Voter Competence in Liquid Democracy

Voter Competence in Liquid Democracy

a thesis submitted
by
Montague Mawere

in partial fulfillment of the requirements
for the degree of
Bachelor of Arts
in
Computer Science And Philosophy

Harvard University
Cambridge, Massachusetts
March 26, 2021

Thesis advisors: Professor Jim Waldo and Meica Magnani Montague Mawere

Voter Competence in Liquid Democracy

Abstract

Liquid democracy is a form of democracy where people can either vote directly on policies or delegate their voting power. It combines features of direct democracy and representative democracy. In doing so, some believe it can realize the strengths and avoid the weaknesses of both alternatives. This thesis uses computational and philosophical debates on liquid democracy to evaluate its relative performance, given different assumptions regarding voter competence. Whether or not liquid democracy can realize these strengths depends on voter competence. Thus, the organizing question will be: under what assumptions regarding human competence does liquid democracy lead to better outcomes than those of direct democracy and representative democracy?

The computational models of democracy build upon the social choice literature on liquid democracy. Single elections have two policies, where one policy is independently preferable to the other. Voter competence is measured by one's likelihood of selecting that better policy. Two further distinctions are considered: personal competence and social competence. Personal competence determines how well people vote individually. Social competence determines how well people select others to vote on their behalf. Interestingly, the simulation results show that liquid democracy outperforms the alternatives when voters have high social competencies but varied personal competencies.

These models and simulations illuminate the critical role that voter competence plays in models of democracy. The models bring to light the fact that the relative success of liquid democracy relies heavily on various understandings of voter competence. The effect of voter competence shown in these simulations depends on even more assumptions about human behaviors, social structures, and the normative goals of democracy. Computational models are subject to many limitations resting on assumptions that must be further justified or shown in empirical research.

The philosophical discussion of this thesis, then, identifies and applies pressure to some of these limitations. It is essential to clarify what is modeled and the method of evaluation. With regards to the former, this thesis only considers democratic voting for policies at scale. With regards to the latter, some evaluate democracy concerning preferences, while others evaluate it concerning the truth. Given these approaches, one could focus on voters' preferences and the truth being maximized in the process or the outcome. These particular models present something closer to the truth-tracking outcome-oriented approach, though this is not incompatible with the other views. An understanding of what qualifies democracy establishes what makes an individual's vote better. Ultimately, notions of voter competence are rooted in normative democratic ideals. The hope is that these discursive efforts will contextualize and guide further investigations of liquid democracy.

Contents

Acknowledgments

My deepest appreciation to my advisors, Professor Jim Waldo and Professor Meica Magnani, whose guidance, proofreading, and insightful comments made this project possible. I am extremely grateful to Professor Ariel Procaccia for conversations concerning this work and to London Lowmanstone for his contribution to the code.

0

Introduction

Democracy is a system of collective decision-making where individuals share equal voting power. The origins of this system are attributed to ancient Greece, where free people of Athens gathered at the agora ("gathering space") and voted on aspects of public life by majority rule. This process is a direct democracy that is still widely used by many organizations and regions. As societies grew large, this form of democracy was not scalable. It is no longer

possible for many nations to have a populace gather in a physical public space. In the 21st-century, information and communication technology can create virtual public spaces akin to the Athenian agora. Social media and online forums are scalable platforms already used for public discourse and could be implemented for voting.

Assuming direct voting could be done securely through technology, a remaining concern about direct democracy is that ordinary citizens are not informed enough to vote on all issues in different domains(economics, public health, infrastructure, etc.). Society has become complex, and it takes a great deal of effort to make informed decisions in different areas of expertise. A solution is to elect representatives who legislate on public policy for a fixed term—relieving the burden of voting from ordinary citizens and still giving them control over policy through electing representatives. Still, under a representative democracy, ordinary citizens have high stakes in certain decisions. In these cases, they could be better off voting themselves instead of lobbying their representative.

Liquid democracy attempts to address these common concerns about modern democracy. Citizens can recall their delegation of voting power and give it to other experts or even vote themselves on any policy issue. Like direct democracy, liquid democracy enables all people to vote directly on policy. And like representative democracy, liquid democracy mobilizes experts to reduce the continual voting burden on ordinary citizens. Voting in liquid democracy contains aspects of both these alternatives. The answer to what degree

society should prefer direct voting to representative voting depends significantly on several assumptions about society and humans. In order to understand this relationship, one must address these questions: What are the goals of voting in a democracy? How well do people vote concerning those goals? Are people able to select others to vote on their behalf? The answers to these questions change which form of democracy is most suitable.

Liquid Democracy

Liquid democracy is a form of democracy where people can either vote directly or delegate their voting power on specific issues. Those who receive increased voting power can re-delegate to someone else. Finally, anyone can instantly recall their delegation. In academic literature, this theory exists under the titles of liquid democracy, proxy voting, or delegative democracy. A formal definition of a liquid democracy is as follows: all eligible community members can: (I) vote directly on policy issues (direct vote); (II) delegate their vote to another member on (1) a single issue, or (2) all policy issues in one or more domains, or (3) all policy issues in all domains (flexible delegation); (III) re-delegate votes received to another member (transitive delegation); (IV) terminate their delegation at any time (instant recall)(Blum & Zuber, 2016). These features create a "liquidity" of voting power as direct and delegated votes can be continuously recast based on current public opinion. Liquid democracy is implemented using information and communication technology to cast votes,

delegate votes on specific issues, and instantly recall votes. Notable computer softwares with some of these features include LiquidFeedback, GoogleVotes, Votorola, Civicracy, and Sovereign.

Thesis Structure

This thesis examines philosophical assumptions implicit in computational models of liquid democracy, direct democracy, and representative democracy. The relative success of liquid democracy is dependent on various understandings of voter competence. Simulations of the models show the effect of voter competence in democratic elections. The primary aim is to describe assumptions regarding voter competence where liquid democracy leads to better decisions than its alternatives. The models are limited by other assumptions on human behaviors, social structures, and the democratic goals of voting. These concerns are discussed after presenting the models and simulation results. This investigation unfolds in three chapters.

The first chapter presents the computational work that extends existing liquid democracy models of truth-tracking binary elections[1]. These models capture elections with two outcomes, where one outcome is perceived as correct. Given this constraint, one can compare the performance of individuals and decision-making processes by their likelihood of selecting the desired outcome. Moreover, one can observe performance trends as the voting

[1] First presented in Kahng et al. (2018) and later used in Caragiannis & Micha (2019).

population increases to determine how these voting systems scale. The liquid democracy model is extended to accounts of direct democracy and representative democracy.

The second chapter presents simulations of these models with different assumptions regarding voter competence. From these results, I note some trends: the delegation of voting power to fixed small groups can be detrimental or beneficial depending on average voter competence; all models perform similarly when the experts grow proportionally to the population; liquid democracy does not strictly outperform its alternatives in any given election simulation; liquid democracy performs best in a series of elections when the relative group competencies vary per election. An ideal scenario for liquid democracy is when voters have high social competencies but varied personal competencies.

In the final chapter, the philosophical discussions characterizes the limitations and content of the models. Firstly, this begins by identifying unrealistic modeling assumptions. The simulations are limited in scope due to these assumptions but offer results worth discussing further. Secondly, democracy encompasses many meanings, so it is crucial to clarify relationships between them. Systems of democratic governments are arbitrarily complex as they are rooted in a lengthy history of societal shifts and the protection of special interests. Voting itself is simpler to model and plays an integral role in determining decisions. Democratic debates can include other government structures beyond merely voting. For the most part, I am concerned with the mechanisms behind voting. However, the discussions are

relevant to all democratic government structures.

Thirdly, I discuss the normative goals of democracy. There are various ideals used to assess the success of democratic systems, each addressing different but related concerns. An epistemic democracy is concerned with whether or not collective decisions are truth-tracking. This ideal relies on the wisdom of crowds: large-scale democracies are valued at least in part for their collective knowledge-producing potential and are defended in relation to this. This notion is not incompatible with other democratic ideals, namely egalitarian ones. An egalitarian democracy requires that people and their preferences are treated equally. One can further assess the success of democratic systems by looking at either the outcomes or the process. When modeling epistemic democracy, one can evaluate whether outcomes align with predetermined truths or whether the process is trustworthy. Similarly, when modeling egalitarian democracy, one can evaluate whether outcomes reflect the voters' preferences' or whether the process treats voters equally. This thesis's models are similar to the epistemic outcome approach. This epistemic ideal is compatible with the other ideals introduced. Finally, I define the types of voter competence used in the models: personal and social competence. That is how well people vote and how well they select others to vote on their behalf. The metrics for determining voter competence are derived from the democratic ideals. These definitions contextualize aspects of the models and simulation results.

Related Work

The first mentions of voting with delegation begin in the 19th century with Dodgson (1884) suggesting a combination of direct democracy and representative democracy to elect candidates to the house of representatives. Later in the 20th century, Miller (1969) and Tullock (1967) proposed using delegation in a more practical sense enabled by technology. Modern advances in communications, social media, and portable computing devices have made liquid democracy more prominent and possible in the 21st century. Behrens (2017) presents a historical overview of the origins of liquid democracy. Ultimately, the contemporary academic literature on liquid democracy falls under three broad disciplines.

The first field is the computational work in computer science, political science and social choice theory. Green-Armytage (2015) begin the modern academic investigation of liquid democracy by presenting proxy voting to reduce performance loss in direct demoracy and representative democracy. Since then, two approaches have been taken in this line of research. Some investigate the potential drawbacks, limitations, and complexities of voting in liquid democracy. For example, Brill & Talmon (2018) and Christoff & Grossi (2017) look at the issues of delegation cycles and rationality. Caragiannis & Micha (2019) and Kahng et al. (2018) evaluate poor performance compared to other models of democracy in non-strategic models, while Bloembergen et al. (2019) do similar work for game-theory strategic ones. Others investigate potentially better delegation features for liquid democracy(Brill

& Talmon, 2018). Escoffier et al. (2019) suggests delegating with preferences over many experts. Gölz et al. (2018) allow for multiple forms of delegation. Colley et al. (2020) enable more complicated delegations that allow for conditioning on results of sub-elections. Boldi et al. (2011) and Kotsialou & Riley (2018) present dampened delegations and breadth-first delegations, respectively.

The second field consists of theoretical work done in normative political philosophy. Blum & Zuber (2016) introduce egalitarian and epistemic ideals as two methods for evaluating liquid democracy. They conclude liquid democracy is more epistemically accurate (better at identifying the truth) and more egalitarian (greater participatory equality) than representative democracy. They identify the potential problems for liquid democracy related to these two ideals and propose solutions. These two democratic ideals provide a starting point for the later discussions of voter competence in this thesis.

The third field is the empirical work studying real-life implementations of liquid democracy in organizations. LiquidFeedback is the most popular software implementation of liquid democracy. Its usage has been studied within the German Pirate Party by Kling et al. (2015). They investigate the unequal distribution of voting power in liquid democracy and the potential danger of high-weight voters or super-voters. They observe super-voters have a two-fold positive impact as they vote according to the majority and stabilize the approval process. De Cindio & Stortone (2013) provide an analysis of LiquidFeedback used in an

initiative in Lombardy, Italy. Hardt & Lopes (2015) gives an account of GoogleVotes, an internal Google software used for employees to vote on dining preferences. Many theoretical drawbacks of liquid democracy have been non-problematic in this empirical research.

It should be kept in mind, then, that liquid democracy is defined and evaluated across different fields. The definitions and standards of evaluation may not cleanly overlap or overlap at all. First, computational social choice theorists use theoretical computer science to create models of liquid democracy elections and study their algorithmic complexity, network structures, and inconsistencies. Second, normative democratic theorists in political philosophy analyze the moral foundations and principles of liquid democracy. Third, empirical researchers analyze data from real-life experiments of liquid democracy softwares within political parties and organizations. This thesis primarily engages with the computational and philosophical work but presents modeling assumptions and discussions that can be justified with empirical evidence.

The computational models build upon those in Caragiannis & Micha (2019) and Kahng et al. (2018). One addition is an account for representative democracy within the same framework. In the existing literature, liquid democracy is compared to either direct democracy or representative democracy, but there is less work on which of the three is, all things considered, preferable. Realistically, representative democracy contains more complex mechanisms to determine the representatives before policy elections. Even without all these

features, modeling representation is essential for this democratic analysis because representative democracy is prevalent in modern government.

The discussions regarding democratic ideals are inspired by the philosophical arguments of Blum & Zuber (2016). They reference democratic ideals as being either concerning equality or the truth. This thesis expands upon those and further breaks those down into process or outcome-oriented approaches. Every computational model mentioned in this section lies within these axes. These distinctions clarify the concepts used in the computational and philosophical literature. With this clearer understanding, models of democracy and voter competence can be expanded.

1

Models

This chapter introduces the voting models for liquid democracy, direct democracy, and representative democracy. The models capture elections with two policies. One can imagine this as a vote for or against a single policy proposal or between two incompatible policy proposals. One policy is assumed to be preferable to the other. There is some

independent reason that this policy is superior, which could be related to the truth, voters' preferences, etc. This assumption is made in order to quantify the performance of a democratic voting mechanism. Voting mechanisms that select these preferred policies more often are considered to be better. This assumption need not apply to every possible election in the real world. There are elections where either option offers no substantive benefits over the other. If an election outcome is this, then the stakes seem lower than what is being considered here. The motivating idea is using democracy to ascertain the best policy in an election. Most intuitively, the best policies are associated with proposals that are more informed by the truth. However, it is worth considering that the best policies also align with voters' preferences. Distinctions between these two democratic ideals are elaborated in the final chapter.

1.1 Overview of Models

Using the assumption that there is a superior policy outcome, I can construct definitions for collective performance and voter competence. Given that there are only two policies, the decision space is represented on a binary scale. Collective performance and voter competence are on this same scale: a probability range from zero to one. Voting mechanisms with high probabilities of selecting the better policy are desirable. Voters with high probabilities of voting for the better policy are considered personally competent. Those who can

select others with high personal competence are socially competent. The models assume all voters have high social competence but allow for variable personal competence. The effect of this is investigated in the simulations. The assumption that all voters have high social competence despite variable personal competence is discussed in the final chapter.

After defining performance and voter competence, I construct and compare the liquid democracy, direct democracy, and representative democracy models. Liquid democracy and direct democracy are directly comparable—a liquid democratic vote without delegations is a direct democratic vote. By constructing a model for liquid democracy that includes a separate mechanism for voters to choose whether to vote or delegate, it is straightforward to extend it to direct democracy by modifying that mechanism to never delegate and to only vote directly[1].

Representative democracy is complicated to consider within a single election model. The average citizen does not vote in policy elections; they vote in other elections to determine which representatives vote in policy elections. Attempting to model that process presents a considerable computational and conceptual challenge to generalize. Instead, I simplify the characteristics of representatives and have them selected accordingly. A prime characteristic is that representatives are well-known; a second is that they are usually more competent. In reality, these two characteristics might be related: more competent people become well-

[1] Modelling both liquid democracy and direct democracy in this way are introduced by Kahng et al. (2018).

known, and well-known people become more competent. For numerous reasons, this is not necessarily true and is not assumed to be within these models. A typical counterexample is that citizens who campaign well become well-known, increasing their likelihood of becoming representatives regardless of their competence. The social networks and exact competencies in these models could be improved. Even so, they capture this asymmetry between popular experts or gurus and their followers. In these models, gurus are the representatives. Only the representatives vote in the model for representative democracy.

1.2 Preliminaries

For these models[2], I consider binary elections with two alternatives T, representing a correct policy and F representing an incorrect policy. Elections are assumed to be binary for simplicity. The epistemic assumption that there exists a ground truth is useful. It entails one policy outcome, T, that is assumed to be better than the other, F. The stakes of these elections are high, and it matters for performance to select T more often than not. Nothing about the nature of T is assumed, only that a better policy decision exists and is denoted by T. The following mechanisms result in collective decisions and are evaluated on their probability of selecting T.

There is a set $V = \{1, \ldots, n\}$ of n voters connected in a social network. The social network is a directed graph $G(V, E)$ with voters as vertices and their relationships as directed

[2] First introduced by Kahng et al. (2018).

edges. A directed edge from voter i to voter j or $(i, j) \in E$ means that voter i knows of voter j. Voters known by voter i are called the neighborhood of voter i. The senses of "knowing" and "neighbor" here are different from their common usages. It is more than physical or virtual acquaintance. A voter knows relevant information on the competence of their neighbors, which they use to make judgments on their relative competence. Each voter $i \in V$ has a personal competence level $p_i \in [0, 1]$ which represents i's confidence in policy T or a probability that i votes correctly for T should they decide to vote. Let $\mathbf{p}$ be the vector of all voters' personal competence.

1.3 Local Delegation Mechanisms

Voters have three potential actions. They can vote directly, delegate their voting power to another voter in their neighborhood, or abstain. Delegations are transitive, which means that received votes can be re-delegated for increased voting weight. I assume that voters are prevented from delegating in a cycle[3] (e.g., voter 1 delegates to voter 2, who delegates to voter 3, who then delegates back to voter 1). The result is a voter graph resembling a forest of trees with edges leading to single roots known as direct voters. Each direct voter i votes with their respective competence p_i and voting weight equal to the sum of their received delegations.

[3] Delegation cycles can pose a serious problem to liquid democracy models(Christoff & Grossi, 2017). But has also been resolved with simple means in practice such as preventing them from occurring as is the case in LiquidFeedback(Behrens et al., 2014).

A local delegation mechanism (LDM) decides how voters choose to vote, delegate or abstain. It is specified to be local because each voters' choice is independent of the global state. For a non-local mechanism, one could imagine a centralized authority strategically coordinating how everyone acts to maximize selecting better policies. However, this does not capture the spirit of democratic voting. In local mechanisms, each voter takes action based upon their local judgment. An LDM uses the network graph $G(V, E)$ and competence vector $\mathbf{p}$ to produce a voter graph indicating who delegates (the tree leaves), who votes directly (the tree roots), and who abstains (removed from the graph). Abstention is only for the convention that non-representatives abstain from voting in representative democracy.

1.4 Liquid Democracy Model

In liquid democracy, voters can either vote directly or transitively delegate their vote within their neighborhood. An idealistic LDM in liquid democracy would have voters delegating only to the most competent voters they know, maximum delegation. This LDM strongly assumes high social competencies because every voter reliably judges themselves and their neighborhood's relative competencies. Additionally, they act according to that judgment and impartially give their voting power to most competent voters. These assumptions are optimistic[4]. One would expect a model of this type to outperform alternatives because the

[4]Caragiannis & Micha (2019) argue for a more realistic delegation mechanism where voters with less than $\frac{1}{2}$ personal competence delegate to less competent voters rather than to more competent voters.

most competent voters amass more voting weight. However, this is not what leads to the best overall performance in all cases.

LD is the liquid democratic process that uses maximum delegation on $G(V, E)$ and competence vector $\mathbf{p}$ to produce an acyclic voter graph. Each root or direct voter i votes with their respective competence p_i and weight equal to their tree size on the voter graph. A collective decision is determined by weighted majority rule. $P_{LD}(G, p)$ denotes the probability that the outcome of this process is policy T.

1.5 Direct Democracy Model

In a direct democracy, voters vote directly and never delegate. Direct democracy can be modeled by altering the LDM of LD[5]. If voters never delegate, then liquid democracy is equivalent to direct democracy. It is straightforward to use the above model of liquid democracy and only change the delegation mechanism. Let direct voting be the voter mechanism where voters never delegate their vote. *DD* is the direct democratic process that uses direct voting on $G(V, E)$ and competence vector $\mathbf{p}$ to produce an edgeless voter graph (i.e., no delegation). Each voter i votes with their respective competence p_i and an equal weight of one. The collective decision is determined by majority rule. $P_{DD}(G, p)$ denotes the probability that the outcome of this process is policy T.

[5] Building three comparable models by modifying only the delegation mechanism was originally done by Caragiannis & Micha (2019).

1.6 Condorcet's Jury Theorem

Condorcet's Jury Theorem (CJT), a computational social choice theory, shows that as group size increases, the probability that direct voting results in the correct policy converge to one, on the assumption that every voters' probability of voting for the correct policy is greater than 0.5 (Grofman et al., 1983). In other words, better-than-random but worse-than-perfect voters converge on perfect collective decision-making as the group size increases when using direct democracy. A corollary is that the largest group of these voters outperforms any subset of those voters.

The CJT proof rests on a few assumptions: universal personal competence assumes that every voter is better than random when voting directly ($\forall i \in V; p_i > \frac{1}{2}$). Voter independence is the assumption that every voter votes independently of any other voter. Voter homogeneity is the assumption that the voters are equally competent ($\forall i, j \in V; p_i = p_j = \overline{\mathbf{p}}$). Independence and homogeneity simplify the probability formula and model. If voters are not independent, it is challenging to model various ways to votes correlate (media, religion, family, geography). If voters are not homogeneous, it complicates the concise probability formula representing every combination of votes causing policy T to win. Universal personal competence is crucial because if voters are all personally incompetent ($\forall i \in V; p_i < \frac{1}{2}$), then the same proof shows the opposite result: the performance worsens as the group size increases.

Condorcet Jury Theorem. Given voter independence and homogeneity, if $\forall i \in V; p_i > \frac{1}{2}$ (universal personal competence), P_{DD} only increases with n (number of voters) and $\lim_{n\to\infty} P_{DD} \to 1$. If $\forall i \in V; p_i < \frac{1}{2}$ (universal personal incompetence), P_{DD} only decreases with n and $\lim_{n\to 0} P_{DD} \to 0$. If $\forall i \in V; p_i = \frac{1}{2}$ (universal personal randomness), $P_{DD} = \frac{1}{2}$ for all V. Finally, $P_{DD} = \sum_{k=\lceil \frac{n}{2} \rceil}^{n} \binom{n}{k} (\overline{\mathbf{p}})^k (1-\overline{\mathbf{p}})^{n-k}$ is the probability formula used for direct democracy under these constraints.

1.7 Star Social Network Graphs

Voter homogeneity is an assumption worth modifying to create more realistic social structures. As mentioned, this assumption simplifies the formula in CJT, but delegation and representation mechanisms typically rely on social networks where some voters are more competent than others. Usually, the motivation is that we want these more competent voters to have more decision-making power. It then makes sense to consider a graph in which we have different classes of voters with different competencies. An example network structure used in liquid democracy literature is a k-center (k+1)-uniform star. Some n_f followers are connected to $k = n_g$ central gurus and one other follower. There are n_p partisans disconnected from gurus and followers. Followers, gurus, and partisans have competencies p_f, p_g, and p_p, respectively. With three types of voters, there is no longer complete homogeneity but homogeneity within each group.

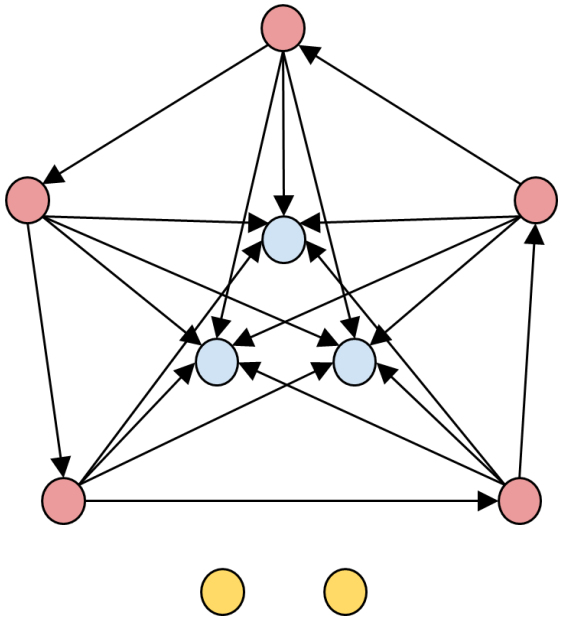

Figure 1.1: Graph $G(V, E)$ is a k-center-(k+1)-uniform star with $k = n_g = 3$ gurus, $n_f = 5$ followers, and $n_p = 2$ partisans. Followers are connected to all k gurus and one other follower.

1.8 Liquid Democracy's Impossibility Result

For direct democracy and CJT, the results still hold though the formula is more challenging to generalize. Under the direct democracy model, the probability of T winning tends towards one when the group size increases. Since the direct democracy model is independent of the network structure, it reduces to voters voting independently. The least competent voters are better than random by universal personal competence, so the collective probability of selecting T is be greater than a homogenous group of those voters with the same competence, which by CJT tends to one with increasing group size. In summary, if G is a k-center-(k+1)-uniform star, then $\lim_{n\to\infty} P_{DD} \to 1$ all the same by CJT. Thus under uni-

versal personal competence, direct democracy still leads to perfect performance on this star social network.

One would expect LD to always be at least as good as DD, but this is not the case. In liquid democracy, voters can always choose never to delegate (i.e., direct voting), and the performance would be identical to direct democracy. Delegation is an optional feature intended to increase performance. Moreover, maximum delegation seems to be an idealistic assumption that voters always delegate to the most competent voters. However, when more competent voters have more voting power, the collective performance is bounded by the competence of that group of voters. This consistently underperforms relative to this direct democracy model as the voter population scales.

For example, imagine a star graph that has $n - 1$ followers with $p_g = \frac{2}{3}$ connected to a single guru with $p_g = \frac{4}{5}$, and no partisans. This social network has many competent followers connected to a single more competent guru. By CJT, $\lim_{n\to\infty} P_{DD} \to 1$ and this holds regardless of whether the population grows with more gurus, followers, or partisans. So direct democracy still tends towards perfect performance as the voter population increases. Voter independence and universal personal competence still cause this effect with multiple voter types. Under these assumptions, the collective performance grows to one on this example star social network. However, $\lim_{n\to\infty} P_{LD} \to \frac{4}{5}$. By maximum delegation, every follower delegates to the single guru with personal competency of $p_g = \frac{4}{5}$. This single guru

decides the election's outcome, so the collective performance is bounded by their personal competence. So *LD* with maximum delegation has a loss of performance compared to *DD* (i.e., *LD* with no delegation).

This example has been generalized by Kahng et al. (2018), who prove that there is no LDM for liquid democracy that is guaranteed to outperform direct democracy, the liquid democracy's impossibility result. Still, it might not be concerning that liquid democracy is not guaranteed a performance gain if it usually outperforms direct democracy. But for the k-center-(k+1)-uniform star, direct democracy outperforms liquid democracy with maximum delegation. So this particular social network structure and delegation mechanism are helpful to extend those negative results into the simulations presented in the next chapter. Under different assumptions regarding voter competence, this results does not hold. Such as if every voter is personally incompetent (i.e., worse than random), then bounding the collective performance to a group of slightly better voters offers performance gains. If this scenario is likely, one could even avoid per election delegations and always delegate collective decisions to a select group of popular experts. The following model of representative democracy captures this concept.

1.9 Representative Democracy Model

With this social network in mind, we can use it to select popular experts (i.e., gurus) as representatives. In a representative democracy, voters elect representatives, and only those representatives vote directly on policy issues. Non-representatives decline their right to vote (i.e., abstain) on policy issues. Generally, candidates campaign before elections presenting their plans and increasing their followers. Representatives are usually selected through voting. They then act as general-purpose experts, voting on a wide range of policy issues in every election for a fixed term. The selection process for representatives and the ways they vote for policies can vary greatly depending on a government's particular system. For simplicity, selected representatives are the gurus or center of the star social networks. Gurus usually have higher personal competencies than their followers.

There are many assumptions about the kinds of people who become societal gurus and whether all those people would win representative elections. This selection process is arbitrary. The relevant feature is that a representative democracy election is a direct democracy election with a select subset of voters. In this case, that subset is well-known and usually above-average personal competence, which is equally as optimistic as previous assumptions on voter independence and social competence. Modeling representation in this way allows for a direct comparison between the three types of democracies in elections. Liquid democracy maximizes delegations to the most competent voter. Direct democracy only employs

a direct vote. Representative democracy fully delegates to the well-connected central voters in a star social network. Let full delegation be the LDM that produces a new voter graph containing only the representatives (or gurus).

RD is the representative democratic process that uses full delegation on $G(V, E)$ to produce a subset of representatives with no edges (i.e., no delegation). Each representative voter i votes with their respective competency p_i and an equal weight of one, which results in a collective decision by majority rule. $P_{RD}(G, p)$ denotes the probability that the outcome of this process is policy T.

1.10 Relative Guru Competence

Under certain assumptions, *LD* and *RD* are nearly identical in performance. When the gurus are more competent than their followers, maximum delegation delegates most votes to the gurus. The representative democracy model always fully delegates voting power to the gurus. So if one assumes that those who become gurus are more competent than their followers, both models perform the identically because gurus amass deciding voting power. Imagine, on some issues followers are more competent than their gurus. In the liquid democracy model, followers would not delegate to gurus in these elections by maximum delegation. Since representative democracy always leaves decisions to the gurus, it could suffer performance losses in these elections.

It is usually assumed that gurus are more competent than their followers, but there are many reasons to deny this. The effect of breaking this assumption is shown by the simulations. The justifications for doing so are discussed in the final chapter. These models enable either gurus, followers, or both to be personally competent. Additionally, they allow the voter group that is most competent to be altered. These allowances are used to investigate their effect on the models' performances. The simulations that follow use the above models with these various constructions of voter competence to compare their relative performance of these democracies. In doing so, I show the assumptions regarding voter competence that cause liquid democracy to outperform the alternative models.

2

Simulations

The simulations of the models ran closely following the design in the previous chapter. The social networks are always k-center-(k+1)-uniform stars with the three voter types of gurus, followers, and partisans. There are $n_g = k$ center gurus, n_f followers connected to all of the gurus and one other follower, and n_p partisans isolated. Gurus, followers, and

partisans vote with p_g, p_f, and p_p, respectively, if they vote. Generally, $p_g > p_p > p_f$, unless specified otherwise. In *LD*, voters use maximum delegation, delegating to any of the most competent voters in their neighborhood. In *DD*, voters use direct voting, always voting directly. In *RD*, voters use full delegation, delegating only to gurus who vote with equal weighting. A star graph and competence vector is generated for each group size, then input into the various democratic voting mechanisms. Each process iterates over each voter, having them vote, delegate, or abstain according to the respective LDM. P_{LD}, P_{DD}, and P_{RD} are calculated by running iterations of identical elections until performance converges. These simulations are open-source, and the code is available at `https://github.com/Micmonta/liquid-democracy-models`.

2.1 Delegations to Fixed Size Groups

The first result is the bounding property that appears when delegating to a fixed size group. Figure 2.1 plots the performance of democratic voting mechanisms (the probability that policy T is selected) as the number of followers increases, but the number of gurus remains constant. This type of population growth aligns with the structure of modern representative democracies with fixed representative seats. The number of gurus, the experts with public attention and full-time devotion to politics, remains fixed while their following grows. In (a), all the voters are personally competent, while in (b), all the voters are person-

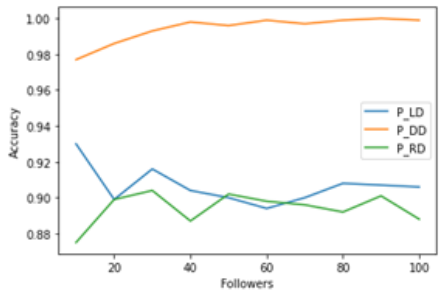

(a) Universal personal competence

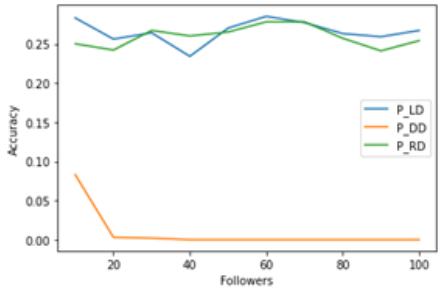

(b) Universal personal incompetence

Figure 2.1: Accuracy of democratic voting mechanisms as the number of followers increases. There is a fixed number of gurus and partisans as the number of followers increases. Specifically, $n_g = 3$ and $n_p = 2$ while n_f is plotted along the x-axis. Gurus have the highest personal competence, followed by partisans, then followers. (a) With universal personal competence, every voter has a greater than $\frac{1}{2}$ probability of voting for T. In particular, $p_g = \frac{4}{5}, p_p = \frac{3}{4}$ and $p_f = \frac{2}{3}$. (b) With universal personal incompetence, every voter has less than $\frac{1}{2}$ probability of voting for T. In particular, $p_g = \frac{1}{3}$, $p_p = \frac{1}{4}$ and $p_f = \frac{1}{5}$.

ally incompetent. The understanding of personal competence is the same as before: being better than random at voting for a correct policy. In the first case, P_{DD} has the highest starting performance and tends towards perfect performance as group size increases. This result reveals the effect of CJT on multiple voter types. The direct democracies' performance remains much higher than the individual competence of any guru, the most competent voter; for these simulations, $p_g = \frac{4}{5}$. Both P_{LD} and P_{RD} are approximately equal in performance and remain so as the group size increases. Their performance is bounded by the competence of the fixed sized group of gurus. Liquid democracy performs slightly better than representative democracy for smaller groups because of the partisan voters and its randomness of delegation assignments.

One could see results in Figure 2.1 (a) and (b) occurring in the same world over different

elections. In some decisions, people are much better than random at selecting between two policies, and in others, they are much worse. That is more realistic than assuming one or the other states of total competence or total incompetence. Using delegation to bound the group's performance is appealing to avoid the worst case of incompetent voters in direct democracy.

This bounding property of delegation to fixed sized groups is often cited in favor of direct democracy over its alternatives. However, it only holds under the assumption of universal personal competence[1]. CJT also shows that the collective performance strictly decreases for worse than random direct voters. The simulations in (b) confirm this and show the improved performance due to the bounding property of delegating to fixed size groups. In this case, gurus are still more competent than their followers. Under these circumstances, the bounded group's performance is beneficial and not detrimental to liquid democracy and representative democracy. Suppose all voters are incompetent for certain decisions. It may be better to have them delegate to a smaller group of voters, even if they are no more competent than the others.

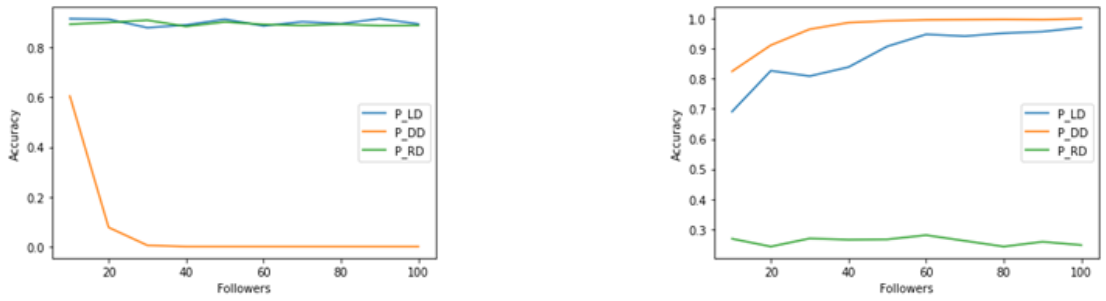

(a) Follower personal incompetence

(b) Guru personal incompetence

Figure 2.2: Accuracy of democratic voting mechanisms as the number of followers increases. There is a fixed number of gurus and partisans as the number of followers increases. Specifically, $n_g = 3$ and $n_p = 2$ while n_f is plotted along the x-axis. (a) With follower personal incompetence, followers have less than $\frac{1}{2}$ probability of voting for T. In particular, $p_g = \frac{4}{5}, p_p = \frac{3}{4}$ and $p_f = \frac{1}{5}$. (b) With guru personal incompetence, gurus have less than $\frac{1}{2}$ probability of voting for T. In particular, $p_g = \frac{1}{3}, p_p = \frac{3}{4}$ and $p_f = \frac{2}{3}$.

2.2 Guru and Follower Incompetence

This section controls for specific groups being incompetent, namely gurus and followers. These two groups play essential roles in direct democracy and representative democracy, respectively. Direct democracy is most appealing when average citizens (the followers) are competent and capable of performing their democratic duties. In comparison, representative democracy is most appealing when popular experts (the gurus) are competent and qualified to decide on others' behalves. Figure 2.2 plots simulation results with one voter type having a worse than random chance of voting for T. Again, note the population only grows with more followers. Hence, follower incompetence has a significant effect when many followers vote, as is the case in direct democracy.

[1] The liquid democracy impossibility result in Kahng et al. (2018) shows that liquid democracy is not guaranteed to outperform direct democracy, only under the assumption of universal personal competence.

Given that followers are the most dominant majority group, direct democracy follows the same CJT trends as in Figure 2.1, tending towards perfect performance when followers are competent (a) and poor performance if they are not (b). Representative democracy uses only the guru group; its performance is again bounded by that group, which is relatively high in (a) and low in (b). Finally, this setup is the ideal case for liquid democracy, especially using maximum delegation, as it performs well in either case. In (a) incompetent followers delegate their votes to more competent gurus, while in (b) competent followers do not delegate their votes to incompetent gurus. Whether incompetent followers can be expected to employ maximum delegation or some similar mechanism to delegate to the more competent gurus is discussed in the next chapter.

Once again, these simulations are more intuitive if imagined in the same world over different elections. For some policy areas, average citizens are not informed enough to make better than random decisions while the experts are. For other areas, it is better to rely on the wisdom of crowds. When average voters are invested in specific elections, the majority may disagree with the experts. In this case, enforcing an expert-decided policy that disagrees with the majority's wishes may lead to unrest. In this circumstance, gurus seem less qualified than their followers in voting for the best policies. In other elections that rely less on voters' preferences and more on facts, those gurus could be much more competent than the average follower.

The conclusion from these results is that the liquid democracy model performs whether the gurus or the followers are incompetent while the other two models do not. In any single election, liquid democracy does not significantly outperform direct or representative democracy. However, in a series of elections that vary between guru and follower incompetence, *LD* would outperform the other models. Liquid democracy's core feature is the potential flexibility to switch between the benefits of direct and representational voting.

2.3 Growth with Proportional Voter Types

In this section, Figure 2.3 plots the results when all voters groups grow proportionally. When the population increases, the number of popular experts (gurus) and outliers (partisans) increases in proportion to the followers. In larger populations, there could be more experts in each policy domain. More representatives or experts could be trained to maintain a fixed ratio between representatives and their base as the population increases. When universal personal competence and voter independence are maintained, all models follow the trend seen in CJT, albeit at different rates. The collective performance converges to perfect decision-making in all models of democracy.

First, these trends imply there may not be significant performance differences between these forms of democracy. Suppose voters are universally personally competent and significantly more so than random. In other words, every voter is reasonably competent. It

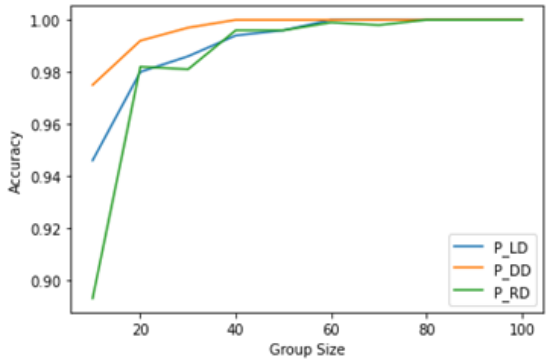

Figure 2.3: Accuracy of democratic voting mechanisms as the group size increases proportionally. The ratio of gurus to followers to partisans remains the same. Specifically, the ratio $n_g : n_p : n_f = 3 : 2 : 5$ is maintained while $n = n_g + n_p + n_f$ is plotted along the x-axis. Gurus have the highest competence, followed by partisans and followers. With universal personal competence, every voter has a greater than $\frac{1}{2}$ probability of voting for T. In particular, $p_g = \frac{4}{5}$, $p_p = \frac{3}{4}$ and $p_f = \frac{2}{3}$.

matters less who votes and what mechanism they use. Second, CJT identifies the addition of independent voters as a factor associated with performance. Direct democracy performs better despite allowing less competent voters to vote. If the expert group of independent voters also grows, then the performance is no longer bounded. The result of this investigation could support proportional representation over fixed size representative seats.

However, I will motivate the fixed size guru assumption. As the population increases, so do the number and complexity of issues, giving rise to new domains of expertise. The number of experts in a given field could remain constant over time as the population grows. There are more gurus overall but more domains such that there is a relatively fixed number of gurus in each domain. However, the number of followers in a given election increases as this group includes every citizen, even the potential gurus in other contexts. It remains

for further research whether proportional or fixed group population growth is desirable and realistic. Regarding this result, it is desirable only under the assumption of universal personal competence.

2.4 Summary of Results

This chapter's simulation results rely on models that assume high social competencies, social network structures, and best policy outcomes to show interesting trends related to voter competence as the population grows. First, when delegating to a fixed size group of experts, the collective performance is bounded by that group. This bounding property is detrimental if voters are universally personally competent but beneficial otherwise. On the other hand, when all voters satisfy the sufficient competence threshold of being better than random and are independent: voters are better off with a direct vote. Conversely, if all voters do not satisfy the sufficient competence threshold, delegation and representation are appealing.

Further, if gurus are periodically less competent than their followers, liquid democracy outperforms the alternatives. In any given election, liquid democracy does not outperform direct democracy or representative democracy[2]. With some basic assumptions on voter competence, direct democracy or representative would be better options. The main ben-

[2] It has been proved in another computational model that liquid democracy never outperforms the best of direct voting and full delegation(Caragiannis & Micha, 2019). Though the delegation mechanism for liquid democracy was notably different than the one here.

efit of liquid democracy is its flexibility allowing voters to either vote directly or delegate voting power. High social competence enables deciding whether to vote directly or delegate reliably. Varied distributions of personal competence between elections increase the importance of that choice.

When the competent representative groups grow proportionally to the population, all modeled democracies converge to similar performance. This result raises an optimization problem for balancing the trade-offs between voter group sizes and competencies required for reliable decision-making. In this final simulation, there is no significant performance difference between these models of democracy. Since proportionally increasing the experts or gurus minimizes the effect of voter competence in these models under the assumption of voter independence.

3

Discussions

This chapter discusses limitations, concerns, and definitions regarding the models and simulation results. Firstly, I explain how some limitations are troubling these models. These are the assumptions that are unrealistic but conceptually challenging to address. These assumptions are on voter independence, the social network structure, the

model of representative democracy, and binary elections. Secondly, I acknowledge the relationship between democratic voting and other government systems. Voting is integral to modern democracies, but other mechanisms can solve the same problems. Thirdly, I discuss normative democratic ideals compatible with these models. Though an epistemic outcome-oriented understanding is most suitable for these models, it is compatible with other democratic ideals. Finally, voter competence is revisited relative to these democratic ideals.

3.1 Unrealistic Social Assumptions

The voter independence assumptions within these models are unrealistic. It has the most significant bearing on direct democracy as performance tends towards the extremes. However, there are many ways in which voters and votes are correlated (family, media, geography). Then, there are voting strategies where people compromise on beliefs to gain higher consensus on fewer policies. Deliberation[1] and public discourse are also prominent components of democracy, with significant opportunities for citizens to influence each other's votes. Modeling voter correlation, the opposite of voter independence, would likely minimize the phenomena of CJT. In other words, collective performance would be bounded by

[1] Deliberation is a special kind of voter correlation which is commonly argued to improve democratic performance. Deliberative democracy can be considered its own form of democracy where deliberation is valued as part of the collective decision-making process. This research could be combined with the democratic models here to account for this kind of voter correlation.

the individual competence of a few influential experts. One might even see the correlation effect of liquid democracy in direct democracy when popular expert voters convince others to vote in a certain way. There numerous ways that votes are correlated that are challenging to model realistically. A straightforward method of modeling correlation is to use the social network graph to reinforce or weaken personal competencies. Even without this, these models show an inverse relationship between performance and voter correlation worth investigating further.

The delegation mechanism in liquid democracy is idealistic. Maximum delegation means that voters delegate only to the most competent voters. This presupposes voters have the skills, awareness, and competencies to perform this reliably. Caragiannis & Micha (2019) suggest an alternative delegation mechanism for liquid democracy whereby voters who are personally incompetent delegate to the less competent voters. Such a delegation mechanism would likely reduce the performance of liquid democracy in these models. In order to simplify this investigation, only one liquid democracy delegation mechanism is used. Maximum delegation best captures the ideals of liquid democracy where voters would only delegate in advantageous ways from their local perspectives. It is surprising how this alone does not outperform direct voting. This thesis presents other conditions that cause liquid democracy with maximum delegation to outperform the alternatives. Investigating more realistic delegation mechanisms in liquid democracy would be better to show strong posi-

tive results. Presenting and evaluating more complex delegations mechanisms is beyond the scope of this thesis.

The social network structures used in the simulations and models could be improved. A better approach would be to use randomly generated(Gölz et al., 2018) or more sophisticated(Bloembergen et al., 2019) social networks. A relevant feature captured by the star social networks is the asymmetry between popularity and competence. In the simulations, gurus are not necessarily more competent than their followers. It might be such that the highly competent voters become well-known and well-known voters become highly competent[2]. However, this relationship is not assumed. Still, the star social network extends examples in the literature that show negative results for liquid democracy. The simulations highlight the conditions behind those negative results and alternative conditions that would change those findings. Though the social networks could be made more realistic, this thesis explores this example thoroughly.

Some simulations would perform equivalently on fully connected network graphs. This approach was avoided to capture the nature of local mechanisms where voters are not connected to every other voter. Connections on these networks are more than being acquainted. The connections represent a reasonably good understanding of other voters' views and competence such that one could make reasonable judgments. This type of con-

[2] Given representatives or gurus are usually equipped with aids, consultants etc. it may be reasonable to assume they are more competent.

nection seems unlikely to occur between all voters at scale. Instead, average voters are connected to the popular experts and few others. In contrast, well-connected experts are not mutually connected to their followers. This scenario is more realistic than a fully connected social network graph.

The social network has the most significant effect on representative democracy by construction. The liquid and representative democracy models are nearly equivalent for most simulations. The significant differences between the two are partisans' votes, and that followers do not delegate to less competent gurus in liquid democracy. Partisans act as additional independent voters, who are beneficial to performance when personally competent and independent. It is also controversial whether followers can reliably choose when to delegate. This latter point is discussed later in this chapter.

There are relevant differences not modeled between liquid democracy and representative democracy. Firstly, liquid democracy allows for domain-specific delegation as opposed to general-purpose delegation. One can pick experts regarding their specific areas of expertise instead of electing representatives that create policy packages. Both approaches to representation have distinct merits. Liquid democracy's approach allows for the relevant experts to vote, while representative democracy's approach allows for overall consistency and the budgeting of priorities. Secondly, liquid democracy allows voters to instantly recall delegations, whereas representatives are elected for fixed terms. Liquid democracy aims to hold experts

continuously accountable to their followers, while representative democracy allows for periodic stability and the forced exchange of power. These additional distinctions within the models would improve the comparison between liquid democracy and representative democracy

Finally, these models only capture binary elections between two policies. It is computationally challenging to extend these models to multiple policy elections. Voters could select acceptable but non-optimal policies. Personal competence would be a more complex to define for multiple voting outcomes. One approach to this problem is to model voter's competence as a vector of probabilities and assign each policy a performance weighting. Many elections and votes operate in a binary fashion, where citizens propose policies that are accepted or rejected. Nevertheless, many decisions operate beyond this binary.

3.2 Voting within Democracy

One helpful distinction for navigating the interdisciplinary literature on liquid democracy is between voting and other government systems. The computational models only capture the voting mechanisms. A voting mechanism refers to the rules that determine winners of a collective decision. Democratic voting mechanisms usually follow a one person, one vote principle. Other standards are plurality rule (the option with the most votes wins), quota rule (the option with more than X% of the votes wins), and majority rule (the option with

more than 50% of the votes wins). In a direct democracy, decisions are made using majority rule on policies. While in a representative democracy, representatives are selected by citizens, and policies are selected by representatives using majority rule (or possibly higher quota rules, such as $\frac{2}{3}$ quota known as supermajority rule). In a liquid democracy, decisions are made using transitive proxy majority rule: the option with direct and transitive proxy votes whose weighted sum is greater than 50% of the total voters wins. The models only capture the voting aspects of democracy.

However, democracy also refers to any government structure where citizens hold equal rights and authority over their governing body. Voting is usually an essential mechanism for reaching collective decisions. This distinction and relationship between a voting mechanism and democracy is not explicitly made but helps avoid confusion. These computational models represent just the voting mechanisms, while many philosophical arguments address broader debates on government structures, where voting is usually integral. Still, other mechanisms besides voting can account for failures in decisions by majority rule: the absolute protection of fundamental rights, limiting the power of representatives, special protection for minorities and vulnerable populations. These additional features can alleviate some of the worries about democratic voting alone. For instance, the US Supreme Court functions to ensure all citizens are treated equally before the law and can overturn policy decisions from democratic elections. Assuming the Supreme Court functions as

intended, equality-based concerns need not be addressed by democratic voting at all. Nevertheless, democratic voting remains the primary means of reaching collective decisions in democracy.

3.3 Evaluating Democracy

The democratic models in this thesis are most compatible with the independent standard epistemic approach. In an epistemic approach, democracy derives value from generating good policies grounded in true beliefs[3]. In other words, citizens in a democracy consider issues, propose solutions, and determine each issue's best policies. The independent standard refers to the justification for the superiority of a policy being independent of the process. A similar assumption is present in the models: that one policy is correct and better than the other. Ideally, society would always decide on policies leading to better outcomes. Realistically, better outcomes are not known with certainty, but we can assume they exist. Democracy serves as a generally good truth-tracker taking various individual beliefs as input and discovering the truth through voting. This account can be less intuitive because democracy is commonly associated with reflecting the preferences of people and not the truth. However, both components seem to be necessary. The best outcomes likely balance voters' preferences and the truth. There are compelling reasons that the democratic process is effec-

[3] Note the distinction between a policy and a belief. Believes track truths in the world. Policies are practical responses to those truths. It is assumed that the best policies respond to true beliefs. In this way the policies are said to be truth-tracking.

tive at truth-tracking because it leverages many sources of information and diverse critical reasoning to enhance the collective ability to make decisions. The independent standard epistemic approach is instrumental because it looks at outcomes, as opposed to the process.

Another epistemic mode of evaluation is epistemic proceduralism(Estlund, 2009). This democratic ideal is partially compatible with the models. It is less demanding than the independent standard definition. The democratic goals for this theory are not about approximating a process-independent truth. In this view, democracy only needs to satisfy some minimal epistemic conditions. Universal personal competence in the models (all voters being better than a random) is an example of such a condition. Unlike the independent standard approach, any outcomes from this process are acceptable. The standards of the procedure justify the outcome. Consider legal rulings within courts. Judges do not necessarily track the truth. Lawyers present compelling evidence for both sides. Impartial bystanders listen in. There are many hours of work, preparation, deliberation, and gathering of evidence. These conditions create an epistemically sound procedure that justifies its outcomes. This epistemic approach is less related to these models because performance is not evaluated in relation to the outcomes. Yet, universal personal competence and the social competence could be good conditions for an epistemic procedure. However, epistemic proceduralism avoids the independent standard assumption, central to the models' definitions.

There are democratic concerns besides epistemic ones, namely egalitarian ones. In the

egalitarian view, democracy derives value from prioritizing the equality of all people. Egalitarian evaluations of democracy are concerned with fairness, equality, the protection of rights, and the aggregation of preferences. On egalitarian procedural grounds, some highlight fairness as something to be achieved by a procedure committed to securing citizens' equality[4]. In social choice theory, the egalitarian procedural or fairness approach investigates important properties of democratic voting. Usually showing that voters and opinions are rational and weighed equally. This type of evaluation could be done on these models, though it is not attempted in this thesis.

On instrumental egalitarian grounds, some argue that democracy derives its value from selecting decisions that maximize voters' preferences. The instrumental egalitarian approach is often related to utilitarianism. Every voter gets some utility if the collective decision aligns with their views. Democracy serves the purpose of a preference aggregator taking individual preferences as inputs and producing a collective preference output that maximizes utility. Theories of utility are on instrumental grounds because they focus on the outcomes as opposed to the process. One can evaluate the relationship between individuals and that collective decision to see how well a process aggregates preferences given certain constraints. For example, suppose most people prefer policy A to its alternatives.

[4]John Rawls' Theory of Justice (2009) offers such an account by describing a society of free citizens with equal rights in an economy that minimizes inequality. His guiding procedure for this account is the original position, a hypothetical state where citizens are blind to facts about themselves. He argues that policies resulting from this process would be just.

In that case, an equality-based preference aggregator should select A. Policy A's decision should remain when an equal number of people prefer A, regardless of the specific individuals. Some conclude certain forms of democracy are better than their alternatives using other such constraints or utility functions. The instrumental egalitarian and epistemic ideals seem exceptionally compatible with each other. A preference-based utility function could serve as part of the independent standard. In other words, the correct policy partially maximizes voters' preferences and partially maximizes truth-tracking. It seems compelling that the best policies would align with both voters' preferences and the truth. Though, unlike other preference aggregation models, the ones in this thesis do not provide metrics to measure performance in relation to voters' preferences alone.

Four democratic understandings have been presented in this section: epistemic instrumentalism, egalitarian instrumentalism, epistemic proceduralism, and egalitarian proceduralism. This thesis's evaluation and definitions are best attributed to the epistemic instrumental approach because of the correct policy assumption and outcome-orientated definitions of performance. Still, these models are not devoid of aspects of the other approaches. One, personal and social competence can serve as sufficient conditions for an epistemic proceduralist. Two, egalitarian procedural or fairness properties can be compared between the models. Three, the best policies likely align with citizens' preferences. Policy performance often requires citizens to cooperate, which is more likely when outcomes match their pref-

erences when possible. The truth-tracking and preference aggregation functions of democracy are compatible and could be defined together in the independent standard framework[5]. Presenting these democratic ideals adds context many models of democracy. The computational literature in liquid democracy is often unclear which one of these approaches, if any, is being taken. This clarification unifies the democracy literature across disciplines. In the next section, I explain voter competence in relation to the democratic ideals.

3.4 Revisiting Voter Competence

The models contain various types of voter competencies. The main two categories are personal competence and social voter competence. Personal competence is how well people vote, and social competence is how well people select others to vote on their behalf. In the models, personal competence is represented as a probability of voting for the correct policy. Social competence is represented by the delegation mechanisms that specify how voters choose to delegate voting power. The meaning of "voting well" is derived from some of the philosophical ideals mentioned in the previous section. In the independent standard epistemic framework, personal voter competency is one's likelihood of voting for the presumed correct outcome. Within the models, universal personal competence is a threshold of personal competence. Social competence is selecting other voters with high personal

[5] Some other models of liquid democracy capture the pure preference aggregation ideals better than those presented here(Christoff & Grossi, 2017).

competence to vote. Maximum delegation in liquid democracy is a mechanism that presupposes high social competence. So the personal and social notions of competence used in the models and simulations are mainly defined relative to the independent standard epistemic framework. Democratic egalitarians might be concerned that focusing on the truth ignores the importance of voters' preferences. However, with modifications, voter competence could be expanded to include preferences in an independent standard approach.

Voter competence concerning preferences aggregation is less intuitive than truth-tracking. Personal competence is how well one votes for their own preferences, and social competence is how well one selects others to vote for their own preferences. The preferential versions of voter competence are from a voter's subjective position, as opposed to objective truths. When done well, these tasks can maximize voters' preferences collective decisions. It is often strongly assumed that voters know their own preferences with great certainty. However, voters can vote against their interests, especially with limited information. In a binary decision, we can suppose voters' preferences are hidden, and their personal competencies are a probability that they vote for their hidden informed preference[6]. There could be inherent accuracy loss with preferential social competence, which entrusts individuals with many others' preferences. Delegation seems less useful under this framework because others are likely less accurate at voting according to another voter's preferences. Neverthe-

[6]This egalitarian approach has been modelled by Bloembergen et al. (2019) using a utility function where voters try to minimize the effort needed to increase their own preferential accuracy by delegating power when best to do so.

less, if the elections' issues are sufficiently complex, informed third-parties may vote better in line with their followers' preferences.

The models can encompass both voters' preferences and the truth. Understanding personal and social competence can reflect either an independent standard outcome in both senses. Voters can be competent individually or delegate to the best expert to maximize these goals. Personally competent voters vote for the best policies. These best policies are the ones that track-truth and align with their individual preferences. Socially competent voters delegate to others who vote on their behalf for policies that track-truth and align with their preferences. Combining the epistemic and egalitarian ideals within the independent standard model makes these models more robust. With this understanding of democratic ideals and voter competence, I will defend two related objections.

Firstly, can voters have strong social competence despite weak personal competence? The simulations assume this to be true and more. All voters are assumed to possess social competence in every case. Maximum delegation relies on high social competence. The skills required for this task are demanding: self-awareness, social-awareness, reasoning, and decision-making. These skills may seem unlikely when voters have low personal competence. However, these two types of competence are not necessarily related. One could have low personal competence in an area but be highly self-aware, socially-aware, and reasonable. This concept relates to a social epistemology problem on selecting experts[7]. There

[7] Mentioned in the philosophical literature by Estlund (2009) as the novice/expert problem.

are multiple reasons that one could expect voters to reliably select experts, despite being uninformed in that domain. Voters can analyze arguments, debates, consensus, qualifications, biases, and past performances of the experts without any area-specific expertise. These sources of information highlight that it is possible to have high social competence with low personal competence. However, it is not immediately convincing that high social competencies can be expected from everyone as is done in the models.

Secondly, why would gurus be periodically less personally competent than their followers? Ideally, those who are political gurus are positioned because of their relatively high competence. The simulations show that voters with high social competencies but extremely varied personal competency presents ideal scenarios for liquid democracy compared to both direct and representative democracy. In other words, liquid democracy performs well when the gurus routinely are worse than their followers. Suppose some elections in a democracy are primarily truth-tracking or preferential. Different conceptions of competence would be used between elections. In the truth-tracking election, it appears more likely that the well-informed gurus possess more truth-tracking personal competence than their followers at selecting the correct outcome. While in preferential elections, it appears likely that the average citizen possesses more preferential personal competence than their guru. Liquid democracy outperforms its alternatives when gurus are occasionally less personally competent than their followers because they do not delegate. A society where elec-

tions are both truth-tracking and preferential motivates this possibility.

4

Conclusion

The first chapter of this thesis introduces the binary election models of liquid democracy, direct democracy, and representative democracy by changing the delegation mechanism. The second chapter describes simulations' results using the models to show some interesting trends as populations grow. The results highlight the intuition that liquid democracy benefits from its flexibility, especially when popular experts are less personally competent

than average voters on occasion. The third chapter discusses some of the limitations and democratic ideals for understanding democracy and voter competence.

Discussing these models' limitations at length adds a unique philosophical contribution to modelling liquid demoracy. The voter independence and the star social network are the most troubling assumptions. Voter independence causes the extreme CJT phenomena. Simultaneously, the direct democracy model does not use the social network in the direct voting mechanism. These cause the performance to tend towards extremes when introducing voter correlation would minimize this trend. The star social network is intuitive but should be substituted with a more convincing representation of society. The model of representative democracy can be significantly improved to differentiate it from that of liquid democracy. It is especially pertinent to press on the different types of social competencies in selecting policy-specific or general-purpose experts. More work should be done to confirm the relationships between voter competence and performance identified in these models.

There are numerous assumptions identified in this thesis that could account for when liquid democracy leads to better outcomes than both direct democracy and representative democracy. Voting is assumed to be essential in determining outcomes. It is assumed that there exist independently correct and preferable outcomes. Both truth-tracking and preferential concerns are compatible for an independent standard assumption. The delegations

mechanisms assume high social competence despite low person competence. The star social networks and different voter groups are only specific examples of social structures. Elections are only binary policy decisions. Voters vote with independence. There are some cases when followers are more competent than gurus and vice versa. Under all these constraints, the simulations show the performance of liquid democracy is relatively highest in a series of elections with high social competencies and varied personal competencies. The primary cause is the flexibility of liquid democracy's delegation mechanisms. This result highlights the effect of voter competence on liquid democratic performance, which can guide further research in democracy.

References

Behrens, J. (2017). The origins of liquid democracy. *The Liquid Democracy Journal*, 5(2), 7–17.

Behrens, J., Kistner, A., Nitsche, A., & Swierczek, B. (2014). *The principles of LiquidFeedback*. Interaktive Demokratie e. V. Berlin.

Bloembergen, D., Grossi, D., & Lackner, M. (2019). On rational delegations in liquid democracy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33 (pp. 1796–1803).: aaai.org.

Blum, C. & Zuber, C. I. (2016). Liquid democracy: Potentials, problems, and perspectives. *The journal of political philosophy*.

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

Brill, M. & Talmon, N. (2018). Pairwise liquid democracy. In *IJCAI*, volume 18 (pp. 137–143).: pdfs.semanticscholar.org.

Caragiannis, I. & Micha, E. (2019). A contribution to the critique of liquid democracy. In *IJCAI* (pp. 116–122).: ijcai.org.

Christoff, Z. & Grossi, D. (2017). Binary voting with delegable proxy: An analysis of liquid democracy.

Colley, R., Grandi, U., & Novaro, A. (2020). Smart voting. *irit.fr*.

De Cindio, F. & Stortone, S. (2013). Experimenting liquidfeedback for online deliberation in civic contexts. In *Electronic Participation* (pp. 147–158).: Springer Berlin Heidelberg.

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

Escoffier, B., Gilbert, H., & Pass-Lanneau, A. (2019). The convergence of iterative delegations in liquid democracy in a social network. In *Algorithmic Game Theory* (pp. 284–297).: Springer International Publishing.

Estlund, D. (2009). *Epistemic Proceduralism and Democratic Authority*, (pp. 15–27). Springer Netherlands.

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

Grofman, B., Owen, G., & Feld, S. L. (1983). Thirteen theorems in search of the truth. *Theory and decision*, 15(3), 261–278.

Gölz, P., Kahng, A., Mackenzie, S., & Procaccia, A. D. (2018). The fluid mechanics of liquid democracy. In *Web and Internet Economics* (pp. 188–202).: Springer International Publishing.

Hardt, S. & Lopes, L. C. R. (2015). Google votes: A liquid democracy experiment on a corporate social network.

Kahng, A., Mackenzie, S., & Procaccia, A. D. (2018). Liquid democracy: An algorithmic perspective. *AAAI 2018*.

Kling, C. C., Kunegis, J., Hartmann, H., Strohmaier, M., & Staab, S. (2015). Voting behaviour and power in online democracy: A study of liquidfeedback in germany's pirate party.

Kotsialou, G. & Riley, L. (2018). Incentivising participation in liquid democracy with breadth-first delegation.

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

Rawls, J. (2009). *A Theory of Justice*. Harvard University Press.

Tullock, G. (1967). *Toward a Mathematics of Politics*. University of Michigan Press.