informs
https://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–19
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Tracking Truth with Liquid Democracy

Adam J. Berinsky,[a] Daniel Halpern,[b] Joseph Y. Halpern,[c] Ali Jadbabaie,[d] Elchanan Mossel,[e] Ariel D. Procaccia,[b] Manon Revel[d,*]

[a] Department of Political Science, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; [b] School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138; [c] Department of Computer Science, Cornell University, Ithaca, New York 14853; [d] Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; [e] Department of Mathematics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

*Corresponding author

Contact: berinsky@mit.edu, https://orcid.org/0000-0001-7827-9396 (AJB); dhalpern@g.harvard.edu (DH); halpern@cs.cornell.edu, https://orcid.org/0000-0002-9229-1663 (JYH); jadbabai@mit.edu, https://orcid.org/0000-0003-1122-3069 (AJ); elmos@mit.edu, https://orcid.org/0000-0001-7812-7886 (EM); arielpro@seas.harvard.edu, https://orcid.org/0000-0003-0318-491X (ADP); mrevel@mit.edu, https://orcid.org/0000-0002-8335-946X (MR)

Received: August 7, 2023
Revised: April 23, 2024; July 11, 2024
Accepted: July 13, 2024
Published Online in Articles in Advance: January 8, 2025

https://doi.org/10.1287/mnsc.2023.02470

Abstract. The dynamics of random transitive delegations on a graph are of particular interest when viewed through the lens of an emerging voting paradigm: *liquid democracy*. This paradigm allows voters to choose between directly voting and transitively delegating their votes to other voters so that those selected cast a vote weighted by the number of delegations that they received. In the epistemic setting, where voters decide on a binary issue for which there is a ground truth, previous work showed that a few voters may amass such a large amount of influence that liquid democracy is less likely to identify the ground truth than direct voting. We quantify the amount of permissible concentration of power and examine more realistic delegation models, showing that they behave well by ensuring that (with high probability) there is a permissible limit on the maximum number of delegations received. Our theoretical results demonstrate that the delegation process is similar to well-known processes on random graphs that are sufficiently bounded for our purposes. Along the way, we prove new bounds on the size of the largest component in an infinite Pólya urn process, which may be of independent interest. In addition, we empirically validate the theoretical results, running six experiments (for a total of $N = 168$ participants, 62 delegation graphs, and over 11,000 votes collected). We find that empirical delegation behaviors meet the conditions for our positive theoretical guarantees. Overall, our work alleviates concerns raised about liquid democracy and bolsters the case for the applicability of this emerging paradigm.

History: Accepted by Martin Bichler, market design, platform, and demand analytics.
Funding: This work was supported by the Michael Hammer Fellowship, the Office of Naval Research [2016 Vannevar Bush Faculty Fellowship, 2020 ONR Vannevar Bush Faculty Fellowship, and Grant N00014-20-1-2488], the Office of Secretary of Defense [Grants ARO MURI W911NF-19-0217 and ARO W911NF-17-1-0592], Simons [Investigator Award 622132], the Open Philanthropy Foundation, and the National Science Foundation [Grants IIS-178108, CCF-1733556, CCF-1918421, CCF-2007080, IIS-1703846, and IIS-2024287]. J. Y. Halpern was supported by a grant from the Cooperative AI Foundation [ARO Grant W911NF-22-1-0061].
Supplemental Material: The online appendix and data files are available at https://doi.org/10.1287/mnsc.2023.02470.

Keywords: liquid democracy • random graph theory • collective-decision making • social choice theory

1. Introduction

Liquid democracy is a voting paradigm that is conceptually situated between *direct democracy*, in which voters have direct influence over decisions, and *representative democracy*, where voters choose delegates who represent them for a period of time. Under liquid democracy, voters have a choice; they can either vote directly on an issue, like in direct democracy, or delegate their vote to another voter, entrusting that voter to vote on their behalf. The defining feature of liquid democracy is that these delegations are *transitive*; if voter 1 delegates to voter 2 and voter 2 delegates to voter 3, then voter 3 votes (or delegates) on behalf of all three voters.

In recent years, liquid democracy has gained prominence around the world. The most impressive example is that of the German Pirate Party, which adopted the LiquidFeedback platform in 2010 (Kling et al. 2015). Other political parties, such as the Net Party in Argentina and Flux in Australia, have run on the wily promise that, once elected, their representatives would be essentially controlled by voters through a liquid democracy platform. Companies are also exploring the use of liquid democracy for corporate governance; Google, for example, has run a proof-of-concept experiment (Hardt and Lopes 2015). Blockchain systems have also been experimenting with related weighted decentralized voting systems (Benhaim et al. 2023, Li et al. 2023).

Practitioners, however, recognize that there is a potential flaw in liquid democracy: namely, the possibility of *concentration of power* in the sense that certain voters amass a relatively large number of delegations, giving them pivotal influence over the final decision. This scenario seems inherently undemocratic, and it is not a mere thought experiment. Indeed, in the LiquidFeedback platform of the German Pirate Party, a linguistics professor at the University of Bamberg received so many delegations that as noted by *Der Spiegel*,[1] his "vote was like a decree."

Kahng et al. (2021) examine liquid democracy's concentration-of-power phenomenon from a theoretical viewpoint and establish a troubling impossibility result in what has been called the *epistemic* setting: that is, one where there is a ground truth.[2] Informally, they demonstrate that even under the strong assumption that voters delegate only to more "competent" voters, any "local mechanism" satisfying minimal conditions will, in certain instances, be subject to concentration of power, leading to relatively low accuracy. More specifically, Kahng et al. (2021) model the problem as a decision problem where voters decide on an issue with two outcomes $\{0, 1\}$, where one is correct (the ground truth) and zero is incorrect. Each of the voters $i \in \{1, \ldots, n\}$ is characterized by a *competence* $p_i \in [0, 1]$. The binary vote V_i of each voter i is drawn independently from a Bernoulli distribution; that is, each voter votes correctly with probability p_i. Under direct democracy, the outcome of the election is determined by a majority vote. The correct outcome is selected if and only if more than half of the voters vote for the correct outcome; that is, it is correct if and only if $\sum_{i=1}^{n} V_i \geq n/2$. Under liquid democracy, there exists a set of weights, weight_i, for each $i \in [n]$, which represent the number of votes that voter i gathered transitively after delegation (if voter i delegates, then $\mathsf{weight}_i = 0$). The outcome of the election is then determined by a weighted majority; it is correct if and only if $\sum_{i=1}^{n} \mathsf{weight}_i V_i \geq n/2$.

Kahng et al. (2021) also introduce the concept of a *delegation mechanism*, which determines whether voters delegate and if so, to whom they delegate. They are especially interested in *local mechanisms*, where the delegation decision of a voter depends only on their local neighborhood according to an underlying social network. They assume that voters delegate only to those with strictly higher competence, which excludes the possibility of cyclic delegations. To evaluate liquid democracy, Kahng et al. (2021) test the intuition that society makes more informed decisions under liquid democracy than under direct democracy (especially given the foregoing assumption about upward delegation). To that end, they define the *gain* of a delegation mechanism to be the difference between the probability that the correct outcome is selected under liquid democracy and the probability that the correct outcome is selected under direct democracy. A delegation mechanism satisfies *positive gain* if its gain is strictly positive in some cases, and it satisfies *do no harm* (DNH) if for all $\varepsilon > 0$, its gain is at least $-\varepsilon$ for sufficiently large instances. Assuming that competence after delegation remains strictly above $1/2$, this will follow from the law of large number that applies to the weighted majority with weights relatively spread out (Häggström et al. 2006). The main result of Kahng et al. (2021) is that local mechanisms can never satisfy these two requirements. Caragiannis and Micha (2019) further strengthen this negative result by showing that there are degenerate instances where local mechanisms perform much worse than either direct democracy or dictatorship (the most extreme concentration of power).[3]

These results undermine the case for liquid democracy; the benefits of delegation appear to be reversed by concentration of power. However, the negative conclusion relies heavily on worst-case modeling assumptions. Our research represents a significant advance as it offers a comprehensive framework that not only captures the worst-case scenarios of previous works but also, provides insights into more intriguing "high-probability" cases. In particular, in this paper, we provide a new theoretical model and extensive experiments that show that liquid democracy will typically satisfy a probabilistic version of *positive gain* and *do no harm* under minimal assumptions.

1.1. Our Contributions and Techniques

Our contributions are the following. First, building on the work of Kahng et al. (2021), we provide a general framework to analyze the stochastic network dynamics of transitive delegations that captures liquid democracy's intricate interactions between local delegation choices and global properties. Second, we identify large classes of delegation models where liquid democracy performs well in that delegations induce a sufficiently small amount of concentration of power and liquid democracy almost surely results in correct outcomes. Along the way, we prove new high-probability bounds on the size of the largest component in an infinite Pólya urn process;[4] this result may be of independent interest. Finally, we conduct the first series of laboratory experiments on liquid democracy that can test epistemic performance. This involved over 11,000 votes from 168 participants in six experimental groups, where each group had pre-existing social ties. Our novel experimental design allows us to compare the performance of direct democracy and liquid democracy as well as to analyze properties of real voter delegation behavior. Importantly, the behaviors that we observe align with one of the models that we introduce, thus lending support to this approach. Taken together, these results exhibit a regime in which liquid democracy displays promising performance. We next elaborate on some of our specific techniques.

1.1.1. Stochastic Delegations. Our point of departure from the existing literature is the way that we model delegation in liquid democracy. To emphasize these differences, instead of calling these delegation functions *mechanisms*, we instead call them delegation *models* as they are intended to capture independent voter behavior rather than prescribing to each voter to whom they must delegate. Our delegation models are defined by $M = (q, \varphi)$, where $q : [0,1] \to [0,1]$ is a function that maps a voter's competence to the probability that they delegate and $\varphi : [0,1]^2 \to \mathbb{R}_{\geq 0}$ maps a pair of competencies to a weight. In this model, each voter i votes directly with probability $1 - q(p_i)$ and conditioned on delegating with probability $q(p_i)$, delegates to voter $j \neq i$ with probability proportional to $\varphi(p_i, p_j)$. These delegation functions do not model explicit reasoning; rather, they model behaviors that may be influenced by tacit knowledge captured by q and φ. A voter does not need to "know" the competence of another voter to decide whether to delegate. Rather, the delegation *probabilities* are influenced by competence as captured by φ; note that delegation cycles are possible, and we take a worst-case approach to dealing with them. If the delegations form a cycle, then all voters in the cycle are assumed to be incorrect (vote 0).[5]

The most significant difference between our model of delegation and that of Kahng et al. (2021) is that in our model, each voter has a chance of delegating to any other voter, whereas in their model, an underlying social network restricts delegation options. Our model captures a connected world where in particular, voters may have heard of experts on various issues, even if they do not know them personally. Although our model eschews an explicit social network, it can be seen as embedded into the delegation process, where the probability that i delegates to j takes into account the probability that i is familiar with j in the first place. Another difference between our model and that of Kahng et al. (2021) is that we model the competencies $p_1, \ldots, p_n$ as being sampled independently from a distribution $\mathcal{D}$. Although this assumption is made mainly for ease of exposition, it allows us to avoid edge cases and obtain robust results.

1.1.2. Delegation Models. Our goal is to identify delegation models that satisfy (probabilistic versions of) positive gain and do no harm. Our first technical contribution, in Section 2.5, is the formulation of general conditions on the model and competence distribution that are sufficient for these properties to hold (Lemma 1). In particular, to achieve the more difficult do no harm property, we present conditions that guarantee that the maximum weight $\mathsf{max\text{-}weight}(G_n)$ accumulated by any voter is sublinear with high probability and that the expected increase in competence after delegation is at least a positive constant times the population size. These conditions prevent extreme concentration of power and ensure that the representatives after delegation are sufficiently better than the entire population to compensate for any concentration of power that does happen.

Although the proof is straightforward, the benefit of this lemma is that it then suffices to identify models and distribution classes that verify these conditions. A delegation model M and a competence distribution $\mathcal{D}$ induce a distribution over delegation instances that generates random graphs in ways that relate to well-known graph processes, which we leverage to analyze our models. Specifically, we introduce three models, all shown to satisfy do no harm and positive gain under *any* continuous distribution over competence levels. The first models, *upward delegation* and *confidence-based delegation*, are interesting but restricted case studies that demonstrate the robustness of our approach. By contrast, the *general continuous delegation* model is, as the name suggests, quite general. Moreover, it is realistic; its predictions are consistent with our experiments.

1.1.2.1. Upward Delegation. In Section 3, we consider a model according to which the probability p of delegation is exogenous and constant across competencies, $q(p_i) = p$, and delegation can occur only to voters with strictly higher competence; the weight that any voter i puts on another voter j is $\varphi(p_i, p_j) = \mathbb{I}_{\{p_j - p_i > 0\}}$. This model captures the fact that there might be some reluctance to delegate regardless of the voter's competence, but it does assume that voters act in the interest of society by only delegating to voters who are more competent than they are.

To generate a random graph induced by such a model, one can add a single voter at a time in order of decreasing competence and allow the voter to either not delegate (with probability $1 - p$) and create their own disconnected component or delegate to the creator of any other component with probability proportional to the size of the component. This works because delegating to any voter in the previous components is possible (because they have strictly higher competence) and would result in the votes being concentrated in the originator of that component by transitivity. Such a process exactly generates a preferential attachment graph with a positive probability of not attaching to the existing components, also called an infinite Pólya urn process (Simon 1955). We can then show that with high probability, no component grows too large so long as $p < 1$ (see Section 1.1.3 for an overview of this step). Further, continuity of the competence distribution ensures that enough lower-competence voters delegate to higher-competence voters to sufficiently increase the average.

1.1.2.2. Confidence-Based Delegation. In Section 4, we consider a model in which voters delegate with probability decreasing in their competencies and choose someone at random when they delegate. That is, the probability $q(p_i)$ that any voter i delegates is decreasing in p_i, and the weight that any voter i gives to any voter j is $\varphi(p_i, p_j) = 1$. In other words, in this model, competence does not affect the probability of receiving delegations, only the probability of delegating.

To generate a random graph induced by such a model, one can begin from a random vertex and study the delegation tree that starts at that vertex. A delegation tree is defined as a branching process, where a node i's "children" are the nodes that delegated to node i. In contrast to classical branching processes, the probability for a child to be born increases as the number of people who already received delegations decreases. Nevertheless, we prove that with high probability, as long as a delegation tree is no larger than $O(\log n)$, our heterogeneous branching process is dominated by a subcritical graph branching process (Alon and Spencer 2016). We can then conclude that no component has size larger than $O(\log n)$ with high probability. Next, we show that the expected competence among the voters who do not delegate is strictly higher than the average competence.

1.1.2.3. General Continuous Delegation. Finally, we consider a general model in Section 5, where the likelihood of delegation is fixed and the weight assigned to each voter when delegating is increasing in their competence. That is, each voter i delegates with probability $q(p_i) = p$, and the weight that voter i places on voter j is $\varphi(p_i, p_j)$, where φ is continuous and increases in its second coordinate. Thus, in this model, the delegation distribution is slightly skewed toward more competent voters.

To generate a random graph induced by such a model, we again consider a branching process, but now, voters j and k place different weights on i per φ. Therefore, voters have a *type* that governs their delegation behavior; this allows us to define a multitype branching process with types that are continuous in $[0, 1]$. The major part of the analysis is a proof that with high probability, as long as the delegation tree is no larger than $O(\log n)$, our heterogeneous branching process is dominated by a subcritical Poisson multitype branching process. In a manner similar to confidence-based delegation, we also show that there is an expected increase in competence after delegation.

1.1.3. Component Sizes in Infinite Pólya Urn Processes. Recall that to prove that upward delegation satisfies do no harm, we show that the largest component in an infinite Pólya urn process is sublinear with high probability (Lemma 2). We briefly expand on the proof as this result was, to the best of our knowledge, not previously known in the random graph literature and may be of independent interest. We begin by focusing on the first t^γ bins (for a suitably chosen γ depending on the attachment probability p) and derive an upper bound on the expected size of these bins. This allows us to use Markov's inequality and union bound over all bins to show that simultaneously all of them are sublinear in size with high probability.

Second, we take care of the remaining bins by observing that each additional bin's growth is isomorphic to a classic Pólya urn process with two bins, whose limiting dynamic follows a Beta distribution. We analyze the rate of convergence, which allows us to give sufficiently strong bounds using Chebyshev's inequality after exactly $t - t^\gamma$ steps, and union bound over all of these bins, concluding that all are sublinear with high probability.

1.1.4. Consistency with Experiments. Lastly, we conduct six experiments to statistically estimate the functions q and φ, and we test the overall effectiveness of liquid democracy. Participants were presented with several yes or no questions on various topics. We call the set of questions related to each topic a *task*. For each task, participants could either choose to vote directly or delegate their vote (for all questions) to another participant. They only saw the questions in a task if they chose to vote directly. In a later phase, they were asked to answers the questions that they had delegated (and not seen) to see how they would have voted. This setup allows us to do a few things. First, it induces a matched-pair design, where for each task and experiment, we can compare the accuracy of voting under liquid and direct democracy. Second, we use the answers to all questions to estimate participants' competencies. Using this information, we study how delegation behavior depends on competence and investigate whether it is consistent with the theoretical findings.

Results suggest that (i) competence is inversely correlated with the chance of delegation and that (ii) the likelihood of delegating to another voter increases with their competence. The results, therefore, support the assumptions and predictions made by the confidence-based and general continuous delegation models. Taken together, these results exhibit a regime in which liquid democracy is overall more likely to pinpoint the truth than direct democracy.

1.2. Related Work

The most closely related paper is that of Kahng et al. (2021), which was discussed in detail above. It is worth noting, however, that they complement their negative result with a positive one; when the mechanism can restrict the maximum number of delegations (transitively) received by any voter to $o(\sqrt{\log n})$, do no harm and positive gain are satisfied. Imposing such a restriction would require a central planner that monitors and controls delegations. Gölz et al. (2018) build on this idea; they study liquid democracy systems where voters may nominate multiple delegates and where a central planner chooses a single delegate in order to minimize the maximum weight of any voter. Similarly, Brill and Talmon (2018) introduce a process that allows voters to specify ordinal preferences over delegation options and possibly restricting or modifying delegations in a centralized way. Caragiannis and Micha (2019) and then, Becker et al. (2021) also consider central planners; they show that for given competencies, the problem of choosing among delegation options to maximize the probability of a correct decision is hard to approximate. In any case, implementing these proposals would require a fundamental rethinking of the practice of liquid democracy. By contrast, our positive results show that *decentralized* delegation models can be inherently self-regulatory, which supports the effectiveness of the current practice of liquid democracy.

More generally, there has been a significant amount of theoretical research on liquid democracy in recent years. To give a few examples, Green-Armytage (2015) studies whether it is rational for voters to delegate their vote from a utilitarian viewpoint. Christoff and Grossi (2017) examine a similar question but in the context of voting on logically interdependent propositions. Bloembergen et al. (2019), Zhang and Grossi (2021), and Dhillon et al. (2023) study liquid democracy from a game-theoretic viewpoint.

Next, our work builds on the random graph literature as our delegation processes are related to well-known stochastic graph processes. Upward delegation can be viewed as a generalization of the preferential attachment model, where agents do not attach to the existing component(s) with a fixed probability. Classical preferential attachment models assume that a new node attaches to an existing node n_0 with probability (parameterized by an *attachment function*) depending on the degree of n_0 (Barabási and Albert 1999, Durrett 2007). In our framework, a new component may be created with fixed probability, a setup introduced by Simon (1955) and usually referred to as an infinite Pólya urn process. Others have studied the distribution of degrees (Drinea et al. 2001), the distribution of the number of components with k people at time t (Chung et al. 2003), and the conditions for the emergence of infinite components (Collevecchio et al. 2013). However, to the best of our knowledge, the existing results do not allow us to derive bounds on the size of the largest component with high probability after a finite amount of time.

In terms of our experiments on liquid democracy, ours is the first paper to conduct experiments with human subjects. Previous papers have studied different aspects of liquid democracy through experiments in corporate (Hardt and Lopes 2015) and political environments Independent of and essentially concurrent with our work, Campbell et al. (2022) tested a game-theoretic formulation of liquid democracy. Unlike our experiments, they used online platforms to gather participants who did not know each other. Participants were assigned a probability of being correct and asked whether they would want to delegate to others, with experts (those with the highest probability of being correct) being publicly known. The delegations were randomly assigned to the predetermined experts in one setup and through the random dot kinematogram task in another one. The group sizes considered were 5 people with 1 expert, 15 people with 3 experts, and 125 people with 25 experts. Although this study reveals interesting connections between individuals' perceived competence and delegation behavior, it cannot investigate how experts are (or are not) identified endogenously through interpersonal knowledge embedded in a social networks because the participants do not know each other.

Last, our work relates to recent advances in managerial studies that consider novel forms of governance, such as corporate governance (e.g., Huang 2023), blockchain technologies (e.g., Benhaim et al. 2023, Li et al. 2023), and prediction markets (e.g., Chen et al. 2008, Atanasov et al. 2017).

2. Model

There is a set of n voters denoted $[n] = \{1, \dots, n\}$. We assume that voters are making a decision on a binary issue with possible answers of zero and one; there is a correct alternative (one) and an incorrect alternative (zero). Each voter i has a *competence level* $p_i \in [0, 1]$, which is the probability that i votes correctly. We denote the vector of competencies by $\vec{p}_n = (p_1, \dots, p_n)$. When n is clear from the context, we sometimes drop it from the notation.

2.1. Delegation Graphs

A *delegation graph* $G_n = ([n], E)$ on n voters is a directed graph with voters as vertices and a directed edge $(i, j) \in E$ denoting that i delegates their vote to j. Again, if n is clear from the context, we occasionally drop it from the notation. The out degree of a vertex in the delegation graph is at most one because each voter can delegate to at most one person. Voters who do not delegate have no outgoing edges. In a delegation graph G_n, the *delegations received* by a voter i, $\mathsf{dels}_i(G_n)$, are defined as the total number of people who (transitively) delegated to i in G_n (i.e., the total number of ancestors of i in G_n). The *weight* of a voter i, $\mathsf{weight}_i(G_n)$, is $\mathsf{dels}_i(G_n) + 1$ (the number of delegations that they received plus their own weight) if i votes directly and zero if i delegates. We define $\mathsf{max\text{-}weight}(G_n) = \max_{i \in [n]} \mathsf{weight}_i(G_n)$ to be the largest weight of any voter and define $\mathsf{total\text{-}weight}(G_n) = \sum_{i=1}^{n} \mathsf{weight}_i(G_n)$. Because each vote is counted at most once, we have that $\mathsf{total\text{-}weight}(G_n) \le n$. However, note that if delegation edges form a cycle, then the weights of the voters on the cycle and voters delegating into the cycle are all set to zero and hence, will not be counted. In particular, this means that total-weight(G_n) may be strictly less than n.[6]

2.2. Delegation Instances

We call the tuple $(\vec{p}_n, G_n)$ a *delegation instance* or simply, an instance on n voters. Let $V_i = 1$ if voter i would vote correctly if i did vote, and $V_i = 0$ otherwise. Fixed competencies $\vec{p}_n$ induce a probability measure $\mathbb{P}_{\vec{p}_n}$ over the n possible binary votes V_i, where $V_i \sim \text{Bern}(p_i)$. Given votes $V_1, \ldots, V_n$, we let X_n^D be the number of correct votes under direct democracy; that is, $X_n^D = \sum_{i=1}^n V_i$. We let $X_{G_n}^F$ be the number of correct votes under liquid democracy with delegation graph G_n; that is, $X_{G_n}^F = \sum_{i=1}^n \text{weight}_i(G_n) \cdot V_i$. The probabilities that direct democracy and liquid democracy are correct are $\mathbb{P}_{\vec{p}_n}[X_n^D > n/2]$ and $\mathbb{P}_{\vec{p}_n}[X_{G_n}^F > n/2]$, respectively.

2.3. Gain of a Delegation Instance

We define the *gain* of an instance as

$$\text{gain}(\vec{p}_n, G_n) = \mathbb{P}_{\vec{p}_n}[X_{G_n}^F > n/2] - \mathbb{P}_{\vec{p}_n}[X_n^D > n/2].$$

In words, it is the difference between the probability that liquid democracy is correct and the probability that majority is correct.

2.4. Randomization over Delegation Instances

In general, we assume that both competencies and delegations are chosen randomly. Each voter's competence p_i is an independent and identically distributed (i.i.d.) sample from a fixed distribution $\mathcal{D}$ with support contained in $[0,1]$. Delegations will be chosen according to a *model* M. A model $M = (q, \varphi)$ is composed of two parts. The first, $q : [0,1] \to [0,1]$, is a function that maps competencies to the probability that the voter delegates. The second, $\varphi : [0,1]^2 \to \mathbb{R}_{\geq 0}$, maps pairs of competencies to a weight. A voter i with competence p_i will choose how to delegate as follows.

- With probability $1 - q(p_i)$, the voter does not delegate.
- With probability $q(p_i)$, i delegates; i places weight $\varphi(p_i, p_j)$ on each voter $j \neq i$ and randomly samples another voter j to delegate to proportional to these weights. In the degenerate case where $\varphi(p_i, p_j) = 0$ for all $j \neq i$, we assume that i does not delegate.

A competence distribution $\mathcal{D}$, a model M, and a number n of voters induce a probability measure $\mathbb{P}_{\mathcal{D},M,n}$ over all instances $(\vec{p}_n, G_n)$ of size n.

We can now redefine the *do no harm* and *positive gain* properties from Kahng et al. (2021) in a probabilistic way.

Definition 1 (Probabilistic Do No Harm). A model M satisfies *probabilistic do no harm* with respect to a class $\mathfrak{D}$ of distributions if for all distributions $\mathcal{D} \in \mathfrak{D}$ and all $\varepsilon, \delta > 0$, there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$,

$$\mathbb{P}_{\mathcal{D},M,n}[\text{gain}(\vec{p}_n, G_n) \geq -\varepsilon] > 1 - \delta.$$

Definition 2 (Probabilistic Positive Gain). A model M satisfies *probabilistic positive gain* with respect to a class $\mathfrak{D}$ of distributions if there exists a distribution $\mathcal{D} \in \mathfrak{D}$ such that for all $\varepsilon, \delta > 0$, there exists $n_0 \in N$ such that for all $n \geq n_0$,[7]

$$\mathbb{P}_{\mathcal{D},M,n}[\text{gain}(\vec{p}_n, G_n) \geq 1 - \varepsilon] > 1 - \delta.$$

2.5. Core Lemma

Next, we give a key lemma, which provides sufficient conditions for a model M to satisfy probabilistic do no harm and probabilistic positive gain with respect to a class $\mathfrak{D}$ of distributions. This lemma will form the basis of all of our later results. Because the result follows from relatively straightforward concentration inequalities, we defer the proof to Online Appendix B.1.

Lemma 1. *If M is a model, $\mathfrak{D}$ is a class of distributions, n is a number of persons, and for all distributions $\mathcal{D} \in \mathfrak{D}$, there is an $\alpha \in (0,1)$ and $C : \mathbb{N} \to \mathbb{N}$ with $C(n) \in o(n)$ such that*

$$\mathbb{P}_{\mathcal{D},M,n}[\text{max-weight}(G_n) \leq C(n)] = 1 - o(1) \tag{1}$$

$$\mathbb{P}_{\mathcal{D},M,n}\left[\sum_{i=1}^n \text{weight}_i(G_n) \cdot p_i - \sum_{i=1}^n p_i \geq 2\alpha n\right] = 1 - o(1); \tag{2}$$

then, M satisfies probabilistic do no harm. If in addition, there exists a distribution $\mathcal{D} \in \mathfrak{D}$ and an $\alpha \in (0,1)$ such that

$$\mathbb{P}_{\mathcal{D},M,n}\left[\sum_{i=1}^n p_i + \alpha n \leq n/2 \leq \sum_{i=1}^n \text{weight}_i(G_n) \cdot p_i - \alpha n\right] = 1 - o(1), \tag{3}$$

then M satisfies probabilistic positive gain.

In words, Condition (1) ensures that as the number of voters grows large, the weighted number of correct votes under liquid democracy will concentrate around its expectation, $\sum_{i=1}^n \text{weight}_i(G_n) \cdot p_i$. Standard concentration results already imply that this holds for direct democracy. Condition (2) ensures that these expectations are sufficiently separated. So, with high probability, liquid democracy will have more correct votes than direct democracy, which is sufficient to guarantee DNH. Finally, Condition (3) ensures that in some cases, the expectations for direct and liquid votes will be below and over half the voters, respectively, which after applying concentration, means that there will likely be a large gain.

In the following sections, we investigate natural delegation models and identify conditions such that the models satisfy probabilistic do no harm and probabilistic positive gain. In all instances, we will invoke

Lemma 1 after showing that its sufficient conditions are satisfied.

3. Strictly Upward Delegation Model

We now turn to the analysis of a simple model that assumes that voters either do not delegate with fixed exogenous probability or do delegate to voters who have a competence greater than their own.

Formally, for a fixed $p \in [0,1]$, we let $M_p^U = (q, \varphi)$ be the model consisting of $q(p_i) = p$ for all $p_i \in [0,1]$, and $\varphi(p_i, p_j) = \mathbb{I}_{\{p_j > p_i\}}$ for all $i, j \in [n]$. That is, voter i delegates with fixed probability p and puts equal weight on all of the more competent voters. In other words, if voter i delegates, then i does so to a more competent voter chosen uniformly at random. Note that a voter with maximal competence will place zero weight on all other voters and hence, is guaranteed not to delegate. We refer to M_p^U as the *upward delegation model* parameterized by p.

Theorem 1 (Upward Delegation Model). *For all $p \in (0,1)$, M_p^U satisfies probabilistic do no harm and probabilistic positive gain with respect to the class $\mathfrak{D}^C$ of all continuous distributions.*

The proof of the theorem relies on novel bounds that we drive on the largest bin size in an infinite Pólya urn process (Simon 1955, Chung et al. 2003). We first formally define the process and present our bound in Lemma 2. A *Pólya urn process with attachment probability* p begins at time $t = 1$ with one ball in one bin. At each time step $t > 1$, a new ball arrives. With probability $1 - p$, a new bin is created, and the new ball is placed in that bin; with probability p, the ball joins an existing bin, and it does so with probability proportional to the number of balls in the bins (i.e., if there are three bins containing one, two, and three balls, respectively, it joins each with probability 1/6, 2/6, and 3/6, respectively). We then have the following.

Lemma 2. *For all $p \in (0,1)$ and $t \ge 1$, let L_t^p be the random variable denoting the maximum number of balls in any bin after running the infinite Pólya urn process with new-bin probability p for t steps. Then, there exists $\delta < 1$ depending only on p such that for all $T \ge 1$, $\Pr[L_T^p \le T^\delta] = 1 - o(1)$.*

Proof. Fix the parameter $p \in (0,1)$. Choose γ to be a constant such that $3/4 < \gamma < 1$; note that $p + (1-p)\gamma < p + (1-p) = 1$. Choose δ (for the lemma statement) such that $p + (1-p)\gamma < \delta < 1$. Notice that we can choose γ and δ such that δ is arbitrarily close to $3/4 + p/4$.

Let $B^{(k)}$ denote the kth bin. Let $U_t^{(k)}$ be the size of $B^{(k)}$ at time t. Because there are at most t bins by time t, notice that $L_t^p = \max(U_t^{(1)}, \ldots, U_t^{(t)})$. In general, our approach will be to analyze bins separately and show that $U_T^{(k)}$ remains below T^δ with high-enough probability so that we can union bound over all possible $k \le T$. That is, we will show

$$\sum_{k=1}^{T} \Pr[U_T^{(k)} > T^\delta] = o(1),$$

which also implies $\Pr[L_T^p > T^\delta] = o(1)$. Hence, it will be useful to consider this process more formally from the perspective of the kth bin, $B^{(k)}$. The kth bin $B^{(k)}$ is "born" at some time $t \ge k$, the kth time in which a ball does not join a pre-existing bin, at which point $U_t^{(k)} = 1$ (prior to this, $U_t^{(k)} = 0$). More specifically, the first bin $B^{(k)}$ is guaranteed to be born at time $t = 1$, and for all other $k > 1$, $B^{(k)}$ will be born at time $t \ge k$ with probability $\binom{t-1}{k-1}(1-p)^k p^{t-k}$, although these exact probabilities will be unimportant for our analysis. Once born, we have the following recurrence on $U_t^{(k)}$ describing that the probability $B^{(k)}$ will be chosen at time t:

$$U_t^{(k)} = \begin{cases} U_{t-1}^{(k)} + 1 & \text{with probability } \dfrac{p \cdot U_{t-1}^{(k)}}{t-1} \\ U_{t-1}^{(k)} & \text{with probability } 1 - \dfrac{p \cdot U_{t-1}^{(k)}}{t-1}. \end{cases}$$

Let $W_t^{(k)}$ be the process for the size of bin that is born at time k. That is, $W_k^{(k)} = 1$, and for $k > t$, $W_t^{(k)}$ follows the exact same recurrence as $U_t^{(k)}$. Note that because the kth bin $B^{(k)}$ can only be born at time k or later, we have that $W_t^{(k)}$ stochastically dominates $U_t^{(k)}$ for all k and t. Hence, it suffices to show that

$$\sum_{k=1}^{T} \Pr[W_T^{(k)} > T^\delta] = o(1). \tag{4}$$

We split our analysis into two parts; the first considers the first T^γ bins, whereas the second considers the last $T - T^\gamma$ bins.

We first show that $\sum_{k=1}^{T^\gamma} \mathbb{P}[W_T^{(k)} > T^\delta] = o(1)$. Note that the expectation of $W_n^{(k)}$

$$\mathbb{E}[W_n^{(k)}] = \frac{\Gamma(n+p)\Gamma(k)}{\Gamma(p+k)\Gamma(n)} \tag{5}$$

for all $k \le n$, where Γ represents the Gamma function. We relegate the argument for Equation (5) to Online Appendix B.2. Using this along with Gautchi's inequality (Gautschi 1959), $(t + p - 1)^p \le \frac{\Gamma(p+t)}{\Gamma(t)} \le (t+p)^p$, to approximate the Γ terms, we can apply Markov's inequality and use algebra to get $\sum_{k=1}^{n^\gamma} \mathbb{P}[W_n^k > n^\delta] = o(1)$. We again relegate these arguments to Online Appendix B.2.

Now, consider the final $T - T^\gamma$ components. We will prove that $\Pr[W_T^{(T^\gamma+1)} > T^\delta] = o(1/T)$. Because $W_T^{(k)}$ stochastically dominates $W_T^{(k')}$ for all $k' \ge k$, this implies that $\Pr[W_T^{(k)} > T^\delta] = o(1/T)$ for all $k \ge T^\gamma + 1$. Hence,

$$\sum_{k=T^\gamma+1}^{T} \Pr[W_T^{(k)} > T^\delta] = o(1).$$

To do this, we compare the $W_t^{(T^\gamma+1)}$ process with another process, V_t. We define $V_0 = 1$, and for $t > 0$, we take V_t to satisfy the following recurrence:

$$V_t = \begin{cases} V_{t-1}+1 & \text{with probability } \dfrac{V_{t-1}}{t+n^\gamma} \\ V_{t-1} & \text{with probability } 1-\dfrac{V_{t-1}}{t+n^\gamma}. \end{cases}$$

This is identical to the W recurrence with t shifted down by $n^\gamma+1$, except without the p factor. Hence, $V_{T-T^\gamma+1}$ clearly stochastically dominates $W_T^{(T^\gamma+1)}$. For convenience in calculation, we will instead focus on bounding V_T, which itself stochastically dominates $V_{T-T^\gamma+1}$.

Next, note that the V_t process is isomorphic to the following classic Pólya urn process. We begin with two bins: one with a single ball and the other with n^γ balls. At each time, a new ball is added to one of the two bins with probability proportional to the bin size. The process V_t is isomorphic to the size of the one-ball urn after t steps. Classic results tell us that for fixed starting bin sizes a and b, as the number of steps grows large, the possible proportion of balls in the a bin follows a Beta(a,b) distribution (Markov 1917, Eggenberger and Pólya 1923, Pólya 1930, Johnson and Kotz 1978, Mahmoud 2009).

The mean and variance of such a Beta distribution would be sufficient to prove our necessary concentration bounds; however, for us, we need results after exactly $T-T^\gamma$ steps, not simply in the limit. Hence, we will be additionally concerned with the speed of convergence to this Beta distribution.

Let $X_T=\frac{V_T}{T}$ and $Z_T\sim \text{Beta}(1,T^\gamma)$. From Janson (2020), we know that the rate of convergence is such that for any $p\ge 1$,

$$\ell_p(X_T,Z_T)=\Theta(1/T), \tag{6}$$

where ℓ_p is the *minimal* L_p *metric* defined as

$$\ell_p(X,Y)=\inf\{\mathbb{E}[|X'-Y'|^p]^{1/p}\,|\,X'\overset{d}{=}X, Y'\overset{d}{=}Y\},$$

which can be thought of as the minimal L_p norm over all possible couplings between X and Y. For our purposes, the only fact about the ℓ_p metric that we will need is that $\ell_p(X,0)=\mathbb{E}[|X|^p]^{1/p}$, where zero is the identically zero random variable. Because ℓ_p is, in fact, a metric, the triangle inequality tells us that $\ell_p(0,X_n)\le \ell_p(0,Z_n)+\ell_p(Z_n,X_n)$, so combining with (6), we have that

$$\mathbb{E}[|X_T|^p]^{1/p}\le \mathbb{E}[|Z_T|^p]^{1/p}+\Theta(1/T) \tag{7}$$

for all $p\ge 1$.

Note that because $Z_T\sim\text{Beta}(1,T^\gamma)$,

$$\mathbb{E}[Z_T]=\frac{1}{1+T^\gamma}=\Theta(T^{-\gamma})$$

and

$$\text{Var}[Z_T]=\frac{T^\gamma}{(2+T^\gamma)(1+T^\gamma)^2}=\Theta(T^{-2\gamma}).$$

Given these results, we are ready to prove that V_T is smaller than T^δ with probability $1-o(1/T)$. Precisely, we want to show that $\Pr[X_T\ge T^{\delta-1}]=o(1)$. By Chebyshev's inequality,

$$\Pr[X_T\ge T^{\delta-1}]\le \frac{\text{Var}[X_T]}{(T^{\delta-1}-\mathbb{E}[X_T])^2}.$$

Inequality (7) with $p=1$ along with the fact that X_T and Z_T are always nonnegative implies that $\mathbb{E}[X_T]\le \mathbb{E}[Z_T]+\Theta(1/T)=O(T^{-\gamma})$. Hence, $T^{\delta-1}-\mathbb{E}[X_T]=\Omega(T^{\delta-1})$ because $\delta-1>-1/2>-\gamma$. We can, therefore, write

$$(T^{\delta-1}-\mathbb{E}[X_T])^2=\Omega(T^{-2(\delta-1)}). \tag{8}$$

Inequality (7) with $p=2$ implies that $\sqrt{\mathbb{E}[X_T^2]}\le\sqrt{\mathbb{E}[Z_T^2]}+\Theta(1/T)$. Hence,

$$\begin{aligned}\mathbb{E}[X_T^2]&\le\left(\Theta(1/T)+\sqrt{\mathbb{E}[Z_T^2]}\right)^2\\ &\le\left(\Theta(1/T)+\sqrt{\mathbb{E}[Z_T]^2+\text{Var}[Z_T]}\right)^2\\ &\le\left(\Theta(1/T)+\sqrt{\Theta(T^{-2\gamma})}\right)^2\\ &=(\Theta(1/T)+\Theta(T^{-\gamma}))^2\\ &=\Theta(T^{-\gamma})^2\\ &=\Theta(T^{-2\gamma}).\end{aligned}$$

Next, note that $\text{Var}[X_T]\le\mathbb{E}[X_T^2]$, so

$$\text{Var}[X_T]=O(T^{-2\gamma}) \tag{9}$$

as well. Combining (8) and (9), we have that

$$\Pr[X_T\ge T^{\delta-1}]\le\frac{\text{Var}[X_T]}{(T^{\delta-1}-\mathbb{E}[X_T])^2}=O(T^{-2\gamma+2(1-\delta)}).$$

Because $-2\gamma+2(1-\delta)<1$, given our assumption that $3/4<\gamma<\delta$, it follows that $\Pr[X_T\ge T^{\delta-1}]=o(1/T)$, which allows us to conclude that

$$\sum_{k=T^\gamma+1}^{T}\Pr[W_T^{(k)}>T^\delta]=o(1).$$

Because we showed earlier that $\sum_{k=1}^{T^\gamma}\Pr[W_T^{(k)}>T^\delta]=o(1)$, we have that

$$\sum_{k=1}^{T}\Pr[W_T^{(k)}>T^\delta]=o(1)$$

as needed. □

We are now ready to prove the theorem about upward delegation.

Proof of Theorem 1. To prove the theorem, we will prove that the upward delegation model with respect to $\mathfrak{D}^C$ satisfies (1), (2), and (3), which implies that the model satisfies probabilistic do no harm and positive gain by Lemma 1. We show (1) here and relegate (2) and (3) to Online Appendix B.3. □

3.1. Upward Delegation Satisfies (1)

To do this, we will simply show that the component sizes in G_n sampled according to $\mathbb{P}_{D,M,n}$ have the same distribution as the bin sizes in a Pólya urn process with attachment probability p, and hence, max-weight(G_n) follows the same distribution as L_n^p. Once we have shown this, (1) follows immediately from Lemma 2 as $n^\delta \in o(n)$.

To that end, fix some sampled competencies $\vec{p}_n$. Recall that each entry p_i in $\vec{p}_n$ is sampled i.i.d. from $\mathcal{D}$, a continuous distribution. Hence, almost surely, no two competencies are equal. From now on, we condition on this probability 1 event. Now, consider sampling the delegation graph G_n. By the design of the model M_p^U, we can consider a random process for generating G_n that is isomorphic to sampling according to $\mathbb{P}_{D,M,n}$ as follows. First, order the competencies $p_{(1)} > p_{(2)} > \cdots > p_{(n)}$ (note that such strict order is possible by our assumption that all competencies are different), and rename the voters such that voter i has competence $p_{(i)}$; then construct G_n iteratively by adding the voters one at a time in decreasing order of competencies: voter 1 at time 1, voter 2 at time 2, and so on.

We start with the voter with the highest competence: voter 1. By the choice of φ, voter 1 places weight 0 on every other voter and hence, by definition, does not delegate. This voter forms the first component in the graph G_n, which we call $C^{(1)}$. Then, we add voter 2, who either delegates to voter 1 joining component $C^{(1)}$ with probability p or starts a new component $C^{(2)}$ with probability $1-p$. Next, we add voter 3. If $2 \in C^{(1)}$ (that is, if voter 2 delegated to voter 1), voter 3 either delegates to voter 1 (either directly or through voter 2 by transitivity) with probability p, or she starts a new component $C^{(2)}$. If $2 \in C^{(2)}$, then voter 3 delegates to voter 1 with probability $p/2$ and is added to $C^{(1)}$, delegates to voter 2 with probability $p/2$ and is added to $C^{(2)}$, or starts a new component $C^{(3)}$. In general, at time t, if there are k existing components $C^{(1)}, \ldots, C^{(k)}$, voter t either joins each component $C(j)$ with probability $\frac{p|C(j)|}{t-1}$ or starts a new component with probability $1-p$. To construct G_n, we run this process for n steps. Notice that this is identical to the Pólya urn process with bins $B^{(k)}$, balls replaced with components $C^{(k)}$, and voters being run for n steps as needed.

4. Confidence-Based Delegation Model

We now explore a model according to which voters delegate with probability that is strictly decreasing (or monotonically decreasing; that is, $x < y$ implies $f(x) > f(y)$) in their competence, and when they do decide to delegate, they do so by picking a voter uniformly at random. This models the case where voters do not need to know anything about their peers' competencies, but they do have some sense of their own competence and delegate accordingly.

Formally, for any q, let $M_q^C = (q, \varphi^1)$, where $\varphi^1(p_i, p_j) = 1$ for all $i, j \in [n]$. Voter i puts equal weight on all of the voters and hence, samples one uniformly at random when they delegate. We refer to M_q^C as the confidence-based delegation model.

Theorem 2 (Confidence-Based Delegation Model). *All models M_q^C with monotonically decreasing q satisfy probabilistic do no harm and probabilistic positive gain with respect to the class $\mathfrak{D}^C$ of all continuous distributions.*

Proof. We show that the confidence-based model satisfies (1) and (2) here, and we relegate showing (3) to Online Appendix B.4. □

4.1. Confidence-Based Delegation Satisfies (1)

Fix some distribution $\mathcal{D} \in \mathfrak{D}^C$. We show that there exists $C(n) \in O(\log n)$ such that (1) holds.

Note that when sampling an instance $(\vec{p}_n, G_n)$, the probability that an arbitrary voter i chooses to delegate is precisely $p := \mathbb{E}_{\mathcal{D}}[q]$. To see this, consider how a voter i chooses whether to delegate; the voter first samples a competence $p_i \sim \mathcal{D}$ and then, samples whether to delegate from Bern($q(p_i)$). Treating this as a single process, it is clear that the overall probability of choosing to delegate is exactly $\mathbb{E}_{\mathcal{D}}[q]$ by integrating out the competence.

Further, because $\mathcal{D}$ is continuous and q is monotonically decreasing, $p \in (0,1)$. When a voter does decide to delegate, the voter does so by picking another voter uniformly at random. Hence, we can consider the marginal distribution of delegation graphs directly (ignoring the competencies). We will show that when sampling a delegation graph for any specific voter i, with probability $1 - o(1/n)$, $\mathsf{dels}_i(G_n) \le C(n)$, which implies $\mathsf{weight}_i(G_n) \le C(n)$. A union bound over all n voters implies max-weight$(G_n) \le C(n)$ with probability $1 - o(1)$.

To that end, we will describe a branching process similar to the well-known *graph branching process* (Alon and Spencer 2016), which has the property that the distribution of its size exactly matches the distribution of $\mathsf{dels}_i(G_n)$ for an arbitrary voter i. We will compare this process with a known graph branching process that has size at most $O(\log n)$ with high probability. We will show that our process is sufficiently dominated such that it too has size at most $O(\log n)$ with high probability. The branching process works as follows. Fix our voter i. We sample which other voters end up in i's "delegation tree" (i.e., its ancestors in G_n) dynamically over a sequence of time steps. As is standard for these processes, all voters V will be one of three types: live, dead, or neutral. Dead voters are those whose "children" (i.e., voters who delegate to them) we have already sampled. Live voters are voters who have decided to delegate but whose children have not yet been sampled. Neutral voters are still in the "pool" and have yet to commit to a delegation. At time 0,

i is a live voter, there are no dead voters, and all other voters $V \setminus \{i\}$ are neutral. At each time step, we take some live voter j, sample which of the neutral voters choose to delegate to j, add these voters as live vertices, and update j as dead. The procedure ends when there are no more live vertices, at which point the number of delegations received by i is simply the total number of dead vertices.

Let us now describe this more formally. Following the notation of Alon and Spencer (2016), let Z_t denote the number of voters who we sample to delegate at time t. Let Y_t be the number of live vertices at time t; we have that $Y_0 = 1$. At time t, we remove one live vertex and add Z_t more, so we have the recursion $Y_t = Y_{t-1} - 1 + Z_t$. We let N_t be the number of neutral vertices at time t. We have that $N_0 = n - 1$ and $N_t = N_{t-1} - Z_t$. Note that after t time steps, there are t dead vertices and Y_t live ones, so this is equivalent to $N_t = n - 1 - t - Y_t$. To sample Z_t, we fix some live voter j and ask how many of the neutral voters chose to delegate to j conditioned on them not delegating to any of the dead voters. Note that when sampling at this step, there are $t - 1$ dead voters, and conditioned on the neutral voters not delegating to the dead ones, the probability that they delegate to any of the other $n - t$ individuals (not including themselves) is exactly $\frac{p}{n-t}$, equally split between them for a total delegation probability of p. Hence, $Z_t \sim \mathrm{Bin}(N_{t-1}, \frac{p}{n-t}) \sim \mathrm{Bin}(n - t - Y_{t-1}, \frac{p}{n-t})$. We denote by $\mathfrak{X}^D_{n,p}$ the random variable that counts the size of this branching process (i.e., the number of time steps until there are no more live vertices). Note that the number of delegations received by any voter has the same distribution as $\mathfrak{X}^D_{n,p}$.

Choose some constant p' such that $p < p' < 1$. We will be comparing the $\mathfrak{X}^D_{n,p}$ with a graph branching process $\mathfrak{X}^G_{n,p'}$. The graph branching process is nearly identical except that the probability that each of the neutral vertexes joins our component is independent of the number of dead vertices and is simply $\frac{p'}{n}$. In other words, $Z_t \sim \mathrm{Bin}(N_{t-1}, \frac{p'}{n})$. A key result about this branching process is the probability of seeing that a component of a certain size ℓ decreases exponentially with ℓ. In other words, there is some constant c such that

$$\mathbb{P}_{\mathcal{D}, M^C_q, n}[\mathfrak{X}^G_{n,p'} \le c \log(n)] = 1 - o(1/n).$$

Take $C(n) = c\log(n)$. Note that as long as t is such that $\frac{p}{n-t} \le \frac{p'}{n}$, the sampling in the delegation branching process is dominated by the sampling in this graph branching process. Hence, as long as $\frac{p}{n-C(n)} \le \frac{p'}{n}$, $\mathbb{P}[\mathfrak{X}^D_{n,p} \le c \log(n)] \ge \mathbb{P}[\mathfrak{X}^G_{n,p'} \le c\log(n)]$. Because $C(n) \in O(\log n)$, this is true for sufficiently large n, so for such n, $\mathbb{P}[\mathfrak{X}^D_{n,p} \le c\log(n)] = 1 - o(1/n)$. By a union bound over all n voters, this implies the desired result.

4.2. Confidence-Based Delegation Satisfies (2)

Let $\overline{q}$ be such that $\overline{q}(x) = 1 - q(x)$, so $\overline{q}$ represents the probability that someone with competence x does not delegate. Notice that $\mathbb{E}_{\mathcal{D}}[\overline{q}]$ is exactly the probability that an arbitrary voter will not delegate. Let $q^+(x) = \overline{q}(x)x$, and let

$$\mu^* = \frac{\mathbb{E}_{\mathcal{D}}[q^+]}{\mathbb{E}_{\mathcal{D}}[\overline{q}]}.$$

Expanding the definition, we see that μ^* is exactly the expected value of a voter's competence, conditioned on them not voting. Let $\mu_{\mathcal{D}}$ be the mean of the competence distribution $\mathcal{D}$. We first show that $\mu^* > \mu_{\mathcal{D}}$. Indeed, because both x and $\overline{q}(x)$ are strictly increasing functions of x, the Fortuin–Kasteleyn–Ginibre inequality (Fortuin et al. 1971) tells us that $\mathbb{E}_{\mathcal{D}}[q^+] > \mathbb{E}_{\mathcal{D}}[\overline{q}] \cdot \mathbb{E}_{\mathcal{D}}[x] = \mathbb{E}_{\mathcal{D}}[\overline{q}] \cdot \mu_{\mathcal{D}}$. This implies that the expected competence conditioned on not delegating is strictly higher than the overall expected competence.

Next, we will show that for any constant $\gamma > 0$, with high probability, both $\sum_{i=1}^n p_i \le (\mu + \gamma)n$ and $\sum_{i=1}^n \mathsf{weight}_i(G)p_i \ge (\mu^* - \gamma)n$. If we choose $\gamma = (\mu^* - \mu)/3$ and $\alpha = \gamma/2$, it follows that with high probability,

$$\sum_{i=1}^n \mathsf{weight}_i(G)p_i - \sum_{i=1}^n p_i \ge 2\alpha n,$$

implying that (2) is satisfied.

Because the p_i's are bounded independent variables, it follows directly from Heoffding's inequality that $\sum_{i=1}^n p_i \le n(\mu + \gamma)$ with high probability, so we now focus on showing $\sum_{i=1}^n \mathsf{weight}_i(G) \cdot p_i \ge (\mu^* - \gamma)n$ with high probability. To do this, we will first show that with high probability, the delegation graph G satisfies $\mathsf{dels}_i(G) \le C(n)$ for all i and $\mathsf{total\text{-}weight}(G) \ge n - C(n)\log^2 n$.

We showed in the earlier part of this proof that $\mathsf{dels}_i(G) \le C(n)$ with high probability. We will now prove that $\mathbb{P}_{\mathcal{D}, M^C_q, n}[\mathsf{total\text{-}weight}(G) \ge n - C(n)\log^2 n \mid \mathsf{dels}_i(G) \le C(n)] = 1 - o(1)$. To do this, we will first bound the number of voters who with high probability, end up in cycles. Fix a voter i and sample i delegation tree. Voter i will only end up in a cycle if i chooses to delegate to someone in this delegation tree. Because we are conditioning on $\mathsf{dels}_i(G) \le C(n)$, the maximum size of this tree is $C(n)$. Hence, the total φ weight that voter i places on someone in the tree is at most $C(n)$, whereas the total weight that voter i place on all voters is $n - 1$. Hence, the probability that i delegates to someone in the tree can be at most $p \cdot C(n)/(n-1)$. Because this is true for each voter i, the expected number of voters in cycles is at most $np\frac{C(n)}{(n-1)} \in O(\log n)$. By Markov's inequality, the probability that more than $\log^2 n$ voters are in cycles is at most $np\frac{C(n)}{(n-1)\log^2 n} = O(1/\log n) = o(1)$.

Next, because we have conditioned on $\mathsf{dels}_i(G) \le C(n)$, no single voter and in particular, no single voter in a cycle can receive more than $C(n)$ delegations. So, conditioned on the high-probability event that there are at most $\log^2 n$ voters in cycles, there are at most $C(n)\log^2 n$ voters who delegate to those in cycles. This implies that $\text{total-weight}(G) \geq n - C(n)\log^2 n + \log^2 n$ with high probability.

We now show that conditioned on the graph satisfying these properties, the instance $(\vec{p}, G)$ satisfies $\sum_{i=1}^{n} \text{weight}_i(G) \cdot p_i \geq n(\mu^* - \gamma)$ with high probability. Note that the competencies satisfy that those that do not delegate are drawn i.i.d. from the distribution of competencies conditioned on not delegating, which has mean μ^*. Fix an arbitrary graph G satisfying the properties. Suppose M is the set of voters who do not delegate. Note that for each $i \in M$, $\text{weight}_i(G) \leq C(n)$ by assumption. Further, $\sum_{i \in M} \text{weight}_i(G) \geq n - C(n)\log^2(n)$. Hence, when we sample the nondelegator p_i's, $\mathbb{E}[\sum_{i \in M} \text{weight}_i(G) \cdot p_i] \geq (n - C(n)\log^2(n)) \cdot \mu^*$. Moreover,

$$\mathrm{Var}\left[\sum_{i \in M} \text{weight}_i(G) \cdot p_i\right] \leq \sum_{i \in M} \text{weight}_i(G)^2 \leq C(n) \cdot n.$$

This follows from the fact that $\mathrm{Var}[p_i] \leq 1$ and that we have fixed the graph G and hence, $\text{weight}_i(G)$ for each i, so these terms can all be viewed as constants. In addition, we know that for each voter i, $\text{weight}_i(G) \leq C(n)$, and $\sum_{i=1}^{n} \text{weight}_i(G) \leq n$. Hence, we can directly apply Chebyshev's inequality:

$$\begin{aligned}
&\mathbb{P}_{\mathcal{D}, M_q^C, n}\left[\sum_{i \in M} \text{weight}_i(G) p_i < n(\mu^* - \gamma)\right] \\
&< \frac{\mathrm{Var}\left[\sum_{i \in M} \text{weight}_i(G) p_i\right]}{\left(\mathbb{E}\left[\sum_{i \in M} \text{weight}_i(G) p_i\right] - n(\mu^* - \gamma)\right)^2} \\
&\leq \frac{nC(n)}{(\gamma n - C(n)\log^2(n)\mu^*)^2} \\
&\in o(1),
\end{aligned}$$

where the final step holds because the numerator is $o(n^2)$ and the denominator is $\Omega(n^2)$. Hence, $\sum_{i \in M} \text{weight}_i(G) p_i \geq n(\mu^* - \gamma)$ with high probability as needed.

To summarize, we have proven that conditioned on $\text{dels}_i(G) \leq C(n)$ for all i and $\text{total-weight}(G) \geq n - C(n)\log^2 n$, $\sum_{i=1}^{n} \text{weight}_i(G) \cdot p_i \geq n(\mu^* - \gamma/3)$ occurs with high probability. Given that conditioned on $\text{dels}_i(G) \leq C(n)$, $\text{total-weight}(G) \geq n - C(n)\log^2 n$ occurs with high probability and $\text{dels}_i(G) \leq C(n)$ occurs with high probability, we can conclude by the chain rule that the intersection of these events holds with high probability. Given that the probability of any of this event is greater than the probability of the intersection, we can conclude that $\sum_{i=1}^{n} \text{weight}_i(G) \cdot p_i \geq n(\mu^* - \gamma/3)$ occurs with probability $1 - o(1)$ as desired.

5. Continuous General Delegation Model

Finally, we study a model in which voters delegate with fixed probability, and they do so by picking a voter according to a continuous increasing delegation function. This is a general model in which delegations can go to either more competent neighbors or less competent neighbors but where more competent voters are more likely to be chosen over less competent ones.

Formally, let $M_{p,\varphi}^S = (q^p, \varphi)$, where q^p is a constant function equal to p: that is, $q^p(x) = p$ for all $x \in [0,1]$, and $\varphi(x,y)$ is nonzero, continuous, and increasing in y. We then have the following.

Theorem 3 (Continuous General Delegation Model). *All models $M_{p,\varphi}^S$ with $p \in (0,1)$ and φ that is nonzero, continuous, and increasing in its second coordinate satisfy probabilistic do no harm and probabilistic positive gain with respect to the class $\mathfrak{D}^C$ of all continuous distributions.*

Proof. A majority of the proof is related to Online Appendix B.5. In the following, we show the beginning of the proof, which describes the setup for proving (1).

Fix $M_{p,\varphi}^S$ and $\mathcal{D} \in \mathfrak{D}^C$. Note that because φ is continuous and always positive on the compact set $[0,1]^2$, φ is in fact uniformly continuous, and there are bounds $L, U \in \mathbb{R}^+$ such that φ is bounded in the interval $[L, U]$. Additionally, we can assume without loss of generality that for all $x \in [0,1]$, $\mathbb{E}_{\mathcal{D}}[\varphi(x,\cdot)] = 1$. Indeed, $\mathbb{E}_{\mathcal{D}}[\varphi(x,\cdot)]$ is a positive, continuous function of x, so replacing φ by $\varphi'(x,y) = \varphi(x,y)/\mathbb{E}_{\mathcal{D}}[\varphi(x,\cdot)]$ induces the same model and satisfies the desired property.

5.1. The Continuous General Delegation Model Satisfies (1)

Our goal is to show there is some $C(n) \in O(\log n)$ such that with high probability, no voter receives more than $C(n)$ delegations. To do this, just as in the proof of Theorem 2, we consider a branching process of the delegations received beginning with some voter i. We will show that under minimal conditions on the sampled competencies (which all occur with high probability), this branching process will be dominated by a well-known *subcritical multitype Poisson branching process* (Bollobás et al. 2007), which has size $O(\log n)$ with high probability.

For a fixed competence vector $\vec{p}_n$, the branching process for the number of delegations received by a voter i works as follows. We keep track of three sets of voters: those who are alive at time t (L_t), those who are dead at time t (D_t), and those who are neutral at time t (N_t). Unlike in the proof of Theorem 2, where it was sufficient to keep track of the number of voters in each category, here we must keep track of the voter identities as well as they do not all delegate with the same probability. At time 0, the only live voter is voter i, and the rest are neutral; so, $L_0 = \{i\}$, $D_0 = \emptyset$, and $N_0 = [n] \setminus \{i\}$. As long as there are still live voters, we sample the next set of delegating voters Z_t in time t by choosing some live voter $j \in R_{t-1}$ and sampling the voter's children. Once

j's children are sampled, j becomes dead, and j's children become live. All voters who did not delegate and were not delegated to remain neutral. The children are sampled independently; the probability that they are included is the probability that they delegate to j conditioned on them not delegating to the dead voters in D_{t-1}. For each voter $k \in N_{t-1}$, k will be included with probability

$$p \cdot \frac{\varphi(p_k, p_j)}{\sum_{k' \in [n] \setminus (D_{t-1} \cup \{k\})} \varphi(p_k, p_{k'})}.$$

This is precisely the probability that k delegates to j conditioned on them not delegating to any voter in D_{t-1}. We continue this process until there are no more live voters, at which point the number of delegations is simply the number of dead voters or equivalently, the total number of time steps. We denote by $\mathfrak{X}^D_{\vec{p}_n, i}$ the size of the branching process parameterized by competencies $\vec{p}_n$ and a voter $i \in [n]$.

Our goal will be to compare $\mathfrak{X}^D_{\vec{p}_n, i}$ with the outcome of a well-known multitype Poisson branching process. In this branching process, there are a fixed finite number k of types of voters.[8] The process itself is parameterized by a $k \times k$ matrix M, where $M_{\tau\tau'}$ is the expected number of children of type τ' that a voter of type τ will have. The process is additionally parameterized by the type $\tau \in [k]$ of the starting voter. The random variable Y_t keeps track of the number of live voters of each type; it is a vector of length k, where the τth entry is the number of live voters of type τ. Hence, $Y_0 = e_\tau$, the (basis) vector with a one in entry τ and an entry of zero for all other types. We sample children by taking an arbitrary live voter of type τ' (the τ' component in Y_{t-1} must be positive, indicating that there is such a voter) and sampling the voter's children Z_t, which is also a vector of length k, each entry indicating the number of children of that type. The vector Z_t is sampled such that the τ'' entry is from the $\text{Pois}(M_{\tau'\tau''})$ distribution. That is, children of different types are sampled independently from a Poisson distribution, with the given expected value. We have the recursion $Y_t = Y_{t-1} + Z_t - e_{\tau'}$.

Note that this means that there is no "pool" of voters to choose from; in fact, it is possible for this process to grow unboundedly large (see Alon and Spencer 2016, section 11.6 for the classical description of the single-type Poisson branching process). Nonetheless, this process will still converge often enough to remain useful. We denote by $\mathfrak{X}^P_{M,\tau}$ the random variable that gives the size of this branching process, parameterized by expected-children matrix M and starting voter type $\tau \in [k]$. Such a branching process is considered *subcritical* if the largest eigenvalue of M is strictly less than one (Bollobás et al. 2007). In such a case, if we begin with a voter of any type $\tau \in [k]$, the probability of the branching process surviving ℓ steps decreases exponentially in ℓ. Hence, there is some c such that for all $\tau \in [k]$,

$$\mathbb{P}[\mathfrak{X}^P_{M,\tau} \le c \log(n)] = 1 - o(1/n).$$

To compare these branching processes, we make a sequence of adjustments to the original branching process that at each step creates a dominating branching process slightly closer in flavor to the multitype Poisson. In the end, we will be left with a subcritical multitype Poisson process that we can bound.

Fix some $\varepsilon > 0$, which is a parameter in all of our steps. Later, we will choose ε to be sufficiently small (specifically, such that $p\frac{(1+\varepsilon)^3}{1-2\varepsilon} < 1$) to ensure that the Poisson branching process is subcritical. To convert from our delegation branching process to the Poisson branching process, we take a voter's type to be the competence (which completely characterizes the delegation behavior). However, to compare with the Poisson process, there must be a finite number of types. Hence, we partition the interval $[0,1]$ into B buckets, each of size $1/B$, such that voters in the same bucket delegate and are delegated to "similarly." We choose B large enough such that all points in $[0,1]^2$ within a distance of $\sqrt{2}/B$ of each other differ in φ by at most $L \cdot \varepsilon$. (Recall that the range of φ is in the interval $[L, U]$.) This is possible because φ is uniformly continuous. Further, this implies that any points $(x,y), (x',y')$ within a square with side length $1/B$ have the property that $\varphi(x,y) \le \varphi(x',y') + L \cdot \varepsilon \le (1+\varepsilon) \cdot \varphi(x',y')$. Note that B depends only on φ and ε and hence, is a constant with respect to the number of voters n.

We say a voter i is of type τ if $\frac{\tau-1}{B} < p_i \le \frac{\tau}{B}$ for $1 \le \tau \le B$ (with a nonstrict inequality for $\tau = 1$, so zero is of type 1). Let $S_\tau = (\frac{\tau-1}{B}, \frac{\tau}{B}]$ be the set of competencies of type τ (except that in the case that $\tau = 1$, we take S_1 to be the closed interval $[0, \frac{1}{B}]$). Let $\pi_\tau = \mathcal{D}[S_\tau]$ be the probability that a voter has type τ. Because the types form a partition of $[0,1]$, we have that $\sum_{\tau=1}^{B} \pi_\tau = 1$.

For any two types τ, τ', we define[9]

$$\varphi'(\tau, \tau') = \sup_{(x,y) \in S_\tau \times S_{\tau'}} \varphi(x,y).$$

We abuse notation by extending φ' to operate directly on competencies in $[0,1]$ by first converting competencies to types and then, applying φ'. Then, φ' has the property that for any $p_i, p_j \in [0,1]$,

$$\varphi(p_i, p_j) \le \varphi'(p_i, p_j) \le (1+\varepsilon)\varphi(p_i, p_j).$$

We have that for all τ, if $x \in S_\tau$, then

$$\sum_{\tau'=1}^{B} \varphi'(\tau, \tau')\pi_{\tau'} = \mathbb{E}_{\mathcal{D}}[\varphi'(x, \cdot)] \le (1+\varepsilon) \cdot \mathbb{E}_{\mathcal{D}}[\varphi(x, \cdot)] = (1+\varepsilon).$$

Hence, we define

$$\tilde{\varphi}(\tau, \tau') = \varphi'(\tau, \tau') \cdot \frac{(1+\varepsilon)}{\sum_{\tau''=1}^{B} \varphi'(\tau, \tau'')\pi_{\tau''}}.$$

We again abuse notation to allow $\tilde{\varphi}$ to operate directly on competencies. We have that $\tilde{\varphi}(x, y) \geq \varphi'(x, y) \geq \varphi(x, y)$ for all competencies $x, y \in [0, 1]$, and further, for all τ, $\sum_{\tau'=1}^{B} \tilde{\varphi}(\tau, \tau')\pi_{\tau'} = 1 + \varepsilon$.

The Poisson branching process that we will eventually compare with is one with B types parameterized by the expected-children matrix M, where

$$M_{\tau\tau'} = p\frac{(1+\varepsilon)^2}{1-2\varepsilon}\tilde{\varphi}(\tau, \tau').$$

First, we show that M has largest eigenvalue strictly less than one (for our choice of ε) so that the branching process will be subcritical. Indeed, M has only positive entries, so we only need to exhibit an eigenvector with all nonnegative entries such that the associated eigenvalue is strictly less than one. The Perron–Frobenius theorem tells us that this eigenvalue must be maximal. □

The remainder of this proof can be found in Online Appendix B.5. At a high level, we give details for proving that the Poisson process is subcritical as well as completing the comparison between the original delegation process and this one. The comparison makes use of the concentration of the number of voters in each bucket. The proofs of (2) and (3) follow a similar structure to confidence based; however, they are quite a bit more intricate because of the interdependencies between competence level and delegation probability.

6. Liquid Democracy in Experiments

In six experiments, we statistically estimate the functions q and φ to assess the real-world implications of our theoretical findings. These experiments rely on a novel design measuring the vote in a nonstrategic, nonincentivized liquid democracy setting while simultaneously estimating voters' competence. Our empirical results consistently exhibit a regime in which liquid democracy enhances collective intelligence, leveraging interpersonal knowledge embedded in social networks and identifying diverse sets of experts. Data and code are available at http://tinyurl.com/osf-liqdem.[10]

6.1. Experimental Design

6.1.1. Experiments and Material. We conducted $E = 6$ experiments after an initial pretest[11] between March 21 and November 27, 2022.[12] In each experiment e, a group of participants[13] performed $|\mathcal{T}_e|$ tasks.[14] Each task consisted of eight questions on the corresponding topic that were primarily taken and adapted from the work of Simoiu et al. (2019).[15] A total of $N = 168$ individuals participated. They hailed from over 30 countries; 33% were female, 1% were nonbinary, 64% were male, and 2% preferred to self-describe. Each experiment e had a number of participants N_e ranging between 14 and 50. A description of the settings and group sizes is presented in Online Appendix C, and the survey material can be found in Online Appendix D.

6.1.2. Survey Flow. Participants began by providing informed consent and inputting their name. Next, they completed the following steps.

First experimental stage. Participants were presented with a task and could either answer a series of questions related to that theme or delegate the task to a peer. For instance, a task read: "You will be shown images of architectural landmarks from around the world, and asked to select the country where the landmark is located" (https://github.com/ManRev/liquiddemocracy/blob/main/material/LD_template.qsf). This was followed by "Do you want to vote yourself or delegate your vote to a trusted peer?" If they chose to vote themselves, they were taken to the eight questions contained in the task. If they chose to delegate, they were asked to select the name of their delegate and then, immediately directed to the next task. Importantly, when deciding whether to delegate, participants did not see the questions.

Second experimental stage. Participants were then asked to answer "additional questions." These were all of the questions corresponding to tasks that they had chosen to delegate in the first stage. We collected these data at the end of the experiment so as not to prime the participants on the exercise.[16]

Finally, optional background questions were asked on the last page. Note that the orders in which tasks, questions within each task, and the "true/false" options appeared were all randomized. The entire flow is summarized in Figure 3 in Online Appendix D.

6.1.3. Data Collected. Let $[N]$ be the set of N participants, and let $[E]$ be the set of E experiments. Each experiment $e \in [E]$ has N_e participants so that $N = \sum_{e\in[E]} N_e$. $[N_e]$ denotes the subset of voters in experiment e and $\mathcal{T}$ is the set of tasks surveyed ($|\mathcal{T}| = 15$). For each task $t \in \mathcal{T}$, there is a set R_t of eight corresponding questions. We let $R = \cup_t R_t$ be the set of all questions. For each participant i, $e(i) \in [E]$ is the experiment that participant i participated in; for each question r, $t(r) \in \mathcal{T}$ is its corresponding task.

In the experiments, we collect (i) the direct vote to each question i answered $v_{i,r} \in \{0, 1\}$, where one means correct and zero means incorrect, and (ii) the binary signal $\delta_{i,t}$ equal to one if i delegated on task t and zero otherwise (note that $\delta_{i,t}$ is constant at the task level), along with which voter they delegated to. From these collected data, we can compute $w_{i,t}$, the weight of voter i on task t. This is i's total weight after adding up all transitive delegations; it is set to zero when i delegates. Figure 1 provides an example of a collected delegation graph.

In rare cases, a delegation could not be included for a couple of possible reasons. The first reason is if a

Figure 1. (Color online) Delegation Graphs for Task T_7 ("You Will Be Given Upcoming European Men Soccer Games and Asked to Predict the Games' Outcomes") from Experiment 6

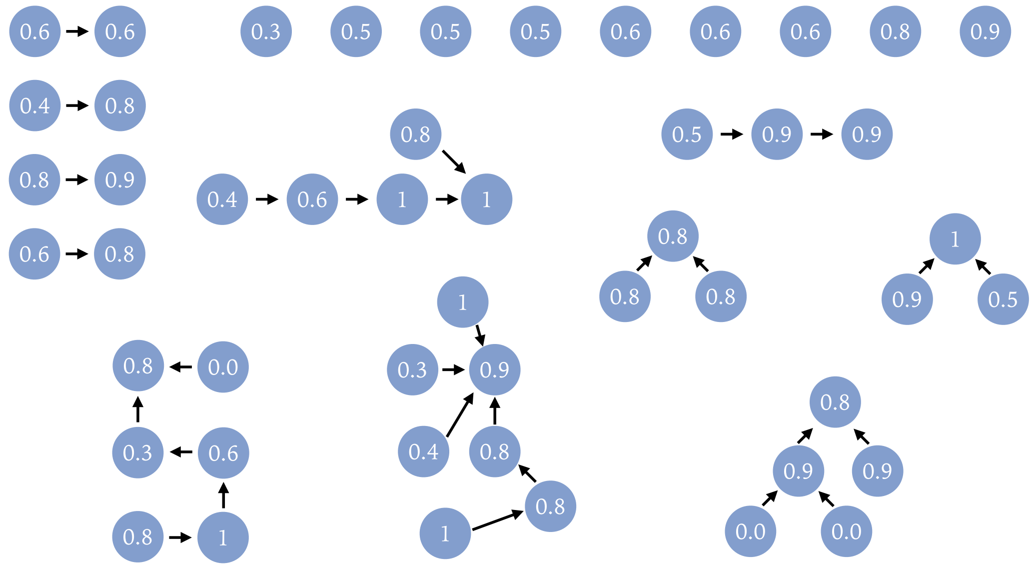

Note. Each node is a voter, and the node's number represents the rounded expertise $\eta_{i,t}$ of a given voter i for task t computed using item response theory; see Online Appendix E.

participant delegated to somebody who did not complete the survey. In this case, we would simply ignore the delegation (assuming that they directly voted). The second reason is in an instance of a cycle (e.g., participant i delegated to participant j who delegated to participant i). These were also ignored (i.e., assumed that no voter on the cycle delegated). In many real-world implementations, such participants would be notified of the cycle and asked to choose a new delegate or vote directly.

6.2. Delegation and Competence Statistics

Over the 1,096 (participant/task) pairs, we observed a total of 505 delegations, meaning that participants delegated 47% of the time (standard deviation (*std*) = 0.49). The rate varied across experiments from 32% (*std* = 0.49) in experiment 2 to 54% (*std* = 0.50) in experiment 5, and the rate varied across tasks from 22% (*std* = 0.50) in task T_8 to 80% (*std* = 0.40) in task T_{15}. Among those who voted directly, 15% received only one delegation besides their own (hence, had weight 2 in the decision), 6% received two delegations, and just about 5% received five or more delegations. However, in one experiment, over half of the votes were delegated to a single participant. Additionally, throughout all of the experiments, we observed only four delegation cycles, and all were only of size 2 (where a delegates to b and then, b delegates back to a). These occurred in Experiment 4 with $N_4 = 27$ and in Experiment 6 with $N_6 = 50$. Examples of additional delegation graphs can be found in Online Appendix D.

6.2.1. Estimating Competence. In order to evaluate how delegation behavior relates to competence, we need to estimate participants' competence. We denote by $\eta_{i,t}$ the estimated competence of participant i in task t. Naively, participants' competence per task could be estimated by averaging the number of correct answers given on all eight questions of that task, $\eta^{\text{naive}}_{i,t} = \frac{\sum_{r \in R_t} v_{i,r}}{|R_t|}$. However, such a computation does not account for the questions' heterogeneity. We thus estimate $\eta_{i,t}$ using the item response theory (IRT) framework (Lalor and Rodriguez 2023), which provides a widely used parametric model to estimate competence $\eta_{i,t}$ and question difficulty from repeated measurements. We explain the parametric estimation in Online Appendix E.[17]

6.2.2. Gender-Based Statistics. Although we might worry that delegation patterns vary across gender because of significant differences in confidence (e.g., Ellis et al. 2016, Sarsons and Xu 2021), we actually find no significant differences in these experiments, neither in measured competence in tasks nor in propensity to delegate. Analysis of variance tests for the propensity to delegate (respectively, competence) across gender show no significant differences with $p = 0.464$ (respectively, $p = 0.112$). Tukey tests for pair-wise mean comparison further validate the absence of significant differences across the different genders (see Online Appendix G).

6.3. Estimating the Probability of Delegating as a Function of Competence

We now turn to estimating q and φ as a function of voters' competence. Recall that $q(\eta)$ represents the probability that somebody of competence η chooses to delegate. We have observations $\delta_{i,t}$ encoding participant i's delegation choice for task t and an estimate $\eta_{i,t}$ of i's competence on task t. We use these to estimate the relationship between competence $\eta_{i,t}$ and the probability of delegating $q(\eta_{i,t})$.

6.3.1. Methods. To estimate q, we fit a logistic model, regressing $\delta_{i,t}$ against $\eta_{i,t}$. The following equation shows the relationship that we wish to fit, where α_0 is the intercept and β^q is the effect size that we measure:

$$\log\left(\frac{\Pr[\delta_{i,t}=1]}{1-\Pr[\delta_{i,t}=1]}\right)=\alpha_0+\beta^q\eta_{i,t}+\varepsilon_i. \qquad (10)$$

To account for potential correlation in the error term within participants' answers, when estimating the parameters in Equation (10), we cluster standard errors (*s.e.*s) at the participant level. We also test for the data normality; results for these test can be found in Online Appendix H.

We repeat the procedure above on data sets filtered by task, this time having a distinct β_t for each task t to measure the task-specific estimates. Additionally, we run these with fixed effects for individuals and tasks to more directly measure the impact of competence (rather than just looking at population trends). Additional details and results can be found in Online Appendix I.

6.3.2. Results. We find $\beta^q=-2.24$, with $s.e.=0.42$, statistics $z=-7.12$, and $p=10^{-7}$. In turn, we estimate that $q(\eta_{i,t})=\widehat{\Pr}[\delta_{i,t}=1]=\frac{1}{1+\exp^{-(-1.39-2.24\times\eta_{i,t})}}$, suggesting that the probability of delegating decreases with competence. We can also test for monotonic dependence through a model-free method using a Pearson correlation test and its associated p-value. We find a correlation coefficient of –0.17 and $p<5\times10^{-8}$.

6.4. Estimating Weight Function Used to Delegate

Recall that in the theoretical model, a voter with competence η_1 delegates to another with competence η_2 with probability proportional to $\varphi(\eta_1,\eta_2)$.

6.4.1. Methods. We first bucket the observed competence levels into B clusters $c_1,\dots,c_B$. We assume that φ is constant across inputs in the same bucket and fit it based on bucket "centers," $\eta_1,\dots\eta_B$, which are simply taken to be the mean values of the competences in each bucket (i.e., $\eta_\ell=\frac{\sum_{i,t:\eta_{i,t}\in c_\ell}\eta_{i,t}}{|\{(i,t)|\eta_{i,t}\in c_\ell\}|}$). This means that we can estimate $\varphi(x,y)$ using the number of delegations from any competence x' to competence y', where x' and y' fall in the same bucket as x and y, respectively. Finally, we determine the Kendall tau rank correlation coefficient between $\varphi(x,y)$ and y with its associated p-value to test for the monotonic relation between φ and its second coordinate.

6.4.1.1. Bucketing Strategies. We discretize the segment $[0,1]$ into B buckets. We do so using several methods (to ensure the robustness of our approach); we describe here the k-means clustering procedure and discuss the rest in Online Appendix J.1.

To bucket using k-means, we optimize for B clusters, $c_1,\dots,c_B$, that minimize $\sum_{k=1}^{B}\sum_{\eta_{i,t}\in c_k}\left(\eta_{i,t}-\frac{\sum_{\eta_{i,t}\in c_k}\eta_{i,t}}{|c_k|}\right)^2$. In words, we compute a partition of the $[0,1]$ segment such that the total squared distance from elements to their cluster centers is minimized. We use the standard k-means clustering algorithm to find the clusters (Hartigan and Wong 1979).

6.4.1.2. Estimation of φ for a Given Delegation Graph. Next, we wish to fit a function φ. For given experiment e and task t, we estimate $\varphi_{e,t}(\eta_\ell,\eta_k)$ for each $\ell,k\in[B]$, so for conciseness, we write $\varphi^\ell_{e,t}(\eta_k):=\varphi_{e,t}(\eta_\ell,\eta_k)$.[18]

Observe that the experiment can then be viewed as a multinomial trial; from the perspective of an ℓ participant, there are B choices to pick from, where the probability of picking a k participant is proportional to all of the $\varphi^\ell_{e,t}(\eta_k)$ for $k\in[B]$ and the number of participants of each competence level. We observe instances of these choices; the z^ℓ_k's are the observed numbers of times that someone of type ℓ delegated to someone of type k. It then suffices to find the maximum likelihood estimators for $\varphi^\ell_{e,t}(\eta_k)$ as a function of z^ℓ_k, n_k, and n_ℓ. We provide more technical details in Online Appendix J.2.

6.4.1.3. Testing for Monotonic Dependence of φ in Its Second Coordinate. Finally, we test for potential monotonic dependence of $\varphi^\ell_{e,t}(\eta_k)$ as a function of η_k, visualizing the $\varphi^\ell_{e,t}(\eta_k)$ in Figure 2 and computing the Kendall tau rank correlation coefficient between $\varphi^\ell_{e,t}(\eta_k)$ and η_k for a fixed ℓ and its associated p-value. The Kendall tau rank correlation coefficient evaluates the similarity between two vectors of rank; its form and significance are detailed in the Online Appendix.

6.4.2. Results. We show here the results for using k-means bucketing with $B=4$.[19] Additional results for other strategies and other numbers of buckets can be found in Online Appendix J.5. Descriptions of these four buckets can be found in Table 1.

In Figure 2, the blue crosses in each column show the $\varphi^\ell_{e,t}(\eta_k)$ for all (e, t) (for a particular η_ℓ in the four plots in the left panel and for all combined in the right panel) as a function η_k. The pink points in Figure 2 represent the average across all experiments and tasks for a given η_k, and the regression lines in Figure 2 correspond to ordinary least square regression on the mean values. We show this pooled both only for those in the same bucket as well as all grouped together.

Figure 2. (Color online) Estimates of $\varphi_{e,t}^{\ell}$ for Each Bucket (Left) and for Buckets Grouped (Right)

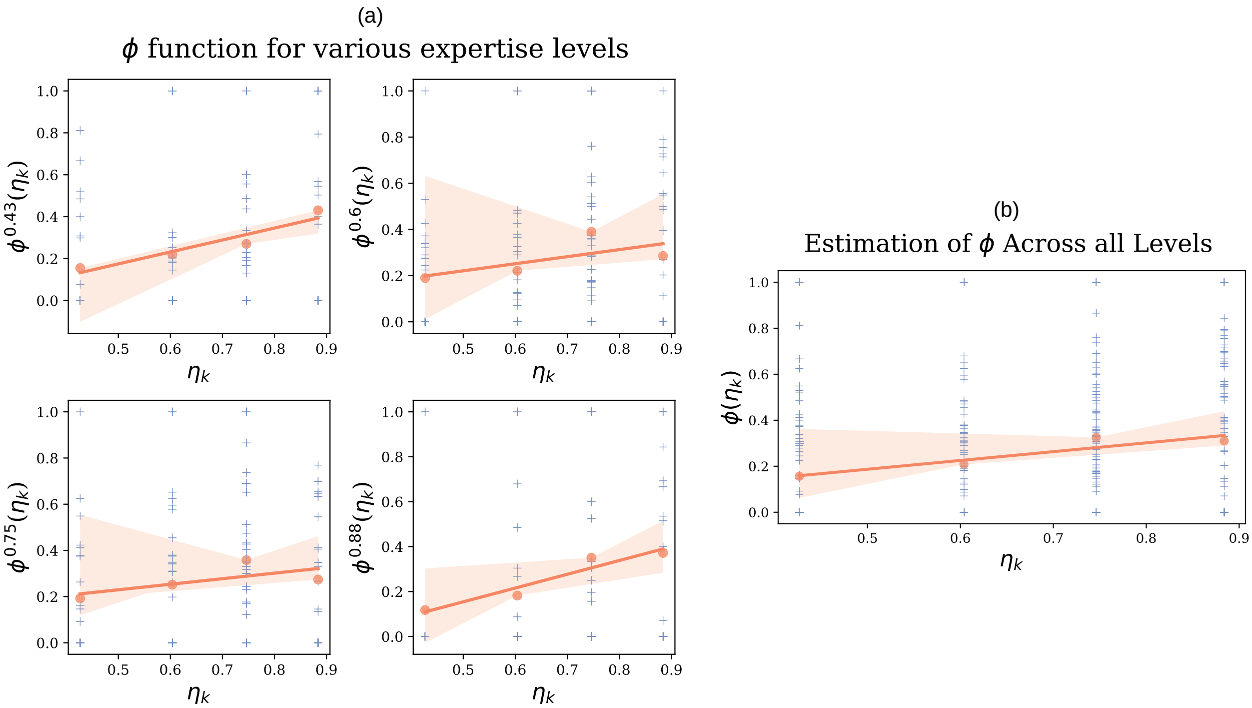

Notes. The crosses show the values computed for $\varphi_{e,t}^{\ell}(\eta_k)$. The dots show the average across all values for that η_k, and the lines correspond to a linear regression over the mean values. We observe increasing trends across the board, with slopes (coefficient of determination) being 0.53(0.90), 0.28(0.46), 0.29(0.47), and 0.60(0.92), respectively, for individual buckets and 0.38(0.85) for the pooled test. The shaded areas represent the 95% confidence intervals. (a) Estimated values separated by input bucket. (b) Estimated values pooled together.

To test the significance of the trends observed in Figure 2, we test whether the Kendall tau rank correlation coefficient between $\varphi_{e,t}^{\ell}(\eta_k)$ and η_k signals significant associations both at the overall level and when fixing ℓ or η_ℓ, the first coordinate in $\varphi_{e,t}(\eta_\ell, \eta_k)$. Table 2 shows the resulting correlation coefficients and significance tests. The trends observed in both Figure 2 and Table 2 confirm that there is a statistically significant increase in $\varphi_{e,t}(\eta_\ell, \eta_k)$ as a function η_ℓ across all expertise levels, confirming that voters behave according to the general continuous delegation model. We further note that these significant trends are valid at the granularity of three of the four buckets (the third bucket c_3 exhibits nonstatistically significant positive Kendall tau rank correlation). We also run the same tests partitioned into tasks. The results can be found in Online Appendix J.4. We last check that these results are not sensitive to the bucketing strategy in Online Appendix J.5.

6.5. Experimental Conclusions

We found that voters' likelihood to delegate decreases with their competence (as suggested by the confidence-based model). In addition, voters are more likely to delegate to someone of increasing competence (as suggested by the general continuous model). Note that in fact, the general continuous model can be generalized to allow monotonically decreasing q as well (indeed, it suffices to consider the expectation of $q(p_i)$ taken over the distribution of competence as the constant probability of voting). Our empirical results are hence consistent with a general continuous delegation model. Unsurprisingly, the upward delegation model that we described based on Caragiannis and Micha (2019) and Kahng et al. (2021) that leads to a catastrophic concentration of power is not consistent with experimental data; voters do not delegate only to higher-competency agents, and we do not observe constant delegation competence.

Table 1. Bucket Descriptions

Bucket	Interval	Mean competence	Proportion of participants (%)
c_1	[0.00–0.514]	0.43	16
c_2	[0.515–0.674]	0.60	32
c_3	[0.677–0.814]	0.75	35
c_4	[0.818–1.00]	0.88	17

Table 2. Summary of Correlation Effects

		For fixed ℓ			
	Overall	c_1	c_2	c_3	c_4
Correlation	0.17****	0.29**	0.12*	0.11	0.28***
p-value	2×10^{-5}	2×10^{-2}	9×10^{-2}	1×10^{-1}	3×10^{-3}

*$p < 0.1$; **$p < 0.05$; ***$p < 0.01$; ****$p < 0.0001$.

6.6. Additional Results

In Online Appendix K, we run additional tests and check other properties of the collected data. These include comparing the frequency of correctness between liquid democracy and direct democracy (Online Appendix K.2), analyzing the increase in competence (Online Appendix K.3), and analyzing the concentration of power using both the maximum weight (Online Appendix K.4) and the power of small coalitions (Online Appendix K.5). Note that the latter results are particularly relevant to the concentration of power in liquid democracy, a topic that was extensively discussed in previous work. Recall that although previous work exhibited scenarios in which a concentration of power occurs, our work is concerned with measuring whether such an extreme concentration of power is likely. In all but 1 of the 32 instances, we found no evidence of a concentration of power. Anecdotally, we further observe that a number roughly the square root of the number of voters controls half of the votes.

7. Discussion

Our paper relies on a set of assumptions and modeling choices that are worth discussing.

First, a prominent feature of our model is that there is no underlying social network; that is, there is no restriction on whom a voter may delegate to. As we explained in Section 1, we believe that this is realistic in a variety of scenarios. But, we can, in fact, extend our results to a model where a directed social network is first sampled, and then, a (q, φ) model is followed. The social network must be sampled such that the neighbors of each voter are chosen uniformly at random, although the number of such neighbors could follow any small-tailed distribution. Intuitively, delegation proportional to weighting the neighbors of i (rather than the entire population) is equivalent to a possibly different weighting over the entire population.[20] An open research direction is to consider graph topologies not covered by these dynamics.

Second, building on Kahng et al. (2021), we assume that there exists a true best alternative. Needless to say, this assumption is necessary if we wish to "defend" liquid democracy against their conclusions. It is also an extremely well-studied assumption that dates back to the eighteenth century (Young 1988). The existence of a ground truth is easily justified in the contexts of prediction markets or corporate governance, where alternative policies can be measured in terms of concrete metrics, like "estimated revenue in five years," and these metrics can be communicated to voters. That said, some decisions inherently rely on other subjective criteria that we do not capture.

Third, again like previous papers (Caragiannis and Micha 2019, Becker et al. 2021, Kahng et al. 2021), we assume that voters vote independently. Admittedly, this is not a realistic assumption; relaxing it, as it was relaxed for the classic Condorcet jury theorem (Häggström et al. 2006, Nitzan and Paroush 2017), is a natural direction for future work.

Fourth, our models do not take strategic behaviors into account. In the same vein, our experiments do not involve explicit incentives, and although we do not have reasons to believe that significant strategic voting occurs, studying it these issues is beyond the scope of this work. It would be of interest to extend our results to a more game-theoretic setting and relate them to work focusing on game-theoretic issues in liquid democracy (Bloembergen et al. 2019, Zhang and Grossi 2021, Dhillon et al. 2023). Along those lines, experiments with monetary incentives would be interesting.

Fifth, our framework for stochastic delegations opens up interesting directions for more research. For instance, one could characterize all of the delegation models satisfying positive gain and do no harm. One may also consider more general local delegation models that would depend on the competences of all voters.

More generally, our work aims to provide a better understanding of a prominent shortcoming of liquid democracy: concentration of power. But, there are others. For example, any voter can see the complete delegation graph under current liquid democracy systems—a feature that helps voters make informed delegation decisions (because one's vote can be transitively delegated). This may lead to voter coercion, however, and the trade-off between transparency and security is poorly understood.

Finally, to summarize, we have introduced a general framework to investigate stochastic dynamics in liquid democracy and proved new conditions for the convergence of weighted majorities; we then identified regimes in which liquid democracy leads to correct outcomes with high probability. In that sense, our work is to liquid democracy what the Condorcet jury theorem is to direct democracy. There are many reasons to be excited about the potential of liquid democracy (Blum and Zuber 2016). We believe that our results provide another such reason, and we hope that our techniques will be useful in continuing to build the theoretical and empirical understanding of this compelling paradigm.

Endnotes

[1] See https://tinyurl.com/y52j6nfs, last accessed January 2, 2025.

[2] The use of the term "epistemic" in this context is well established in the social choice literature (List and Goodin 2001, Pivato 2012).

[3] The former constructs an instance where even with arbitrarily many voters, a constant number will receive a majority of the delegations. The group has an average competence above 1/2. The probability that liquid democracy gives the right answer can be upper bounded by a constant strictly below one, whereas direct democracy is correct with probability approaching one. In the latter

case, the voters' numbers and relative competence are chosen so that liquid democracy almost always gives the incorrect answer (as does direct voting), whereas dictatorship is correct with a constant probability.

[4] An infinite Pólya urn process models an urn process where each new ball picks its urn with a probability proportional to the size of the urn or creates its own urn with constant probability.

[5] In LiquidFeedback, delegation cycles are, in fact, ignored.

[6] This is a worst-case approach, where cycles can only hurt the performance of liquid democracy because this assumption is equivalent to assuming that all voters on the cycles vote incorrectly.

[7] Note that positive gain and do no harm relate to the notion of concentration of the weighted sum $\sum_{i=1}^{n} \mathsf{weight}_i V_i$. Indeed, the probability of direct democracy being correct approaches one as n increases when the average competence is strictly above $1/2$. As a result, do no harm is satisfied by a delegation model exactly when the probability that liquid democracy is correct also approaches one. This happens when the competence after delegation remains strictly above $1/2$, and the weighted sum $\sum_{i=1}^{n} \mathsf{weight}_i V_i$ concentrates. Positive gain also holds if there exists a setup where the average group competence is strictly below $1/2$, the average competence after delegation remains strictly above $1/2$, and the weighted sum $\sum_{i=1}^{n} \mathsf{weight}_i V_i$ concentrates. In turn, these established benchmarks are directly mapped to existing results in social choice theory on the convergence of weighted majorities Häggström et al. (2006).

[8] In the literature, these are often called *particles*, but to be consistent with our other branching processes, we call them voters here.

[9] Note that because φ is increasing in its second coordinate, one can actually write $\tilde{\varphi}(\tau, \tau') = \sup_{x \in S_\tau} \varphi(x, \frac{\tau'}{B})$.

[10] The full link is https://osf.io/skxwg/?view_only=3671d431bcfd4a9cb94ded5aa86a0a95.

[11] A description of the setup and the results from the prestudy can be found in Online Appendix L, and initial results can be found in Revel et al. (2022).

[12] Our protocol E-3948 was approved and exempted by the university's Committee on the Use of Humans as Experimental Subjects.

[13] Note that liquid democracy depends on the potential for beneficial delegation. It is, therefore, necessary to work with participants who have at least a passing familiarity with each other. Experiments were conducted in places such as classrooms and company workshops, where pre-existing group structures guaranteed such conditions. Although significant preparation was needed to ensure correct experimental setups for these environments, this design did have the benefit of producing high-quality data with few missing entries and minimal drop out.

[14] $|\mathcal{T}_e| = 4$ except for experiment $e = 6$, in which $|\mathcal{T}_e| = 12$; the final experiment was conducted over a longer period of time, allowing more tasks to be completed.

[15] To be consistent with the theoretical setup under study, we converted all categorical questions into binary questions. For example, for a question from Simoiu et al. (2019) of the form "Where is this famous landmark from? (see in https://github.com/stanford-policylab/wisdom-of-crowds/blob/master/data/original/tasks.csv.zip)" with four options (Italy, Tibet, Greece, or Brazil), we selected a possible answer (e.g., Brazil) to reformulate the question: "Is this famous landmark from Brazil?" In more detail, we first randomly selected which questions would be correct (to avoid the sense that most questions are incorrect) and then, for the incorrect ones, drew a wrong option at random. We found multiple inconsistencies in the Simoiu et al. (2019) data that we corrected; also, the prediction questions pertained to events that had passed, so these were replaced with new ones.

[16] We validated this approach with a robustness check on the time spent by participants as a function of how often they delegated (see Online Appendix K.1).

[17] Although $\eta_{i,t}^{\text{naive}}$ takes on one of nine values (multiples of $1/8$), $\eta_{i,t}$ (computed using IRT) is a continuous variable that can take on arbitrary values in $\mathbb{R}$. We normalize so that $\eta_{i,t} \in [0,1]$ and assume this to be the competence, the probability of being correct. Note that these different methods yield a correlation between $\eta_{i,t}^{\text{naive}}$ and $\eta_{i,t}$ of more than 94%.

[18] Note that because the number of participants in each bucket changes for different experiments/tasks, it is difficult to fit a single function. Instead, we first found the most likely φ to have generated each experiment/task and then, combined these to find an overall best fit.

[19] This was the optimal number found using the *Kneedle algorithm* (Satopaa et al. 2011).

[20] This extension does not carry over to undirected networks because if voters have a small number of neighbors, we would expect many two cycles to form after delegation, which under the worst-case cycle approach, may not be canceled out by the overall increase in competence.

References

Alon N, Spencer JH (2016) *The Probablistic Method* (John Wiley & Sons, Hoboken, NJ).

Atanasov P, Rescober P, Stone E, Swift SA, Servan-Schreiber E, Tetlock P, Ungar L, Mellers B (2017) Distilling the wisdom of crowds: Prediction markets vs. prediction polls. *Management Sci.* 63(3):691–706.

Barabási AL, Albert R (1999) Emergence of scaling in random networks. *Science* 286(5439):509–512.

Becker R, D'angelo G, Delfaraz E, Gilbert H (2021) Unveiling the truth in liquid democracy with misinformed voters. *Proc. 7th Internat. Conf. Algorithmic Decision Theory (ADT)* (Springer, Cham, Switzerland), 132–146.

Benhaim A, Falk BH, Tsoukalas G (2023) Scaling blockchains: Can committee-based consensus help? *Management Sci.* 69(11):6525–6539.

Bloembergen D, Grossi D, Lackner M (2019) On rational delegations in liquid democracy. *Proc. 33rd AAAI Conf. Artificial Intelligence (AAAI)* (AAAI Press, Palo Alto, CA), 1796–1803.

Blum C, Zuber CI (2016) Liquid democracy: Potentials, problems, and perspectives. *J. Political Philos.* 24(2):162–182.

Bollobás B, Janson S, Riordan O (2007) The phase transition in inhomogeneous random graphs. *Random Structures Algorithms* 31(1):3–122.

Brill M, Talmon N (2018) Pairwise liquid democracy. *Proc. 27th Internat. Joint Conf. Artificial Intelligence (IJCAI)* (AAAI Press, Palo Alto, CA), 137–143.

Campbell J, Casella A, de Lara L, Mooers VR, Ravindran D (2022) Liquid democracy. Two experiments on delegation in voting. Preprint, submitted December 19, https://arxiv.org/abs/2212.09715v1.

Caragiannis I, Micha E (2019) A contribution to the critique of liquid democracy. *Proc. 28th Internat. Joint Conf. Artificial Intelligence (IJCAI)* (ijcai.org), 116–122.

Chen MK, Ingersoll JE Jr, Kaplan EH (2008) Modeling a presidential prediction market. *Management Sci.* 54(8):1381–1394.

Christoff Z, Grossi D (2017) Binary voting with delegable proxy: An analysis of liquid democracy. *Proc. 16th Conf. Theoret. Aspects Rationality Knowledge (TARK)* 251:134–150.

Chung F, Handjani S, Jungreis D (2003) Generalizations of Polya's urn problem. *Ann. Combinatorics* 7(2):141–153.

Collevecchio A, Cotar C, LiCalzi M (2013) On a preferential attachment and generalized Pólya's urn model. *Ann. Appl. Probab.* 23(3):1219–1253.

Dhillon A, Kotsialou G, Ravindran D, Xefteris D (2023) Information aggregation with delegation of votes. Preprint, submitted June 6, https://arxiv.org/abs/2306.03960.

Drinea E, Enachescu M, Mitzenmacher MD (2001) Variations on random graph models for the web. Technical report, Harvard Computer Science Group, Cambridge, MA.

Durrett R (2007) *Random Graph Dynamics* (Cambridge University Press, Cambridge, UK).

Eggenberger F, Pólya G (1923) Über die statistik verketteter vorgänge. *ZAMM* 3(4):279–289.

Ellis J, Fosdick BK, Rasmussen C (2016) Women 1.5 times more likely to leave stem pipeline after calculus compared with men: Lack of mathematical confidence a potential culprit. *PLoS One* 11(7):e0157447.

Fortuin CM, Kasteleyn PW, Ginibre J (1971) Correlation inequalities on some partially ordered sets. *Comm. Math. Phys.* 22(2): 89–103.

Gautschi W (1959) Some elementary inequalities relating to the gamma and incomplete gamma function. *J. Math. Phys.* 38(1):77–81.

Gölz P, Kahng A, Mackenzie S, Procaccia AD (2018) The fluid mechanics of liquid democracy. *Proc. 14th Conf. Web Internet Econom. (WINE)* (Springer Nature, Cham, Switzerland), 188–202.

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

Häggström O, Kalai G, Mossel E (2006) A law of large numbers for weighted majority. *Adv. Appl. Math.* 37(1):112–123.

Hardt S, Lopes LCR (2015) Google votes: A liquid democracy experiment on a corporate social network. Technical Disclosure Commons. Accessed December 27, 2024, https://www.tdcommons.org/dpubs_series/79.

Hartigan JA, Wong MA (1979) A k-means clustering algorithm. *J. Roy. Statist. Soc. Ser. C Appl. Statist.* 28(1):100–108.

Huang J (2023) Thy neighbor's vote: Peer effects in proxy voting. *Management Sci.* 69(7):4169–4189.

Janson S (2020) Rate of convergence for traditional Pólya urns. *J. Appl. Probab.* 57(4):1029–1044.

Johnson NL, Kotz S (1978) *Urn Models and Their Application: An Approach to Modern Discrete Probability Theory* (Wiley, New York).

Kahng A, Mackenzie S, Procaccia AD (2021) Liquid democracy: An algorithmic perspective. *J. Artificial Intelligence Res.* 70:1223–1252.

Kling CC, 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. *Proc. 9th Internat. AAAI Conf. Web Social Media (ICWSM)* (AAAI Press, Palo Alto, CA).

Lalor JP, Rodriguez P (2023) PY-IRT: A scalable item response theory library for python. *INFORMS J. Comput.* 35(1):5–13.

Li C, Xu R, Duan L (2023) Liquid democracy in DPOs blockchains. *Proc. 5th ACM Internat. Sympos. Blockchain Secure Critical Infrastructure* (ACM, New York), 25–33.

List C, Goodin RE (2001) Epistemic democracy: Generalizing the Condorcet jury theorem. *J. Polit. Philos.* 9(3):277–306.

Mahmoud H (2009) *Pólya Urn Models* (CRC Press, Boca Raton, FL).

Markov AA (1917) Sur quelques formules limites du calcul des probabilités. *Bulletin de l'Académie des Sciences* 11(3):177–186.

Nitzan S, Paroush J (2017) Collective decision making and jury theorems. Parisi F, ed. *The Oxford Handbook of Law and Economics: Volume 1: Methodology and Concepts* (Oxford Academic, Oxford, UK), 494–516.

Pivato M (2012) A statistical approach to epistemic democracy. *Episteme* 9(2):115–137.

Pólya G (1930) Sur quelques points de la théorie des probabilités. *Annales de l'institut Henri Poincaré* 1(2):117–161.

Revel M, Halpern D, Berinsky A, Jadbabaie A (2022) Liquid democracy in practice: An empirical analysis of its epistemic performance. *Proc. 2nd ACM Conf. Equity Access Algorithms Mechanisms Optimization (EAAMO)* (ACM, New York).

Sarsons H, Xu G (2021) Confidence men? Evidence on confidence and gender among top economists. *AEA Papers Proc.* (American Economic Association, Nashville, TN), vol. 111, 65–68.

Satopaa V, Albrecht J, Irwin D, Raghavan B (2011) Finding a "Kneedle" in a haystack: Detecting knee points in system behavior. *2011 31st Internat. Conf. Distributed Comput. Systems Workshops* (IEEE, Piscataway, NJ), 166–171.

Simoiu C, Sumanth C, Mysore A, Goel S (2019) Studying the "wisdom of crowds" at scale. *Proc. AAAI Conf. Human Comput. Crowdsourcing* 7(1):171–179.

Simon HA (1955) On a class of skew distribution functions. *Biometrika* 42(3/4):425–440.

Young HP (1988) Condorcet's theory of voting. *Amer. Political Sci. Rev.* 82(4):1231–1244.

Zhang Y, Grossi D (2021) Power in liquid democracy. *Proc. 35th AAAI Conf. Artificial Intelligence (AAAI)* (AAAI Press, Palo Alto, CA), 5822–5830.