6] Francisco M Bersetche. Generalizing liquid democracy to multi-agent delegation: A voting power measure and equilibrium analysis. *arXiv preprint arXiv:2209.14128*, 2022.

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

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

[9] Markus Brill, Théo Delemazure, Anne-Marie George, Martin Lackner, and Ulrike Schmidt-Kraepelin. Liquid democracy with ranked delegations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 4884–4891, 2022.

[10] Ioannis Caragiannis and Evi Micha. A contribution to the critique of liquid democracy. In *IJCAI*, pages 116–122, 2019.

[11] Zoé Christoff and Davide Grossi. Binary voting with delegable proxy: An analysis of liquid democracy. In *Proceedings of the Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 2017.

[12] Gal Cohensius, Shie Manor, Reshef Meir, Eli Meirom, and Ariel Orda. Proxy voting for better outcomes. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 858–866, 2017.

[13] R. Colley, T. Delemazure, and H. Gilbert. Measuring a priori voting power in liquid democracy. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2607–2615, 2023.

[14] Rachael Colley, Umberto Grandi, and Arianna Novaro. Unravelling multi-agent ranked delegations. *Autonomous Agents and Multi-Agent Systems*, 36(1):1–35, 2022.

[15] Palash Dey, Arnab Maiti, and Amatya Sharma. On parameterized complexity of liquid democracy. In *Proceedings of the Conference on Algorithms and Discrete Applied Mathematics*, pages 83–94, 2021.

[16] Bruno Escoffier, Hugo Gilbert, and Adèle Pass-Lanneau. The convergence of iterative delegations in liquid democracy in a social network. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*, pages 284–297, 2019.

[17] Bruno Escoffier, Hugo Gilbert, and Adèle Pass-Lanneau. Iterative delegations in liquid democracy with restricted preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1926–1933, 2020.

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

[19] Pallavi Jain, Krzysztof Sornat, and Nimrod Talmon. Preserving consistency for liquid knapsack voting. In *Proceedings of the European Conference on Multi-Agent Systems (EUMAS)*, pages 221–238, 2022.

[20] Anson Kahng, Simon Mackenzie, and Ariel Procaccia. Liquid democracy: An algorithmic perspective. *Journal of Artificial Intelligence Research*, 70:1223–1252, 2021.

[21] Matthias Köppe, Martin Koutecký, Krzysztof Sornat, and Nimrod Talmon. Fine-grained liquid

democracy for cumulative ballots. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, page 1029–1037. IFAAMAS, 2024.

[22] Grammateia Kotsialou and Luke Riley. Incentivising participation in liquid democracy with breadth-first delegation. In *Proceedings of the International Conference on Autonomous Agents and MultiAgent Systems*, pages 638–644, 2020.

[23] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. *Computer science review*, 3 (2):65–69, 2009.

[24] Evangelos Markakis and Georgios Papasotiropoulos. As Time Goes By: Adding a Temporal Dimension to Resolve Delegations in Liquid Democracy. In *Proceedings of the International Conference on Algorithmic Decision Theory (ADT)*, pages 48–63, 2025.

[25] Jonathan A Noel, Mashbat Suzuki, and Adrian Vetta. Pirates in wonderland: Liquid democracy has bicriteria guarantees. In *Proceedings of the International Symposium on Algorithmic Game Theory (SAGT)*, pages 391–405, 2021.

[26] Georgios Papasotiropoulos and Ulrike Schmidt-Kraepelin. Liquid democracy: An annotated reading list. *ACM SIGecom Exchanges*, 22(2):130–135, 2025.

[27] Giannis Tyrovolas, Andrei Constantinescu, and Edith Elkind. Unravelling expressive delegations: Complexity and normative analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 9918–9925, 2024.

[28] Markus Utke and Ulrike Schmidt-Kraepelin. Anonymous and copy-robust delegations for liquid democracy. *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 36: 69441–69463, 2023.

[29] Hiroshi Yamakawa, Michiko Yoshida, and Motohiro Tsuchiya. Toward delegated democracy: Vote by yourself, or trust your network. *Proceedings of the International Journal of Human and Social Sciences*, 1(2):146–150, 2007.

[30] Yuzhe Zhang and Davide Grossi. Power in liquid democracy. In *Proceedings of the AAAI conference on Artificial Intelligence*, volume 35, pages 5822–5830, 2021.

Markus Brill
University of Warwick
markus.brill@warwick.ac.uk

Rachael Colley
University of Glasgow
rachael.colley@glasgow.ac.uk

Anne-Marie George
University of Oslo
annemage@ifi.uio.no

Grzegorz Lisowski
AGH University
University of Groningen
glisowski@agh.edu.pl

Georgios Papasotiropoulos
University of Warsaw
gpapasotiropoulos@gmail.com

Ulrike Schmidt-Kraepelin
TU Eindhoven
u.schmidt.kraepelin@tue.nl