How Should We Vote? A Comparison of Voting Systems within Social Networks

Shiri Alouf-Heffetz[1,2] , **Ben Armstrong**[3] , **Kate Larson**[3] and **Nimrod Talmon**[2]
[1]TwoFiveOne Research
[2]Ben-Gurion University
[3]University of Waterloo
shirihe@post.bgu.ac.il, ben.armstrong@uwaterloo.ca, kate.larson@uwaterloo.ca, talmonn@bgu.ac.il

Abstract

Voting is a central methodology for eliciting and combining agents' preferences and information across many applications. Just as there are numerous voting rules exhibiting different properties, we also see many different voting *systems*. In this paper we investigate how different voting systems perform as a function of the characteristics of the underlying voting population and social network. In particular, we compare direct democracy, liquid democracy, and sortition in a *ground truth* voting context. Through simulations – using both real and artificially generated social networks – we illustrate how voter competency distributions and levels of direct participation affect group accuracy differently in each voting mechanism. Our results can be used to guide the selection of a suitable voting system based on the characteristics of a particular voting setting.

1 Introduction

One person, one vote; is this the best way to run elections? This paper explores whether new methods can lead to high quality election outcomes while requiring less effort from voters. Paradigms of voting such as *liquid democracy* and *sortition* aim to apply modern understanding to centuries-old social choice methods. Liquid democracy [Blum and Zuber, 2016] tries to make use of expert knowledge by allowing voters to transitively delegate their vote on issues they feel less informed of or less interested in. Sortition [Dowlen, 2017] involves selecting a subset of voters to participate in the election to reduce the effort required of the total voter population while maintaining desired properties such as fairness or representation of subgroups [Flanigan *et al.*, 2021]. In both systems only a subset of voters are fully active in the election, in contrast to direct democracy where all voters participate equally.

We examine these paradigms within a binary *ground truth* setting. That is, voters are asked to choose between two options – one correct and one incorrect – and all wish to choose the correct option. An election might be any question being asked of a group where an objectively measurable truth exists, such as a trial where a jury is asked whether a defendant is innocent or guilty [Forsyth, 1852], or an investors meeting where attendants are asked whether a stock will have a higher or lower price 30 days in the future. In some cases, voters may be unsure of, e.g., whether a stock price will rise, and will prefer to rely on someone whose expertise they trust by delegating their vote. In other cases, an organization selecting participants in an election from a larger pool may know that some voters bring more expertise than others and will try to select such voters, as in jury selection.

These settings are examples of group decision-making. In many such settings – particularly in rapidly growing online communities – attention is a scarce resource. Communities need to find methods of coping with the increased burden of governance as they grow. Liquid democracy and sortition allow the community to balance the number of individuals making participating with the attention the group is willing to put towards governance.

Voters in the above example situations have their own knowledge, understanding, and biases that affect the likelihood that they will select the correct outcome. It is natural to model this by viewing each participant as able to contribute some noisy evaluation of the ground truth. Understanding the dynamics that affect the accuracy of an entire group's decision are key to both anyone running an election and anyone participating in such an election.

Some papers in the literature have provided primarily theoretical results showing, e.g., the conditions under which liquid democracy *can* reliably improve the decision-making ability of a group [Kahng *et al.*, 2018], or that perfectly optimizing delegations to maximize the ability of a group to identify ground truth is NP-Hard [Caragiannis and Micha, 2019]. Here we focus on understanding the practical factors affecting the accuracy of the decision-making ability of a group that employs either liquid democracy or sortition. To allow for a more intuitive understanding, and due to the difficulty of analytically analysing the situations we are interested in, we focus our efforts on computer-based simulation results and expand upon our previous work [Alouf-Heffetz *et al.*, 2022].

The primary question we seek to answer is: *What factors affect the quality of the group decision when using liquid democracy or sortition?* To answer this question, we explore a model in which voters exist within a social network. In modern times, the exponential increase in size and prevalence of social networks allows us to consider the social structure as a given fact for many social choice settings. Additionally, we use a single parameter that denotes the voters' *competency*, which is their intrinsic probability of selecting the correct outcome. Within this model we compare the ability of direct democracy, sortition, and liquid democracy to uncover ground truth.

Our results demonstrate that achieving better accuracy than direct democracy is achievable as long as there is some method to approximate the relative competence of voters, in which case, it is possible to run a highly accurate election where only a fraction of voters must actively participate. Our experiments found consistent results across both artificial and real social networks. Furthermore, we show there is no clear correlation between the properties of a social network and group accuracy. Lastly, we find that a version of simulated annealing for Liquid Democracy approaches the optimal distributions of delegations among voters. This was not easy to foresee as finding optimal delegations is NP-Hard, therefore it seems that the annealing algorithms rather quickly find states that yield much higher accuracy than other forms of liquid democracy.

The remainder of the paper is organized as follows. This section finishes with a look at the related work while Section 2 describes our model in greater depth. Section 3 introduces our experimental design including the specific voting rules we explore and the parameters we analyze. We describe the experiments that we run and discuss their results in Section 4. Section 5 concludes and briefly describes possible future work.

1.1 Related Work

Voting systems. Forms of direct democracy and sortition [Sintomer, 2018] were originally practiced well over 2000 years ago in Athens [Tridimas, 2017]. Liquid democracy is comparatively recent, having roots in Charles Dodgson's work in the 19th century [Dodgson, 1884]. Most recent works on these voting systems in the AI and computational social choice communities has focused on worst case performance scenarios or improving their fairness properties (see [Caragiannis and Micha, 2019], [Escoffier *et al.*, 2018], and [Kahng *et al.*, 2018]).

Liquid democracy. In recent years liquid democracy has been adapted for use by several political parties. This has been most well documented in the case of the LiquidFeedback software [Kling *et al.*, 2015]; a recent overview is also found in [Paulin, 2020]. Liquid democracy [Blum and Zuber, 2016] is usually studied for settings with a ground truth, in which its worst-case performance is theoretically analyzed. Kahng *et al*. [2018] have shown that, to guarantee a superior group accuracy from liquid democracy than from direct democracy, voters must have access to non-local information. Caragiannis and Micha [2019] have shown that finding optimal delegations is NP-Hard. These hardness results were strengthened by Becker *et al*. [2021], who presented simulations showing that, under particular settings, various delegation methods are able to lead to a more accurate result than direct democracy. They also apply an optimization process to delegations to achieve a similar result as our simulated annealing, albeit without exploration of the similarity of the result to optimal weights. Bloembergen *et al*. [2019] have considered a game-theoretic model of liquid democracy where each voter has a "type" and competency is construed as the ability of a voter to identify and communicate their own type.

Sortition. Sortition is typically studied in a setting focused on aspects other than decision accuracy. In particular, Benade *et al*. [2019] have shown that sortition guarantees fair representation of subsets of a population. Similarly, Flanigan *et al*. [2021] developed several algorithms that select sortition participants to optimally balance between equal selection probability for each voter and accurate representation of the underlying population. We primarily focus on sortition as a voting system that is easy to implement but, since it requires only a subset of the voters to actively participate in the election, minimizes the collective effort of voting.

Voter competence. Relatively little work has focused on understanding the effects of voter competency distributions on the outcome of elections. Grofman [1983] studied direct democracy in a binary ground truth setting; several of their theorems can be viewed as extensions of the Condorcet Jury Theorem. Nitzan and Paroush [2017] summarize the long history of research into jury theorems.

2 Model

Competence networks. We consider a basic social choice setting with n agents, $V = \{v_1, \ldots, v_n\}$, and two alternatives $A = \{a^+, a^-\}$. Our focus is on a *ground truth* voting scenario where a^+ is the objectively correct outcome that the agents collectively aim to elect. However, each voter $v_i \in V$ has a competency level $q_i \in [0, 1]$ that corresponds to the probability v_i would vote "correctly" (i.e., would vote for a^+).[1]

The agents are connected via an underlying social network, thus we have a set of undirected edges E, representing the connections between the agents. Combining the ingredients of our setting, our basic mathematical object is a so-called *competence network*, which is an undirected labeled graph $G = (V, E)$, where $V = \{v_1, \ldots, v_n\}$, and each vertex $v_i \in V$ acts as voter and is labeled with its competence value q_i. We thus use *agents*, *voters*, and *vertices* interchangeably.

Voting Systems. A *voting system* defines a process that operates on a competence network and defines how the agent groups make their joint binary decision. Formally speaking, a voting system takes as input a competence network and outputs a winner, which is either the correct alternative a^+ or the incorrect alternative a^-. As we always have exactly two alternatives we use a weighted majority voting system. Each voter that votes in an election commits their weight to a single alternative and the alternative with the most weight is the winner.

[1]Note that we do not require $q_i \geqslant 0.5$. Caragiannis *et al.* [2019] point out that $q_i \to 0$ might indicate a strongly held incorrect belief. It is possible that a voter with $q_i < 0.5$ prefers to delegate to a voter with even lower competency, thus inducing an echo-chamber dynamic. However, we agree with the view proposed by Becker *et al.* [2021] that there exist objective indications of competency that all voters recognize.

Elections. An *election* is a competence network paired with a voting system. When the election is run the voting system is used on the competence network and results in a winning alternative. We refer to the probability that a^+ will be chosen as the *accuracy* of the election.

3 Experimental Design

In this section we describe our experimental setup and the details of each voting system we use. In particular, the steps we follow in each experiment are as follows:

1. We create the underlying network that connects the agents.
2. We assign competence levels to the voters in the network.
3. We choose and use a voting system.
4. As our setting is probabilistic, we estimate the outcome of the election.

3.1 Voting Systems

We consider several variants of sortition and liquid democracy, as well as direct democracy, which we use as a benchmark. Note that in direct democracy all voters actively participate in the election by voting. In contrast, in sortition and liquid democracy only a subset of the voters actively vote.

Definition 1. An *active voter* is one whose choice of alternative is considered when computing the winner of the election. The set of active voters is denoted V^{active}.

Direct democracy. In *direct democracy*, all voters actively vote. That is, in direct democracy, every voter v_i votes with equal weight for either a^+ or a^-, where v_i chooses a^+ with probability equal to q_i. The winning alternative is the one that receives the most votes. Throughout the paper, ties are broken in favour of a^+.

Sortition. In sortition, only a subset of the voters are active and all those are given equal weight while all other voters do not participate in the election at all and do not affect the outcome. There are many variants of sortition, corresponding to the method by which the set of active voters is selected. Here we consider sortition as it might be used for a task such as jury selection, where some individuals could be considered more competent than others, rather than the sortition used in citizen's assemblies where the goal is to most accurately represent the underlying population. Sortition works best if the set of active voters are exactly those voters who are the most competent. As it may not be possible to accurately estimate the competence levels of voters, we consider the following, parameterized variant.

Definition 2. *ρ-noisy sortition* is a method of sortition in which all voters are sorted in descending order of competence. Then, beginning from the most competent voter, in each iteration $i \in [1, ..., n]$, with probability $\rho \in [0, 1]$, the i^{th} voter is swapped with one of the voters from the $\{i + 1, .., n\}$, chosen uniformly at random. Finally, the first k voters are selected as the set of active voters (and other voters are disregarded).

Observe that, for $\rho = 0$, the ρ-noisy sortition variant is equivalent to selecting the k most accurate voters (i.e., pure meritocracy); while, for $\rho = 1$, it is equivalent to selecting k voters uniformly at random (i.e., pure sortition).

Liquid democracy. In liquid democracy, like in sortition, there is a set of active voters; however the other voters are not disregarded. Each voter chooses whether to actively vote or to delegate to another voter of their choice, and these delegations are transitive (e.g. If v_1 delegates their vote to v_2 then v_2 can vote with a weight of 2 or can instead delegate their vote as well as v_1's vote to v_3.). Each active voter has a set of voters who delegate (either directly or transitively) to them. In contrast to sortition, the voting weight of each active voters is not the same, but is equal to the number of delegations they receive (either directly or transitively), plus one (for their own vote).

In our simulations, each voter has an equal chance of being chosen to be an active voter, based upon an experiment parameter. For the delegations of the nonactive voters, we use the following variants:

- *Liquid Better*: For each nonactive voter v_i, we select one of their neighbours uniformly at random, among the neighbours with higher competence than v_i. That is, v_i delegates to a random voter from the set $\{j \in N_G^+(i) \,|q_j > q_i, (i, j) \in E\}$ where $N_G^+(i) = \{j \in V \mid (i, j) \in E\ ,\ q_j > q_i\}$.
- *Liquid Max*: For each nonactive voter v_i, we select their delegate to be their most competent neighbour.

In liquid democracy computing the delegations that result in the highest probability of the resulting group decision being correct is NP-hard [Caragiannis and Micha, 2019]. For this reason, we also consider a heuristic approach that takes a competence network as input and aims to find the accuracy-maximizing partition of voters into active/nonactive sets and set of delegations for those nonactive voters. Our heuristic is based on the popular local search heuristic of simulated annealing (SA).

In our implementation of SA, we begin with either existing delegations (Experiments 1 and 2), or a state with no delegations (Experiment 4), then make a single, randomly chosen, local delegation change at each time step. If the delegation leads to a higher group accuracy that state is accepted, otherwise the state is accepted probabilistically based on the difference between accuracies of the old and new states. If s_i is the state after i steps of annealing and $s_{current}$ the current state, the probability of accepting a new state is $exp(-(s_{current} - s_i)/T)$. In all our experiments we hold T constant at 1.0 and use the *basinhopping* function within the SciPy library to implement simulated annealing [Virtanen *et al.*, 2020]. For a detailed overview of simulated annealing see the overview in [Dowsland and Thompson, 2012] or [Wales and others, 2003] for more detail on basin-hopping.

3.2 Network Generation

We use both real-world and artificially-generated networks.

Real-world networks. We have explored an array of real-world networks from the SNAP database [Leskovec and

Krevl, 2014] and included a real-world instance that is a sampled network of Facebook users with 534 nodes and 9626 edges.

Artificial networks. We use the following two classical probabilistic models for generating artificial networks, which are often used to replicate properties of real-world social networks (see, e.g., the overview of Kleinberg [2010]): (1) *Erdős–Rényi (ER)*: ER networks are parameterized by two parameters—the number n of nodes and the connection probability p—and are generated by initializing an empty network of n nodes and inserting each possible edge with the fixed connection probability p, independently [Erdos *et al.*, 1960]; and (2) *Barabási–Albert (BA)*: BA networks are parameterized by one parameter m and are generated via a preferential attachment process where nodes are added one at a time and connected to m of the existing nodes where the probability of an existing node to be connected to the current node is proportional to its degree [Albert and Barabási, 2002].

3.3 Competence Assignment

Given a network $G = (V, E)$ generated via one of our network generation processes, we use different probability distributions to assign competence levels to the voters $V = \{v_1, \ldots, v_n\}$ in the network.

In particular, in our experiments we select some distribution $\mathcal{D}$ from which we draw voter competencies, so that for each $v_i \in V$ we sample a value $q_i \sim \mathcal{D}$. That is, the competence levels are chosen independently for each voter. In particular, we consider the following families of distributions:

- uniform distributions;
- truncated normal distributions;[2] and
- truncated exponential distributions.[3]

3.4 Accuracy Calculation

Recall that the accuracy of an election is the probability that the correct alternative a^+ will be selected. The accuracy of a given election can be computed using dynamic programming with a table of size $O\left(n^2\right)$ [Becker *et al.*, 2021]. However, our preliminary experiments with this algorithm found it expensive in memory and time (particularly with well connected networks of several hundred nodes or more). Therefore in our experiment we use an approximation algorithm, based on Monte Carlo simulation [Mooney, 1997]. Concretely, we proceed by performing 1000 iterations. In each iteration we sample the votes of the voters using their competence values (so that a voter v_i with competence value q_i is voting for a^+ with probability q_i), and then use the voting system with these votes. Finally, we take the fraction of iterations resulting in a^+ winning as our estimate of the accuracy of the election. This method proved to be computationally efficient and highly accurate in estimating the accuracy of our elections.

[2]We use the SciPy implementation of the truncated normal distribution [Burkardt, 2014].

[3]Since the exponential distribution does not provide an upper bound on sampled values whenever we sample a competency value greater than 1 we map the value to 1. This leads to a mean value slightly lower than the original distribution, i.e $\mu' = \frac{1}{\lambda} - \frac{e^{-\lambda}}{\lambda}$.

4 Results and Discussion

Here we describe the specific experiments we have performed. For each experiment, we describe its experimental design, show its results, and discuss them. The goal of our experiments is to identify the conditions under which each voting system is most useful. That is: are there particular competency distributions, network structures, or levels of voter participation where any system give higher or lower accuracy?

Our simulation code is written in Python and run using Python 3.8. We use Matplotlib [Hunter, 2007] to visualize our data, and Scikit [Virtanen *et al.*, 2020] for our statistical analysis. In all our experiments we study voters with an average competency of 0.5 in order to demonstrate the ability of liquid democracy and sortition to utilize expert knowledge.

4.1 Experiment 1: Fraction of Active Voters

In this experiment we concentrate on how the number of active voters affects the accuracy of the group decision. We consider networks of 100 and 1000 voters, all with an average degree of approximately 20. For BA networks, we set $m = 10$ and in ER networks $p = 0.20202$ ($p = 0.02002$) for networks of 100 (1000) voters. Voter competencies are sampled as follows for each distribution:

(i) Uniform - $\forall i : q_i \sim U\,(0.3, 0.7)$

(ii) Gaussian - $\forall i : q_i \sim \mathcal{N}\,(0.5, 0.1)$

(iii) Exponential - $\forall i : q_i \sim Exp\,(2)$

We use $\rho = 0.1$ for our ρ-noisy sortition model and report the average result of 10 trials.

In each individual plot of Figure 1, which shows the results of this experiment, the X-axis shows the portion of *active voters* used, ranging from 10% to 90%. The Y-axis shows the resulting group accuracy.

Discussion. Whenever the fraction of active voters is low (corresponding to the left side of the plots in Figure 1), the possible variance of delegation is higher as there are more nonactive voters. Interestingly, whenever more than half of the voters delegate (i.e., when the fraction of active voters is less than half) the accuracy achieved by *Liquid Better* is very close to that of *Liquid Max*. This means that even though *Liquid Better* only requires the nonactive voters to find a delegate who is more competent than themselves, it still is able to elicit the "experts" of the network. Observe also that both *Liquid Better* and *Liquid Max* achieve group accuracy levels very close to that of the SA solution. Moreover, all methods are significantly better than direct democracy. Lastly, note that due to the asymmetry of exponential distribution the differences described above are magnified. More than half of the voters have a competence below 0.5 so direct democracy approaches an accuracy of 0 when the number of voters increases and the room for improvement from delegation or sortition increases dramatically.

4.2 Experiment 2: Variance of Competencies

Here we consider the effect that the variance of voter competence has on the group accuracy. In particular, we keep the mean competency constant at 0.5 and slowly increase the

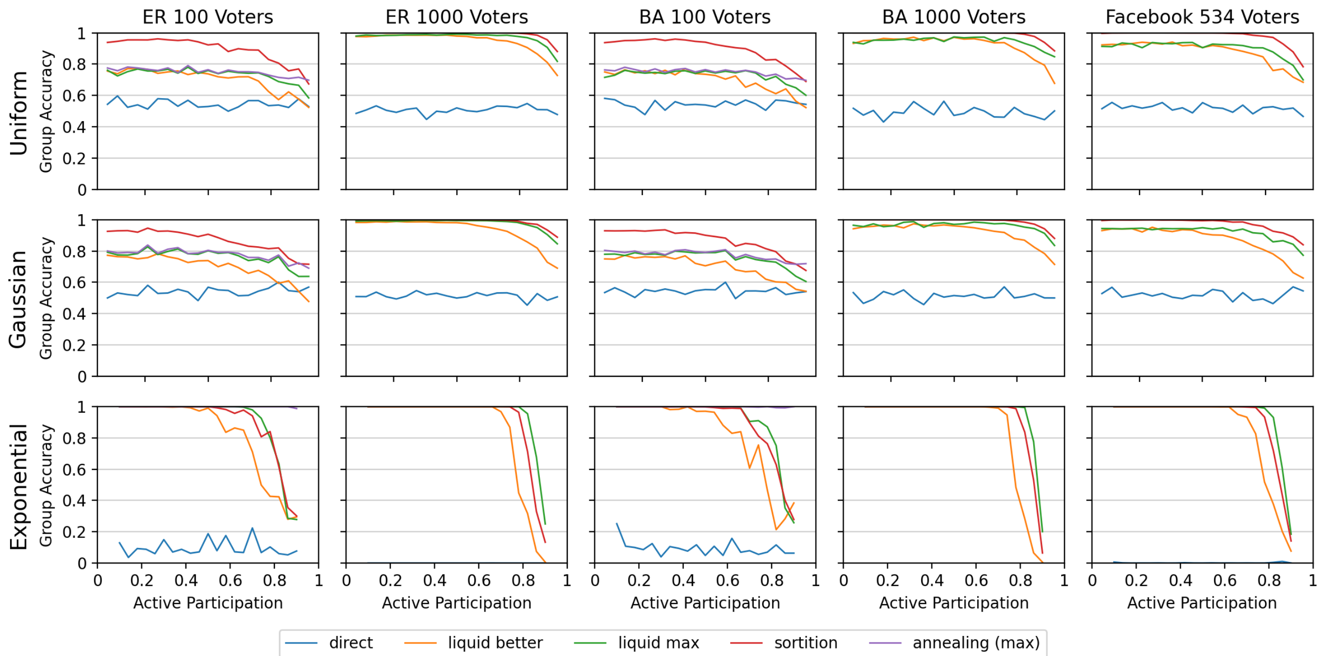

Figure 1: Results of Experiment 1 showing the decrease in accuracy (y-axis) as a higher fraction of voters become active participants (x-axis). The relative accuracy of each voting method is constant between network size and type but accuracy changes between competency distributions. When direct democracy is not visible it has an accuracy of approximately 0. Note that annealing results are only shown for 100 voters.

variance. The results of this experiment are shown in Figure 2. Note that, as the exponential distribution accepts only a single parameter, that affects both the mean and the variance, in this experiment we consider only the uniform and the Gaussian distributions. We consider 20 uniform distributions with upper and lower bounds evenly-spaced between $U(0.475, 0.525)$ to $U(0, 1)$, and 20 Gaussian distributions with $\sigma \in \{0.05, 0.5\}$. We average all results over 10 trials. In all cases 90% of voters are active.

Discussion. Figure 2 shows, particularly with more voters present, that having even just 10% of voters delegating or left out by sortition can lead to a dramatic increase in accuracy. This suggests that indirect voting methods are very well suited to situations where expert knowledge exists; as the higher the competence variance is the higher the number of voters with very high competency.

4.3 Experiment 3: Optimizing Delegations and Active Voter Weight

Here we investigate the performance of simulated annealing over time. In particular, we first run the heuristic and record the accuracy it obtains as more iterations are being made.

We also complement our analysis by comparing the specific vote weights that result from using the heuristic to a theoretical result from the literature. In particular, Grofman *et al.* [1983, Theorem XIII] consider weighted direct democracy and show that in a binary ground-truth setting such as in ours, in order to maximize the accuracy of direct democracy, the weight of each voter v_i with competence level q_i shall be proportional to $log(\frac{q_i}{1-q_i})$. We thus set to investigate whether our SA heuristic find delegations that result in vote weights of the active voters that are getting closer to the theoretical optimum of Grofman *et al.* To this end, we proceed by examining the set of active voters selected by SA, each 10 iterations of SA, as follows:

(i) Calculate the normalized optimal weight $\frac{1}{n}log(\frac{q_i}{1-q_i})\forall i \in V^{active}$. Sort the results in descending order and denote them by w^{opt}.

(ii) Define w^* as the normalized actual weights of all active voters induced by SA and sort them in descending order.

(iii) Finally, calculate the $L2$ distance between the two weight vectors, $L2 = \sqrt{\Sigma_{i \in V^{active}} \left(w_i^{opt} - w_i^*\right)^2}$.

The results of this experiment are shown in Figure 3

Discussion. By comparing to the optimal weights described by Grofman *et al.*, this experiment indicates how well annealing does at finding optimal delegations. When the L2 distance from optimal weights is low the delegations are more closely approximating the optimal weight distribution among active voters. And, indeed, we do see in Figure 3 that the best solution found by annealing approaches to the optimal weights of Grofman *et al.* [1983]. Interestingly, the figure also shows that there are delegations found by annealing with high accuracy that are very dissimilar to the optimal weight distribution.

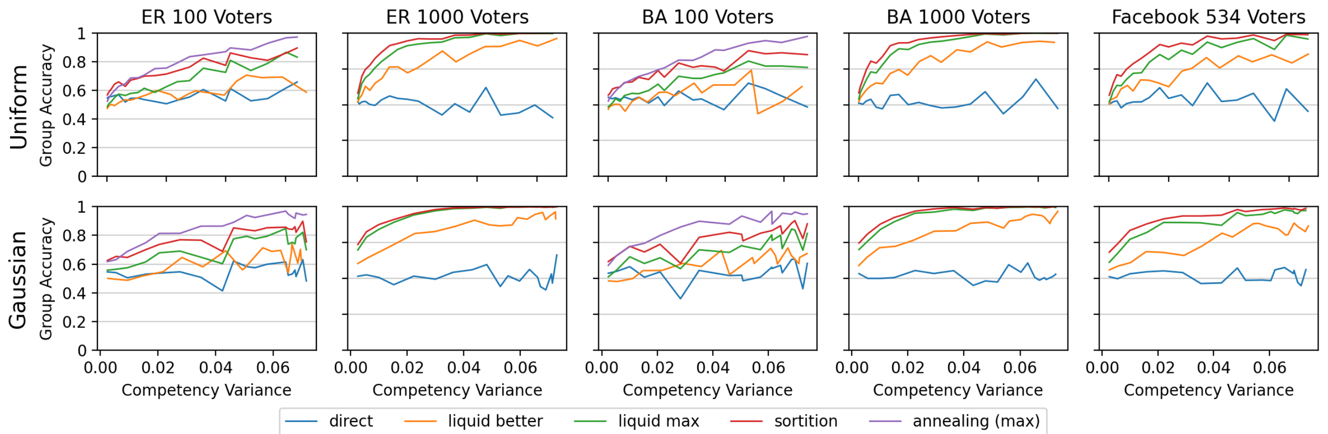

Figure 2: Results of Experiment 2 showing the increase in accuracy from delegation and sortition when competency distributions have larger variance. Note that annealing results are only shown for networks with 100 voters.

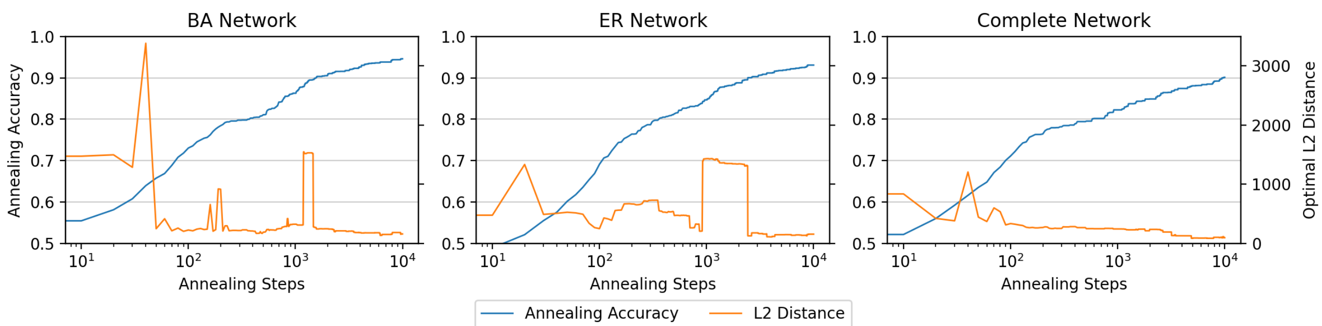

Figure 3: Results of simulated annealing in Experiment 3. Annealing quickly reaches high accuracy while the L2 distance from the optimal weight distribution shrinks.

4.4 Experiment 4: Sortition

To better understand the effect of different values of ρ in our model of sortition we ran an experiment considering several competency distributions over a wide range of ρ. These included two Gaussian distributions with $\sigma = 0.1$ and 0.5, two uniform distributions ranging from $0.4 - 0.6$ and $0.1 - 0.9$, and an exponential distribution created as in previous experiments. As in previous experiments, all distributions have a mean competency of 0.5. We ran simulations with these values and examined the outcome of sortition when selecting 10%, 25%, and 50% of the voters to actively participate. Figure 4 presents the results of this experiment.

Discussion. In general, the results are approximately as expected: As ρ increases, the most competent voters are more likely to be swapped for less competent voters and accuracy decreases. Interestingly, different distributions respond differently to changes in the fraction of active voters. For example, when ρ is low, the low variance uniform distribution is more accurate when 50% of voters are active than when 10% are active. In contrast, as ρ increases, the exponential distribution loses accuracy more quickly with more active voters. These differences reflect the relative difference between competent and incompetent voters in each distribution. An additional voter from the low variance uniform distribution (which has a minimum competency of 0.4) can do much less "harm" than an additional voter from the exponential distribution (which has a minimum competency of 0).

4.5 Experiment 4: Network Properties

We now set to examine how different network properties affect accuracy. To this end, we have fixed the competency distribution and the chance of active participation and ran elections on each network type. Then, we calculated the following properties: mean degree, mean neighbour degree, connectivity, clustering coefficient, radius, and diameter. Figure 5 show the comparison between the average degree of voters in an ER network and group accuracy, using a Gaussian competency distribution with $\mu = 0.5$, $\sigma = 0.2$. We considered 99 ER networks with attachment parameters uniformly distribution between 0.1 and 0.99 and each data point represents the mean of 30 elections. We also performed similar experiments on BA networks; all experiments for each property found results similar to those in Figure 5.

Discussion. None of the network properties we considered exhibited a strong relationship with group accuracy. We only observed a very small drop in the accuracy of liquid democracy in some extremely sparse networks. Of course, as with

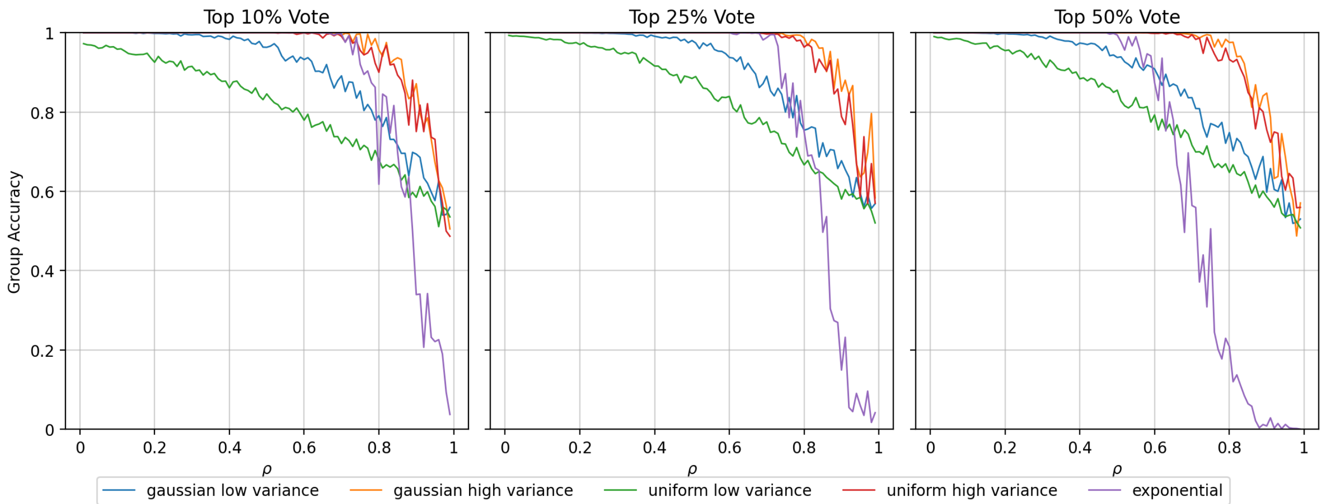

Figure 4: Results of an experiment varying the value of ρ in ρ-sortition. The three plots show group accuracy across three different levels of voter participation for five competency distributions, each with a mean of 0.5.

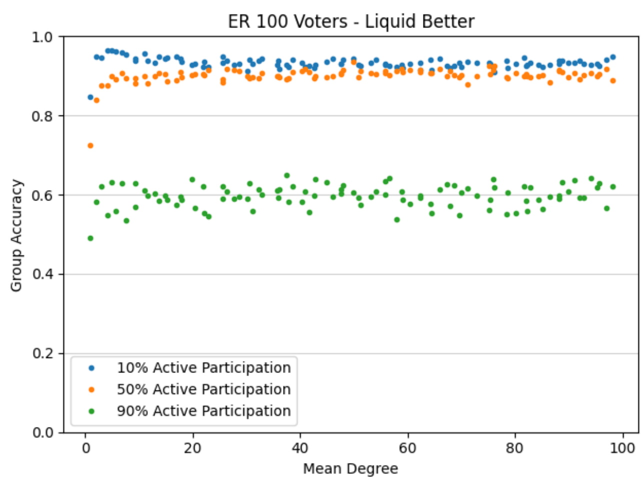

Figure 5: The accuracy of the Liquid Better delegation mechanism on Erdős–Rényi networks with a range of attachment parameters. Group accuracy does not tend to change based on the average degree of voters.

previous experiments, sampling from different competency distributions on the same network affects the accuracy but such variations appear unrelated to any network parameter within the family of ER and BA networks we examined. We believe we did not observe any relationship between network properties and accuracy because only a bare minimum of connectivity is required for delegation to be maximally effective. This strongly suggests that any real social network structure would not be a hindrance to the benefits of liquid democracy.

4.6 Summary

The overarching message from our experiments is that if voter competency can be identified then direct democracy is unlikely to be the most accurate decision-making mechanism. When a coordinator is able to identify competency with high accuracy sortition is an excellent tool and liquid democracy performs well even when voters are only able to recognize another voter's relative competence. Liquid democracy has proven itself robust to variations in network structure as long as the network is sufficiently connected and we have also seen its ability to effectively utilize expert knowledge when a majority of voters have low competency..

5 Conclusion

We started this paper with the question "One person, one vote; is this the best way to run an election?". In the context studied here we believe that our findings conclusively show that the answer to this question is "No". All of our results combine to show that liquid democracy and sortition are both highly effective tools for uncovering ground truth. Furthermore, we believe that our experiments also serve to demonstrate that these social choice paradigms are well suited for empirical use across many settings.

Our work opens up many avenues for future research. Our results were empirical and formalizing the relationship, for example, of the trade-off between accuracy and the portion of active voters could provide deeper insight into the problem. Furthermore, we assumed that voters with $q_i < 0.5$ will choose more competent delegates. However, in the face of disinformation this may not be the case, and exploring the resulting echo-chamber dynamics and methods for combating such an issue may prove important.

References

[Albert and Barabási, 2002] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47, 2002.

[Alouf-Heffetz *et al.*, 2022] Shiri Alouf-Heffetz, Ben Armstrong, Nimrod Talmon, and Kate Larson. Seeking ground truth: A simulation-based exploration of liquid democracy and sortition on social networks. *Fourth Games, Agents, and Incentives Workshop*, 2022.

[Becker *et al.*, 2021] Ruben Becker, Gianlorenzo D'Angelo, Esmaeil Delfaraz, and Hugo Gilbert. When can liquid democracy unveil the truth? *CoRR*, abs/2104.01828, 2021.

[Benadè *et al.*, 2019] Gerdus Benadè, Paul Gölz, and Ariel D Procaccia. No stratification without representation. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 281–314, 2019.

[Bloembergen *et al.*, 2019] 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.

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

[Burkardt, 2014] John Burkardt. The truncated normal distribution. *Department of Scientific Computing Website, Florida State University*, pages 1–35, 2014.

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

[Dodgson, 1884] Charkes L. Dodgson. *The principles of parliamentary representation*. Harrison and Sons, 1884.

[Dowlen, 2017] Oliver Dowlen. *The political potential of sortition: A study of the random selection of citizens for public office*, volume 4. Andrews UK Limited, 2017.

[Dowsland and Thompson, 2012] Kathryn Anne Dowsland and Jonathan Thompson. Simulated annealing. *Handbook of natural computing*, pages 1623–1655, 2012.

[Erdos *et al.*, 1960] Paul Erdos, Alfréd Rényi, et al. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60, 1960.

[Escoffier *et al.*, 2018] Bruno Escoffier, Hugo Gilbert, and Adèle Pass-Lanneau. The convergence of iterative delegations in liquid democracy. *CoRR abs/1809.04362*, 2018.

[Flanigan *et al.*, 2021] Bailey Flanigan, Paul Gölz, Anupam Gupta, Brett Hennig, and Ariel D Procaccia. Fair algorithms for selecting citizens' assemblies. *Nature*, 596(7873):548–552, 2021.

[Forsyth, 1852] William Forsyth. *History of trial by jury*. JW Parker, 1852.

[Grofman *et al.*, 1983] Bernard Grofman, Guillermo Owen, and Scott L Feld. Thirteen theorems in search of the truth. *Theory and decision*, 15(3):261–278, 1983.

[Hunter, 2007] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.

[Kahng *et al.*, 2018] Anson Kahng, Simon Mackenzie, and Ariel D Procaccia. Liquid democracy: An algorithmic perspective. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[Kleinberg, 2010] Jon Kleinberg. Ect volume 26 issue 5 cover and back matter. *Econometric Theory*, 26(5):b1–b4, 2010.

[Kling *et al.*, 2015] Christoph Carl Kling, Jérôme Kunegis, Heinrich Hartmann, Markus Strohmaier, and Steffen Staab. Voting behaviour and power in online democracy: A study of liquidfeedback in germany's pirate party. *arXiv preprint arXiv:1503.07723*, 2015.

[Leskovec and Krevl, 2014] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. http://snap.stanford.edu/data, June 2014.

[Mooney, 1997] Christopher Z Mooney. *Monte carlo simulation*. Sage, 1997.

[Nitzan and Paroush, 2017] Shmuel Nitzan and Jacob Paroush. Collective decision-making and jury theorems. In *The Oxford Handbook of Law and Economics*. 2017.

[Paulin, 2020] Alois Paulin. An overview of ten years of liquid democracy research. In *The 21st Annual International Conference on Digital Government Research*, pages 116–121, 2020.

[Sintomer, 2018] Yves Sintomer. From deliberative to radical democracy? sortition and politics in the twenty-first century. *Politics & Society*, 46(3):337–357, 2018.

[Tridimas, 2017] George Tridimas. Constitutional choice in ancient athens: The evolution of the frequency of decision making. *Constitutional Political Economy*, 28(3):209–230, 2017.

[Virtanen *et al.*, 2020] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

[Wales and others, 2003] David Wales et al. *Energy landscapes: Applications to clusters, biomolecules and glasses*. Cambridge University Press, 2003.