Cycles in Liquid Democracy: A Game-Theoretic Justification

Markus Brill[1], **Rachael Colley**[2], **Anne-Marie George**[3], **Grzegorz Lisowski**[4], **Georgios Papasotiropoulos**[5] and **Ulrike Schmidt-Kraepelin**[6]

[1]University of Warwick, [2]University of Glasgow, [3]University of Oslo, [4]AGH University, [5]University of Warsaw, [6]Eindhoven University of Technology

Abstract

Liquid democracy is a voting scheme in which individuals either vote directly or delegate their voting power to others. A common critique in the academic literature 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 long-term strategic delegation behavior. 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 our theoretical findings, we perform computational experiments using best-response dynamics.

1 Introduction

Liquid democracy is a flexible voting system that allows voters to either vote directly or *delegate* their vote to another participant, who can vote on their behalf. Delegations are *transitive*, meaning delegated votes can be delegated further, creating *delegation chains*. The participant at the end of a chain casts a ballot 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 mode of participation individually [Blum and Zuber, 2016].

Over the past decade, academic interest in liquid democracy has grown rapidly (see Section 1.1). However, practitioners have noted that some parts of this literature overlook or misinterpret key aspects of liquid democracy as it is implemented in practice [Behrens, 2015; Behrens *et al.*, 2014, 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 an one-time event. Specifically, this divergence has led to differing views on the issue of *delegation cycles*. Namely, in the literature *delegating* and *voting* are frequently modeled as mutually exclusive options. Consequently, a *delegation cycle* (i.e., a situation where a voter i delegates to j, and this delegation eventually returns to i through a chain of delegations) is seen as problematic because none of the voters in the cycle eventually casts a vote, leaving their collective voting weight unused. In contrast, Behrens *et al.* [2022] propose an alternative interpretation: 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, delegation cycles are not only unproblematic but may even be intentionally created by voters who trust each other. For instance, a group of like-minded participants may create a delegation cycle so that all their voting weight is used as long as at least one of them is casting a vote.

To bring theory closer to practice, we develop the *default delegation model* capturing long-term delegation decisions by voters. We are interested in scenarios where an electorate must make multiple decisions, referring to each such decision as an *election*. In our setting, voters declare *default delegations*, to be used when they do not cast a ballot. Participation in each election is probabilistic: voters have different probabilities for taking part in elections and there is no *a priori* distinction between *delegating* and *casting* voters. As a result, a voter's *utility* from an election depends not only on the delegation decisions of others but also on the voters casting a vote for the considered election. In this setting, 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. Our utility model enables us to analyze the behavior of voters who strategically aim to optimize their expected utility over a long-term horizon.

We provide theoretical evidence supporting the practical observation that delegation cycles naturally arise among rational users of liquid democracy platforms. More concretely, we study the (i) *existence*, (ii) *structure* and (iii) *quality* of (pure) Nash equilibria in the default delegation model and demonstrate that, under mild assumptions, delegation cycles are necessarily being formed. All 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 has shown a growing interest in liquid democracy, with various models and methodologies emerging rather than converging on a standardized framework. Below, we discuss studies that are related to ours in terms of modeling choices and methodological approach.

Central Role of Cyclic Delegations. A key branch of the literature addresses delegation cycles, typically viewing them as undesirable and proposing solutions to eliminate them [Brill *et al.*, 2022; Christoff and Grossi, 2017; Colley *et al.*, 2022; Dey *et al.*, 2021; Jain *et al.*, 2022; Köppe *et al.*, 2022; Kotsialou and Riley, 2020; Markakis and Papasotiropoulos, 2025; Tyrovolas *et al.*, 2024; Utke and Schmidt-Kraepelin, 2023]. All 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 [2025], like ours, was directly motivated by the study of Behrens *et al.* [2022]. Their work focuses on a temporal framework as well, but differs in the modeling choice: it examines delegation updates over discrete time-steps, whereas in our case, the temporal aspect is captured through the probabilistic model for ballot casting.

One-Dimensional Spatial Models. In order to model voters' preferences over their potential representatives, we consider voters as positioned in an one-dimensional metric space. This is a common modeling scenario (often representing political alignment) which has been explored in settings of delegative voting by Anshelevich *et al.* [2021], Cohensius *et al.* [2017], Escoffier *et al.* [2020], Green-Armytage [2015] and Yamakawa *et al.* [2007], though without the temporal aspect.

Strategic Delegation Behavior. A prominent line of research on liquid democracy frameworks focuses on the game-theoretic perspective. Indicatively, we refer to the works of Armstrong and Larson [2021], Bersetche [2022], Bloembergen *et al.* [2019], Escoffier *et al.* [2019], Noel *et al.* [2021] and Zhang and Grossi [2021] who share the main examined solution concept (i.e., Nash equilibrium) with our study. These assume that voters can always opt to vote directly, ensuring each non-casting voter is ultimately represented by another, whereas we model representation through a probability distribution over different ultimate delegates and ballot loss.

1.2 Our Contribution

A central contribution of our work is conceptual: we present the novel *default delegation model* for liquid democracy that captures and explains the long-term strategic delegation behavior of participants of real-life liquid democracy systems. We analyze this model from a game-theoretic perspective and address the following considerations.

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 equilibria. In more general settings, computational experiments reveal that the vast majority of components contain cycles. The width of these cycles appear 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 to obtain a close-to-optimal social welfare.

Omitted proofs and further details can be found in the Technical Appendix.

2 The Default Delegation Model

We consider a finite set V of *voters* using a liquid democracy platform. Each voter $i \in V$ nominates a default delegate. We assume that voters' default delegations remain fixed across a given series of elections. For each of those elections, each voter $i \in V$ may either cast a vote or abstain. If a voter does *not* cast a vote, their voting power is passed to their default delegate, continuing transitively until a casting voter is reached. This casting voter is referred to as voter i's *ultimate delegate*. If none of the voters in the chain of default delegations casts a vote, then voter i has no ultimate delegate for that election, and their voting power is lost. While default delegates remain fixed across elections, the set of casting voters changes from election to election, resulting in different ultimate delegates for non-casting voters.

Default Delegations. For each voter $i \in V$, 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. Consider an individual election and let $X \subseteq V$ denote the set of voters casting a vote in this election. Then, 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 a given election. Rather, when choosing a default 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 in each election with a fixed probability $p_i \in [0, 1]$, which remains constant across all elections. Moreover, whether or not a voter casts a vote is determined independently for each election. 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 model reflects the idea that voters are not aware of (or engaged in) every single election. The probability p_i can be interpreted as the fraction of elections in which voter i typically casts a vote. 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) $d(y_k) \in \{y_0, y_1, y_2, \cdots, y_k\}$.

The ultimate delegate of voter i 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 from a single election 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 an 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 across all elections. Let $\boldsymbol{x} = (x_i)_{i \in V}$ denote the *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 single election is defined as their tolerance minus the distance to their ultimate delegate, or 0 if the ultimate delegate is undefined. In 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 voter i is a random variable (distributed according to (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*) 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 $\boldsymbol{d}_{SW}$.

We conclude this section with an example.

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

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

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

$\boldsymbol{p}$:	0.8	0.3	0.2	0.3	0.1	0.3
$\boldsymbol{x}$:	0.2	0.3	0.4	0.5	0.6	0.8

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.

In order to compute voters' acceptability sets and utility, let us assume that $\tau_i = 0.25$ for all $i \in V$. Then, e.g., $\mathcal{A}_D(\boldsymbol{x}, \boldsymbol{\tau}) = \{B, C, D, E\}$. Furthermore, we can calculate the expected utility of voters w.r.t. delegation profile $\boldsymbol{d}$. For instance, $u_A(\boldsymbol{d}) = 0.8 \cdot 0.25 + 0.2 \cdot 0.3 \cdot (0.25 - 0.1) + 0.2 \cdot 0.7 \cdot 0.2 \cdot (0.25 - 0.2) \approx 0.21$ 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. This can be verified with the help of Table 1.*

	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 profile $\boldsymbol{d}$ 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, $\mathbf{d}$ is not the only NE of this instance. It can be verified that $\mathbf{d}'$, the delegation graph $G_{\mathbf{d}'}$ of which appears in the following figure, is also a NE.

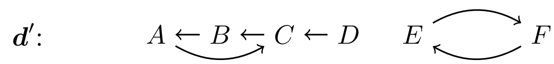

We note that $G_{\mathbf{d}'}$ has two (weakly) connected components, in contrast to $G_{\mathbf{d}}$. Moreover, $\mathbf{d}$ and $\mathbf{d}'$ differ in terms of social welfare, casting voters' weights and expected number of votes that are lost (a metric we are referring to in Appendix B.4).

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. We provide details about best-response dynamics and these experiments in Appendix C.1.

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. *Consider an instance $\mathcal{I}$ defined as follows: $V = \{A, B, C, D\}$, $\boldsymbol{x} = (0, 0.05, 0.1, 0.5)$, $\boldsymbol{\tau} = (1, 0, 0.2, 0)$, $\boldsymbol{p} = (0.4, 0.05, 0.2, 0.4)$, and suppose that it admits a NE $\mathbf{d}$. Let us start by observing that in $\mathbf{d}$, the voters B and D must choose to delegate to themselves as $\tau_B = \tau_D = 0$, i.e., delegating to any other voter lowers their utility compared to self-delegation. Hence, since $\tau_C = 0.2$, $dist(C, D) = 0.4$ and $d(D) = D$, we conclude that C does not delegate to D in $\mathbf{d}$. Moreover, we observe that in $\mathbf{d}$, 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 following table we show the utilities of A and C in all profiles that were not ruled out by the previous reasoning.*

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

Specifically, the table depicts the normal form representation of the game induced by $\mathcal{I}$, where rows correspond to the possible choices of voter C and columns to the choices of voter A, for specifying $\mathbf{d}$. It is routine to check that none of the possible delegation profiles $\mathbf{d}$ is a Nash equilibrium.

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

3.1 Special Cases

We discuss three special cases of the default delegation model and draw a complete picture on 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 lead to guaranteeing the existence of NE.

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. *For the following restrictions, there exists an instance $\mathcal{I}$ for which no Nash equilibrium exists:*

(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 or 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. Furthermore, voter C can be considered a leftist as $\mathcal{A}_C(\mathcal{I})$ only includes voters to their left. Thus, we observe that if a profile has both leftists and rightists, a NE is not guaranteed to exist. In contrast, we show in Theorem 3 that any instance with only leftists or only rightists contains a NE.

For the sake of concreteness, we define an example for 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 where 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, effectively we 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 [Cohensius *et al.*, 2017; Anshelevich *et al.*, 2021], a variant of liquid democracy in which voters are divided into two groups — delegating voters and casting voters — with ballots being delegated only from the former to the latter, forming delegation chains on 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. That is, 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 existence of NE.

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.*

Proof Sketch. The proof begins by assuming, for contradiction, that a weakly connected component W of $G_{\boldsymbol{d}}$ with at least two vertices has no cycle. 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. The uniqueness follows directly from the fact that each vertex in the component has out-degree 1. □

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

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 $G_{\mathbf{d}}$, where $\mathbf{d}$ is a Nash equilibrium of an instance that is not of mutual acceptance or where deterministic voters exist.*

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

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

If mutual acceptance does not hold, 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 $\mathbf{d}$ of $\mathcal{I}$. Consider a weakly connected component W of $G_{\mathbf{d}}$ 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 $G_{\mathbf{d}}$, $\mathcal{L}(W)$ 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 Appendix B.1 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 (see Appendix C.3 for details). 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 4.5 across instances and grows only slightly even with 200 voters. 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 now turn our focus to evaluating the quality of equilibria. We follow a Price-of-Anarchy approach, comparing the social welfare of Nash equilibria to the best possible social welfare [Koutsoupias and Papadimitriou, 2009]. Before that, we compare the structure of social-welfare-maximizing delegation graphs to that of Nash equilibria, observing an interesting contrast.

Observation 9. *There exist mutual acceptance instances without deterministic voters in which the delegation graph maximizing social welfare does not contain a cycle. For example, consider the symmetric instance with $V = \{A, B, C\}$, $x = (0.12, 0.5, 0.88)$, $p = (0.1, 0.9, 0.1)$, and where $\tau_i = 0.4$ for all $i \in V$. Social welfare is maximized if both A and C delegate to B, who self-loops.*

We define the notion of *Price of Anarchy (PoA)* of an instance $\mathcal{I}$ in the standard way:

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

where $\boldsymbol{d}_{\mathtt{SW}}(\mathcal{I})$ is the profile maximizing social welfare and $\boldsymbol{d}_{\mathtt{NE}}(\mathcal{I})$ is the profile achieving the *lowest* social welfare among the Nash equilibria of $\mathcal{I}$. The following result demonstrates that this ratio can be arbitrarily large.

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

Proof sketch. We prove the statement by describing a family of instances, parameterized by ε and n. In the limit for $\varepsilon \to 0$ and $n \to \infty$, the worst Nash equilibrium $\boldsymbol{d}$ satisfies $SW(\boldsymbol{d}) \to 0$ and there is a 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 1. Note that voters are equidistant: $\text{dist}(i, i+1) = \varepsilon/n, \forall i \in [n]$.

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

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

Next, consider the profile $\boldsymbol{d}'$ with $d'(i) = i + 1$ for $i \leqslant n$ and $d'(n+1) = n+1$. It can be shown that $u_i(\boldsymbol{d}') \xrightarrow{n \to \infty} 0$ for all $i \geqslant 2$. We now compute $u_1(\boldsymbol{d}')$ which is equal to the following expression.

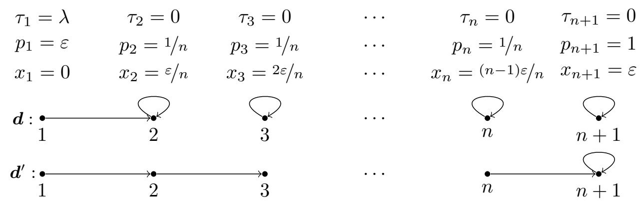

Figure 1: Illustration of the instance $\mathcal{I}_{\varepsilon,n}$ with $n+1$ voters from the proof of Theorem 10, alongside the examined delegation profiles.

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

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 strongly negative result concerning the quality of NE. Note, however, that the constructed instances have certain characteristics, such as a voter with a very low voting probability ($p_1 \to 0$), 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 considered 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 quality of equilibria by focusing on the difference rather than the ratio may lead to less negative conclusions. We define the *additive Price of Anarchy* of an instance $\mathcal{I}$ as

$$PoA^+(\mathcal{I}) = SW(\boldsymbol{d}_{\texttt{SW}}(\mathcal{I})) - SW(\boldsymbol{d}_{\texttt{NE}}(\mathcal{I})),$$

and proceed with the following positive results on both the multiplicative and the additive Price of Anarchy.

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 Appendix C.5, 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 Appendix B.2, 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 for large instances, we approximate its welfare by the sum of each voter's expected utility under their optimal delegation profile, denoted by $\texttt{ODP}(\mathcal{I})$. Formally, $\texttt{ODP}(\mathcal{I}) = \sum_{i\in V} u_i(\mathbf{d}^{\mathbf{i}*})$, where $\mathbf{d}^{\mathbf{i}*}$ is a delegation profile maximizing the utility of voter i. In Appendix B.3, 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, $SW(\boldsymbol{d}_{SW}) \leqslant \texttt{ODP}(\mathcal{I})$.

	Number of Voters				$\tau_i \in [0, \tau_{\max}]$ with $\tau_{\max} =$		
	20	50	100	200	1	0.75	0.5
$\boldsymbol{d}_{BR}$	97.5%	98.8%	99.3%	99.7%	98.8%	98.8%	98.5%
$\boldsymbol{d}_{\texttt{dir}}$	50.9%	50.8%	51.0%	49.9%	50.4%	50.7%	51.5%

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

We generated 30 instances for each number of voters $n \in \{20, 50, 100, 200\}$, with values for $\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\tau}$ chosen uniformly at random. For $n = 50$, we furthermore tested 20 tolerance vectors $\boldsymbol{\tau}$, scaling each by 0.75 and 0.5 to assess the effect of different tolerance levels (full details on the instances can be found in Appendix C.2). For each instance, we computed $\texttt{ODP}(\mathcal{I})$ 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 $\boldsymbol{d}_{\texttt{dir}}$ ("direct voting") in which every voter self-loops.

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$97,5% of $\texttt{ODP}(\mathcal{I})$) and gets closer to $\texttt{ODP}(\mathcal{I})$ as n increases. Given that $\texttt{ODP}(\mathcal{I})$ is an upper bound on 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 long-term 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. A natural starting point could consist in placing voters into a two-dimensional Euclidean space. Preliminary experiments reveal that the structure of Nash equilibria becomes more complicated in that setting. Additionally, considering alternative (e.g., normalized) utility functions could yield further insights. Finally, while our work primarily focused on cycles, the role of *long delegation paths* remains an important and underexplored aspect of liquid democracy. These paths are often seen as undesirable due to their potential to erode trust in ultimate delegates. Potentially, this issue could be explored with the help of a model similar to ours.

References

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.

Ben Armstrong and Kate Larson. On the limited applicability of liquid democracy, 2021. Appeared at *Games, Agents, and Incentives Workshop (GAIW)*.

Jan Behrens, Axel Kistner, Andreas Nitsche, and Björn Swierczek. *The principles of LiquidFeedback*. Interacktive Demokratie, 2014.

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.

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

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

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.

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

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.

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.

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.

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

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.

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.

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.

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

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.

Matthias Köppe, Martin Kouteckỳ, Krzysztof Sornat, and Nimrod Talmon. Fine-grained liquid democracy for cumulative ballots. *arXiv preprint arXiv:2208.14441*, 2022.

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.

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

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.

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.

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.

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.

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.

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

Cycles in Liquid Democracy: A Game-Theoretic Justification

Technical Appendix

A Omitted Proofs

Proof of Theorem 1

(i) Suppose that the voters in V are split into two disjoint groups: The casting voters, i.e., those in the set $X = \{i \in V \mid p_i = 1\}$, and the delegators, i.e., those in the set $V \setminus X = \{i \in V \mid p_i = 0\}$. We set $d(i) = i = \text{argmin}_{j \in \mathcal{A}_i(\boldsymbol{x},\boldsymbol{\tau}) \cap X}(|x_i - x_j|)$ for all voters in X. For each $i \in V \setminus X$, we also set

$$d(i) = \underset{j \in \mathcal{A}_i(\boldsymbol{x},\boldsymbol{\tau}) \cap X}{\text{argmin}} (|x_i - x_j|),$$

i.e., they delegate to their closest casting voter in their acceptability set. If $\mathcal{A}_i(\boldsymbol{x},\boldsymbol{\tau}) \cap X = \emptyset$, then we simply set $d(i) = i$. We proceed to show that $\boldsymbol{d}$ is a Nash equilibrium. No voter $i \in X$ has a profitable deviation, as their expected utility will always be equal to τ_i regardless of the delegation they choose, due to $p_i = 1$. For voters $i \in V \setminus X$, we consider two cases, i.e., when $d(i) = i$ and when $d(i) \neq i$.

In the first case, namely when $d(i) = i$, if i deviates, they will finally be represented by a voter in X, receiving a utility of 0, or by a voter in $V \setminus X$. In the latter case they will necessarily be represented by a voter not in $\mathcal{A}_i(\boldsymbol{x},\boldsymbol{\tau})$, because $X \cap \mathcal{A}_i(\boldsymbol{x},\boldsymbol{\tau}) = \emptyset$. Thus, their expected payoff would be non-positive in both cases, and consequently it will be no more than the utility they receive with $d(i) = i$, which gives a utility of 0.

In the second case, when $d(i) \neq i$, delegating to any other voter in $\mathcal{A}_i(\boldsymbol{x},\boldsymbol{\tau}) \cap X$ will not increase their expected utility, because $d(i)$ is the closest casting voter to i. Similarly, delegating to a voter $j \in V \setminus X$ or to a voter $j \in X \setminus \mathcal{A}_i(\boldsymbol{x},\boldsymbol{\tau})$ might result in i being represented by a voter k satisfying $k = \text{argmin}_{j \in \mathcal{A}_i(\boldsymbol{x},\boldsymbol{\tau}) \cap X}(|x_i - x_k|)$, or not, depending on where j delegates. In both cases, this is again not a profitable deviation from $\boldsymbol{d}$ for i.

(ii) We let $V = \{A, B\}$ and we denote by $d = |x_B - x_A|$. First, we examine the expected utility for voter A, assuming that $d(B) = B$. If $d(A) = B$, then $u_A(d(A), d(B)) = \tau_A p_A + (1 - p_A) p_B (\tau_A - d)$, and if $d(A) = A$, then $u_A(d(A), d(B)) = \tau_A p_A$. Crucially, the utility of A remains exactly the same even when $d(B) = A$. Therefore, in an equilibrium, voter A decides to delegate to B if $(1 - p_A) p_B (\tau_A - d) \geqslant 0$, or equivalently if $\tau_A \geqslant d$, and self-loops otherwise, regardless of what strategy voter B selects. The analogous arguments can determine the delegation of voter B in an equilibrium.

(iii) We call the set of the three voters $V = \{1, 2, 3\}$ and we assume that the voters are named in increasing order of their positions in the line. We can assume that $p_i > 0$ for all $i \in \{1, 2, 3\}$. Otherwise, if there is a deterministic voter i for which this condition does not hold, we do the following: first, we exclude voter i from consideration and then we find a NE in the remaining instance that has at most two voters, which can be done as previously shown. The exclusion of i is safe because it is never profitable for a voter to delegate to i regardless of the choice of voter i. Finally, having fixed the delegation for each voter j such that $p_j > 0$, we fix a delegation for i by determining the optimal delegation for them, given the choices of others.

At a high level, the proof for finding a Nash equilibrium $\boldsymbol{d}$, goes as follows: First, we fix a delegation for voter 2 according to some criterion and, depending on this choice, we then fix the delegation for a (specific) second voter. Finally, we fix the delegation for the remaining voter as well. We will prove that by following this process and the specified criteria we achieve to specify a delegation profile that is a Nash equilibrium. Specifically, we will show that no voter will have an incentive to deviate after fixing their choice, no matter the choices of the rest.

Let us compute the utility that voter 2 will get from each possible delegation, assuming that $d(i) = i$ for $i \in \{1, 3\}$. Then, for each such i, it holds $u_2(1, i, 3) = p_2 \tau_2 + (1 - p_2) p_i (\tau_2 - \text{dist}(2, i))$. We set $d(2) = \text{argmax}_{i \in \{1,3\}} \{p_i(\tau_2 - \text{dist}(2, i))\}$, assuming that $\max_{i \in \{1,3\}} \{p_i(\tau_2 - \text{dist}(2, i))\} > 0$. The case where $\max_{i \in \{1,3\}} \{p_i(\tau_2 - \text{dist}(2, i))\} \leqslant 0$ will be addressed separately later.

Assume without loss of generality that $\text{argmax}_{i \in \{1,3\}} \{p_i(\tau_2 - \text{dist}(2, i))\} = 3$ and so that $d(2) = 3$. In that case, we continue with fixing the delegation for voter 1, and finally for voter 3 (we reverse the order of examination if $d(2) = 1$), and we claim, and we will soon prove, that regardless of those, voter 2 will not have an incentive to deviate from $d(2) = 3$. We begin with examining voter 1, i.e., the voter that voter 2 did not delegate to. We compute the utility that voter 1 would get if $d(2) = 3 = d(3)$, and based on this, we will fix $d(1)$. It holds that:

$$\begin{aligned}
u_1(1, 3, 3) &= p_1 \tau_1 \\
u_1(2, 3, 3) &= p_1 \tau_1 + (1 - p_1) p_2 (\tau_1 - \text{dist}(2, 1)) + (1 - p_1)(1 - p_2) p_3 (\tau_1 - \text{dist}(3, 1)) \\
u_1(3, 3, 3) &= p_1 \tau_1 + (1 - p_1) p_3 (\tau_1 - \text{dist}(3, 1)).
\end{aligned}$$

Since $\text{dist}(2,1) < \text{dist}(3,1)$ we also have that

$$\begin{aligned} u_1(2,3,3) &> p_1\tau_1 + (1-p_1)p_2(\tau_1 - \text{dist}(3,1)) + (1-p_1)(1-p_2)p_3(\tau_1 - \text{dist}(3,1)) \\ &= p_1\tau_1 + (1-p_1)(\tau_1 - \text{dist}(3,1))p_3 = u_1(3,3,3). \end{aligned}$$

Hence, we can assume that $d(1) \neq 3$. Then, we can determine the $\text{argmax}_{i\in\{1,2\}}\{u_1(i,3,3)\}$, and set $d(i)$ to be equal to that value. Note that no matter which the determined delegation choice for voter 1 is, it does not affect the utility of voter 2, so, for the time-being, voter 2 does not have an incentive to deviate.

We continue with fixing a delegation for voter 3 and we now have two cases to consider, i.e., $d(1) = 1$ and $d(1) = 2$.

- We first examine the case where $d(1) = 1$. Then, by mutual acceptance, $u_3(d(1), d(2), 1) < u_3(d(1), d(2), 3)$, or otherwise voter 1 would prefer a delegation to voter 3 than $d(1) = 1$. Similarly, $u_3(d(1), d(2), 2) > u_3(d(1), d(2), 3)$, or otherwise voter 2 would also not prefer a delegation to voter 3. Consequently, we always fix $d(3) = 2$ in this case. It is easy to check that no voter has an incentive to deviate, or otherwise voters would not have fixed the corresponding choices in the previous steps of the procedure.
- Suppose now that $d(1) = 2$. There are two possible options for voter 3 to consider: delegating to voter 1 and hence forming a 3-cycle, or delegating to voter 2 and hence forming a 2-cycle. We select for voter 3 the option among the two that maximizes their own utility, given that $d(2) = 3$ and $d(1) = 2$. We will show that in both cases neither voter 1 nor voter 2 have an incentive to deviate. More precisely, say first that $\boldsymbol{d}$ induces a 3-cycle, and voter 1 wants to deviate. This can only be when $\tau_1 < dist(1,j)$ for some $j \in \{2,3\}$. This cannot hold for $j = 2$ or otherwise voter 1 wouldn't have delegated to 2 beforehand and it cannot also hold for $j = 3$ or otherwise voter 3 wouldn't have delegated to 1. The argument for rejecting a potential deviation of voter 2 is similar. Now say that $\boldsymbol{d}$ induces a 2-cycle between voters 2 and 3 and $d(1) = 2$. The fact that at a previous step voter 2 selected to delegate to 3 means that a deviation of voter 2 will not be profitable for them. It remains to show that voter 1 doesn't have a profitable deviation either. To that end, we will prove that $u_1(\boldsymbol{d}_{-1}, 2) > u_1(\boldsymbol{d}_{-1}, 3)$. This is true due to the following equivalent expressions:

$$\begin{gathered} (\tau_1 - \text{dist}(1,2))p_2 + (\tau_1 - \text{dist}(1,3))p_3(1-p_2) > (\tau_1 - \text{dist}(1,3))p_3 + (\tau_1 - \text{dist}(1,2))p_2(1-p_3) \Leftrightarrow \\ (\tau_1 - \text{dist}(1,2))p_2p_3 > (\tau_1 - \text{dist}(1,3))p_3p_2 \Leftrightarrow \\ \text{dist}(1,2) < \text{dist}(1,3). \end{gathered}$$

Thus, the sequence in which voters are examined, along with the criteria that led to the specification of their delegations, resulted in a Nash equilibrium.

It remains to consider the case where $\max_{i\in\{1,3\}}\{p_i(\tau_2 - \text{dist}(2,i))\} \leqslant 0$. Equivalently, $(\tau_2 - \text{dist}(2,i)) \leqslant 0$, for $i \in \{1,3\}$. If this is the case, we set $d(i) = i$ for every $i \in \{1,2,3\}$. Then, each voter i experiences a utility of $p_i\tau_i$. Consider a possible deviation for voter 2, say by delegating to some voter $i \in \{1,3\}$. Then their new utility becomes $p_2\tau_2 + (1-p_2)p_i(\tau_2 - \text{dist}(2,i)) \leqslant p_2\tau_2$. Therefore, such a deviation is not profitable for voter 2. The same argument holds for all the remaining possible deviations, since, by mutual acceptance, if voter $1 \notin \mathcal{A}_2(\boldsymbol{x},\boldsymbol{\tau})$ and $3 \notin \mathcal{A}_2(\boldsymbol{x},\boldsymbol{\tau})$, then also $2 \notin \mathcal{A}_1(\boldsymbol{x},\boldsymbol{\tau})$ and $2 \notin \mathcal{A}_3(\boldsymbol{x},\boldsymbol{\tau})$, so both voters 1 and 3 would prefer to loop rather than delegating to voter 2, let alone to voter 3 and 1 respectively. □

Proof of Theorem 2

(i) Let us first show that there are instances without a NE even with only three voters present. Take an instance with three voters, A, B, and C, and say that $\boldsymbol{x} = (0.1, 0.9, 0.95)$, $\boldsymbol{p} = (0.02, 0.02, 0.1)$, and $\boldsymbol{\tau} = (0.9, 0.1, 0.9)$. Assume that there exists a NE profile $\boldsymbol{d}$ in that instance. First, notice that B does not choose to delegate to A in $\boldsymbol{d}$, as $\text{dist}(A,B) > \tau_B$. Second, we observe that A and C do not self-delegate, as $\text{dist}(A,B) < \tau_A$ and $\text{dist}(C,B) < \tau_C$. Figure 2 depicts the utilities of the voters in all the remaining strategy profiles. It is routine to check that there is no NE in this instance.

(ii) Consider the following instance $\mathcal{I}$:

$$\begin{aligned} p &= (0.05, 0.49, 0.1, 1.0), \\ x &= (0.0, 0.05, 0.15, 0.3), \\ \tau &= 0.2. \end{aligned}$$

We name the voters of the instance as $V = \{0,1,2,3\}$. We first observe that it is sufficient to focus on delegation profiles where voter 3 loops. This is because if there is a Nash equilibrium where voter 3 doesn't loop, there is another delegation profile where they loop and that it is a Nash equilibrium as well. All possible delegation profiles $\boldsymbol{d}$ in which $d(3) = 3$ are depicted in (the first column of) Table 3. For each such $\boldsymbol{d}$ we identify a voter $i \in V$ for which there exists a profitable deviation, proving that no Nash equilibrium exists in $\mathcal{I}$. □

	$d(B) = C$	$d(B) = B$
$d(A) = B$	0.0247, -0.0104, 0.1082	0.0199, 0.002, 0.1058
$d(A) \to C$	0.0229, -0.0094, 0.0909	0.0229, 0.002, 0.0909

	$d(B) = C$	$d(B) = B$
$d(A) = B$	0.0247, 0.0069, 0.1053	0.0199, 0.002, 0.1053
$d(A) = C$	0.0246, 0.0069, 0.1053	0.0246, 0.002, 0.1053

Figure 2: Normal form representation of the game induced by the instance described in the proof of Theorem 2. The table on the left corresponds to C delegating to A, and the one on the right to C delegating to B. Rows correspond to the choices of voter A and columns to the choices of voter B. The numbers in each cell show the utility of A, B, and C respectively, for the delegation profile under examination.

Proof of Theorem 3

We focus on rightists; the case for leftists is analogous. The proof proceeds by a greedy method. We consider the voters in order of decreasing position, to be denoted as $v_1, v_2, \cdots, v_n$. Voter v_n, which corresponds to the voter at position $x_{\max} = \text{argmax}_{i \in V}\{x_i\}$, will self-loop in an equilibrium, simply due to the fact that, as long as $p_n \neq 1$, any delegation will result in being represented (with some probability) by someone in their left, which gives a negative utility to a rightist voter. If $p_n = 1$, self-looping, trivially, doesn't admit a profitable deviation.

Consider now voter v_{n-1}. Being a rightist, v_{n-1} can only delegate to themselves or to v_n towards getting a non-negative utility. The decision of this voter isn't affected by delegations to them or edges to voters on their left. Therefore, the delegation choices of voters $v_1, \cdots, v_{n-2}$, which will be fixed in a later step of the procedure, cannot influence the decision of v_{n-1}. Thus, for v_{n-1}, we only need to check whether delegating to v_n or to themselves produces more expected utility (knowing that v_n loops) and fix the corresponding choice.

We proceed similarly for the remaining voters in order. Notice that the decision of v_{n-1} impacts only the decisions of voters positioned before v_{n-1}, who will be examined next. In the specified order, when it is time to consider a voter with a lower position, they will choose the best delegation option available to them, taking into account the already-fixed delegations of voters positioned on their right and the fact that they cannot delegate to anyone positioned on their left. Once such a delegation is chosen, it becomes fixed, as changes in the delegations of voters positioned on their left will not affect their expected utility and make them interested in deviating. □

Proof of Theorem 4

We will present a procedure that identifies a Nash equilibrium $\boldsymbol{d}$, in a given instance of the proxy voting setting. Initially we call all voters as *unclassified*. Additionally, we will refer to the following sets of voters: *represented* and *representatives*, initially both empty.

For each unclassified voter i in an arbitrary order, we determine the voter j minimizing the quantity $p_j(\tau_i - \text{dist}(i, j))$ among all j such that j doesn't yet belong to the set of represented voters (i included). Note that if j is represented, a delegation from i to j would create a chain of 3 voters, hence such a delegation is infeasible. Then, we fix $d(i) = j$ and we remove both i and j from the set of unclassified voters, labeling i as a represented and j as a representative. Subsequently, we repeat for another unclassified voter. When all unclassified voters have been examined, we move to the second phase of the algorithm.

At the end of the described first phase, all voters are either represented or representatives but representatives do not have an outgoing edge yet in $G_{\boldsymbol{d}}$ (or they have a self-loop). For those without a loop, it remains to determine their delegation. Then, we consider each such a voter in arbitrary order. Say we examine voter i, the only choices for i are to delegate to one of the voters in $\{k \in V \mid d(k) = i\}$, i.e. those who already delegated to i or to self-loop, due to the constraint of the proxy voting setting. Among those options, i picks the delegation that maximizes their utility, which again corresponds to the voter $j \in \{k \in V \mid d(k) = i\} \cup \{i\}$, maximizing the quantity $p_j(\tau_i - \text{dist}(i, j))$.

To prove that the constructed delegation forms indeed a Nash equilibrium, we consider a voter i and we assume that they prefer to deviate, towards a contradiction. First, we examine a voter i that self-delegates. By construction, there are no feasible choices for i that would improve their utility. Assume now that i is a voter that is in the set of represented voters and didn't self-loop. Then, i is delegating to the one giving them the maximal utility, at the moment we considered i in the first round. It

Delegation Profile	Deviating voter	Current Delegation	Profitable Deviation	Utilities for Deviating voter
$[0, 0, 0, 3]$	0	0	1	[0.0100, 0.0798, 0.0148, -0.0850]
$[0, 0, 1, 3]$	0	0	1	[0.0100, 0.0798, 0.0776, -0.0850]
$[0, 0, 2, 3]$	0	0	1	[0.0100, 0.0798, 0.0148, -0.0850]
$[0, 0, 3, 3]$	0	0	1	[0.0100, 0.0798, -0.0707, -0.0850]
$[0, 1, 0, 3]$	0	0	1	[0.0100, 0.0798, 0.0148, -0.0850]
$[0, 1, 1, 3]$	0	0	1	[0.0100, 0.0798, 0.0776, -0.0850]
$[0, 1, 2, 3]$	0	0	1	[0.0100, 0.0798, 0.0148, -0.0850]
$[0, 1, 3, 3]$	0	0	1	[0.0100, 0.0798, -0.0707, -0.0850]
$[0, 2, 0, 3]$	0	0	1	[0.0100, 0.0822, 0.0148, -0.0850]
$[0, 2, 1, 3]$	0	0	1	[0.0100, 0.0822, 0.0776, -0.0850]
$[0, 2, 2, 3]$	0	0	1	[0.0100, 0.0822, 0.0148, -0.0850]
$[0, 2, 3, 3]$	0	0	1	[0.0100, 0.0386, -0.0707, -0.0850]
$[0, 3, 0, 3]$	0	0	1	[0.0100, 0.0314, 0.0148, -0.0850]
$[0, 3, 1, 3]$	0	0	2	[0.0100, 0.0314, 0.0340, -0.0850]
$[0, 3, 2, 3]$	0	0	1	[0.0100, 0.0314, 0.0148, -0.0850]
$[0, 3, 3, 3]$	0	0	1	[0.0100, 0.0314, -0.0707, -0.0850]
$[1, 0, 0, 3]$	1	0	2	[0.1018, 0.0980, 0.1065, 0.0725]
$[1, 0, 1, 3]$	1	0	2	[0.1018, 0.0980, 0.1031, 0.0725]
$[1, 0, 2, 3]$	1	0	2	[0.1018, 0.0980, 0.1031, 0.0725]
$[1, 0, 3, 3]$	2	3	1	[0.0641, 0.0652, 0.0200, 0.0650]
$[1, 1, 0, 3]$	1	1	2	[0.1018, 0.0980, 0.1065, 0.0725]
$[1, 1, 1, 3]$	1	1	2	[0.1018, 0.0980, 0.1031, 0.0725]
$[1, 1, 2, 3]$	1	1	2	[0.1018, 0.0980, 0.1031, 0.0725]
$[1, 1, 3, 3]$	1	1	0	[0.1018, 0.0980, 0.0802, 0.0725]
$[1, 2, 0, 3]$	2	0	3	[0.0641, 0.0641, 0.0200, 0.0650]
$[1, 2, 1, 3]$	2	1	3	[0.0641, 0.0641, 0.0200, 0.0650]
$[1, 2, 2, 3]$	2	2	3	[0.0641, 0.0641, 0.0200, 0.0650]
$[1, 2, 3, 3]$	1	2	0	[0.1018, 0.0980, 0.0802, 0.0725]
$[1, 3, 0, 3]$	1	3	2	[0.1018, 0.0980, 0.1065, 0.0725]
$[1, 3, 1, 3]$	0	1	2	[0.0100, 0.0314, 0.0340, -0.0850]
$[1, 3, 2, 3]$	1	3	2	[0.1018, 0.0980, 0.1031, 0.0725]
$[1, 3, 3, 3]$	1	3	0	[0.1018, 0.0980, 0.0802, 0.0725]
$[2, 0, 0, 3]$	0	2	1	[0.0100, 0.0798, 0.0148, -0.0850]
$[2, 0, 1, 3]$	0	2	1	[0.0100, 0.0798, 0.0776, -0.0850]
$[2, 0, 2, 3]$	0	2	1	[0.0100, 0.0798, 0.0148, -0.0850]
$[2, 0, 3, 3]$	0	2	1	[0.0100, 0.0798, -0.0707, -0.0850]

Table 3: For each delegation profile $\boldsymbol{d}$ satisfying $d(3) = 3$ (column 1), we identify a voter $i \in N$ (column 2) who can improve their expected utility by unilaterally changing their delegation from $d(i)$ (column 3) to $d(i)'$ (column 4). Column 5 shows the expected utility voter i would obtain by delegating to each voter in V, while keeping the rest voters' delegations unchanged.

holds that i can't delegate to anyone that delegates further, so every other voter in the set of represented voters isn't a feasible choice. Among the representatives, i chose the best for themselves option. Suppose now that i is in the set of representatives and didn't self-loop. Then, simply, no deviation to a voter that doesn't delegate to i is feasible for i, and among those, i has been assigned to their preferred one. □

Proof of Theorem 5

Consider a weakly connected component of a Nash equilibrium $\boldsymbol{d}$ of $\mathcal{I}$ that has at least 2 vertices, or, in other words, at least one (directed) edge. Towards a contradiction, suppose that it does not have a directed cycle. Then, its undirected variant must be a tree, meaning it has a sink vertex that self-delegates. Let this vertex correspond to a voter i, and let j be a voter such that $d(j) = i$. The utility of i under $\boldsymbol{d}$ equals $\tau_i p_i$.

If voter i were to delegate back to voter j, then i would obtain a utility of $\tau_i p_i + (1 - p_i)(\tau_i - \text{dist}(i, j))p_j$. Since the delegation from i to j was not selected in the equilibrium $\boldsymbol{d}$ (otherwise i would not have been a sink), it must hold that

$$\tau_i p_i + (1 - p_i)(\tau_i - \text{dist}(i, j))p_j \leqslant \tau_i p_i \Leftrightarrow (1 - p_i)(\tau_i - \text{dist}(i, j))p_j \leqslant 0.$$

Therefore, it must hold that $p_i = 1$, or $\tau_i \leqslant \text{dist}(i, j)$, or $p_j = 0$. The first and third cases do not hold by assumption. Consequently, it is necessarily true that

$$\tau_i - \text{dist}(i, j) \leqslant 0. \tag{3}$$

Since j delegates to i at $\boldsymbol{d}$, the utility of j equals $\tau_j p_j + (1 - p_j)(\tau_j - \text{dist}(i, j))p_i$, and this value must be no less than $\tau_j p_j$, or j would prefer to self-loop. Therefore,

$$(1 - p_j)(\tau_j - \text{dist}(i, j))p_i \geqslant 0 \Rightarrow \tau_j - \text{dist}(i, j) \geqslant 0.$$

By assumption, $\tau_j - \text{dist}(i, j) \neq 0$, and hence, $\tau_j - \text{dist}(i, j) > 0$. Consequently, $i \in \mathcal{A}_j(\boldsymbol{x}, \boldsymbol{\tau})$, and by the assumption of mutual acceptance between voters, $j \in \mathcal{A}_i(\boldsymbol{x}, \boldsymbol{\tau})$. Therefore, $\tau_i - \text{dist}(i, j) > 0$, which contradicts Expression (3). Thus, the considered weakly connected component must have at least one directed cycle.

We will now prove that no more than a single cycle may appear in a weakly connected component. Consider such a component with at least one cycle and let it contain k vertices. Since each vertex has out-degree 1, the component also has k edges. Removing now one edge from the cycle leaves the component connected, with k vertices and $k - 1$ edges. This structure is necessarily a tree, hence, the original component can be viewed as a tree plus one additional edge, and adding a single edge to a tree creates no more than a single cycle, proving the claim. □

Proof of Observation 6

We will show that there exists an instance $\mathcal{I}$ with a Nash equilibrium $\boldsymbol{d}^*$ such that there is a weakly connected component of at least two vertices in $G_{\boldsymbol{d}^*}$ that does not have a cycle. The instance $\mathcal{I}$ will have at least one of the following properties: (1) for some voter $i \in V$, we have that $p_i = 1$, (2) for some voter $j \in V$, we have that $p_j = 0$, (3) there is a pair of voters $(i, j) \in V \times V$ such that $i \in \mathcal{A}_j(\boldsymbol{x}, \boldsymbol{\tau})$ but $j \notin \mathcal{A}_i(\boldsymbol{x}, \boldsymbol{\tau})$.

Consider an instance $\mathcal{I}$ with only two voters, namely i and j, where $i \in \mathcal{A}_j(\boldsymbol{x}, \boldsymbol{\tau})$. We call $\boldsymbol{d}^*$ the profile in which i self-delegates and j delegates to i. Observe that no cycle appears in the connected component of size 2 of $G_{\boldsymbol{d}^*}$. Note that j receives a utility under $\boldsymbol{d}^*$ that is at least as large as under $(\boldsymbol{d}^*_{-j}, j)$, if any of the conditions (1)-(3) hold. It remains to prove that for i it is not better to delegate to j instead. It is first immediate that if $p_i = 1$, then i receives the same expected utility by self-delegating as by delegating to j. Also, if $p_j = 0$, then delegating to j provides the same utility to i as self-delegation. Then, if $j \notin \mathcal{A}_i(\boldsymbol{x}, \boldsymbol{\tau})$, i receives strictly lower utility in $(\boldsymbol{d}^*_{-i}, j)$ than in $\boldsymbol{d}^*$. Therefore, in all of the examined cases $\boldsymbol{d}^*$ is a Nash equilibrium. □

Proof of Theorem 7

Suppose that $\boldsymbol{d}$ is a Nash equilibrium and that $G_{\boldsymbol{d}}$ contains a weakly connected component W of more than one vertex that does not have a directed cycle. Then there exists a voter i in W that is self-delegating under $\boldsymbol{d}$ and a voter j such that $d(j) = i$. If $p_j \neq 0$ and $p_i \neq 1$, the proof of Theorem 5 works. So, we first assume that $p_i = 1$. We claim that there is another NE profile $\boldsymbol{d}'$, where $d(j)' = i$ and $d(i)' = j$. The fact that $p_i = 1$ means that voter i as well as every voter delegating directly or indirectly to voter i under $\boldsymbol{d}$ is not affected by whether i self-delegates or not. The same holds trivially for the rest of the voters, so $\boldsymbol{d}' = (\boldsymbol{d}_{-i}, j)$ is a Nash equilibrium as well. The proof is similar for the case where for the voter j it holds that $p_j = 0$. Specifically, we claim again that there is another NE $\boldsymbol{d}'$ such that $d'_j = i$ and $d'_i = j$. The fact that $p_j = 0$ means that the change of delegation for voter i from self-looping to j will not affect the utility of any voter as j will never cast a ballot. Therefore, the expected utility of every voter is the same under both delegation profiles. □

Proof of Theorem 8

Take a weakly connected component W of a graph $G_{\boldsymbol{d}}$ where $\boldsymbol{d}$ is a Nash equilibrium of an instance $\mathcal{I}$ satisfying the conditions of the statement. For a voter $v \in \mathcal{L}(W) \cup \mathcal{R}(W)$ and a voter $w \in \mathcal{C}(W)$ we denote as $\mathcal{Q}^v_w(w')$ the probability that a vertex $w' \in \mathcal{C}(W)$ corresponds to the ultimate delegate of v, (to be called $\mathcal{P}^v_w(w')$) times $(1 - p_v)$ in the case that v delegates to w. We note that $\mathcal{P}^v_w$ is actually the same for every voter $v \in \mathcal{L}(W) \cup \mathcal{R}(W)$ and hence we will drop the superscript.

The following lemma is a direct consequence of the definition of expected utility together with the fact that no voter in $\mathcal{C}(W)$ delegated to a voter outside of $\mathcal{C}(W)$ according to $\boldsymbol{d}$.

Lemma 12. *For a voter $v \in \mathcal{L}(W) \cup \mathcal{R}(W)$ that delegates to a voter $w \in \mathcal{C}(W)$, the expected utility of v is $p_v\tau + \sum_{i \in \mathcal{C}(W)} \mathcal{Q}^v_w(i) \cdot (\tau - \mathit{dist}(v, i)) = p_v\tau + (1 - p_v)\sum_{i \in \mathcal{C}(W)} \mathcal{P}_w(i) \cdot (\tau - \mathit{dist}(v, i))$.*

Suppose, without loss of generality, that there are two distinct voters $v, v' \in \mathcal{L}(W)$ that are delegating to a vertex in $\mathcal{C}(W)$, under $\boldsymbol{d}$. Assume further that $x_v < x_{v'}$. First, notice that since both v and v' are on the left side of all of the voters in $\mathcal{C}(W)$, from Lemma 12, we have that the set of voters in $\mathcal{C}(W)$ maximizing the expected utility if delegated to is the same for v and v' and, since $\boldsymbol{d}$ is a NE, both of them delegate to such a voter, say w.

We will show now that $\boldsymbol{d}$ is not a NE by showing that v would benefit from switching their delegation to v' instead of w. In that case, the utility of v would be:

$$\tau p_v + (1 - p_v)\left(p_{v'}(\tau - \text{dist}(v, v')) + \sum_{i \in \mathcal{C}(W)} \mathcal{P}_w(i)(1 - p_{v'})(\tau - \text{dist}(v, i))\right)$$

Delegation Profile	**Social Welfare**
$[A, A, B]$	0.4532
$[A, A, C]$	0.4402
$[A, B, B]$	0.4562
$[A, B, C]$	0.4400
$[A, C, B]$	0.4564
$[A, C, C]$	0.4402
$[B, A, B]$	0.4694
$[B, A, C]$	0.4564
$[B, B, B]$	**0.4724**
$[B, B, C]$	0.4562
$[B, C, B]$	0.4693
$[B, C, C]$	0.4532

Table 4: Social Welfare achieved by profiles in which A doesn't delegate to C and vice versa, in the instance examined in the proof of Observation 9. The maximizing value of social welfare appears in bold.

Now, recall that the expected utility of v under $\boldsymbol{d}$ amounts to

$$\tau p_v + (1 - p_v) \sum_{i \in \mathcal{C}(W)} \mathcal{P}_w(i)(\tau - \text{dist}(v, i)).$$

It suffices to show that

$$p_{v'}(\tau - \text{dist}(v, v')) + \sum_{i \in \mathcal{C}(W)} \mathcal{P}_w(i)(1 - p_{v'})(\tau - \text{dist}(v, i)) > \sum_{i \in \mathcal{C}(W)} \mathcal{P}_w(i)(\tau - \text{dist}(v, i)) \Leftrightarrow$$

$$(\tau - \text{dist}(v, v')) > \sum_{i \in \mathcal{C}(W)} \mathcal{P}_w(i)(\tau - \text{dist}(v, i)),$$

which is true because

$$\sum_{i \in \mathcal{C}(W)} \mathcal{P}_w(i)(\tau - \text{dist}(v, i)) < \sum_{i \in \mathcal{C}(W)} \mathcal{P}_w(i)(\tau - \text{dist}(v, v')) = (\tau - \text{dist}(v, v')) \sum_{i \in \mathcal{C}(W)} \mathcal{P}_w(i) \leqslant \tau - \text{dist}(v, v'),$$

where the last inequality holds as $\sum_{i \in \mathcal{C}(W)} \mathcal{P}_w(i) \leqslant 1$.

Finally, because W is a connected component containing exactly one cycle (by Theorem 5), it directly follows that $\mathcal{L}(W)$ and $\mathcal{R}(W)$ induce directed (in-)trees, which are directed towards $\mathcal{C}(W)$. □

Proof of Observation 9

Proof. Consider the following symmetric instance of three voters, namely A, B, and C, with $\tau = 0.4$:

$$x = (0.12, 0.5, 0.88),$$
$$p = (0.1, 0.9, 0.1).$$

Notice first that in a profile that maximizes total utility A and C do not delegate to each other, as $\text{dist}(A, C) > 0.4$, and hence changing a delegation of one of them to B would strictly improve the total utility. The social welfare of the rest possible delegation profiles is shown in Table 4. It follows that $\boldsymbol{d} = [B, B, B]$ maximizes total utility; its delegation graph appears below.

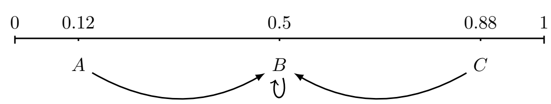

Obviously, $G_{\boldsymbol{d}}$ does not have a cycle. □

Proof of Theorem 10

We will prove the statement by describing an instance in which the worst Nash equilibrium profile in terms of social welfare $\boldsymbol{d}$ satisfies $SW(\boldsymbol{d}) \to 0$ and the social welfare maximizing delegation $\boldsymbol{d}'$ satisfies $SW(\boldsymbol{d}') \to \frac{e-1}{e}\lambda$, where λ corresponds to some fixed value and e is Euler's number.

Fix a value $\lambda > 0$ as well as a value $\varepsilon \in [0, 1]$ which is infinitesimal and satisfies $\lambda > \varepsilon/n$. Consider an instance of $n + 1$ voters, where for $i \in [n]$, voter i is positioned at a distance of ε/n before v_{i+1}. For voter v_1 it holds that $\tau_1 = \lambda$ and $p_1 = \varepsilon$, for v_{n+1} it holds that $\tau_{n+1} = 0$ and $p_{n+1} = 1$, whereas for the rest voters it holds that they vote with probability $1/n$ and their tolerance value is equal to zero. For an illustration see the image below:

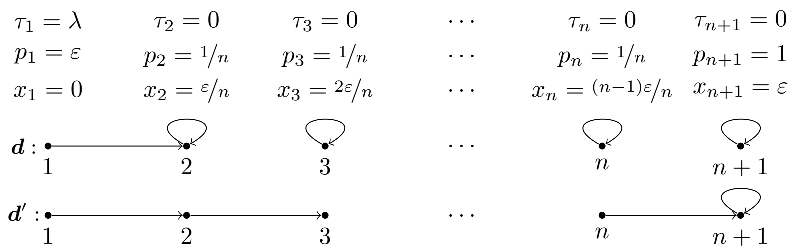

The profile $\boldsymbol{d}$ where v_1 delegates to v_2 and all others delegate to themselves is a Nash equilibrium. This holds as voters of $\lambda_i = 0$ do not want to delegate anywhere other than to themselves and v_1 prefers delegating to v_2 rather than anyone else, if all after v_1 vote. We will now compute the total utility gained in this profile. It holds that $u_i(\boldsymbol{d}) = 0$, for each $i \in \{2, 3, \cdots, n+1\}$, and hence

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

Next, consider $\boldsymbol{d}'$ such that $\boldsymbol{d}'_{v_i} = v_{i+1}$ for $i \in \{1, \cdots, n\}$, and $\boldsymbol{d}'_{v_{n+1}} = v_{n+1}$. Clearly, this is not an equilibrium, since all but v_1 have an incentive to deviate and self-loop in order to increase their utility (to $\lambda_i p_i = 0$, as their current utility is negative). In order to show that PoA is unbounded, it suffices to show that the total utility obtained from this profile is a constant, we will argue that it approaches (approximately) 0.632λ.

We start by computing the utility for v_1. This will be

$$\lambda\varepsilon + (1-\varepsilon)(\lambda - \frac{\varepsilon}{n})\frac{1}{n} + (1-\varepsilon)(\lambda - \frac{2\varepsilon}{n})(1-\frac{1}{n})\frac{1}{n} + (1-\varepsilon)(\lambda - \frac{3\varepsilon}{n})(1-\frac{1}{n})^2\frac{1}{n} + \cdots + (1-\varepsilon)(\lambda - \frac{(n)\varepsilon}{n})(1-\frac{1}{n})^{n-1}\frac{1}{n}$$

Assume that $\varepsilon \to 0$. Then this utility becomes

$$\lambda\frac{1}{n}\sum_{i=0}^{n-1}(1-\frac{1}{n})^i.$$

It now suffices to compute the limit when n goes to infinity, for the expression above. Using the geometric series $\sum_{i=0}^{t} r^i = \frac{1-r^{t+1}}{1-r}$ and the fact that $\lim_{n\to\infty}(1-\frac{1}{n})^n = \frac{1}{e}$ we have the following:

$$\lim_{n\to\infty}\lambda\frac{1}{n}\sum_{i=0}^{n-1}(1-\frac{1}{n})^i = \lim_{n\to\infty}\lambda\frac{1}{n}\frac{1-(1-1/n)^n}{1-(1-1/n)} =$$

$$\lim_{n\to\infty}\lambda\frac{1}{n}\frac{1-(1-1/n)^n}{1/n} = \lim_{n\to\infty}\lambda(1-(1-1/n)^n) =$$

$$\lambda - \lambda\lim_{n\to\infty}(1-1/n)^n = \lambda - \frac{\lambda}{e} = \lambda(\frac{e-1}{e})$$

Similarly, the utility of v_2 can be expressed as (being independent of λ):

$$0\cdot\frac{1}{n} + (1-\frac{1}{n})(0-\frac{\varepsilon}{n})\frac{1}{n} + (1-\frac{1}{n})^2(0-\frac{2\varepsilon}{n})\frac{1}{n} + \cdots + (1-\frac{1}{n})^{n-1}(0-\frac{(n-1)\varepsilon}{n})\frac{1}{n} \xrightarrow[\varepsilon\to 0]{n\to\infty} 0.$$

The utility for the rest voters can be computed along the same lines to v_2. As a result, the total utility in the best delegation profile is at least (only a bit below) $\lambda\frac{e-1}{e}$ and this concludes the proof. □

Proof of Theorem 11

We begin by noting that the maximum utility a voter i can get is τ_i. This is because, by Theorem 14 (see Appendix B.3), we know how each voter's best delegation path looks like and since distances from a voter are decreasing further along the path, the best case for i is that their vote is being cast with probability 1 at a distance of 0.

Now we will argue that the minimum utility a voter i can get in an equilibrium profile is $\tau_i p_i$. Consider a profile $\boldsymbol{d}$ that is a Nash equilibrium and towards a contradiction, say that for a voter i it holds $u_i(\boldsymbol{d}) < \tau_i p_i$. But then $d(i) \neq i$. We now compute the utility that voter i will experience under the profile $(\boldsymbol{d}_{-i}, i)$. This is simply equal to $\tau_i p_i$, contradicting the fact that the considered profile $\boldsymbol{d}$ is a NE.

0.8008	0.8107	0.8191	0.8279	**0.8287**	0.8239	0.8218	0.8194
0.7377	0.1320	0.2448	0.4958	0.5023	0.4579	0.4377	0.4163
0.7046	**0.7246**	0.1408	0.5382	0.5446	0.5007	0.4807	0.4601
0.5936	0.6114	0.6330	0.2904	0.3398	0.6062	**0.6530**	0.6417
0.4353	0.4597	0.4891	**0.6786**	0.0792	0.5250	0.6063	0.5933
0.6946	0.7001	0.7075	0.7567	**0.7630**	0.6072	0.7310	0.7396
0.5363	0.5473	0.5616	0.5220	0.5447	**0.7651**	0.4136	0.4556
0.1525	0.1653	0.1839	0.5445	0.5378	0.5838	**0.6047**	0.1056

Table 5: Representation of utilities for potential deviations for the instance proving the first two statements from Appendix B.1. Specifically, cell (i, j) corresponds to the utility of voter i after delegating to voter j, assuming that the rest of the delegations are given by $\boldsymbol{d}_{-i}$. The expected utility for each voter under $\boldsymbol{d}$ (maximal per row) appears in bold.

0.0420	**0.7240**	0.7226	0.4468	0.3501	0.3438	0.3366
0.6481	0.6384	0.6434	**0.7478**	0.7234	0.7218	0.7200
0.6797	**0.6799**	0.0252	0.6030	0.5032	0.4968	0.4894
0.7242	0.7231	0.7250	0.6300	**0.8052**	0.8036	0.8017
0.6366	0.6383	0.6397	0.6990	0.4452	0.8036	**0.8199**
0.8113	0.8117	0.8118	0.8305	**0.8316**	0.7812	0.8236
0.7221	0.7235	0.7240	0.7819	0.7291	**0.8297**	0.6132

Table 6: Representation of utilities for potential deviations for the instance proving the third statement from Appendix B.1. Specifically, cell (i, j) corresponds to the utility of voter i after delegating to voter j, assuming that the rest of the delegations are given by $\boldsymbol{d}_{-i}$. The expected utility for each voter under $\boldsymbol{d}$ (maximal per row) appears in bold.

As a result, the ratio between the maximum achievable utility and the minimum utility in an equilibrium is at most

$$\frac{\sum_{i\in V}\tau_i}{\sum_{i\in V}\tau_i p_i} \leqslant \frac{\sum_{i\in V}\tau_i}{\sum_{i\in V}\tau_i\cdot\min_{i\in V}\{p_i\}} = \frac{1}{\min_{i\in V}\{p_i\}}.$$

We now focus on measuring the upper bound of PoA^+. Following the same arguments than before, we have that the difference between the welfare in the socially optimal solution and the one in the worst in terms of utility Nash equilibrium is at most

$$\sum_{i\in V}\tau_i - \sum_{i\in V}\tau_i p_i \leqslant \sum_{i\in V}\tau_i - \sum_{i\in V}\tau_i \min_{i\in V}\{p_i\} = (1-\min_{i\in V}\{p_i\})\sum_{i\in V}\tau_i,$$

which proves the statement. □

B Additional Concepts and Results

B.1 Further Insights On the Structure of Nash Equilibria

Consider a mutual acceptance instance $\mathcal{I}$ with no deterministic voters, a profile $\boldsymbol{d}$ that is a Nash equilibrium of $\mathcal{I}$ and a weakly connected component W of $G_{\boldsymbol{d}}$ that consists of more than a single vertex. Then, the following hold:

- If there is a vertex in $\mathcal{C}(W)$ having in-deg $= 2$ this is not necessarily the left-most or right-most of $\mathcal{C}(W)$.
- Paths appearing in $G_{\boldsymbol{d}}$ are not necessarily following an order of increasing or decreasing position.
- Vertices in $\mathcal{L}(W)$ or $\mathcal{R}(W)$ do not necessarily form a path in $G_{\boldsymbol{d}}$.

The following instance $\mathcal{I}$, where $V = \{1, 2, \cdots, 8\}$ proves the first two statements, specifically that the entry points of a cycle are not necessarily the left-most and right-most vertices and that a path ending to a cycle need not to be between consecutive vertices:

$$x = (0.1, 0.15, 0.2, 0.4, 0.46, 0.63, 0.66, 0.88),$$
$$p = (0.91, 0.15, 0.16, 0.33, 0.09, 0.69, 0.47, 0.12),$$
$$\tau = 0.88.$$

The delegation profile $\boldsymbol{d} = (5, 1, 2, 7, 4, 5, 6, 7)$ is a Nash equilibrium for $\mathcal{I}$ as Table 5 shows. More precisely, in Table 5 we can see that no voter has a profitable deviation from $\boldsymbol{d}$. The graph $G_{\boldsymbol{d}}$ is depicted in Figure 3 and its structure proves the corresponding statements.

The third statement, namely that the vertices in $\mathcal{L}(W)$ or $\mathcal{R}(W)$ do not necessarily form a path in $G_{\boldsymbol{d}}$, where $\boldsymbol{d}$ is a NE, holds due to the following instance, where $V = \{1, 2, \cdots, 7\}$:

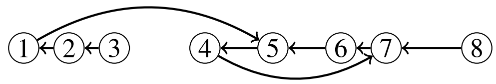

Figure 3: Delegation graph of the instance proving the first two statements from Appendix B.1.

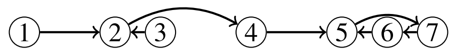

Figure 4: Delegation graph of the instance proving the first two statements from Appendix B.1.

$$x = (0.37, 0.4, 0.54, 0.75, 0.87, 0.89, 0.9),$$
$$p = (0.05, 0.76, 0.03, 0.75, 0.53, 0.93, 0.73),$$
$$\tau = 0.84.$$

Consider the profile $\boldsymbol{d} = (2, 4, 2, 5, 7, 5, 6)$, the delegation graph $G_{\boldsymbol{d}}$ of which is depicted in Figure 4.
Table 6 proves that no voter has a profitable deviation from $\boldsymbol{d}$. Notice that vertices in the left of the cycle form a tree.

B.2 Further Insights on the Structure of Optimal Delegation Profiles

Building on the negative result of Observation 9, we now identify the condition under which a weakly connected component of $G_{\boldsymbol{d}_{SW}}$ contains a cycle. The next natural question is whether a single cycle passing through all vertices is the optimal solution in such cases, as it is for the delegation profile minimizing vote loss (see Appendices B.4 and C.5 for details on this metric). We address this to the negative.

Theorem 13. *In a mutual acceptance instance $\mathcal{I}$ without deterministic voters, if $i \in \mathcal{A}_j(\boldsymbol{x}, \boldsymbol{\tau})$ for every pair i, j of voters of a weakly connected component of a socially optimal delegation profile, then, the component has a cycle. However, it is not the case that a cycle including every vertex corresponds to the delegation profile maximizing social welfare.*

Proof. Towards a contradiction, assume that for an instance $\mathcal{I}$, there is a weakly connected component of at least 2 vertices that doesn't have a directed cycle, therefore, there is a voter i such that $d(i) = i$ in the utility maximizing solution $\boldsymbol{d}$. Since i belongs to the component, there is another voter, say j, such that $d(j) = i$. Consider also the profile $\boldsymbol{d}' = (\boldsymbol{d}_{-i}, j)$. The delegation graph $G_{\boldsymbol{d}'}$ contains a cycle in the component of i and j. We will prove that, under $\boldsymbol{d}'$, the total utility of the voters of the component will be strictly greater than under $\boldsymbol{d}$, while the utility of the rest of the voters remains the same.

By the fact that i and j both belong to the component, and hence, $j \in \mathcal{A}_i(\boldsymbol{x}, \boldsymbol{\tau})$ it holds that $u_i(\boldsymbol{d}') = p_i\tau_i + (1 - p_i)p_j(\tau_i - \text{dist}(i, j)) > p_i\tau_i = u_i(\boldsymbol{d})$. Consider now any other arbitrary voter of the examined weakly connected component of $G_{\boldsymbol{d}}$, say $k \neq i$. It holds that there is a path from k to i in $G_{\boldsymbol{d}}$, and say that $\mathcal{M}$ corresponds to the set of vertices in that path (k and i included). The change of $d(i)$ from i to j will increase the utility of k by $\prod_{t \in \mathcal{M}}(1 - p_t)p_j(\tau_k - \text{dist}(k, j))$, because with a probability equal to the probability that none in $\mathcal{M}$ will vote and j will vote, voter k will be represented by j. By assumptions, this quantity is strictly positive. Therefore, every voter in the component will have a strictly greater utility under $\boldsymbol{d}'$ compared to their utility under $\boldsymbol{d}$.

Finally, it is not hard to observe that in the following symmetric instance, the graph $G_{\boldsymbol{d}}$, where $\boldsymbol{d}$ is the social welfare maximizing delegation profile, consists of two weakly connected components: One including the first two and another including the last two vertices.

$$x = [0.05, 0.06, 0.94, 0.95],$$
$$p = [0.5, 0.5, 0.5, 0.5],$$
$$\tau = 0.1.$$

This is because, due to each voter's tolerance value, a delegation from one of the first two voters to one of the last two, or vice versa, will not appear in the social welfare maximizing profile. □

B.3 Individually Optimal Delegation Profiles

For an individual voter $i \in V$, we define i's optimal delegation profile as the profile that maximizes i's expected utility. This profile will be denoted by $\boldsymbol{d}^{i*}$. It represents the choice i would make, among all possible profiles, if they could fix all voters' delegations to solely maximize their own expected utility. This concept serves as a key metric for evaluating the quality of a delegation profile in terms of social welfare (e.g., whether or not it constitutes a Nash equilibrium). Interestingly, the cyclic structure plays a crucial role here as well. Specifically, we prove that for any voter $i \in V$, there is no profile that results in strictly higher utility for i than one where the corresponding delegation graph includes a cycle consisting of all voters in

$\mathcal{A}_i(\boldsymbol{x}, \tau)$, ordered by increasing distance to i, and closing back to i. The utility i obtains from such a profile is equal to the utility they achieve when the edge returning to i is removed.

The main result in this regard is stated below. The individually optimal delegation profile for a voter is referred to as i's optimal delegation path. In principle, specifying this path is sufficient, as delegations involving voters outside this path do not influence i's utility. Specifically, we show that the most beneficial strategy for a voter is to sequentially delegate their vote to the closest voter, continuing step-by-step to the farthest voter within the τ distance.

Theorem 14. *For any $i \in V$ and instance $\mathcal{I}$, we have $\mathbb{E}(i, \boldsymbol{d}) \leqslant \mathbb{E}(i, \boldsymbol{d}^{i*})$, where*

- *$\boldsymbol{d}$ is some delegation profile of $\mathcal{I}$;*
- *$\boldsymbol{d}^{i*}$ is the profile for which $G_{\boldsymbol{d}^{i*}}$ includes a path starting in i that traverses all voters within $\mathcal{A}_i(\boldsymbol{x}, \boldsymbol{\tau})$ distance of i in the order of their distance to i.*

That is, $\boldsymbol{d}^{i*}$ is an optimal delegation profile for voter $i \in V$.

Before proving the result, to illustrate its statement, we revisit our running example.

Example 1 Continued. *By Theorem 14, it holds that no better delegation profile exists for voter D, than those including either the delegation cycle $\{C, E, B, D\}$ or $\{E, C, B, D\}$. In both cases, voter D has an optimal expected utility of* 0.11196.

Now we move to proving Theorem 14. We highlight that by providing a constructive proof for an optimal solution with respect to a single voter, our result contrasts with strategic choice selection in other areas of social choice theory, where the analogous problem is intractable.

Proof of Theorem 14. The proof is a direct consequence of a series of the three lemmas that follow. The first lemma shows that any delegation path of a voter v along which voters are sorted by ascending distance to v achieves higher expected utility than the same delegation path where the order of voters is permuted. For a voter v we denote by $\mathbb{E}(v, \pi)$ the (expected) utility that v gains from a delegation profile the graph of which has π as the maximal path from v.

Lemma 15. *Let $\pi = (v_1, \cdots, v_k)$ be a delegation path for a voter v_1. Assume there exists an index i such that v_{i+1} is closer to v_1 than i, i.e., dist$(v_1, v_{i+1}) <$ dist(v_1, v_i). Then the expected utility v_1 gains from π is lower than the expected utility from path $\pi' = (v_1, \cdots, v_{i-1}, v_{i+1}, i, v_{i+2}, \cdots, v_k)$ where the delegation path first visits v_{i+1} and then i.*

Proof. We compare the expected utility of v_1 w.r.t. π

$$\mathbb{E}(v_1, \pi) = \sum_{\ell=1}^{k} (\tau_1 - \text{dist}(v_1, v_\ell)) p_{v_\ell} \prod_{m=1}^{\ell-1} (1 - p_{v_m})$$

to the expected utility of v_1 w.r.t. π'

$$\begin{aligned} \mathbb{E}(v_1, \pi') = \quad & [\sum_{\ell=1}^{i-1} (\tau_1 - \text{dist}(v_1, v_\ell)) p_{v_\ell} \prod_{m=1}^{\ell-1} (1 - p_{v_m})] + (\tau_1 - \text{dist}(v_1, v_{i+1})) p_{v_{i+1}} \prod_{m=1}^{i-1} (1 - p_{v_m}) \\ & + (\tau_1 - \text{dist}(v_1, v_i)) p_{v_i} (1 - p_{v_{i+1}}) \prod_{m=1}^{i-1} (1 - p_{v_m}) + \sum_{\ell=i+2}^{k} (\tau_1 - \text{dist}(v_1, v_\ell)) p_{v_\ell} \prod_{m=1}^{\ell-1} (1 - p_{v_m}). \end{aligned}$$

These expected utilities only differ in terms that include the utilities of i and v_{i+1}. Further, these terms have common factors. Thus, we have:

$$\begin{gathered} \mathbb{E}(v_1, \pi) < \mathbb{E}(v_1, \pi') \Leftrightarrow \\ (\tau_1 - \text{dist}(v_1, v_i)) p_{v_i} + (\tau_1 - \text{dist}(v_1, v_{i+1})) p_{v_{i+1}} (1 - p_{v_i}) < (\tau_1 - \text{dist}(v_1, v_{i+1})) p_{v_{i+1}} + (\tau_1 - \text{dist}(v_1, v_i)) p_{v_i} (1 - p_{v_{i+1}}) \Leftrightarrow \\ \tau_1 p_{v_i} - \text{dist}(v_1, v_i) p_{v_i} + \tau_1 p_{v_{i+1}} - \tau_1 p_{v_{i+1}} p_{v_i} - \text{dist}(v_1, v_{i+1}) p_{v_{i+1}} + \text{dist}(v_1, v_{i+1}) p_{v_{i+1}} p_{v_i} < \\ \tau_1 p_{v_{i+1}} - \text{dist}(v_1, v_{i+1}) p_{v_{i+1}} + \tau_1 p_{v_i} - \tau_1 p_{v_i} p_{v_{i+1}} - \text{dist}(v_1, v_i) p_{v_i} + \text{dist}(v_1, v_i) p_{v_i} p_{v_{i+1}} \Leftrightarrow \\ \text{dist}(v_1, v_{i+1}) p_{v_{i+1}} p_{v_i} < \text{dist}(v_1, v_i) p_{v_i} p_{v_{i+1}} \Leftrightarrow \\ \text{dist}(v_1, v_{i+1}) < \text{dist}(v_1, v_i) \end{gathered}$$

□

Note that Lemma 15 holds even for paths that may include voters outside their τ-threshold, i.e., contribute negative utility. However, as we see next, eliminating such voters from (the end of) a delegation path increases the expected utility.

Lemma 16. *Let $\pi = (v_1, \cdots, v_k)$ be a delegation path for a voter v_1 and let $\pi' = (v_1, \cdots, v_{k-1})$ be the shortened path by deleting the last voter v_k. Then the expected utility of v_1 w.r.t. delegation path π is lower than that for path π', i.e., $\mathbb{E}(v_1, \pi) < \mathbb{E}(v_1, \pi')$, if and only if $\tau_1 -$ dist$(v_1, v_k) < 0$.*

Proof. We compare the expected utility of v_1 w.r.t. π and π':

$$\mathbb{E}(v_1, \pi) < \mathbb{E}(v_1, \pi') \Leftrightarrow$$
$$\sum_{\ell=1}^{k}(\tau_1 - \text{dist}(v_1, v_\ell))p_{v_\ell} \prod_{m=1}^{\ell-1}(1-p_{v_m}) < \sum_{\ell=1}^{k-1}(\tau_1 - \text{dist}(v_1, v_\ell))p_{v_\ell} \prod_{m=1}^{\ell-1}(1-p_{v_m}) \Leftrightarrow$$
$$(\tau_1 - \text{dist}(v_1, v_k))p_{v_k} \prod_{m=1}^{k-1}(1-p_{v_m}) < 0 \Leftrightarrow \tau_1 - \text{dist}(v_1, v_k) < 0,$$

hence the result follows. □

Thus, to increase the expected utility, we can first sort voters along a delegation path in order of their distance and then repeatedly shorten the path to only contain voters with a distance smaller or equal τ. We now show that inserting a voter into a path increases the expected utility if the subsequent voters have a larger distance.

Lemma 17. *Let $\pi = (v_1, \cdots, v_k)$ be a delegation path for a voter v_1. Assume there exists an index i such that dist$(v_1, v_i) <$ dist(v_1, v_ℓ) for all $\ell \in [i+1, k]$. Let $\pi' = (v_1, \cdots, v_{i-1}, v_{i+1}, \cdots, v_k)$ where i is deleted. Then the expected utility of v_1 w.r.t. delegation path π is higher than that for path π', i.e., $\mathbb{E}(v_1, \pi) > \mathbb{E}(v_1, \pi')$.*

Proof. We compare the expected utility of v_1 w.r.t. π

$$\mathbb{E}(v_1, \pi) = \sum_{\ell=1}^{k}(\tau_1 - \text{dist}(v_1, v_\ell))p_{v_\ell} \prod_{m=1}^{\ell-1}(1-p_{v_m})$$

to the expected utility of v_1 w.r.t. π'

$$\mathbb{E}(v_1, \pi') = \sum_{\ell=1}^{i-1}(\tau_1 - \text{dist}(v_1, v_\ell))p_{v_\ell} \prod_{m=1}^{\ell-1}(1-p_{v_m}) + \sum_{\ell=i+1}^{k}(\tau_1 - \text{dist}(v_1, v_\ell))p_{v_\ell} \prod_{m=1}^{i-1}(1-p_{v_m}) \prod_{m=i+1}^{\ell-1}(1-p_{v_m})$$

These expected utilities are only differing in the term that includes utilities of i and some terms that have an additional factor $(1-p_i)$. Thus, we have:

$$\mathbb{E}(v_1, \pi) > \mathbb{E}(v_1, \pi') \Leftrightarrow$$
$$\sum_{\ell=i}^{k}(\tau_1 - \text{dist}(v_1, v_\ell))p_{v_\ell} \prod_{m=1}^{\ell-1}(1-p_{v_m}) > \sum_{\ell=i+1}^{k}(\tau_1 - \text{dist}(v_1, v_\ell))p_{v_\ell} \prod_{m=1}^{i-1}(1-p_{v_m}) \prod_{m=i+1}^{\ell-1}(1-p_{v_m}) \Leftrightarrow$$
$$\sum_{\ell=i}^{k}(\tau_1 - \text{dist}(v_1, v_\ell))p_{v_\ell} \prod_{m=i}^{\ell-1}(1-p_{v_m}) > \sum_{\ell=i+1}^{k}(\tau_1 - \text{dist}(v_1, v_\ell))p_{v_\ell} \prod_{m=i+1}^{\ell-1}(1-p_{v_m}) \Leftrightarrow$$
$$(\tau_1 - \text{dist}(v_1, i))p_i + (1-p_i) > \sum_{\ell=i+1}^{k}(\tau_1 - \text{dist}(v_1, v_\ell))p_{v_\ell} \prod_{m=i+1}^{\ell-1}(1-p_{v_m}) \Leftrightarrow$$
$$(\tau_1 - \text{dist}(v_1, i))p_i > (1-(1-p_i)) \sum_{\ell=i+1}^{k}(\tau_1 - \text{dist}(v_1, v_\ell))p_{v_\ell} \prod_{m=i+1}^{\ell-1}(1-p_{v_m}) \Leftrightarrow$$
$$\tau_1 - \text{dist}(v_1, i) > \sum_{\ell=i+1}^{k}(\tau_1 - \text{dist}(v_1, v_\ell))p_{v_\ell} \prod_{m=i+1}^{\ell-1}(1-p_{v_m})$$

To see that the last inequality holds we first show the following:

$$
\begin{aligned}
&\sum_{\ell=i+1}^{k} p_{v_\ell} \prod_{m=i+1}^{\ell-1}(1-p_{v_m}) = p_{v_{i+1}} + (1-p_{v_{i+1}}) \sum_{\ell=i+2}^{k} p_{v_\ell} \prod_{m=i+2}^{\ell-1}(1-p_{v_m}) = \\
&p_{v_{i+1}} + (1-p_{v_{i+1}})(p_{v_{i+2}} + (1-p_{v_{i+2}}) \sum_{\ell=i+3}^{k} p_{v_\ell} \prod_{m=i+3}^{\ell-1}(1-p_{v_m})) = \\
&\dots \\
&p_{v_{i+1}} + (1-p_{v_{i+1}})(p_{v_{i+2}} + (1-p_{v_{i+2}})(\cdots(p_{v_{k-1}} + (1-p_{v_{k-1}})p_{v_k}))) \leqslant \\
&p_{v_{i+1}} + (1-p_{v_{i+1}})(p_{v_{i+2}} + (1-p_{v_{i+2}})(\cdots(p_{v_{k-1}} + (1-p_{v_{k-1}}) \cdot 1)))
\end{aligned}
$$

The previous bound was obtained by bounded the last factor: $p_{v_k} \leqslant 1$. This leaves this multiplicative term to be: $p_{v_{k-1}} + (1 - p_{v_{k-1}}) = 1$. Observe that this step will now repeat throughout the terms for decreased subscripts. We see the final steps of this recursive process in the following lines.

$$
\dots = p_{v_{i+1}} + (1-p_{v_{i+1}})(p_{v_{i+2}} + (1-p_{v_{i+2}}) \cdot 1) = p_{v_{i+1}} + (1-p_{v_{i+1}}) \cdot 1 = 1
$$

Together with our assumption that $\text{dist}(v_1, v_i) < \text{dist}(v_1, v_\ell)$ for all $i < \ell$, this yields the desired inequality:

$$
\begin{aligned}
&\sum_{\ell=i+1}^{k} (\tau_1 - \text{dist}(v_1, v_\ell)) p_{v_\ell} \prod_{m=i+1}^{\ell-1}(1-p_{v_m}) < \sum_{\ell=i+1}^{k} (\tau_1 - \text{dist}(v_1, i)) p_{v_\ell} \prod_{m=i+1}^{\ell-1}(1-p_{v_m}) = \\
&(\tau_1 - \text{dist}(v_1, i)) \sum_{\ell=i+1}^{k} p_{v_\ell} \prod_{m=i+1}^{\ell-1}(1-p_{v_m}) \leqslant \tau_1 - \text{dist}(v_1, i).
\end{aligned}
$$

□

Combining the three lemmas above one can show that for an individual voter v, the maximal expected utility can be achieved by a delegation path that traverses all voters within τ_v-distance and in order by their distance to v. □

B.3.1 Providing an Upper Bound for the Optimal Social Welfare

From Theorem 14 we can directly get the following result.

Corollary 18. *For any instance $\mathcal{I} = \langle \boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\tau} \rangle$, if $\boldsymbol{d}_{SW} = \text{argmax}_{\boldsymbol{d} \in V \times \cdots \times V} SW(\boldsymbol{d})$ is an optimal profile and $\boldsymbol{d}^{i*}$ is the delegation profile with voter i's highest expected utility, then $SW(\boldsymbol{d}_{SW}) \leqslant \sum_{i \in V} u_i(\boldsymbol{d}^{i*})$.*

The result of Corollary 18 provides an upper bound on the highest possible social welfare for any delegation profile in an instance. Let $\texttt{ODP}(\mathcal{I})$ denote the sum of expected utilities across the *optimal delegation profiles* of all voters in instance $\mathcal{I}$, i.e., $\texttt{ODP}(\mathcal{I}) = \sum_{i \in V} u_i(\mathbf{d}^{\mathbf{i}*})$. This approximation is employed in experiments involving socially optimal profiles, where finding the exact solution is computationally expensive (see Section 5). One natural question is how far away is $\texttt{ODP}(\mathcal{I})$ from $SW(\boldsymbol{d}_{SW})$. We conducted some simulations on small random instances where the delegation profile with the highest social welfare can be effectively found by brute force. Our preliminary experiments showed that $\texttt{ODP}(\mathcal{I})$ is just a few percentiles larger than its corresponding $SW(\boldsymbol{d}_{SW})$. Details can be found in Appendix C.4.

B.4 Expected Number of Votes Cast

The concept of lost votes is crucial in liquid democracy research, as it has been a key motivation behind the scheme since its inception. In our work, among others, we are evaluating delegation profiles, both theoretically and experimentally, based on their potential to mitigate the loss of voting power over all elections held in the system. By *votes lost*, we refer to the expected number of votes not cast due to voters abstaining from voting in a given election, which in turn results in the loss of the voting power of those who delegated to them as well. For simplicity, consider a profile $\boldsymbol{d}$ where each weakly connected component of $G_{\boldsymbol{d}}$ contains a cycle. Fix such a delegation cycle $\mathcal{C} = (v_1, \cdots, v_k)$. Since the out-degree of every vertex in $G_{\boldsymbol{d}}$ is 1, $\mathcal{C}$ has no outgoing edges. Consequently, no votes from voters corresponding to vertices in the component of $\mathcal{C}$ are lost as long as at least one voter in $\mathcal{C}$ casts a ballot. Equivalently, votes within the connected component can only be lost if all voters in $\mathcal{C}$ abstain. The expected number of votes lost thus depends on the probability that the delegation cycle remains unbroken and the expected number of incoming votes. A similar analysis applies to components that terminate in self-loops. It holds that

$$
\mathbb{E}[\#\text{ votes lost}] = \sum_{v_i \in C} \Pi_{i=1,\cdots,k}(1-p_i) \cdot \left(k + \sum_{i=1,\cdots,k} \mathbb{E}[\#\text{ votes (from outside C) delegated to } v_i] \right).
$$

Note that here, the term $\mathbb{E}[\#\text{ votes delegated to } i]$ is a computation over a tree in $G_{\boldsymbol{d}}$. That is, we can compute the expected number of votes delegated to some voter recursively via the expected number of votes delegated to their direct predecessors.

Algorithm 1 BR Protocol

Input: $\boldsymbol{p}$, $\boldsymbol{x}$, $\boldsymbol{\tau}$ and n
Initialize a random delegation vector $\boldsymbol{d}$
$it = 0$
while $it \leqslant n$ **do**
  $\boldsymbol{d}^1 = \boldsymbol{d}$
  for each $i \in V$ **do**
    $it = it + 1$
    for each $j \in V$ **do**
      $\boldsymbol{d}' = (\boldsymbol{d}^1_{-i}, j)$
      $exp(j) = u_i(\boldsymbol{d}')$
    if $\max_{j \in V}(exp(j)) > u_i(\boldsymbol{d}^1)$ **then**
      $\boldsymbol{d}(i) = \arg\max_{j \in V}(exp(j))$
      $it = 0$
    else if $it \geqslant n$ **then**
      return $\boldsymbol{d}$

$$\mathbb{E}[\#\text{ votes delegated to } v] = \sum_{w \in V : d(w) = v} (1 - p_w) \cdot (1 + \mathbb{E}[\#\text{ votes delegated to } w])$$

This is computationally feasible via a recursion starting from voters with no in-delegations. Observe that, finally, the expected number of votes cast is $n - \mathbb{E}[\#\text{ votes lost}]$. Given the previous definitions, the following structural result regarding the minimization of lost voting power is straightforward.

Proposition 19. *The delegation graph of the profile that minimizes the number of lost votes consists of a single cycle passing through every vertex in arbitrary order.*

C Experimental Analyses of Our Model

In this section of the Appendix, we detail the various experimental analyses that were conducted to support the theoretical results of our model. Note that there are files in the supplementary material corresponding to each of the following sections

C.1 Best-Response Dynamics

With a definition of Nash equilibria in our model, we now need a procedure to find them when they exist. We naturally consider a best-response dynamic, which works by repeatedly checking if any voter has a better delegation choice based on the current profile of delegations. A detailed pseudocode of the best-response dynamic (BR) is being presented as Algorithm 1. The process starts with an arbitrary profile of delegations, and then updates the delegations of voters one by one. Each voter updates their delegation to their best response when one exists, with ties broken arbitrarily. We will refer to each instance where the protocol checks if a voter has a best response as a *round*. The algorithm will stop only after n rounds without any voter finding a best response. This ensures that the resulting delegations form a Nash equilibrium, as no voter can improve their expected utility by unilaterally changing their delegation.

Experimental Analysis of Symmetric Instances. We ran the BR protocol on 20,000 instances with symmetric $\boldsymbol{\tau}$ such that the parameters were chosen uniformly at random as follows:

- n chosen randomly from $\{1, \cdots, 100\}$
- $\tau \in [0, \frac{2}{3}]$ rounded to two decimals.
- $\boldsymbol{x} \in [0, 1]^n$ rounded to two decimals.
- $\boldsymbol{p} \in [0, 1]^n$ rounded to two decimals.

We then ran the BR protocol starting from a random initial profile of delegations. For each of these instances, the BR protocol found a NE.

Example 3. *Consider six voters* $V = \{A, B, C, D, E, F\}$ *whose opinions can be placed on a line such that* $\boldsymbol{x} = (0.2, 0.25, 0.4, 0.4, 0.6, 0.8)$*, probabilities* $\boldsymbol{p} = (0.5, 0.5, 0.9, 0.3, 0.5, 0.3)$*, and tolerance* $\boldsymbol{\tau} = (0.25, 0.25, 0.25, 0.25, 0.25, 0.25)$.

We will perform the best response protocol on our running example, first given in Example 1. We assume the randomly initiated delegations are $\boldsymbol{d} = (A, D, E, B, F, E)$*. Following the BR protocol, we start with voter* A *and check if there is a BR.* A*'s BR is delegating to voter* B*. We then check for voter* B*, who has a BR to delegate to* A*. We then inspect* C *and see that*

No. voters	20	50	100	200
No. SCC	3.7	7.00	12.93	20.23
No. cycles	2.66	4.40	6.97	10.70
No. self loops	1.03	2.6	5.57	9.53
No. WCC of size 1	0.63	2.5	5.57	9.33
No. paths into SCC	3.4	6.6	10.87	19
Av. width of cycle	0.06	0.05	0.03	0.02
Av. width of WCC	0.23	0.18	0.11	0.08
Av. size of cycle	3.78	4.60	5.17	5.69

Table 7: Reporting various average measurements of the structure of its Nash equilibria taken over the 30 instances of each size (i.e., $V \in \{20, 50, 100, 200\}$). All values are rounded to 2 decimal places. SCC denotes strongly connected components (i.e., cycles and self-loops), and WCC denotes weakly connected components (i.e., an SCC and the paths entering the SCC).

their BR is to delegate to D. This continues until we arrive at a NE $\mathbf{d}'' = (B, A, D, C, C, F)$ after 6 iterations. In these first 6 steps, the counter it is set to 0 each time a voter updates their delegations with a BR. The protocol will continue for another 6 iterations with no more updates (as we have reached a NE). The protocol terminates when $it = 6$. Moreover, we note that for our example, there are only two NE, $\mathbf{d}'$ and $\mathbf{d}''$.

C.2 Creating General Instances of our Model

Varying the number of voters. Our first set of instances varies the number of voters $n \in \{20, 50, 100, 200\}$. For each value of n, we created 30 instances, randomly selecting $\boldsymbol{x} \in [0,1]^n$ to three decimal places, $\boldsymbol{p} \in [0,1]^n$ rounded to two decimal places, and $\boldsymbol{\tau} \in [0,1]^n$ rounded to one decimal place. Then for each instance $\mathcal{I} = (\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\tau})$, we find a delegation profile $\boldsymbol{d}_{BR}$ that are Nash equilibria with the best response protocol (given as Algorithm 1, in Appendix C.1).

Varying the size of $\boldsymbol{\tau}$. The second set of instances assesses the impact of the size of tolerance vectors in the various measures we look at in the remainder of the experiments. We take the 30 previous instances when $n = 50$ and then create 20 tolerance vectors $\boldsymbol{\tau} \in [0,1]^n$ and we will denote these vectors as $\boldsymbol{\tau}_{\leqslant 1}$ (no restriction on the number of decimal places). We then modify each of the 20 vectors $\boldsymbol{\tau}_{\leqslant 1}$ by scaling each $\boldsymbol{\tau}_{\leqslant 1}(i)$ by 0.75 and 0.5, i.e., $\boldsymbol{\tau}_{\leqslant 0.75}(i) = 0.75 \times \boldsymbol{\tau}_{\leqslant 1}(i)$ and $\boldsymbol{\tau}_{\leqslant 0.5}(i) = 0.5 \times \boldsymbol{\tau}_{\leqslant 1}(i)$. Resulting in 20 vectors $\boldsymbol{\tau}_{\leqslant 0.75}$ and 20 vectors $\boldsymbol{\tau}_{\leqslant 0.5}$. Therefore, 600 instances of $\mathcal{I} = (\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\tau}_{\leqslant 1})$, $\mathcal{I} = (\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\tau}_{\leqslant 0.75})$, and $\mathcal{I} = (\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\tau}_{\leqslant 0.5})$. For each of these instances, we again use the best response protocol to find a delegation profile that is a NE. These instances are given in the file `Instances_varying_tau.xlxs`.

Delegation profiles for different voting models. When analyzing the expected number of votes lost (Appendix C.5) and the proportion of SW achieved (Section 5), we wanted to compare an arbitrary NE $\boldsymbol{d}_{BR}$ with delegation profiles that reflect different voting models. The first is acyclic liquid democracy `acyc` to replicate the models of liquid democracy in which cycles are not permitted. We create a delegation profile $\boldsymbol{d}_{\texttt{acyc}}$ for each of the instances mentioned previously in the section that modifies the corresponding $\boldsymbol{d}_{BR}$ by breaking each cycle at a random point and replacing the delegation within the cycle with a self-loop. The second delegation profile models the direct democracy voting model without delegations. We model this with the delegation profile $\boldsymbol{d}_{\texttt{dir}}$ where every voter delegates to themselves.

C.3 Experimental Analysis of the Structure of NE

We complement our theoretical results with simulations illustrating how delegation graphs of Nash equilibria, found using our best-response dynamic (Appendix C.1), can appear in synthetic instances. We create instances at random of various sizes: $n \in \{20, 50, 100, 200\}$. For each instance size, we create 30 instances comprised of $\boldsymbol{x}, \boldsymbol{p}$ and $\boldsymbol{\tau}$ chosen uniformly at random such that for each $i \in V$, $p_i, \tau_i \in [0,1]$ (rounded to 2 d.p.), and $x_i \in [0,1]$ (rounded to 3 d.p.).In total, we have 120 instances for which we ran our best-response protocol and found a NE, analyzing its structure according to average measurements for various metrics (see Table 7). First, observe that all of the values present in the table are the average values over the 30 instances with the same number of voters. For example, each of the 30 instances when $n = 20$, on average, has 2.66 have cycles, and when $n = 50$, the average width of a cycle, averaged over all 30 instances, is 0.05. We observe that components are large in the number of voters included. Their size compared to the total number of voters also grows as the latter increases. Both the width of cycles and components (i.e., maximum distance between voters therein) decrease as the number of voters increases, with voters in the same component, especially in cycles, having very close positions. The proportion of voters who prefer self-delegation over participating in a cycle remains steady at approximately 5%, similar to the proportion in weakly connected components of size 1.

The second way we analyze the structure of a NE within the default delegation model is via the impact of the randomly chosen values of $\boldsymbol{\tau}$. Thus, we restrict the following experiments to our previous 30 instances when $n = 50$, and we then study the impact of scaling the tolerance vectors by 0.75 and 0.5. We take 20 vectors $\boldsymbol{\tau}_{\leqslant 1} \in [0,1]^n$ chosen uniformly at random for each pair $\boldsymbol{x}, \boldsymbol{p}$. We let $\boldsymbol{\tau}_{\leqslant 0.75} = 0.75 \times \boldsymbol{\tau}$ and $\boldsymbol{\tau}_{\leqslant 0.5} = 0.5 \times \boldsymbol{\tau}$, scaling each value in the vector by either 0.75 or 0.5. For

	$\boldsymbol{\tau}_{\leqslant 1}$	$\boldsymbol{\tau}_{\leqslant 0.75}$	$\boldsymbol{\tau}_{\leqslant 0.5}$
Av. tolerance	0.50	0.37	0.25
No. SCC	5.63	6.35	7.69
No. cycles	4.55	4.96	5.77
No. self loops	1.08	1.39	1.92
No. WCC of size 1	0.98	1.23	1.60
Av. no. paths into SCC	6.68	7.05	7.69
Av. width of SCC	0.043	0.039	0.032
Av. width of WCC	0.17	0.15	0.12
Av. size of cycle	4.59	4.35	3.99

Table 8: Reporting various average measurements of the structure each of the 600 equilibria found for each tolerance vector $\boldsymbol{\tau}_{\leqslant k}$ with $k \in \{0.5, 0.75, 1\}$. Note that SCC denotes strongly connected components (i.e., cycles and self-loops), and WCC denotes weakly connected components (i.e., an SCC and the paths entering the SCC).

	$\boldsymbol{d}_{BR}$	$\boldsymbol{d}_{\texttt{acyc}}$	$\boldsymbol{d}_{\texttt{dir}}$
20 voters	0.953	0.859	0.497
50 voters	0.966	0.900	0.494
100 voters	0.969	0.916	0.504
200 voters	0.976	0.940	0.501
$\boldsymbol{\tau}_{\leqslant 1}$	0.983	0.913	0.494
$\boldsymbol{\tau}_{\leqslant 0.75}$	0.978	0.902	0.494
$\boldsymbol{\tau}_{\leqslant 0.5}$	0.965	0.878	0.494

Table 9: The average percentage of the votes cast across the 30 instances when varying $n \in \{20, 50, 100, 200\}$ and 600 instances when varying $\boldsymbol{\tau}_{\leqslant k}$ for $k \in \{1, 0.75, 0.5\}$. The three models we compare the percentage of votes cast are the equilibria within our default delegation model (NE), acyclic liquid democracy (acyc), and direct democracy (dir).

$\boldsymbol{x}, \boldsymbol{p}$ and each $\boldsymbol{\tau}_{\leqslant k}$ with $k \in \{1, 0.75, 0.5\}$, a NE is found via our best response protocol (see Algorithm 1 in Appendix C.1). We then take the average measurements over the 600 equilibria found from a certain kind of $\boldsymbol{\tau}$. These values can be found in Table 8.

We see that the number of connected components increases while the values in the tolerance measures decrease, and we see that the number of connected components ending in a self-loop increases slightly as well. Moreover, the number of connected components of size one, i.e., voters delegating to themselves without receiving any delegations, (slightly) increases as the tolerance vector decreases. Another measure we looked at was the width of a cycle and of a weakly connected component, i.e., the largest distance between any two voters within it. The width of cycles and weakly connected components and the size of cycles all decrease slowly as the tolerance vectors decrease, consistently maintaining notably small widths.

C.4 Experimental Analyses of Approximating SW

Our experimental setup created 400 instances for each $n \in \{5, 6, 7, 8\}$, where each $i \in V$ had a distinct position $x_i \in [0, 1]$ rounded to three decimal places where no two voters are at the same position, a randomly chosen $\tau_i \in (0, 1)$ without rounding, and a $p_i \in (0, 1)$ rounded to two decimal places. For each of these instances, we identified the delegation profile $\boldsymbol{d}_{SW}$ that maximized $SW(\boldsymbol{d}_{SW})$ and computed $\texttt{ODP}(\mathcal{I})$. We then examined the values of ${SW(\boldsymbol{d}_{SW})}/{n}$ and ${\texttt{ODP}(\mathcal{I})}/{n}$, which allowed us to calculate the increase in the average expected utility of a voter in $\boldsymbol{d}_{SW}$ compared to their optimal path. This led to ${\texttt{ODP}(\mathcal{I})}/{n}$ being 3.0%, 3.0%, 2.8%, and 2.6% higher than ${SW(\boldsymbol{d}_{SW})}/{n}$ for $n = 5, 6, 7, 8$, respectively. Thus, our upper bound $\texttt{ODP}(\mathcal{I})$ is close to the highest possible social welfare $SW(\boldsymbol{d}_{SW})$. However, there appears to be a trend of decreasing distance as the number of voters grows.

C.5 Expected Number of Votes Cast in Equilibria

We now revisit our instances with varying n and $\boldsymbol{\tau}_{\leqslant k}$ (from Appendix C.2), which include 30 and 600 instances for each variant, respectively. To contextualize our model, we compare the expected percentage of votes cast under three scenarios: a NE the default delegation model NE, acyclic liquid democracy acyc, and direct democracy dir. We use the delegation profiles for the three models described in Appendix C.2. We computed the average expected percentage of votes cast in each instance for all three models, and the results are presented in Table 9.

When varying the number of voters, we observe that in the default delegation model, a random NE achieves a very high percentage of votes being cast—starting at 95% for 20 voters and steadily increasing as the number of voters grows. A similar pattern is seen in the acyc setting, though the increase is sharper, and, even at its peak (with 200 voters) the percentage of votes cast in the acyc setting remains below the lowest value achieved in the default delegation model (achieved for 20 voters). In stark contrast, the direct democracy setting consistently loses around half of the votes of votes, which corresponds to the average probability of delegating. These results highlight the significant advantage of our framework in preserving voting power. When varying the values of $\boldsymbol{\tau}_{\leqslant k}$ for $k \in \{1, 0.75, 0.5\}$ leads to 98.3%, 97.8%, and 96.5% expected votes cast in the default delegation model.