Tracking Truth with Liquid Democracy

Adam J. Berinsky
Department of Political Science, MIT

Daniel Halpern
School of Engineering and Applied Sciences, Harvard University

Joseph Y. Halpern
Department of Computer Science, Cornell University

Ali Jadbabaie
Institute for Data, Systems, and Society, MIT

Elchanan Mossel
Department of Mathematics, MIT

Ariel Procaccia
School of Engineering and Applied Sciences, Harvard University

Manon Revel
Institute for Data, Systems, and Society, MIT
Corresponding author: mrevel@mit.edu

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 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 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 $11k$ 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.

Key words: 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 them 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. model the problem as a decision problem where voters decide on an issue with two outcomes, $\{0, 1\}$, where 1 is correct (the ground truth) and 0 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

---

[1] See https://tinyurl.com/y52j6nfs.

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

the number of votes that voter i gathered transitively after delegation (if voter i delegates, then weight$_i = 0$). The outcome of the election is then determined by a weighted majority; it is correct if and only if $\sum_{i=1}^{n} \text{weight}_i V_i \geq n/2$.

Kahng et al. 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. 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 the correct outcome is selected under liquid democracy and the probability 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* 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. 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.

---

[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 liquid democracy gives the right answer can be upper bounded by a constant strictly below 1, while 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), while dictatorship is correct with a constant probability.

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 lab 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 and liquid democracy, as well as to analyze properties of real voter delegation behavior. Importantly, the behaviors we observe align with one of the models 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 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 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 φnote 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

---

[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.

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 $\mathcal{D}$. While 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 2.1, 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 the maximum weight max-weight(G_n) accumulated by any voter is sub-linear with high probability and 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.

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 does assume that voters act in the interest of society by only delegating to voters that 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 the size of the component. This works because delegating to any voter in the previous components is possible (since 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.

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 sub-critical 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 that do not delegate is strictly higher than the average competence.

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 towards 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 multi-type 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 sub-critical Poisson multi-type 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 sub-linear 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 bins'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 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 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 (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, though, 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 a central planner chooses a single delegate in order to minimize the maximum weight of any voter. Similarly, Brill and Talmon (2018) introduces 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 sfixed 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 cooncurrent 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 pre-determined experts in one set-up, and through the random dot kinematogram task in another one. The group sizes considered are 5 people with one expert, 15 people with 3 experts and 125 people with 25 experts. While 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, since 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 $[n] = \{1, \ldots, n\}$. We assume voters are making a decision on a binary issue with possible answers 0 and 1; there is a correct alternative (1) and an incorrect alternative (0). 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 $\vec{p}_n = (p_1, \ldots, p_n)$. When n is clear from the context, we sometimes drop it from the notation.

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 context, we occasionally drop it from the notation. The outdegree of a vertex in the delegation graph is at most 1 since each voter can delegate to at most one person. Voters that do not delegate have no outgoing edges. In a delegation graph G_n, the *delegations received* by a voter i, $\mathsf{dels}_i(G_n)$, is defined as the total number of people that (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 delegation they received plus their own weight) if i votes directly, and 0 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 $\mathsf{total\text{-}weight}(G_n) = \sum_{i=1}^{n} \mathsf{weight}_i(G_n)$. Since each vote is counted at most once, we have that $\mathsf{total\text{-}weight}(G_n) \leq n$. However, note that if delegation edges form a cycle, then the weight of the voters on the cycle and voters delegating into the cycle are all set to 0 and hence will not be counted. In particular, this means that $\mathsf{total\text{-}weight}(G_n)$ may be strictly less than n.[6]

---

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

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 $V_i \sim$ 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, $X_{G_n}^F = \sum_{i=1}^n \mathsf{weight}_i(G_n) \cdot V_i$. The probability 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.

Gain of a delegation instance We define the *gain* of an instance as

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

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

Randomization over delegation instances In general, we assume that both competencies and delegations are chosen randomly. Each voter's competence p_i is sampled i.i.d. 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 $1 - q(p_i)$ they do 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 sample 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 (DNH)* and *positive gain (PG)* 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 $n \geq n_0$,

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

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 $n \geq n_0$,

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

7

2.1. 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. Since the result follows from relatively straightforward concentration inequalities, we defer the proof to Appendix B.1.

LEMMA 1. *If M is a model, $\mathfrak{D}$ a class of distributions, n 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*

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

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

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

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

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, $\sum_{i=1}^{n}\mathsf{weight}_i(G_n) \cdot p_i$. Standard concentration results already imply 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 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.

---

[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 1 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 1. 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 $1/2$, and 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).

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 delegate to voters that 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 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 0 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 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 timestep $t > 1$, 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 1, 2, and 3 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 $t \geq 1$, 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 $T \geq 1$, $\Pr[L_T^p \leq T^\delta] = 1 - o(1)$.*

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 k-th bin. Let $U_t^{(k)}$ be the size of $B^{(k)}$ at time t. Since 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 $k \leq T$. 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 $t \geq k$, 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 $t \geq k$ 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 the probability $B^{(k)}$ will be chosen at time t:

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

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 since 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 consider the first T^γ bins, while 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 $k \leq n$, where Γ represents the Gamma function. We relegate the argument for Equation (5) to Appendix B.2. Using this along with Gautchi's inequality (Gautschi 1959), $(t+p-1)^p \leq \frac{\Gamma(p+t)}{\Gamma(t)} \leq (t+p)^p$, 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 Appendix B.2.

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

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

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

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

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 $p \geq 1$

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

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

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

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 we will need is that $\ell_p(X, 0) = \mathbb{E}[|X|^p]^{1/p}$ where 0 is the identically 0 random variable. Since ℓ_p is in fact a metric, the triangle inequality tells us that $\ell_p(0, X_n) \leq \ell_p(0, Z_n) + \ell_p(Z_n, X_n)$, so, combining with (6), we have that

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

for all $p \geq 1$.

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

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

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 $\Pr[X_T \geq T^{\delta-1}] = o(1)$. By Chebyshev's inequality,

$$\Pr[X_T \geq T^{\delta-1}] \leq \frac{\mathrm{Var}[X_T]}{(T^{\delta-1} - \mathbb{E}[X_T])^2}.$$

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

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

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

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

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

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

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

$$\Pr[X_T \geq T^{\delta-1}] \leq \frac{\mathrm{Var}[X_T]}{(T^{\delta-1} - \mathbb{E}[X_T])^2} = O\left(T^{-2\gamma + 2(1-\delta)}\right).$$

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

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

Since 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. $\quad\square$

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), (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 Appendix B.3

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 $\mathsf{max\text{-}weight}(G_n)$ 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 2 delegated to 1), 3 either delegates to 1 (either directly or through 2 by transitivity) with probability p or she starts a new component $C^{(2)}$. If $2 \in C^{(2)}$, then 3 either delegates to 1 with probability $p/2$ and is added to $C^{(1)}$, or delegates to 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 $B^{(k)}$ and balls replaced with components $C^{(k)}$ and voters being run for n steps, as needed. $\square$

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 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 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 satisfy (1) and (2) here, and relegate showing (3) to Appendix B.4.

Confidence-Based Delegation satisfies (1)

Fix some distribution $\mathcal{D} \in \mathfrak{D}^C$. We show 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 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: they first sample a competence $p_i \sim \mathcal{D}$ and then sample whether or not to delegate from $\text{Bern}(q(p_i))$. 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, since $\mathcal{D}$ is continuous and q is monotonically decreasing, $p \in (0, 1)$. When a voter does decide to delegate, they do 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 $1 - o(1/n)$, $\text{dels}_i(G_n) \leq C(n)$, which implies $\text{weight}_i(G_n) \leq C(n)$. A union bound over all n voters implies $\text{max-weight}(G_n) \leq C(n)$ 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 $\text{dels}_i(G_n)$ for an arbitrary voter i. We will compare this process to a known graph branching process that has size at most $O(\log n)$ with high probability. We will show 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 zero, 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 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 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 $Z_t \sim \text{Bin}(N_{t-1}, \frac{p}{n-t}) \sim \text{Bin}(n - t - Y_{t-1}, \frac{p}{n-t})$. 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 $\mathfrak{X}^D_{n,p}$ to a graph branching process $\mathfrak{X}^G_{n,p'}$. The graph branching process is nearly identical, except the probability each of the neutral vertex joins our component is independent of the number of dead vertices and is simply $\frac{p'}{n}$. In other words, $Z_t \sim \text{Bin}(N_{t-1}, \frac{p'}{n})$. A key result about this branching process is the probability of seeing a component of a certain size ℓ decreases exponentially with ℓ. In other words, there is some constant c such that

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

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

Confidence Based Delegation satisfies (2)

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

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

Expanding the definition, we see that μ^* is exactly the expected value of a voter's competence, conditioned on them not voting. Let $\mu_{\mathcal{D}}$ the mean of the competence distribution $\mathcal{D}$. We first show that $\mu^* > \mu_{\mathcal{D}}$. Indeed, since both x and $\bar{q}(x)$ are strictly increasing functions of x, the Fortuin–Kasteleyn–Ginibre (FKG) inequality (Fortuin et al. 1971) tells us that $\mathbb{E}_{\mathcal{D}}[q^+] > \mathbb{E}_{\mathcal{D}}[\bar{q}] \cdot \mathbb{E}_{\mathcal{D}}[x] = \mathbb{E}_{\mathcal{D}}[\bar{q}] \cdot \mu_{\mathcal{D}}$. 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 $\sum_{i=1}^n p_i \leq (\mu + \gamma)n$ and $\sum_{i=1}^n \mathsf{weight}_i(G)p_i \geq (\mu^* - \gamma)n$. If we choose $\gamma = (\mu^* - \mu)/3$ and $\alpha = \gamma/2$, it follows that, with high probability,

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

implying that (2) is satisfied.

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

We showed in the earlier part of this proof that $\mathsf{dels}_i(G) \leq C(n)$ with high probability. We will now prove that $\mathbb{P}_{\mathcal{D}, M_{\bar{q}}^C, n}[\mathsf{total\text{-}weight}(G) \geq n - C(n)\log^2 n \mid \mathsf{dels}_i(G) \leq C(n)] = 1 - o(1)$. To do this, we will first bound the number of voters that, with high probability, end up in cycles. Fix a voter i and sample i's delegation tree. Voter i will only end up in a cycle if i chooses to delegate to someone in this delegation tree. Since we are conditioning on $\mathsf{dels}_i(G) \leq C(n)$, 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 $C(n)$, while the total weight they place on all voters is $n - 1$. Hence, the probability that i delegates to someone in their tree can be at most $p \cdot C(n)/(n-1)$. Since 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, since we have conditioned on $\mathsf{dels}_i(G) \leq C(n)$, 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 that delegate to those in cycles. This implies that $\mathsf{total\text{-}weight}(G) \geq n - C(n)\log^2 n + \log^2 n$ with high probability.

We now show that, conditioned on the graph satisfying these properties, the instance $(\vec{p}, G)$ satisfies $\sum_{i=1}^{n} \mathsf{weight}_i(G) \cdot p_i \geq n(\mu^* - \gamma)$ with high probability. Note that the competencies satisfy that those that don't 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 that do not delegate. Note that for each $i \in M$, $\mathsf{weight}_i(G) \leq C(n)$, by assumption. Further $\sum_{i \in M} \mathsf{weight}_i(G) \geq n - C(n) \log^2(n)$. Hence, when we sample the non-delegator p_is, $\mathbb{E}[\sum_{i \in M} \mathsf{weight}_i(G) \cdot p_i] \geq (n - C(n) \log^2(n)) \cdot \mu^*$. Moreover,

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

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

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

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

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

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 either go to more 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 non-zero, 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 non-zero, 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 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 and $\mathcal{D} \in \mathfrak{D}^C$. Note that since 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.

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 multi-type 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 that are live at time t (L_t), those dead at time t (D_t), and those 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 zero, 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 its children. Once j's children are sampled, j becomes dead, and j's children become live. All voters that did not delegate and were not delegated to remain neutral. The children are sampled independently; the probability they are included is the probability 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 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 $\mathfrak{X}^D_{\vec{p}_n,i}$ to the outcome of a well-known multi-type 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' 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 1 in entry τ and an entry 0 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 its 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 *sub-critical* if the largest eigenvalue of M is strictly less than 1 (Bollobás et al. 2007). In such a case, if we begin with 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 $\tau \in [k]$,

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

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 multi-type Poisson. In the end, we will be left with a sub-critical multi-type 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 sub-critical. To convert from our delegation branching process to the Poisson branching process, we take

---

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

a voter's type to be their competence (which completely characterizes their delegation behavior). However, to compare to 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 since φ is uniformly continuous. Further, this implies any points $(x, y), (x', y')$ within a square with side length $1/B$ have the property that $\varphi(x, y) \leq \varphi(x', y') + L \cdot \varepsilon \leq (1 + \varepsilon) \cdot \varphi(x', y')$. 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 $\frac{\tau-1}{B} < p_i \leq \frac{\tau}{B}$ for $1 \leq \tau \leq B$ (with a non-strict inequality for $\tau = 1$, so 0 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. Since 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

$$\varphi'(\tau, \tau') = \sup_{(x,y) \in S_\tau \times S_{\tau'}} \varphi(x, y).\text{[9]}$$

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 $p_i, p_j \in [0, 1]$,

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

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

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

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 we will eventually compare to is one with B types parameterized by the expected-children matrix M, where

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

---

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

First, we show that M has largest eigenvalue strictly less than 1 (for our choice of ε), so that the branching process will be subcritical. Indeed, M has only positive entries, so we need only exhibit an eigenvector with all nonnegative entries such that the associated eigenvalue is strictly less than 1. The Perron-Frobenius theorem tells us this eigenvalue must be maximal.

The remainder of this proof can be found in Appendix B.5. At a high level, we give details for proving the Poisson process is subcritical, as well as completing the comparison of 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 due to the inter dependencies 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 n anon-strategic, non-incentivized liquid democracy stting, while simultaneously estimating voters' competence. Our empirical result are consistent exhibits 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 pre-test[11] between March 21st and November 27th, 2022.[12] In each experiment e, a group of participants[13] performed $|\mathcal{T}_e|$ tasks[14] Each task consisted of 8 questions on the corresponding topic that were primarily taken and adapted from the work of Simoiu et al. (2019).[15] A total of $N = 168$

---

[10] Full link: `https://osf.io/skxwg/?view_only=3671d431bcfd4a9cb94ded5aa86a0a95`

[11] A description of the setup and results from the pre-study can be found in 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 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 that 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. While significant preparation was needed to ensure correct experimental set-ups 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?" with four options (Italy, Tibet, Greece, or Brazil), we selected a possible answer (e.g., Brazil) to reformulate the question as: "Is

individuals participated. They hailed from over 30 countries; 33% were female, 1% were non-binary, 64% were male, and 2% preferred to self-describe. Each experiment e had a number of participants N_e ranging between 14 to 50. A description of the settings and group sizes are presented in Appendix C, and the survey material can be found in 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," 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 8 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 or not to delegate, participants did not see the questions.

Second experimental stage: Participants were then asked to answer "additional questions." These were all the questions corresponding to tasks they had chosen to delegate in the first stage. We collected this 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 order 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 Section D.

6.1.3.   Data Collected Let $[N]$ be the set of N participants and let $[E]$, 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 8 corresponding questions. We let $R = \bigcup_t R_t$ be the set of all questions. For each participant i, $e(i) \in [E]$ is the experiment they participated in; for each question r, $t(r) \in \mathcal{T}$ is its corresponding task.
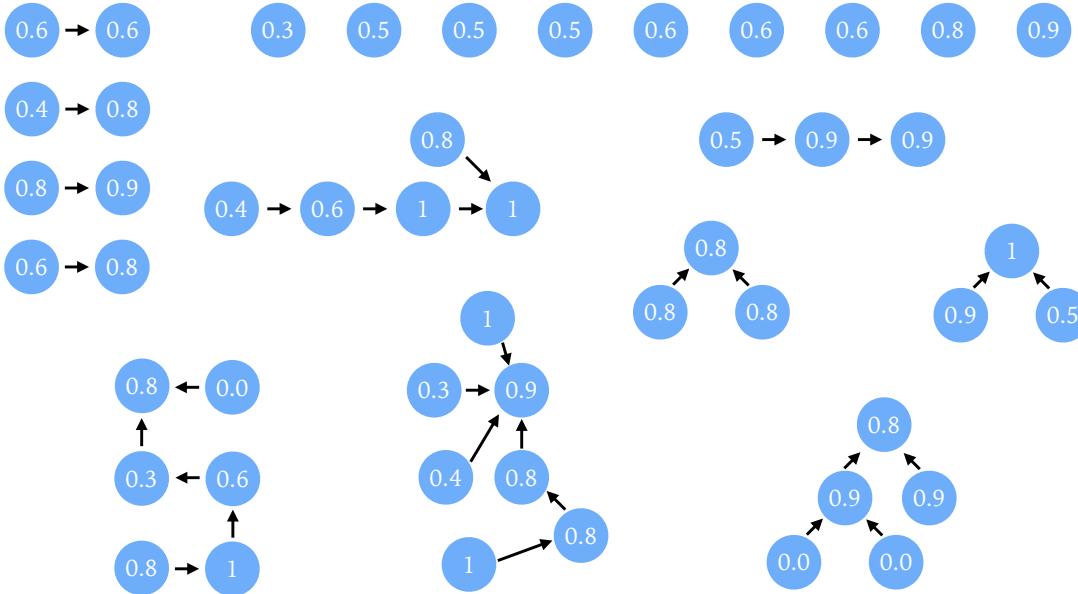
In the experiments, we collect (i) the direct vote to each question i answered $v_{i,r} \in \{0,1\}$ where 1 means correct and 0 means incorrect and (ii) the binary signal $\delta_{i,t}$ equal to 1 if i delegated on task t and 0 otherwise (note that $\delta_{i,t}$ is constant at the task level), along with which voter they

---

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, and 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 Appendix K.1).

delegated to. From this 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 0 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. First, if a participant delegated to somebody who did not complete the survey. In this case, we would simply ignore the delegation (assuming they directly voted). Second, 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.


Figure 1   **Delegation graphs for task** T_7 **("You will be given upcoming European men soccer games and asked to predict the games' outcome.") from Experiment** 6. **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 Section E.**

6.2.   Delegation and Competence Statistics

Over the 1096 (participant/task) pairs, we observed a total of 505 delegations meaning participants delegated 47% of the time ($std = 0.49$) The rate varied across experiments from 32% ($std = 0.49$) in experiment 2 to 54% ($std = 0.50$) in experiment 5 and 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 the votes

were delegated to a single participant. Additionally, throughout all the experiments, we observed only four delegation cycles, and all were only of size two (where a delegates to b, and 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 Appendix F.

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 8 questions of that task, $\eta_{i,t}^{\text{naive}} = \frac{\sum_{r \in R_t} v_{i,r}}{|R_t|}$. However, such a computation does not account for the questions' heterogeneity. We thus estimate $\eta_{i,t}$ using the Item Response Theory framework (IRT) (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 Appendix E.[17]

6.2.2. Gender-based statistics While we might worry that delegation patterns vary across gender due to 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. ANOVA tests for the propensity to delegate (resp. competence) across gender shows no significant differences with $p = 0.464$ (resp. $p = 0.112$). Tukey tests for pairwise mean comparison further validate the absence of significant differences across the different genders (see 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 we wish to fit, where α_0 is the intercept and β^q is the effect size we measure:

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

---

[17] While $\eta_{i,t}^{\text{naive}}$ 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%.

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

6.3.2. Results We find $\beta^q = -2.24$, with standard error $s.e. = 0.42$, statistics $z = -7.12$ and p-value $p = 10^{-7}$. 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 $p < 5 \times 10^{-8}$.

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 $c_1, \ldots, c_B$. We assume that φ is constant across inputs in the same bucket, and fit it based on bucket "centers," $\eta_1, \ldots \eta_B$, 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 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 monotonic relation between φ and its second coordinate.

Bucketing strategies. We discritize 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 Appendix J.1.

To bucket using k-means, we optimize for B clusters, $c_1, \ldots, c_B$, 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 et al. 1979).

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 $\varphi_{e,t}^\ell(\eta_k) := \varphi_{e,t}(\eta_\ell, \eta_k)$.[18]

---

[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.

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 the $\varphi_{e,t}^{\ell}(\eta_k)$ for $k \in [B]$ and the number of participants of each competence level. We observe instances of these choices, the z_k^{ℓ} that are the observed number of times that someone of type ℓ delegated to someone of type k. It then suffices to find the maximum likelihood estimators for $\varphi_{e,t}^{\ell}(\eta_k)$ as a function of z_k^{ℓ}, n_k and n_ℓ. We provide more technical details in Appendix J.2.

Testing for monotonic dependence of φ in its second coordinate. Finally, we test for potential monotonic dependence of $\varphi_{e,t}^{\ell}(\eta_k)$ as a function of η_k visualizing the $\varphi_{e,t}^{\ell}(\eta_k)$ in Figure 2 and computing the Kendall tau rank correlation coefficient between $\varphi_{e,t}^{\ell}(\eta_k)$ 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 is detailed in 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 Appendix J.5. Descriptions of these four buckets can be found in Table 1.
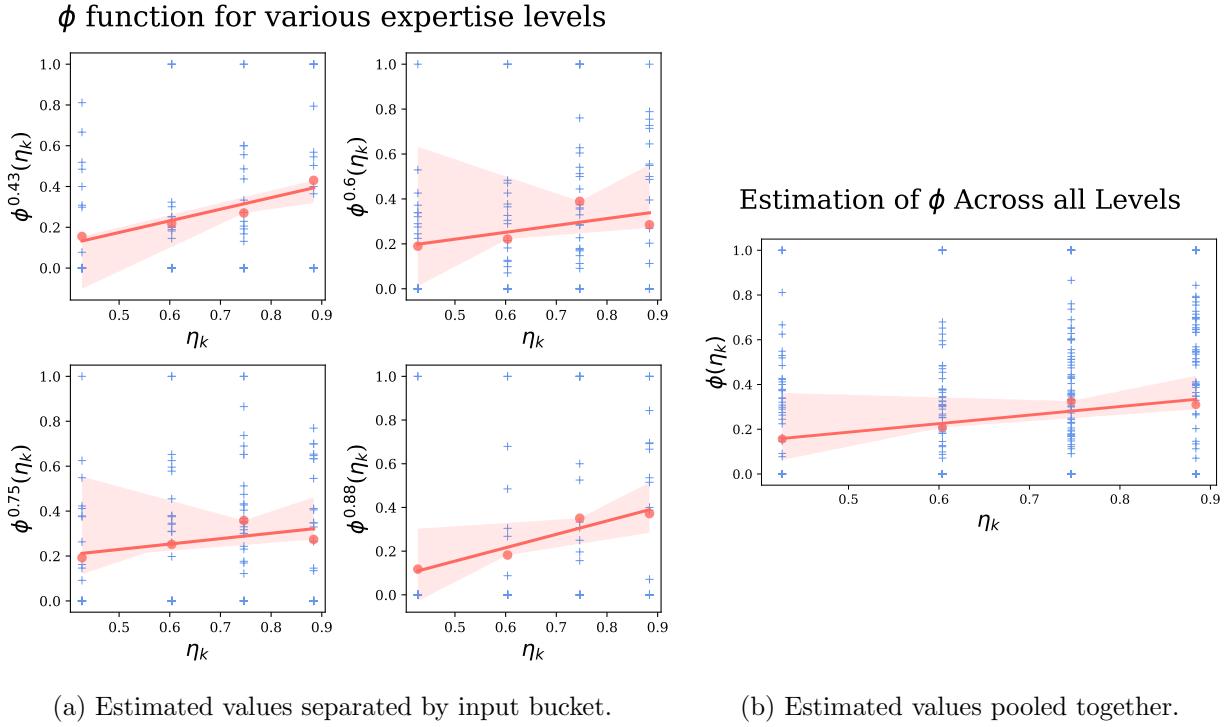
Table 1   **Bucket Descriptions**

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%

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

---

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

ϕ function for various expertise levels

(a) Estimated values separated by input bucket.


(b) Estimated values pooled together.

Figure 2 **Pooled estimates of $\varphi_{e,t}^{\ell}$, both for each bucket individually, and grouped together. The blue crosses show the values computed for $\varphi_{e,t}^{\ell}(\eta_k)$. The pink dots show the average across all values for that η_k, and the pink lines correspond to a linear regression over the mean values. We observe increasing trends across the board, with slope (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 area represents the 95% confidence interval.**

(the third bucket c_3 exhibits non-statistically significant positive Kendall tau rank correlation.) We also run the same tests partitioned into tasks. The results can be found in Appendix J.4. We last check that these results are not sensitive to the bucketing strategy in J.5.

Table 2   Summary of correlation effects

	Overall	For fixed ℓ			
		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}
Note:		*p<0.1; **p<0.05; ***p<0.01; ****p<0.0001			

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 we described based on Caragiannis and Micha (2019), 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-competence agents, and we do not observe constant delegation competence.

6.6. Additional Results

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

---

[20] This extension does not carry over to undirected networks, since if voters have a small number of neighbors, we would expect many 2-cycles to form after delegation, which, under the worst-case cycle approach, may not be canceled out by the overall increase in competence.

their conclusions. It is also an extremely well-studied assumption that dates back to the 18th 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 (Kahng et al. 2021, Caragiannis and Micha 2019, Becker 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, while 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 open interesting directions for more research. For instance, one could characterize all 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 tradeoff 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 hope that our techniques will be useful in continuing to build the theoretical and empirical understanding of this compelling paradigm.

References

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

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 science* 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. *Proceedings of the 7th International Conference on Algorithmic Decision Theory (ADT)*, 132–146.

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

Bloembergen D, Grossi D, Lackner M (2019) On rational delegations in liquid democracy. *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 1796–1803.

Blum C, Zuber CI (2016) Liquid democracy: Potentials, problems, and perspectives. *Journal of Political Philosophy* 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. *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 137–143.

Campbell J, Casella A, de Lara L, Mooers VR, Ravindran D (2022) Liquid democracy. two experiments on delegation in voting. Technical report, National Bureau of Economic Research.

Caragiannis I, Micha E (2019) A contribution to the critique of liquid democracy. *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 116–122.

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

Christoff Z, Grossi D (2017) Binary voting with delegable proxy: An analysis of liquid democracy. *Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 134–150.

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

Collevecchio A, Cotar C, LiCalzi M (2013) On a preferential attachment and generalized pólya's urn model. *The Annals of Applied Probability* 23(3):1219–1253.

Dhillon A, Kotsialou G, Ravindran D, Xefteris D (2023) Information aggregation with delegation of votes. *arXiv preprint arXiv:2306.03960* .

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

Durrett R (2007) *Random graph dynamics* (Cambridge University Press).

Eggenberger F, Pólya G (1923) Über die statistik verketteter vorgänge. *Zeitschrift für Angewandte Mathematik und Mechanik* 3(4):279–289.

Ellis J, Fosdick BK, Rasmussen C (2016) Women 1.5 times more likely to leave stem pipeline after calculus compared to 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. *Communications in Mathematical Physics* 22(2):89–103.

Gautschi W (1959) Some elementary inequalities relating to the gamma and incomplete gamma function. *Journal of Mathematics and Physics* 38(1):77–81.

Gölz P, Kahng A, Mackenzie S, Procaccia AD (2018) The fluid mechanics of liquid democracy. *Proceedings of the 14th Conference on Web and Internet Economics (WINE)*, 188–202.

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

Häggström O, Kalai G, Mossel E (2006) A law of large numbers for weighted majority. *Advances in Applied Mathematics* 37(1):112–123.

Hardt S, Lopes LCR (2015) Google votes: A liquid democracy experiment on a corporate social network. *Technical Disclosure Commons* .

Hartigan JA, Wong MA, et al. (1979) A k-means clustering algorithm. *Applied statistics* 28(1):100–108.

Hoeffding W (1963) Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58(301):13–30.

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

Janson S (2020) Rate of convergence for traditional pólya urns. *Journal of Applied Probability* 57(4):1029–1044.

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

Kahng A, Mackenzie S, Procaccia AD (2021) Liquid democracy: An algorithmic perspective. *Journal of Artificial Intelligence Research* 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. *Proceedings of the 9th International AAAI Conference on Web and Social Media (ICWSM)*.

Lalor JP, Rodriguez P (2023) py-irt: A scalable item response theory library for python. *INFORMS Journal on Computing* 35(1):5–13.

Li C, Xu R, Duan L (2023) Liquid democracy in dpos blockchains. *Proceedings of the 5th ACM International Symposium on Blockchain and Secure Critical Infrastructure*, 25–33.

List C, Goodin RE (2001) Epistemic democracy: Generalizing the Condorcet Jury Theorem. *Journal of Political Philosophy* 9(3):277–306.

Mahmoud H (2009) *Pólya Urn Models* (CRC Press).

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

Natesan P, Nandakumar R, Minka T, Rubright JD (2016) Bayesian prior choice in irt estimation using mcmc and variational bayes. *Frontiers in psychology* 7:1422.

Nitzan S, Paroush J (2017) Collective decision making and jury theorems. *The Oxford Handbook of Law and Economics* 1.

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. *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*.

Sarsons H, Xu G (2021) Confidence men? evidence on confidence and gender among top economists. *AEA Papers and Proceedings*, volume 111, 65–68 (American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203).

Satopaa V, Albrecht J, Irwin D, Raghavan B (2011) Finding a" kneedle" in a haystack: Detecting knee points in system behavior. *2011 31st international conference on distributed computing systems workshops*, 166–171 (IEEE).

Simoiu C, Sumanth C, Mysore A, Goel S (2019) Studying the "wisdom of crowds" at scale. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 7, 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. *The American Political Science Review* 82(4):1231–1244.

Zhang Y, Grossi D (2021) Power in liquid democracy. *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 5822–5830.

Tracking Truth with Liquid Democracy - Appendix

Adam J. Berinsky, Department of Political Science, MIT

Daniel Halpern, School of Engineering and Applied Sciences, Harvard University

Joseph Y. Halpern, Department of Computer Science, Cornell University

Ali Jadbabaie, Institute for Data, Systems, and Society, MIT

Elchanan Mossel, Department of Mathematics, MIT

Ariel Procaccia, School of Engineering and Applied Sciences, Harvard University

Manon Revel, Institute for Data, Systems, and Society, MIT
(Corresponding author: mrevel@mit.edu)

Appendix

A.   Hoeffding's Inequality

Throughout many of the proofs, we will make use of the following well-known concentration inequality (Hoeffding 1963):

LEMMA 3 (**Hoeffding's Inequality**). *Let $Z_1, \cdots, Z_n$ be independent, bounded random variables with $Z_i \in [a, b]$ for all i, where $-\infty < a \leq b < \infty$. Then*

$$\mathbb{P}\left[\frac{1}{n}\sum_{i=1}^{n} Z_i - \mathbb{E}[Z_i] \geq t\right] \leq \exp\left(-\frac{2nt^2}{(b-a)^2}\right)$$

and

$$\mathbb{P}\left[\frac{1}{n}\sum_{i=1}^{n} Z_i - \mathbb{E}[Z_i] \leq -t\right] \leq \exp\left(-\frac{2nt^2}{(b-a)^2}\right)$$

for all $t \geq 0$.

B. Missing Proofs

B.1. Proof of Lemma 1

We establish the two properties separately.

Probabilistic do-no-harm: We first show that a model M that satisfies conditions (1) and (2) satisfies probabilistic do no harm. Fix an arbitrary competence distribution $\mathcal{D} \in \mathfrak{D}$ and let α and C be such that (1) and (2) are satisfied. Without loss of generality, suppose that $C(n) \leq n$ for all n, as replacing any larger values of $C(n)$ with n will not affect (1) (since $\mathsf{max\text{-}weight}(G_n) \leq n$ for all graphs G_n on n vertices). Fix $\varepsilon, \delta > 0$. We must identify some n_0 such that for all $n \geq n_0$, $\mathbb{P}_{\mathcal{D},M,n}[\mathrm{gain}(\vec{p}_n, G_n) \geq -\varepsilon] > 1 - \delta$.

We will begin by showing there exists $n_1 \in \mathbb{N}$ such that for all instances $(\vec{p}_n, G_n)$ on $n \geq n_1$ voters, if both

$$\mathsf{max\text{-}weight}(G_n) \leq C(n) \text{ and} \tag{11}$$

$$\sum_{i=1}^{n} \mathsf{weight}_i(G_n) \cdot p_i - \sum_{i=1}^{n} p_i \geq 2\alpha n, \tag{12}$$

then

$$\mathrm{gain}(\vec{p}_n, G_n) \geq -\varepsilon. \tag{13}$$

Since (11) and (12) each hold with probability $1 - o(1)$ by (1) and (2), for sufficiently large n, say $n \geq n_2$, they will each occur with probability at least $1 - \delta/2$. Hence, by a union bound, for all $n \geq n_2$, they both occur with probability at least $1 - \delta$. By taking $n_0 = \max(n_1, n_2)$, this implies that probabilistic do no harm is satisfied.

We now prove that, for sufficiently large n, (11) and (12) imply (13). First, we will show that

$$\mathrm{gain}(\vec{p}_n, G_n) \geq -\mathbb{P}_{\vec{p}_n}[X_n^D > X_{G_n}^F]. \tag{14}$$

Indeed, we have that

$$\mathbb{P}_{\vec{p}_n}[X_n^D > n/2] = \mathbb{P}_{\vec{p}_n}[X_n^D > n/2, X_{G_n}^F > n/2] + \mathbb{P}_{\vec{p}_n}[X_n^D > n/2, X_{G_n}^F \leq n/2]$$

$$\leq \mathbb{P}_{\vec{p}_n}[X_{G_n}^F > n/2] + \mathbb{P}_{\vec{p}_n}[X_n^D > X_{G_n}^F]$$

where the first transition holds by the law of total probability, and the second because the corresponding events are contained in each other. That is,

$$\left\{ X_n^D > n/2, X_{G_n}^F > n/2 \right\} \subseteq \left\{ X_{G_n}^F > n/2 \right\}$$

and

$$\left\{ X_n^D > n/2, X_{G_n}^F \leq n/2 \right\} \subseteq \left\{ X_n^D > X_{G_n}^F \right\}.$$

Re-arranging the terms above yields (14).

Hence, for our purpose, it suffices to show that (11) and (12) imply $\mathbb{P}_{\vec{p}_n}\left[X_n^D > X_{G_n}^F \right] \leq \varepsilon$. Intuitively, we will use (12) to show the expected value of X_n^D is well below the expected value of $X_{G_n}^F$. Then we will show both X_n^D and $X_{G_n}^F$ concentrate well around their means, where for the latter we will need (11). Together, these observations imply that $X_{G_n}^F > X_n^D$ with high probability.

Fix an instance $(\vec{p}_n, G_n)$ on n voters satisfying (11) and (12). We will show that for sufficiently large n,

$$\mathbb{P}_{\vec{p}_n}\left[X_n^D < \sum_{i=1}^n p_i + \alpha n\right] > 1 - \varepsilon/2 \tag{15}$$

and

$$\mathbb{P}_{\vec{p}_n}\left[X_{G_n}^F > \sum_{i=1}^n \mathsf{weight}_i(G_n) \cdot p_i - \alpha n\right] > 1 - \varepsilon/2. \tag{16}$$

Note that since (12) holds for this instance, $\sum_{i=1}^n p_i + \alpha n \leq \sum_{i=1}^n \mathsf{weight}_i(G_n) \cdot p_i - \alpha n$. Therefore, when both events whose probability is considered in (15) and (16) hold, $X_n^D \leq X_n^F$. Hence,

$$\mathbb{P}_{\vec{p}_n}[X_n^D \leq X_{G_n}^F] \geq \mathbb{P}_{\vec{p}_n}\left[X_n^D < \sum_{i=1}^n p_i + \alpha n, X_{G_n}^F > \sum_{i=1}^n \mathsf{weight}_i(G_n) \cdot p_i - \alpha n\right] > 1 - \varepsilon$$

where the last inequality holds by a union bound. This implies that $\mathbb{P}_{\vec{p}_n}[X_n^D \leq X_{G_n}^F] < \varepsilon$, as needed.

It remains to be shown that (15) and (16) hold for sufficiently large n. For (15), this follows directly from Hoeffding's inequality (Lemma 3 in Appendix A). To prove (16), first note that, as shown in Kahng et al. (2021),

$$\mathrm{Var}_{\vec{p}_n}\left[X_{G_n}^F\right] = \sum_{i=1}^n \mathsf{weight}_i(G_n)^2 \cdot p_i(1 - p_i)$$

$$\leq \frac{1}{4} \cdot \sum_{i=1}^n \mathsf{weight}_i(G_n)^2$$

$$\leq \frac{1}{4} \cdot \sum_{i=1}^{\lceil n/C(n) \rceil} C(n)^2$$

$$< nC(n) \in o(n^2),$$

where the first inequality holds because $p(1 - p)$ is upper bounded by $1/4$, the second because $\sum_{i=1}^n \mathsf{weight}_i(G_n) \leq n$ with each $\mathsf{weight}_i(G_n) \leq C(n)$ so the value is maximized by setting as many terms to $C(n)$ as possible, and the final inequality holds because $C(n) \leq n$.

Hence, by Chebyshev's inequality,

$$\mathbb{P}_{\vec{p}_n}\left[X_{G_n}^F \leq \mathbb{E}_{\vec{p}_n}\left[X_{G_n}^F\right] - \alpha n\right] \leq \frac{\mathrm{Var}_{\vec{p}_n}[X_{G_n}^F]}{(\alpha n)^2}.$$

This bound is $o(1)$ because the numerator is $o(n^2)$ and the denominators is $\Omega(n^2)$. This implies that for sufficiently large n, it will be strictly less than $\varepsilon/2$, so (16) holds.

Probabilistic positive gain: Fix a distribution $\mathcal{D} \in \mathfrak{D}$ and an $\alpha \in (0, 1)$ such that (3) holds. We want to show that M satisfies probabilistic positive gain. Since $\mathcal{D} \in \mathfrak{D}$, it also satisfies (1) for some C. We show below that there exists an n_3 such that all instances $(\vec{p}_n, G_n)$ with $n \geq n_3$ voters satisfying (11) for which $\sum_{i=1}^n p_i + \alpha n \leq n/2 \leq \sum_{i=1}^n \mathsf{weight}_i(G_n) \cdot p_i - \alpha n$, we have that $\mathrm{gain}(\vec{p}_n, G_n) \geq 1 - \varepsilon$. As with the DNH part of the proof, since the events of (1) and (3) each hold with probability $1 - o(1)$, for sufficiently large n, say $n \geq n_4$, they each occur with probability at least $1 - \delta/2$. Hence, by a union bound, for all $n \geq n_4$, they both occur with probability $1 - \delta$. For $n_0 = \max(n_3, n_4)$, probabilistic positive gain is satisfied.

It remains to show that that if (1) and (3) hold for a specific instance $(\vec{p}_n, G_n)$, then $\text{gain}(\vec{p}_n, G_n) \geq 1 - \varepsilon$ for sufficiently large n. Since $\mathcal{D} \in \mathfrak{D}$, (15) and (16) are both satisfied for sufficiently large n. When

$$\sum_{i=1}^{n} p_i + \alpha n \leq n/2 \leq \sum_{i=1}^{n} p_i - \text{weight}_i(G_n) \cdot \alpha n$$

is satisfied as well, we get that $\mathbb{P}_{\vec{p}_n}\left[X_n^D > n/2\right] < \varepsilon/2$ and $\mathbb{P}_{\vec{p}_n}\left[X_{G_n}^L > n/2\right] > 1 - \varepsilon/2$, so $\text{gain}(\vec{p}_n, G_n) > 1 - \varepsilon$ is immediate.   Q.E.D.

B.2.   Missing Details from the Proof of Lemma 2

The first missing details was the proof of Equation (5),

$$\mathbb{E}[W_T^{(k)}] = \frac{\Gamma(T+p)\Gamma(k)}{\Gamma(p+k)\Gamma(T)}$$

for all $k \leq T$, where Γ represents the Gamma function. The second was showing

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

Recall that $W_k^{(k)} = 1$ and we have the following recurrence for all $t > k$:

$$W_t^{(k)} = \begin{cases} W_{t-1}^{(k)} + 1 & \text{with probability } \frac{p \cdot W_{t-1}^{(k)}}{t-1} \\ W_{t-1}^{(k)} & \text{with probability } 1 - \frac{p \cdot W_{t-1}^{(k)}}{t-1}. \end{cases}$$

By the tower property of expectation, for all $t \geq k+1$,

$$\mathbb{E}[W_t^{(k)}] = \mathbb{E}[\mathbb{E}[W_t^{(k)} \mid W_{t-1}^{(k)}]]$$
$$= \mathbb{E}[W_{t-1}^{(k)}(1 + \frac{p}{t-1})]$$
$$= \mathbb{E}[W_{t-1}^{(k)}](1 + \frac{p}{t-1}).$$

Thus, by a straightforward induction argument and the fact that $\mathbb{E}[W_k^{(k)}] = 1$,

$$\mathbb{E}[W_T^{(k)}] = \mathbb{E}[W_k^{(k)}] \prod_{i=k}^{T-1}(1 + \frac{p}{i}) = \prod_{i=k}^{T-1}(1 + \frac{p}{i}).$$

Expanding this, we have

$$\prod_{i=k}^{T-1}(1 + \frac{p}{i}) = \prod_{i=k}^{T-1} \frac{i+p}{i}$$
$$= \frac{1}{\prod_{i=k}^{T-1} i} \cdot \prod_{i=k}^{T-1}(i+p)$$
$$= \frac{(k-1)!}{(T-1)!} \cdot \frac{\prod_{i=0}^{T-1}(i+p)}{\prod_{i=0}^{k-1}(i+p)}$$
$$= \frac{(k-1)!}{(T-1)!} \frac{\frac{\Gamma(p+T)}{\Gamma(p)}}{\frac{\Gamma(k+p)}{\Gamma(p)}}$$
$$= \frac{\Gamma(T+p)\Gamma(k)}{\Gamma(p+k)\Gamma(T)},$$

where the fourth equality holds because $\Gamma(x+1) = x\Gamma(x)$ for all $x \in \mathbb{R}$, and the last uses the fact that $\Gamma(n) = (n-1)!$ for all $n \in \mathbb{N}$. This proves (5).

For the second missing detail, can now use Markov's inequality to show that for all k,

$$\Pr\left[W_T^{(k)} > T^\delta\right] \leq \frac{\mathbb{E}[W_T^{(k)}]}{T^\delta} = \frac{1}{T^\delta} \cdot \frac{\Gamma(T+p)}{\Gamma(T)} \cdot \frac{\Gamma(k)}{\Gamma(k+p)}.$$

Hence,

$$\sum_{k=1}^{T^\gamma} \Pr[W_T^{(k)} > T^\delta] \leq \sum_{k=1}^{T^\gamma} \frac{1}{T^\delta} \cdot \frac{\Gamma(T+p)}{\Gamma(T)} \cdot \frac{\Gamma(k)}{\Gamma(k+p)} = \frac{1}{T^\delta} \cdot \frac{\Gamma(T+p)}{\Gamma(T)} \cdot \sum_{k=1}^{T^\gamma} \frac{\Gamma(k)}{\Gamma(k+p)}.$$

What remains to be shown is that

$$\frac{1}{T^\delta} \cdot \frac{\Gamma(T+p)}{\Gamma(T)} \cdot \sum_{k=1}^{T^\gamma} \frac{\Gamma(k)}{\Gamma(k+p)} = o(1).$$

To do this, we will use Gautschi's inequality (Gautschi 1959) which states that for all $x > 0$, since $p \in (0,1)$,

$$(x+p-1)^p \leq \frac{\Gamma(p+x)}{\Gamma(x)} \leq (x+p)^p$$

We then have that

$$\frac{1}{T^\delta} \cdot \frac{\Gamma(T+p)}{\Gamma(T)} \cdot \sum_{k=1}^{T^\gamma} \frac{\Gamma(k)}{\Gamma(k+p)} \leq \frac{(T+p)^p}{T^\delta} \cdot \sum_{k=1}^{T^\gamma} \frac{1}{(k+p-1)^p}$$

$$= \frac{(T+p)^p}{T^\delta} \cdot \left(\frac{1}{p^p} + \frac{1}{(1+p)^p} + \sum_{k=3}^{T^\gamma} \frac{1}{(k+p-1)^p}\right)$$

$$\leq \frac{(T+p)^p}{T^\delta} \cdot \left(\frac{1}{p^p} + \frac{1}{(1+p)^p} + \sum_{k=3}^{T^\gamma} \frac{1}{(k-1)^p}\right)$$

$$= \frac{(T+p)^p}{T^\delta} \cdot \left(\frac{1}{p^p} + \frac{1}{(1+p)^p} + \sum_{k=2}^{T^\gamma-1} \frac{1}{k^p}\right)$$

$$\leq \frac{(T+p)^p}{T^\delta} \cdot \left(\frac{1}{p^p} + \frac{1}{(1+p)^p} + \sum_{k=2}^{T^\gamma} \frac{1}{k^p}\right)$$

$$\leq \frac{(T+p)^p}{T^\delta} \cdot \left(\frac{1}{p^p} + \frac{1}{(1+p)^p} + \int_1^{T^\gamma} \frac{1}{x^p}\,dx\right)$$

$$= \frac{(T+p)^p}{T^\delta} \cdot \left(\frac{1}{p^p} + \frac{1}{(1+p)^p} + \frac{x^{1-p}}{1-p}\Big|_1^{T^\gamma}\right)$$

$$= \frac{(T+p)^p}{T^\delta} \cdot \left(\frac{T^{\gamma(1-p)}}{1-p} + \frac{1}{p^p} + \frac{1}{(1+p)^p} - \frac{1}{1-p}\right).$$

Notice that asymptotically, this upper bound is $O(T^{-\delta+p+\gamma\cdot(1-p)})$. By our choice of δ, $\delta > p + \gamma \cdot (1-p)$, so this implies that it is is $o(1)$, as desired.

B.3. Missing Details from the Proof of Theorem 1

The missing details were showing that the Upward Delegation Mechansim satisfied (2) and (3).

Upward Delegation satisfies (2)

We will show there exists $\alpha \in (0,1)$ such that $\sum_{i=1}^n \mathsf{weight}_i(G_n) \cdot p_i - \sum_{i=1}^n p_i \geq 2\alpha n$ with high probability, so (2) is satisfied. Note that in the present scheme, cycles are impossible, so do need to worry about ignored voters.

Since $\mathcal{D}$ is a continuous distribution, there exists $a < b$ such that $\pi_a := \mathcal{D}[\{p : p < a\}] > 0$ and $\pi_b := \mathcal{D}[\{p : p > b\}] > 0$. Let $N_{a,n}(\vec{p}_n)$ be the number of voters in $\vec{p}_n$ with competence $p_i < a$ and $N_{b,n}(\vec{p}_n)$ be the

number of voters with competence $p_i > b$. When we sample competencies, since each is chosen independently, $N_{a,n} \sim \text{Bin}(n, \pi_a)$ and $N_{b,n} \sim \text{Bin}(n, \pi_b)$. By Hoeffding's inequality (Lemma 3) and the union bound, with probability $1 - o(1)$, there will be at least $\pi_a/2 \cdot n$ voters with competence $p_i < a$ and $\pi_b/2 \cdot n$ voters with competence $p_i > b$. Indeed,

$$
\begin{aligned}
\mathcal{D}^n[N_{a,n} > \frac{n\pi_a}{2}, N_{b,n} > \frac{n\pi_b}{2}] &= 1 - \mathcal{D}^n[\{N_{a,n} \le \frac{n\pi_a}{2}\} \cup \{N_{b,n} \le \frac{n\pi_b}{2}\}] \\
&\ge 1 - (\mathcal{D}^n[N_{a,n} \le \frac{n\pi_a}{2}] + \mathcal{D}^n[N_{b,n} \le \frac{n\pi_b}{2}]) \\
&\ge 1 - \exp(-\frac{n\pi_a^2}{2}) - \exp(-\frac{n\pi_b^2}{2}),
\end{aligned}
\tag{17}
$$

where the first line comes from De Morgan's law, the second from the union bound, and the last from Heoffding's inequality (Lemma 3).

Conditioned on this occurring, each voter with competence $p_i < a$ has probability at least $p\pi_b/2$ of delegating to a voter with competence at least b. As they each decide to do this independently, the number $N_{ab,n}$ of n voters deciding to do this stochastically dominates a random variable following the $\text{Bin}(\pi_a/2 \cdot n, p \cdot \pi_b/2)$ distribution. We can again apply Hoeffding's inequality to conclude that with probability $1 - o(1)$, at least $\pi_a \cdot \pi_b \cdot p/8 \cdot n$ voters do so. Indeed,

$$
\begin{aligned}
\mathcal{D}[N_{ab,n} > \frac{np\pi_a\pi_b}{8} \mid N_{a,n} > \frac{n\pi_a}{2}, N_{b,n} > \frac{n\pi_b}{2}] &\ge \mathcal{D}[\text{Bin}(\frac{n\pi_a}{2}, \frac{p\pi_b}{2}) > \frac{np\pi_a\pi_b}{8}] \\
&\ge 1 - \exp(-\frac{np^2\pi_a\pi_b^2}{4}),
\end{aligned}
\tag{18}
$$

where the first inequality holds because $N_{ab,n}$ stochastically dominates the corresponding binomial random variable and the second holds by Hoeffding's inequality. Finally, using (17) and (18), we have

$$
\begin{aligned}
\mathcal{D}[N_{ab,n} > \frac{np\pi_a\pi_b}{8}] &\ge \mathcal{D}[N_{ab,n} > \frac{np\pi_a\pi_b}{8} \mid N_{a,n} > \frac{n\pi_a}{2}, N_{b,n} > \frac{n\pi_b}{2}] \\
&\quad \cdot \mathcal{D}[N_{a,n} > \frac{n\pi_a}{2}, N_{b,n} > \frac{n\pi_b}{2}] \\
&\ge 1 - o(1).
\end{aligned}
$$

Under these upward delegation models, delegations can only increase the total competence of all voters. Hence,

$$
\sum_{i=1}^{n} \text{dels}_i(G_n) \cdot p_i - \sum_{i=1}^{n} p_i \ge (b-a)N_{ab.n}.
$$

Each of these $\pi_a \cdot \pi_b \cdot p/8 \cdot n$ voters results in a competence increase of at least $b - a$. Hence, under these high probability events, the total competence increase is at least $(b - a) \cdot \pi_a \cdot \pi_b \cdot p/8 \cdot n$. Indeed, since $\mathcal{D}[N_{ab,n} > \frac{np\pi_a\pi_b}{8}] = 1 - o(1)$, this implies $\mathcal{D}[\sum_{i=1}^{n} \text{dels}_i(G_n) \cdot p_i - \sum_{i=1}^{n} p_i > \frac{np\pi_a\pi_b}{8}] = 1 - o(1)$. By choosing $\alpha = \frac{p\pi_a\pi_b}{8}(b - a)$, we see that there is an $\alpha \cdot n$ increase in competence with high probability, as needed.

We have proved that for any continuous distribution $\mathcal{D}$, and for well-behaved realizations of $\vec{p}_n$ which occur with high probability, the graph generated from the random delegation process yields an increase in the expected sum of the votes of at least $\alpha \cdot n$. We can then conclude that M_p^U satisfies Equation (2) with respect to the class of continuous distributions.

Upward Delegation satisfies (3)

We now show that there exists a distribution $\mathcal{D}$ such that $\sum_{i=1}^{n} p_i + \alpha n \leq n/2 \leq \sum_{i=1}^{n} \mathsf{weight}_i(G_n) \cdot p_i - \alpha n$ with probability $1 - o(1)$ for some $\alpha > 0$. This implies that the model satisfies probabilistic positive gain by Lemma 1, and will conclude the proof.

We take $\mathcal{D}$ to be $\mathcal{D}_\eta$, the uniform distribution $\mathcal{U}[0, 1 - 2\eta]$ for some small $0 < \eta < p/512$. Let $\alpha = \eta/2$. Clearly, $\mu_{\mathcal{D}_\eta}$, the mean of D_η, is $1/2 - \eta$. Since each $p_i \overset{i.i.d.}{\sim} \mathcal{D}_\eta$, the p_is are bounded independent random variables with mean $1/2 - \eta$, so Hoeffding's inequality directly implies that $\sum_{i=1}^{n} p_i \leq n/2 - n\eta/2 = n/2 - n\alpha$ with high probability.

Now consider $\mathcal{E}_F$, the event consisting of instances $(\vec{p}_n, G_n)$ such that $\sum_{i=1}^{n} \mathsf{weight}_i(G_n) \cdot p_i \geq n/2 + n\alpha$. We denote by $\mathcal{E}_D$ the event that $\sum_{i=1}^{n} p_i \geq n/2 - 3n\eta/2$. The same reasoning as before implies that $\mathrm{Pr}_{\mathcal{D}_\eta, M_p^U, n}(\mathcal{E}_D) = 1 - o(1)$.

Let $a = 1/4 - \eta/2$ and $b = 1/2 - \eta$, so we have that $\pi_a := \mathcal{D}_\eta[p_i < a] = 1/4$ and $\pi_b := \mathcal{D}_\eta[p_i > b] = 1/2$. We proved in the preceding derivation that $\sum_{i=1}^{n} \mathsf{weight}_i(G_n) \cdot p_i - \sum_{i=1}^{n} p_i > \frac{np\pi_1\pi_b}{8} = \frac{np}{64}(1 - \eta)$ with high probability. Hence, if both this and $\mathcal{E}_D$ occur, which is the case with high probability, by the union bound, it follows that $\sum_{i=1}^{n} \mathsf{weight}_i(G_n) \cdot p_i > n/2 + n(\frac{p}{128}(1 - \eta) - 3\eta/2)$ with high probability.

Since $\eta < \frac{p}{512} < 1/2$, we have that

$$\frac{p}{128}(1 - \eta) - 3\eta/2 > \frac{p}{256} - 3\eta/2 > 2\eta - 3\eta/2 = \eta/2 = \alpha,$$

and we can conclude that $\mathcal{E}_F$ occurs with high probability. Hence, M_p^U satisfies Equation (3).

B.4. Missing Details from the Proof of Theorem 2

We show that the Confidence-based Model satisfies (3).

Confidence-Based Delegation satisfies (3)

We finally show there exists a distribution $\mathcal{D}$ such that $\sum_{i=1}^{n} p_i + \alpha n \leq n/2 \leq \sum_{i=1}^{n} \mathsf{weight}_i(G_n) \cdot p_i - \alpha n$ with probability $1 - o(1)$. This implies that the model M_q^C satisfies probabilistic positive gain by Lemma 1.

Using the notation of the analogous proof in Section 3, let $\mathcal{D}_\eta = \mathcal{U}[0, 1 - 2\eta]$ for $\eta \in [0, 1/2)$. Note that as a function of η, $\frac{\mathbb{E}_{\mathcal{D}_\eta}[q^+]}{\mathbb{E}_{\mathcal{D}_\eta}[\bar{q}]}$, the expected competence conditioned on not delegating, is continuous. Moreover, if $\eta = 0$, then

$$\frac{\mathbb{E}_{\mathcal{D}_0}[q^+]}{\mathbb{E}_{\mathcal{D}_0}[\bar{q}]} > \mu_{\mathcal{D}_0} = 1/2.$$

Hence, for sufficiently small $\eta > 0$,

$$\frac{\mathbb{E}_{\mathcal{D}_\eta}[q^+]}{\mathbb{E}_{\mathcal{D}_\eta}[\bar{q}]} > 1/2 > \mu_{\mathcal{D}_\eta}.$$

We choose $\mathcal{D}_\eta$ to be our distribution for this choice of η. As in the previous section, let $\mu_{\mathcal{D}_\eta}^* = \frac{\mathbb{E}_{\mathcal{D}_\eta}[q^+]}{\mathbb{E}[\bar{q}]}$. Note that $\mu_{\mathcal{D}_\eta} = 1/2 - \eta$. Let $\gamma = \min(\frac{1/2 - \mu_{\mathcal{D}_\eta}}{2}, \frac{\mu_{\mathcal{D}_\eta}^* - 1/2}{2})$ and $\alpha = \gamma$. By the earlier argument for (2), we have that with high probability

$$\sum_{i=1}^{n} p_i \leq n(\mu + \gamma) \leq n/2 - \alpha n$$

and

$$\sum_{i=1}^{n} \mathsf{weight}_i(G) p_i \geq n(\mu^* - \gamma) \geq n/2 + \alpha n.$$

By the union bound, we have that both occur simultaneously with high probability, so (3) holds.

B.5. Missing Details from the Proof of Theorem 3

We first show the remainder of the General Continuous Model satisfying (1).

The eigenvector we consider is $\vec{\pi} = (\pi_1, \ldots, \pi_B)$ (which has nonnegative entries, as each π_τ is a probability). We show it has eigenvalue $p\frac{(1+\varepsilon)^3}{1-2\varepsilon}$, strictly less than 1 due to our choice of ε. We show it has eigenvalue $p\frac{(1+\varepsilon)^3}{1-2\varepsilon}$, strictly less than 1 due to our choice of ε. Indeed, we have that

$$(M\vec{\pi})_\tau = \sum_{\tau'=1}^{B} \pi_\tau \tilde{\varphi}(\tau, \tau')\pi_{\tau'} = \pi_\tau p\frac{(1+\varepsilon)^3}{1-2\varepsilon}$$

by the definition of $\tilde{\varphi}$. Hence, $\vec{\pi}$ is our desired eigenvector.

Since $\mathfrak{X}_{M,\tau}^P$ is sub-critical for all τ, we have that there is some c such that for all $\tau \in [B]$, $\mathbb{P}[\mathfrak{X}_{M,\tau}^P \leq c\log(n)] = 1 - o(1/n)$. We take $C(n) = c\log(n)$.

Now we consider our branching process, $\mathfrak{X}_{\vec{p},i}^D$. To make the comparison, we will need some minimal concentration properties. We first show that the sampled competencies $\vec{p}$ satisfy these properties with high probability, and then show that, conditioned on these properties, the branching process $\mathfrak{X}_{\vec{p},i}^D$ is easily comparable to a Poisson process. The properties are the following:

1. For each voter $i \in [n]$, $\sum_{j \neq i} \varphi(p_i, p_j) \geq (1-\varepsilon) \cdot n$.

2. For each type $\tau \in [B]$, the number of voters of type τ, $|\{i \mid p_i \in S_\tau\}| \leq (1+\varepsilon)\pi_\tau n$.

For the first property, fix the competence p_i of a single voter i. Then when sampling the p_js, $\sum_{j \neq i} \varphi(p_i, p_j)$ is the sum $n-1$ independent variables, all in the interval $[L, U]$, with mean 1. Hence, by Hoeffding's inequality, for all competencies c, $\mathcal{D}^n[\sum_{j \neq i} \varphi(p_i, p_j) \geq (1-\varepsilon)n \mid p_i = c] = 1 - o(1/n)$, where the $o(1/n)$ term is independent of c. By the law of total probability, this implies that even when p_i is sampled as well, the $1 - o(1/n)$ bound continues to hold. By a union bound over all n voters, this holds for everybody with probability $1 - o(1)$.

For the second property, note that the number of voters of type τ follows a $\mathrm{Bin}(n, \pi_\tau)$ distribution. A simple application of Hoeffding's inequality implies that for this τ, $|\{i \mid p_i \in S_\tau\}| \leq (1+\varepsilon)\pi_\tau n$ (note that this holds even in the extreme cases where $\pi_\tau = 0$ or $\pi_\tau = 1$). As the number B of types is fixed and independent of n, a union bound over all B types implies this holds for all τ with probability $1 - o(1)$.

Now fix some voter competencies $\vec{p}$ such that both properties hold. We will first upper bound the probability a voter of type τ delegates to a voter of type τ'. Hence, we can compare our branching process to one with these larger probabilities, and this will only dominate our original process.

To that end, since $|D_{t-1}| = t-1 \leq t$ (recall that D_{t-1} consists of the dead voters at time $t-1$), using the first property, we have that for all $i \in [n]$,

$$\sum_{j \in [n] \setminus (D_{t-1} \cup \{i\})} \varphi(p_i, p_j) \geq (1-\varepsilon)n - U \cdot t.$$

Hence, as long as $t \leq \varepsilon n/U$, $\sum_{j \in [n] \setminus (D_{t-1} \cup \{i\})} \varphi(p_i, p_j) \geq (1-2\varepsilon)n$.

Including the fact that $\varphi(p_i, p_j) \leq \tilde{\varphi}(p_i, p_j)$ for all p_i and p_j, we have that for all time steps $t \leq \varepsilon n/U$,

$$p \cdot \frac{\varphi(p_i, p_j)}{\sum_{k' \in [n] \setminus (D_{t-1} \cup \{k\})} \varphi(p_i, p_{k'})} \leq \frac{p}{n} \cdot \frac{\tilde{\varphi}(p_i, p_j)}{(1-2\varepsilon)}.$$

Note that for sufficiently large n, $C(n) \leq \varepsilon n / U$, so from now on we restrict ourselves to such n.

Further, note that by the second property, there will never be more than $(1+\varepsilon)\pi_\tau n$ neutral voters of type τ. Hence, if we take a voter of type τ' at time step $t \leq C(n)$, the number of children it will have of type τ will be stochastically dominated by a $\text{Bin}((1+\varepsilon)\pi_\tau n, \frac{p}{n} \cdot \frac{\tilde{\varphi}(p_i, p_j)}{(1-2\varepsilon)})$, and this is independent for each τ. As n grows large, this distribution approaches a $\text{Pois}(p\frac{(1+\varepsilon)}{1-2\varepsilon}\tilde{\varphi}(\tau, \tau'))$. In particular, this means that for sufficiently large n, it will be stochastically dominated by a $\text{Pois}(p\frac{(1+\varepsilon)^2}{1-2\varepsilon}\tilde{\varphi}(\tau, \tau'))$ distribution (note the extra $(1+\varepsilon)$ factor). Hence, if voter i is of type τ, up to time $t \leq C(n)$, $\mathfrak{X}_{\vec{p},i}^D$ is dominated by $\mathfrak{X}_{M,\tau}^P$, so

$$\mathbb{P}_{\mathcal{D}, M_{\vec{p},\varphi}^S \cdot n}[\mathfrak{X}_{\vec{p},i}^D \geq C(n)] \geq \mathbb{P}_{\mathcal{D}, M_{\vec{p},\varphi}^S \cdot n}[\mathfrak{X}_{M,\tau}^P \geq C(n)] = 1 - o(1/n).$$

A union bound over all n voters tells us this is true for all voters simultaneously with probability $1 - o(1)$, as needed.

The Continuous General Delegation Model satisfies (2).

To show (2) holds, we first show the following.

Let $\mu_\mathcal{D}$ be the mean of the competence distribution $\mathcal{D}$. For a fixed x, let $\varphi_x^+(y)$ be the function $\varphi(x,y) \cdot y$. We show that there is some $c > 0$ such that for all $x \in [0,1]$,

$$\mathbb{E}_\mathcal{D}[\varphi_x^+] \geq \mu_\mathcal{D} + c. \tag{19}$$

Indeed, if we view $\mathbb{E}_\mathcal{D}[\varphi_x^+]$ as a function of x for $x \in [0,1]$, first note that it is a continuous function on a compact set, and hence it attains its minimum. Further, for all $x \in [0,1]$, since $\varphi(x,y)$ and y are both increasing functions of y, by the FKG inequality Fortuin et al. (1971),

$$\mathbb{E}_\mathcal{D}[\varphi^+] > \mathbb{E}_\mathcal{D}[\varphi(x,\cdot)] \cdot \mu_\mathcal{D} = \mu_D,$$

since, by assumption, $\mathbb{E}_\mathcal{D}[\varphi(x,\cdot)] = 1$. Hence, this attained minimum must be strictly larger than μ, implying (19).

Since $\varphi(x,y)$ is normalized so that $\mathbb{E}_{y \sim \mathcal{D}}[\varphi(x,y)] = 1$, $\mathbb{E}_{y \sim \mathcal{D}}[\varphi(x,y) \cdot y]$ is the expected competence of the voter to whom someone of competence x delegates to (prior to other competencies being drawn). Hence, (19) tells us that "on average", all voters (regardless of competence) tend to delegate to those with competence strictly above the mean. Ideally, we would choose $\alpha \approx c/2$ and hope that some concentration result tells us that the weighted competencies after delegation will be strictly above $\mu + c/2$ (the mean of all competencies will be close to μ by standard concentration results). However, proving this concentration result is surprisingly subtle, as there are many dependencies between different voter delegations. Indeed, if one voter with high competence and many delegations chooses to delegate "downwards" (that is, to someone with very low competence), this can cancel out all of the "expected" progress we had made thus far. Hence, the rest of this proof involves proving concentration *does* in fact hold. We prove this by breaking up the process of sampling instances into much more manageable pieces, where, in each, as long as nothing goes "too" wrong, concentration will hold.

In particular, we will prove that for all $\gamma > 0$, with high probability,

$$\sum_{i=1}^{n} \mathsf{weight}_i(G) \cdot p_i - \sum_{i=1}^{n} p_i \geq (c(1-p) - \gamma)n. \tag{20}$$

Fix such a γ. As in the previous part, fix $\varepsilon > 0$ which will paramaterize our steps. We will later choose ε sufficiently small to get our desired result (precisely ε such that $6\varepsilon + \varepsilon^2 < \gamma$). By choosing $\gamma < c(1-p)$, this value is positive, so we can choose $\alpha = \frac{c(1-p)-\gamma}{2}$ which proves Equation (2)

To that end, we define a sequence of six sampling steps that together are equivalent to the standard sampling process with respect to $\mathcal{D}$ and $M_{p,\varphi}^S$. In each step, we will show that with high probability, nothing "goes wrong", and conditioned on nothing going wrong in all these steps, we will get the α improvement that we desire. The six steps are as follows:

1. Sample a set $M \subseteq [n]$ of voters that choose not to delegate. Each voter is included independently with probability p.

2. Sample competencies p_i for $i \in [n] \setminus M$. Each p_i is sampled i.i.d. from $\mathcal{D}$.

3. Sample competencies p_j for $j \in M$. Each p_j is sampled i.i.d. from $\mathcal{D}$.

4. Sample a set $R \subseteq [n] \setminus M$ of delegators that delegate to those in M. Each voter $i \in [n] \setminus M$ is included independently with probability $\frac{\sum_{j \in M} \varphi(p_i, p_j)}{\sum_{j \in [n] \setminus \{i\}} \varphi(p_i, p_j)}$, that is, the total φ weight they put on voters in M divided by the total φ weight they put on all voters.

5. Sample delegations of voters in $[n] \setminus (M \cup R)$. At this point, we are conditioning on such voters delegating, and when they do delegate, they do so to voters in $[n] \setminus M$. Hence, for each $i \in [n] \setminus (M \cup R)$, they delegate to $j \in [n] \setminus (M \cup \{i\})$ with probability $\frac{\varphi(p_i, p_j)}{\sum_{j' \in [n] \setminus (M \cup \{i\})} \varphi(p_i, p_{j'})}$.

6. Sample delegations of voters in R. At this point, we are conditioning on such voters delegating to those in M. Hence, for each $i \in R$, they choose to delegate to $j \in M$ with probability $\frac{\varphi(p_i, p_j)}{\sum_{j' \in M} \varphi(p_i, p_{j'})}$.

We now analyze each step, describing what could "go wrong". Let $\mathcal{E}_1, \ldots, \mathcal{E}_6$ be the events that nothing goes wrong in each of the corresponding steps. We define these events formally below. Our goal is to show that $\mathbb{P}_{\mathcal{D}, M_{p,\varphi}^S, n}[\mathcal{E}_1 \cap \cdots \cap \mathcal{E}_6] = 1 - o(1)$.

• Let $\mathcal{E}_1$ be the event that $(p - \varepsilon) \cdot n \leq |M| \leq (p + \varepsilon) \cdot n$. Note that M is the sum of n independent Bernouilli random variables with success probability p. It follows directly from a union bound over both variants of Hoeffding's inequality that

$$\mathbb{P}_{\mathcal{D}, M_{p,\varphi}^S, n}[(p - \varepsilon) \cdot n \leq |M| \leq (p + \varepsilon) \cdot n] = 1 - o(1).$$

• Let $\mathcal{E}_2$ be the event that $\sum_{i \in [n] \setminus M} p_i \leq n(\mu + \varepsilon)(1 - p + \varepsilon)$. Note that $\sum_{i \in [n] \setminus M} p_i$ is the sum of $n - |M|$ i.i.d. random variables with mean μ. Conditioning on event $\mathcal{E}_1$, $|M|$ is lower bounded by $n(p - \varepsilon)$, implying that $n - |M| \leq n(1 - p + \varepsilon)$ as well. It follows from Lemma 3 that

$$\mathbb{P}_{\mathcal{D}, M_{p,\varphi}^S, n}\left[\sum_{i \in [n] \setminus M} p_i \leq n(\mu + \varepsilon)(1 - p + \varepsilon) \,\middle|\, \mathcal{E}_1 \right] = 1 - o(1)$$

which, combined with $\mathbb{P}_{\mathcal{D}, M_{p,\varphi}^S, n}[\mathcal{E}_1] = 1 - o(1)$, proves that $\mathcal{E}_1 \cap \mathcal{E}_2$ occurs with probability $1 - o(1)$.

• Let $\mathcal{E}_3$ be the event consisting of all instances $(\vec{p}, G)$ such that

$$\frac{\sum_{j \in M} \varphi(p_i, p_j) \cdot p_j}{\sum_{j \in M} \varphi(p_i, p_j)} \geq \frac{(1 - \varepsilon)}{(1 + \varepsilon)}(\mu + c)$$

for all $i \in [n] \setminus M$.

We show $\mathcal{E}_3$ occurs with high probability conditional on $\mathcal{E}_1$ and $\mathcal{E}_1$ (conditioning on $\mathcal{E}_2$ is unnecessary. but makes the final statement easier). Fix a set of voters M and p_i for $i \in [n] \setminus M$ satisfying $\mathcal{E}_1$ and $\mathcal{E}_2$. For each $i \in [n] \setminus M$, we will show that with probability $1 - o(1/n)$, when we sample the p_js for $j \in M$, they satisfy

$$\sum_{j \in M} \varphi(p_i, p_j) \le |M|(1 + \varepsilon) \tag{21}$$

and

$$\sum_{j \in M} \varphi(p_i, p_j) \cdot p_j \ge |M|(1 - \varepsilon)(\mu + c). \tag{22}$$

(21) follows from the fact that $\sum_{j \in M} \varphi(p_i, p_j)$ is the sum of $|M|$ bounded independent random variables with mean $\mathbb{E}_{y \sim \mathcal{D}}[\varphi(p_i, y)] = 1$. By Hoeffding's inequlaity, since $|M|$ is linear in n, $\sum_{j \in M} \varphi(p_i, p_j)$ is at most $|M|(1 + \varepsilon)$ with probability $1 - o(1/n)$.

(22) follows from the fact that $\sum_{j \in M} \varphi(p_i, p_j) \cdot p_j$ is also the sum of $|M|$ bounded independent random variables with mean $\mathbb{E}_{\mathcal{D}}[\varphi_{p_i}^+]$. Again, since we have conditioned on $\mathcal{E}_1$, $|M|$ is lower bounded by $(p - \varepsilon)n$, which by Hoeffding's inequality implies that $\sum_{j \in M} \varphi(p_i, p_j)p_j$ is at least $|M|(1 - \varepsilon)\mathbb{E}_{\mathcal{D}}[\varphi_{p_i}]$ with probability $1 - o(1/n)$.

Finally, we can conclude via a union bound that $\frac{\sum_{j \in M} \varphi(p_i, p_j) \cdot p_j}{\sum_{j \in M} \varphi(p_i, p_j)} \ge \frac{(1 - \varepsilon)}{(1 + \varepsilon)}(\mu + c)$ with probability $1 - o(1/n)$ for any $i \in [n] \setminus M$. Hence, by another union bound over the at most n voters $i \in [n] \setminus M$, $\frac{\sum_{j \in M} \varphi(p_i, p_j) \cdot p_j}{\sum_{j \in M} \varphi(p_i, p_j)} \ge \frac{(1 - \varepsilon)}{(1 + \varepsilon)}(\mu + c)$ for all $i \in [n] \setminus M$ with high probability.

By the law of total probability, $\mathcal{E}_3$ conditioned on $\mathcal{E}_1$ and $\mathcal{E}_2$ occurs with probability $1 - o(1)$, which proves that $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3$ occurs with probability $1 - o(1)$ by the chain rule.

- Let $\mathcal{E}_4$ be the entire sample space. Nothing can "go wrong" during this sampling step. So trivially, $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3 \cap \mathcal{E}_4$ occurs with probability $1 - o(1)$.

- Let $\mathcal{E}_5$ be the event that $\mathsf{dels}_i(G) \le C(n)$ for all $i \in [n] \setminus M$ and $\mathsf{total\text{-}weight}(G) \ge n - C(n)^2 \log(n)$ in the subgraph G sampled (i.e., with delegations only from voters not in R or M). We will show $\mathcal{E}_5$ occurs with high probability even when we sample a full delegation graph (that is, samples delegations for all voters), which implies it continues to hold even when we sample only some delegations (recall that at this step we have only sampled delegations from voters in $[n] \setminus (M \cup R)$).

The proof of this is very similar to the one in Theorem 2, with one extra step to allow for different φ weights.

It was proved in the previous part of this proof that, for all voters i, we have that $\mathsf{dels}_i(G) \le C(n)$ with probability $1 - o(1)$ (not conditioned on anything) when we sample entire delegation graphs, so we can safely condition on this fact. We now prove that $\mathbb{P}_{\mathcal{D}, M_{p, \varphi}^S, n}[\mathsf{total\text{-}weight}(G_n) \ge n - O(\log^3 n) \mid \mathsf{dels}_i(G) \le C(n)] = 1 - o(1)$.

We begin by bounding the number of voters that end up in cycles. Fix some voter i, and let us begin by sampling their delegation tree.

Since we are conditioning on the tree having size at most $C(n)$, the most weight that voter i can place on all of the voters in i's delegation tree is $U \cdot C(n)$. The minimum weight that i can place on all voters is

$L(n-1)$. Hence, the probability that i delegates to someone in i's tree conditional the delegation tree having size at most $C(n)$ is at most $p \cdot \frac{U \cdot C(n)}{L \cdot (n-1)}$. Since i was arbitrary, this implies that the expected number of voters in cycles can be at most $n \cdot p \cdot \frac{U \cdot C(n)}{L \cdot (n-1)} \in O(\log n)$.

Applying Markov's inequality just as in the analogous proof in the previous section, the probability that more than $\log^2 n$ voters are in cycles is at most $np \frac{UC(n)}{L(n-1)\log^2 n} = O(1/\log n) = o(1)$. Further, the total number of people that could delegate to voters in cycles is at most $C(n)$ times the number of voters in cycles. Hence, with probability $1 - o(1)$, there are at most $C(n) \cdot \log^2 n$ voters delegating to those in cycles. This implies the desired bound. Hence, we have proved that $\mathbb{P}_{\mathcal{D}, M_{p,\varphi}^S, n}[\mathcal{E}_5] = 1 - o(1)$. Since we have already shown that $\mathbb{P}_{\mathcal{D}, M_{p,\varphi}^S, n}[\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3 \cap \mathcal{E}_4] = 1 - o(1)$, a union bound implies that $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3 \cap \mathcal{E}_4 \cap \mathcal{E}_5$ occurs with probability $1 - o(1)$ as well.

• We now consider the sixth step. To define $\mathcal{E}_6$, we need some new notation. Fix competencies $\vec{p}$ and a partial delegation graph G such that $(\vec{p}, G)$ is in the first five events. We define Q_i for $i \in R$ to be the random variable representing the competence of the voter to whom i delegates. Since we know i delegates to a voter in M, note that

$$Q_i(G) = p_j \text{ with probability } \frac{\varphi(p_i, p_j)}{\sum_{j' \in M} \varphi(p_i, p_{j'})} \text{ for all } j \in M.$$

Let $\mathcal{E}_6$ be the event consisting of all instance $(\vec{p}, G)$ such that that $\sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i(G) \geq \frac{(1-\varepsilon)^2}{1+\varepsilon}(\mu + c)(1 - p - 2\varepsilon) \cdot n$. We show that $\mathbb{P}_{\mathcal{D}, M_{p,\varphi}^S, n}[\mathcal{E}_6 \mid \mathcal{E}_1 \cap \cdots \cap \mathcal{E}_5] = 1 - o(1)$. This, combined with the the fact $\mathbb{P}_{\mathcal{D}, M_{p,\varphi}^S, n}[\mathcal{E}_1 \cap \cdots \cap \mathcal{E}_5] = 1 - o(1)$ (shown earlier), implies that $\mathbb{P}_{\mathcal{D}, M_{p,\varphi}^S, n}[\mathcal{E}_1 \cap \cdots \cap \mathcal{E}_6] = 1 - o(1)$. It follows from the definition of Q_i that

$$\mathbb{E}[Q_i] = \sum_{j \in M} \frac{\varphi(p_i, p_j)}{\sum_{j' \in M} \varphi(p_i, p_{j'})} \cdot p_j = \frac{\sum_{j \in M} \varphi(p_i, p_j) \cdot p_j}{\sum_{j \in M} \varphi(p_i, p_j)}.$$

By conditioning on $\mathcal{E}_3$, we have that $\mathbb{E}[Q_i] \geq \frac{(1-\varepsilon)}{(1+\varepsilon)}(\mu + c)$ for each $i \in R$. Hence, $\mathbb{E}[\sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i] \geq (n - |M| - C(n)^2 \log(n)) \cdot \frac{1+\varepsilon}{1-\varepsilon} \cdot (\mu + c)$, since we are conditioning on $\mathcal{E}_3$ and $\mathcal{E}_5$. Further, for sufficiently large n, $C(n)^2 \log(n) \leq \varepsilon n$; since we are conditioning on $\mathcal{E}_1$, $|M| \leq (p + \varepsilon)n$, so we have that for sufficiently large n,

$$\mathbb{E}[\sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i] \geq (1 - p - 2\varepsilon) \cdot \frac{1+\varepsilon}{1-\varepsilon} \cdot (\mu + c) \cdot n \in \Omega(n).$$

Next, consider $\mathrm{Var}[\sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i]$. Since each Q_i takes on values in $[0,1]$, $\mathrm{Var}[Q_i] \leq 1$. Further, each summand is independent, as each Q_i is independent and we have fixed G, so we can can view $\mathsf{dels}_i(G)$ as a constant. Hence, $\mathrm{Var}[\sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i] \leq \sum_{i \in R} \mathsf{dels}_i(G)^2 \in o(n^2)$ since, for all i, $\mathsf{dels}_i(G) \leq C(n) \in O(\log n)$ and $\sum_i \mathsf{dels}_i(G) \leq n$. Hence,

$$\mathbb{P}_{\mathcal{D}, M_{p,\varphi}^S, n}[\sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i < \frac{(1-\varepsilon)^2}{1+\varepsilon}(\mu + c)(1 - p - 2\varepsilon) \cdot n]$$
$$\leq \mathbb{P}_{\mathcal{D}, M_{p,\varphi}^S, n}[\sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i < (1-\varepsilon)\mathbb{E}[\sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i]]$$
$$\leq \frac{\mathrm{Var}[\sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i]}{\varepsilon^2 \cdot \mathbb{E}[\sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i]^2} \in o(1)$$

where the second inequality is due to Chebyshev's inequality, which is $o(1)$ because the numerator is $o(n^2)$ and the denominator is $\Omega(n^2)$. This implies the desired result.

Finally, we show that for all instance $(\vec{p}, G) \in \mathcal{E}_1 \cap \cdots \cap \mathcal{E}_6$, (2) holds, and hence so does (20). We have that $\sum_{i=1}^{n} w_i(G) \cdot p_i = \sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i(G) + \sum_{j \in M} p_j$, because in G each voter $i \in R$ delegates all of their $\mathsf{dels}_i(G)$ votes to the voter in M with competence $Q_i(G)$. Hence, $\sum_{i=1}^{n} w_i(G) \cdot p_i - \sum_{i=1}^{n} p_i = \sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i(G) - \sum_{i \in [n] \setminus (M \cup R)} p_i$. Since $(\vec{p}, G) \in \mathcal{E}_2$, we have that $\sum_{i \in [n] \setminus M} p_i \leq n(\mu + \varepsilon)(1 - p + \varepsilon)$. Since $(\vec{p}, G) \in \mathcal{E}_6$, we have that $\sum_{i \in R} \mathsf{dels}_i(G) \cdot Q_i(G) \geq \frac{(1-\varepsilon)^2}{1+\varepsilon} (\mu + c)(1 - p - 2\varepsilon) \cdot n$. Hence, this difference is at least

$$((\mu + c)(1 - p - 2\varepsilon) - (\mu + \varepsilon)(1 - p + \varepsilon))n \geq (c(1-p) - 3\varepsilon\mu - 2\varepsilon c - (1-p)\varepsilon - \varepsilon^2)n$$

$$\geq (c(1-p) - 6\varepsilon - \varepsilon^2)n$$

where the second inequality holds because, $c, (1-p), \mu \leq 1$. By choosing ε such that $6\varepsilon + \varepsilon^2 \leq \gamma$ ($\varepsilon = \min(\gamma/7, 1)$ will do), (20) follows.

The Continuous General Delegation Model Satisfies (3)

We now show that there exists a distribution $\mathcal{D}$ and $\alpha > 0$ such that $\sum_{i=1}^{n} p_i + \alpha n \leq n/2 \leq \sum_{i=1}^{n} \mathsf{weight}_i(G_n) \cdot p_i - \alpha n$ with probability $1 - o(1)$. This implies that the model $M_{p,\varphi}^S, n$ satisfies probabilistic positive gain by Lemma 1.

As in earlier arguments, let $\mathcal{D}_\eta = \mathcal{U}[0, 1 - 2\eta]$ for $\eta \in [0, 1/2)$. Note that

$$f(\eta) = \inf_{x \in [0,1]} \left\{ \mathbb{E}_{D_\eta}[\varphi_x^+] \right\} \cdot (1-p) - 3\eta/2$$

is a continuous function of η.

Moreover, $f(0) > 0$. Hence, for sufficiently small $\eta > 0$, $f(\eta) > 0$.

Consider $\mathcal{D}_\eta$ for some $\eta > 0$ such that $f(\eta) > 0$. Let $\alpha = \min(\eta/2, f(\eta)/2)$. Since $\mu_{\mathcal{D}_\eta} = 1/2 - \eta$, by Hoeffding's inequality, $\sum_{i=1}^{n} p_i \leq (1/2 - \eta/2)n \leq n/2 - \alpha n$ with high probability.

Next, note that we can choose $c = \inf_{x \in [0,1]} \left\{ \mathbb{E}_{D_\eta}[\varphi_x^+] \right\}$ in order to satisfy (19). Hence, by choosing $\gamma = f(\eta)/2$, it follows from (20) that

$$\sum_{i=1}^{n} \mathsf{weight}_i(G) \cdot p_i - \sum_{i=1}^{n} p_i \geq (c(1-p) - f(\eta)/2)n = (3\eta/2 + f(\eta)/2)n \geq (3\eta/2 + \alpha)n$$

with high probability. Further, by Hoeffding's inequality, $\sum_{i=1}^{n} p_i \geq (1/2 - 3\eta/2)n$ with high probability, so by the union bound applied to these inequalities,

$$\sum_{i=1}^{n} \mathsf{weight}_i(G) \cdot p_i \geq n/2 + \alpha n$$

with high probability, as needed.
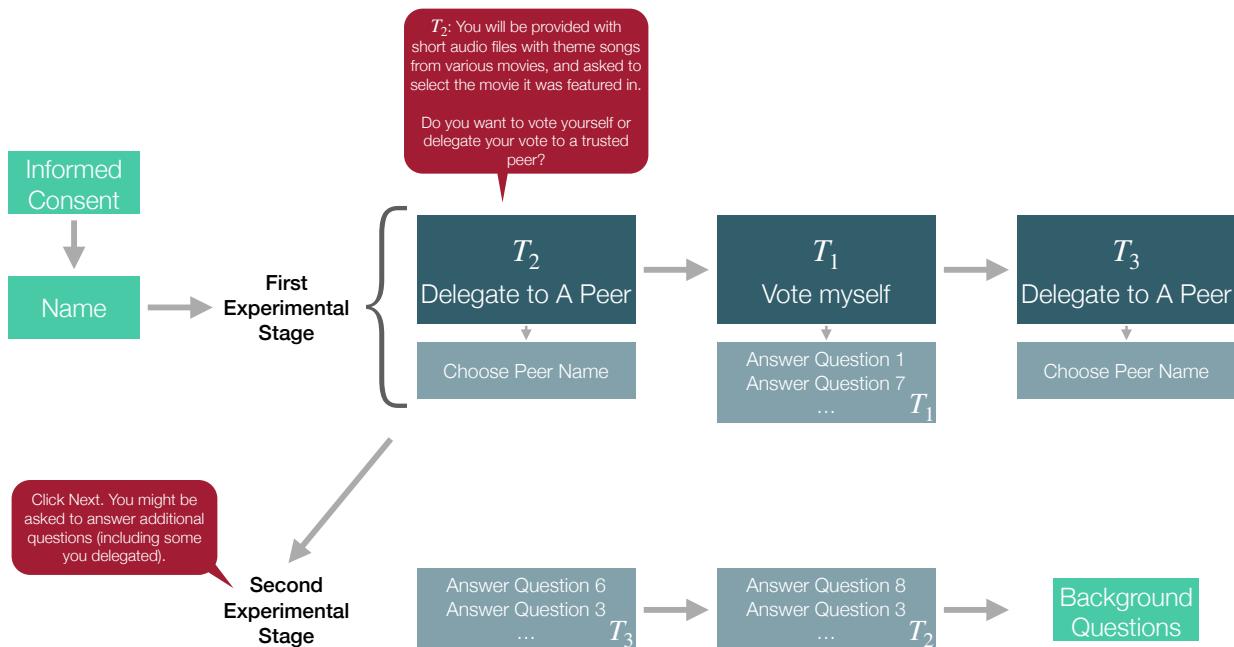
C. Groups Characteristics

Table 3   Qualitative group descriptions and sizes from regular experiment

Group ID	Group Description	Group Size
1	Company Employees Present at a Workshop	14
2	Undergraduate Students Present in Class	22
3	Research Department Meeting	19
4	Company Employees Present at a Workshop	27
5	Participants at an Academic Conference	36
6	Participants at an Academic Conference	50

D.   Survey Materials and Flow

This section describes the participant experience of the survey. Figure 3 shows an example of the survey flow. The green boxes represent the pre and post-survey steps (providing informed consent, name, and optional background questions). In the first stage, participants performed tasks, deciding to either delegate (providing a name) or vote (answering the eight question). The upper red block exemplifies a task prompt (in which the options "delegate" and "vote" also appear in a random order). In the second stage, participants answer additional questions (those they delegated) and optional background questions. Table 4 shows all the prompts given to participants for each of the tasks. Finally, Figure 4 shows screenshots of the survey from the participant's perspective.
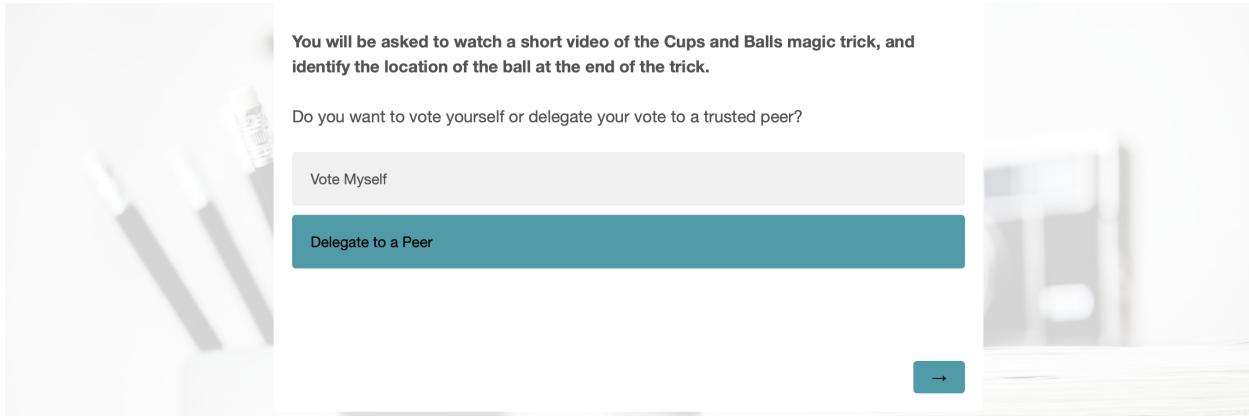

**Figure 3**   **Example of a survey flow with three tasks.**
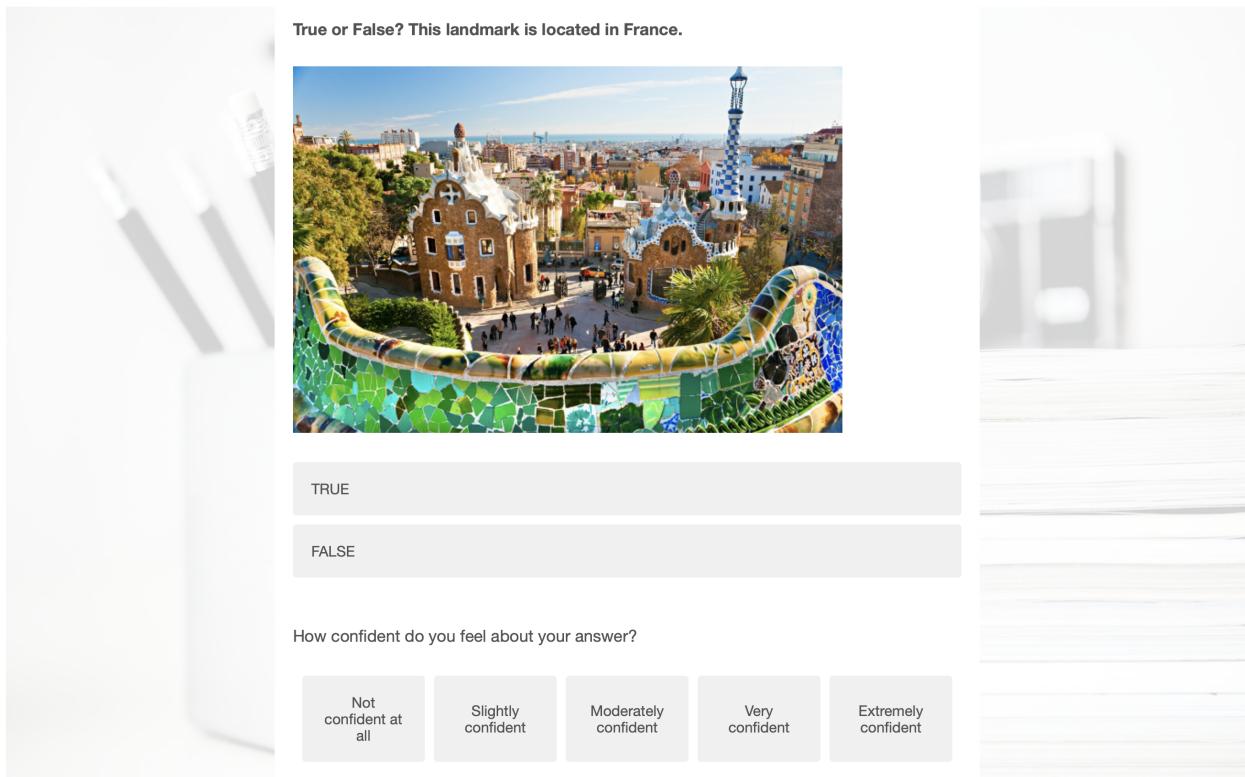
The entire survey can be found in the OSF repository.

Table 4   Prompts Presented at the beginning of each task

ID	Task Prompts	Corresponding Experiment(s)
T_1	You will be shown images of architectural landmarks from around the world, and asked to select the country where the landmark is located.	1, 2, 3, 4, 5, 6
T_2	You will be provided with short audio files with theme songs from various movies, and asked to select the movie it was featured in.	1, 2, 3, 4, 5, 6
T_3	You will be given English idioms, and asked to identify their meaning. An idiom is a group of words that have a meaning not deducible from those of the individual words (e.g., rain cats and dogs, see the light).	1, 2, 3, 4, 5, 6
T_4	You will be given upcoming sport events (soccer and tennis games), and asked to predict the games' outcome?	1
T_5	You will be given the names of tennis players, and asked to predict which round they will make it to in the Tennis French Open (Roland Garros), taking place in May-June 2022?	2, 3, 4
T_6	You will be given the names of tennis players (women and men), and asked to predict which round they will make it to in the ongoing Wimbledon Tennis Tournament (The Championships, Wimbledon), taking place between June 27 and July 10, 2022.	5
T_7	You will be given upcoming European men soccer games and asked to predict the games' outcome.	6
T_8	You will be shown images of flags from around the world, and asked to identify their country of origin.	6
T_9	You will be shown 20 images of famous buildings from around the world, and asked to estimate the year in which the building was completed.	6
T_{10}	You will be shown images of constellations and asked to identify them.	6
T_{11}	You will be given headlines, and asked to identify the magazine that published the article, between The Economist and WIRED.	6
T_{12}	You will be given words, and asked to identify the correct synonym corresponding to each word.	6
T_{13}	You will be asked to listen to audio clips of classical compositions, and asked to identify the composer.	6
T_{14}	You will be given the names of American states, and be asked to predict whether the majority of Congress members elected in that state will be Republican or Democrat.	6
T_{15}	You will be given upcoming NBA games and asked to predict the games' outcome.	6

(a) Screenshot of a task prompt along with a choice to vote or delegate.

(b) Screenshot of a specific question.

Figure 4 **Screenshots from the survey of both task prompts and specific questions.**

E. Item Response Rate Theory and Technical Details in Support of Section 6.2.1

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. A naive way to compute participants' competence per task would be to average their number of correct answers given on all 8 questions of that task, $\eta_{i,t}^{\text{naive}} = \frac{\sum_{r \in R_t} v_{i,r}}{|R_t|}$. However, such a computation does not account for the varying difficulty of the questions.

Instead, we estimate $\eta_{i,t}$ using the Item Response Theory framework (IRT) (Lalor and Rodriguez 2023), which provides a widely used parametric model to estimate competence and question difficulty from repeated measurements. To do this, we fit a one-parameter logistic model to estimate person i's latent ability to do task t, $\eta_{i,t}$, as well as question r's latent difficulty, θ_r, where the probability that person i is correct on question r depends on the person's ability at task t and the question's difficulty.[21] Specifically, we assume a generative process of

$$\Pr[v_{i,r} = 1 | \eta_{i,t}, \theta_r] = \frac{1}{1 + \exp^{-(\eta_{i,t} - \theta_r)}}$$
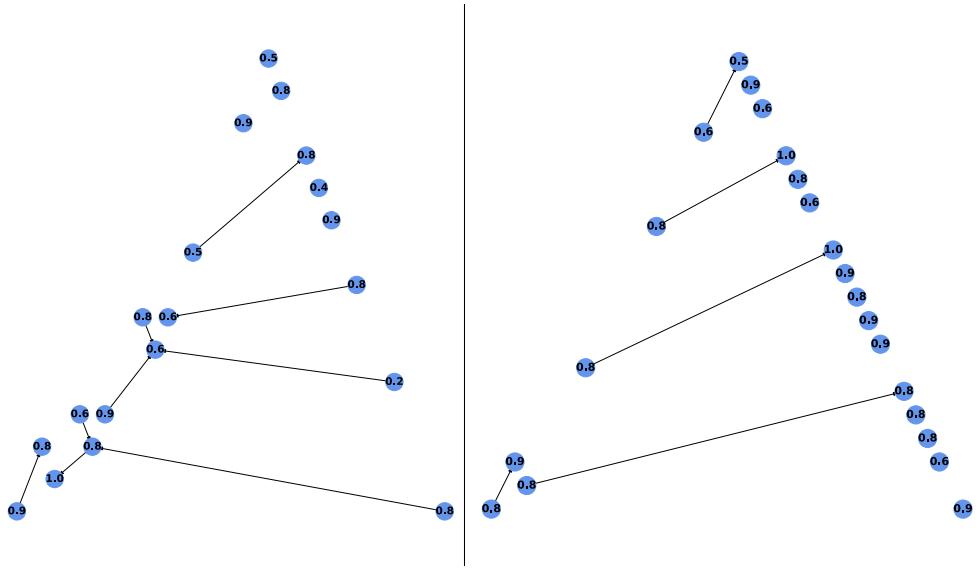
and fit $\eta_{i,t}$ and θ_r to be consistent with the observed answers $v_{i,t}$. We fit these parameters in the canonical way using the Python package py-irt. See Natesan et al. (2016) for more details.

While $\eta_{i,t}^{\text{naive}}$ takes on one of nine values (multiples of 1/8), $\eta_{i,t}$ 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. Even using these different methods, the correlation between $\eta_{i,t}^{\text{naive}}$ and $\eta_{i,t}$ is above 94%.

---
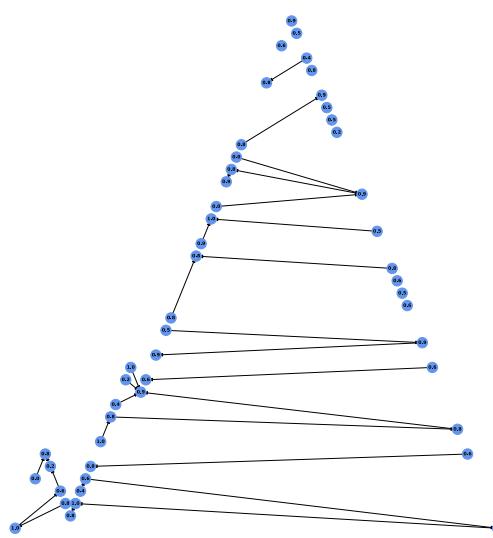
[21] We also fit a three-parameter logistic model estimating $\Pr[v_{i,r} = 1 | \eta_{i,t}, \theta_r, c_r, a_r] = c_r + \frac{1 - c_r}{1 + \exp^{-a_r(\eta_{i,t} - \theta_r)}}$, with c_r, the effect of guessing on question r and a_r, the degree to which question r' differentiate between participants. The resulting competence levels are highly correlated and we stick with the one-parameter model as a result.

F.   Delegation Graph Examples
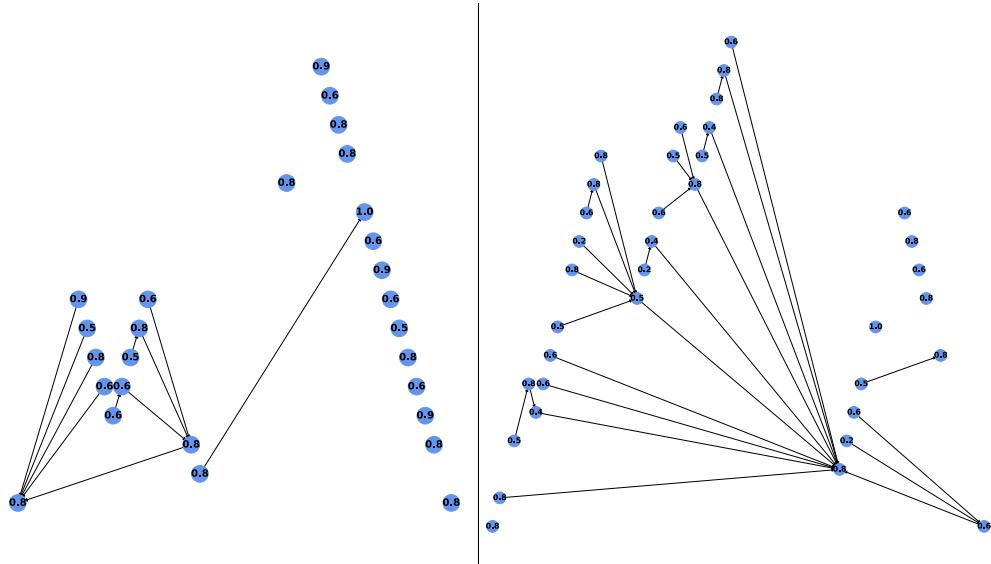
Figures 5 to 8 show representative examples of delegation graphs across different experiments and tasks. Nodes are labeled by the percentage of correct answers for that task, i.e., $\eta_{i,t}^{\mathrm{naive}}$. Figures 5 and 6 display examples of relatively little concentration of power with delegations reaching participants with relatively high competence, Figure 7 displays some of the worst concentration of power observed across the 32 instances and Figure 8 illustrates an example of relatively little concentration of power with delegations reaching participants with relatively low competence.
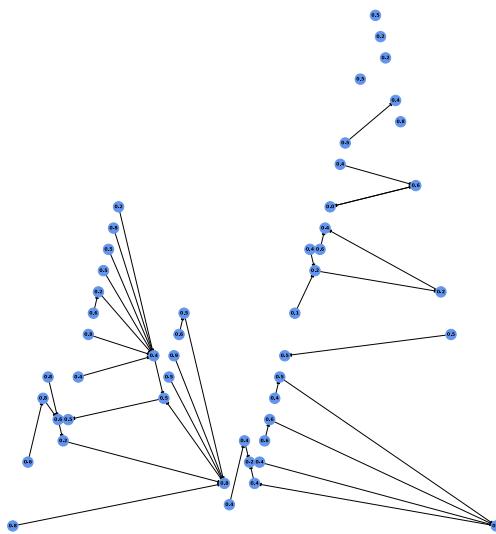
Figure 5   Delegation graphs for task T_3 from Experiment 2 (left) and task T_2 from Experiment 3 (right).

Figure 6   Delegation graph from experiment 6 and task T_7.
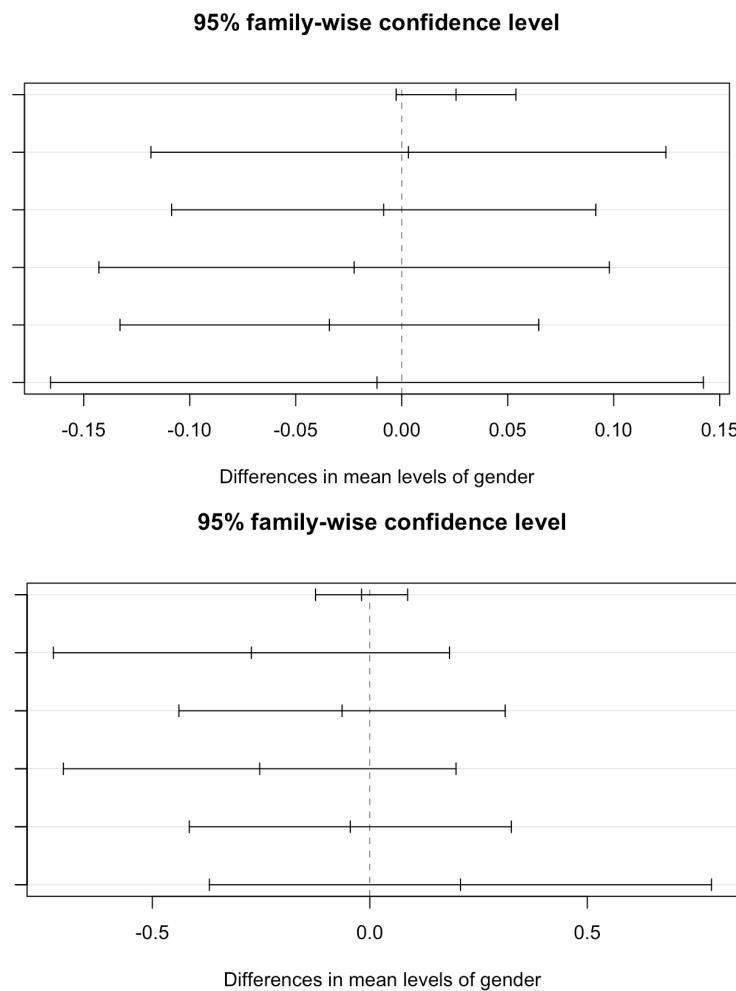
Figure 7 Delegation graphs for task T_1 from Experiment 4 (left) and task T_5 from Experiment 5 (right).

Figure 8 Delegation graph from experiment 6 and task T_{15}.

G.   Pairwise Tukey Tests

95% family-wise confidence level

Differences in mean levels of gender

95% family-wise confidence level

Differences in mean levels of gender
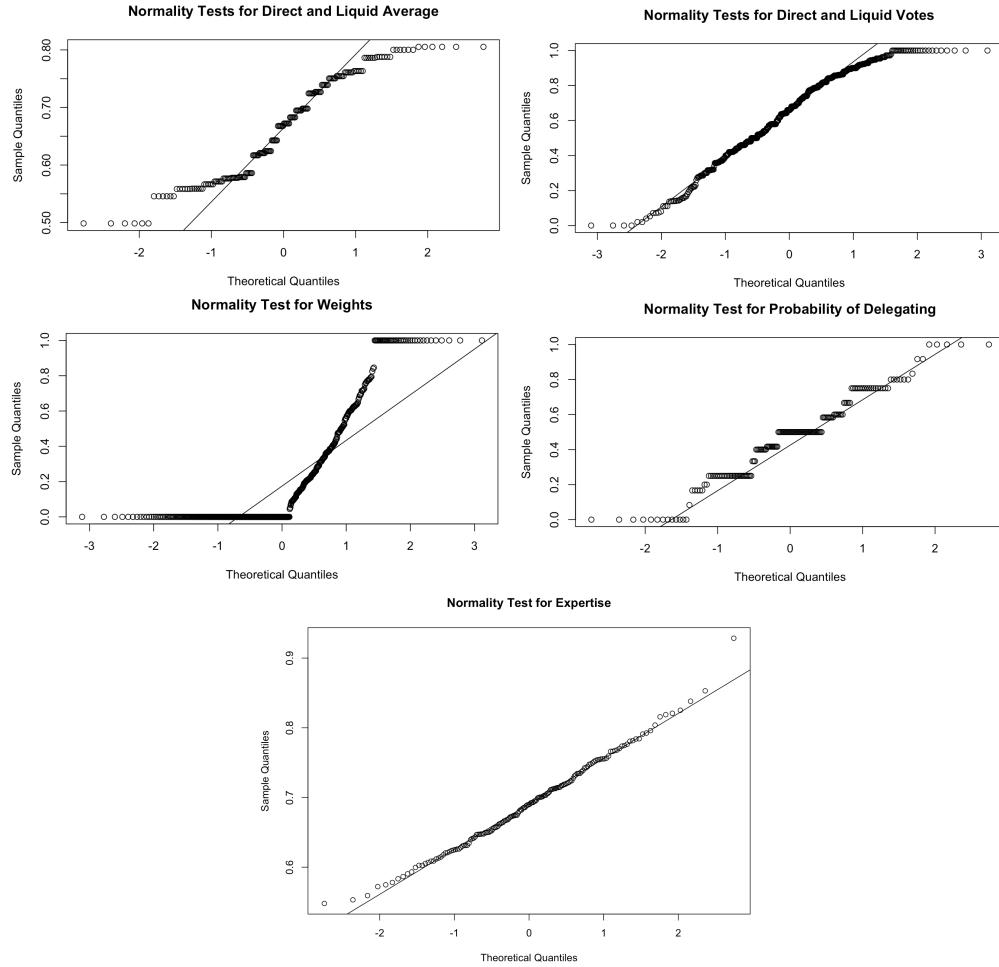
Figure 9   Pairwise Tukey test across different gender regarding competence (top) and propensity to delegate (bottom). Pairwise test shows that competence and propensity to delegate is indistinguishable across gender.

H.   Normality Assumptions for Regressions

We check that the variables used in the different regressions are indeed normal, per the tests' assumptions.

Figure 10   Normality Tests

From top to bottom and left to right, normality tests for the average competence in direct and liquid democracies, the average vote outcome in direct and liquid democracies, the estimated values of φ, the probability of delegating and the participants' competence.

I.   Additional tests for Section 6.3

We run a few additional tests for understanding the relationship between the propensity to delegate and competence. First, we repeat the same analysis we did in Section 6.3, but partitioning by task first. Results can be found in Tables 5 and 6. Note that tasks 8–15 only appeared in experiment 6, so fewer data points were available, resulting in much larger standard errors.

Table 5   Estimates of β_q for tasks 1–7.

	Overall Model	Tasks Models						
		(Tasks)						
		(T_1)	(T_2)	(T_3)	(T_4)	(T_5)	(T_6)	(T_7)
Effect Size β^q	-2.23^{****}	-2.05	-1.96	-4.06^{***}	2.33	0.21	-8.04^{*}	-1.50
	(0.42)	(1.58)	(1.50)	(1.25)	(4.05)	(1.57)	(4.88)	(0.23)
Clustered S.E.	i	NA	NA	NA	NA	NA	NA	NA
Note:						*p<0.1; **p<0.05; ***p<0.01		

Table 6   Estimates of β_q for tasks 8–15.

	Tasks Models							
	(Tasks)							
	(T_8)	(T_9)	(T_{10})	(T_{11})	(T_{12})	(T_{13})	(T_{14})	(T_{15})
Effect Size β^q	-5.13	1.57	-3.87	-2.43	-2.50	-3.62^{*}	3.70	-1.58
	(5.09)	(2.15)	(2.66)	(2.38)	(2.01)	(1.88)	(2.89)	(1.57)
Clustered S.E.	NA	NA	NA	NA	NA	NA	NA	NA
Note:						*p<0.1; **p<0.05; ***p<0.01		

Next, we consider adding fixed effects to the model. This helps us understand the direct impact of competence, even when keeping fixed an individual (as their competence varies across tasks) and the task (which could intrinsically lead to different propensities of delegation). More formally, we fit the following generalized linear model

$$\log\left(\frac{\Pr[\delta_{i,t}=1]}{1-\Pr[\delta_{i,t}=1]}\right) = \alpha_0 + \alpha_i + \alpha_t + \beta^q \eta_{i,t} + \varepsilon_i \qquad (23)$$

where α_0 is the population average across participants and tasks, α_i is the fixed effect for person i and α_t is the fixed effect for task t, and β^q is the coefficient for competence. Results can be found in Table 7.

The results show that participant-specific characteristics (such as confidence) are at play in delegation decisions. We also see that within each task, the probability of delegating decreases more slowly with competence than in the model without fixed effect. This implies that task-specific characteristics (such as difficulty) are also at play in delegation decisions.

Table 7   **Estimates for β^q accounting for fixed-effects.**

	Overall Model	With Fixed Effects		
	Equation (10)	Equation (23)		
Effect Size β^q	-2.23^{****}	-2.90^{****}	-1.67^{****}	-1.86^{****}
	(0.42)	(0.60)	(0.35)	(0.45)
Fixed Effects	NA	i	t	t, i
Clustered S.E.	i	i	i	i
Note:				****p$<$0.0001

J.   Additional tests and details for Section 6.4

J.1.   Alternative Bucketing Strategies

We experiment with varying bucketing methods, to ensure that the results are robust. The strategies (other than k-means) are the following. An example illustrating the differences can be found in Figure 11.

Equal cut: With equal cut, we divide the $[0,1]$ line in B buckets $c_\ell = [(\ell-1)/B, \ell/B]$ for $\ell \in \{1, \ldots, B\}$ of equal size. We vary the number of buckets B from 3 to 10 to ensure robustness of the approach. When buckets are empty, we compute the weights on the existing types and re-normalize in the final stage.

Quantile cut: We cut the $[0,1]$ line in B quantiles c_ℓ for $\ell \in [b]$ so that the number of competence values in each bucket is the same and we take the mean competence in the designated bucket to be the representative competence, η_ℓ. That is, $\eta_\ell = \frac{\sum_{i,t} \eta_{i,t} \mathbb{I}[\eta_{i,t} \in c_\ell]}{\sum_{i,t} \mathbb{I}[\eta_{i,t} \in c_\ell]}$. In short, each competence $\eta_{i,t}$ is assigned to a bucket η_ℓ such that $\ell \leq \eta_{i,t} < l + 1/b$. We vary the number of buckets B to ensure robustness of the approach (see Appendix ??).

Gaussian Mixture Model: We assume that the competence level $\eta_{i,t}$ is drawn from B Gaussian distributions that we intend to reconstruct. To do so, we maximize the log-likelihood $\log \Pr(\vec{\eta}) = \sum_{i,t} \sum_k^b \log \left(\pi(k) \mathcal{N}(\eta_{i,t} | \eta_k, \sigma^2) \right)$, where $\vec{\eta}$ is the vector of $\eta_{i,t}$, $\pi(k)$ is the probability of being in Gaussian k, η_k, σ^2 are the mean and variance of the $k-$th Gaussian and $\mathcal{N}(x | \mu, s^2)$ denotes the probability density function of a Gaussian with mean μ and standard deviation s evaluated at x. This optimization cannot be solved in closed-form, and we use the Expectation-Maximization (EM) algorithm to estimate the Gaussian means η_k, as well as the marginal probabilities $\Pr(k | \eta_{i,k})$.[22] Last, we find the number of Gaussian B that maximizes the likelihood's cross-validation estimate. In turn, we obtain an assignment of competence $\eta_{i,t}$ to B Gaussian and B Gaussian's mean, that we denote η_ℓ for the $\ell-$th Gaussian. Each competence $\eta_{i,t}$ is assigned to a Gaussian η_ℓ based on $\Pr(k | \eta_{i,k})$.

J.2.   Maximum Likelihood Estimation for $\varphi_{e,t}^\ell$

Fix a delegation graph, corresponding to experiment e and task t. We are interested in reconstructing the most likely estimate of $\varphi_{e,t}^\ell(\eta_k)$ (we will drop e and t from the notation when they are clear from context). Let z_k^ℓ be the observed number of times that someone of type ℓ delegated to someone of type k. Let n_ℓ be the number of people of type ℓ, and let $\tilde{n}_\ell$ be the number of people of type ℓ that delegated. Note that φ is always invariant to scaling, so for consistency across experiments, we will always normalize SO that $\sum_k \varphi_{e,t}^\ell(\eta_k) = 1$.
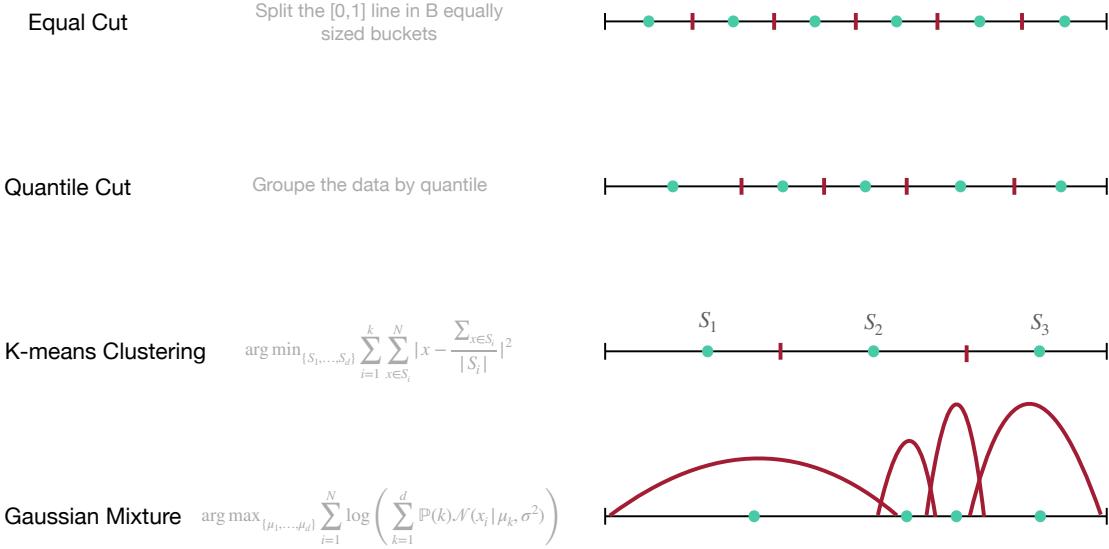
PROPOSITION 1. *The maximum likelihood estimators for* $\varphi_{e,t}^\ell(\eta_1), \ldots, \varphi_{e,t}^\ell(\eta_B)$ *are*[23]

$$\varphi_{e,t}^\ell(\eta_k) = \begin{cases} \frac{z_k^\ell}{n_k} & \text{if } k \neq \ell \\ \frac{z_k^\ell}{n_k - 1} & \text{if } k = \ell \end{cases}$$

up to arbitrary normalization constants.

---

[22] See details here: https://people.csail.mit.edu/rameshvs/content/gmm-em.pdf.

[23] Or 0 in if the denominator is 0.

Equal Cut — Split the [0,1] line in B equally sized buckets

Quantile Cut — Groupe the data by quantile

K-means Clustering — $\arg\min_{\{S_1,\ldots,S_d\}} \sum_{i=1}^{k} \sum_{x \in S_i} |x - \frac{\sum_{x \in S_i}}{|S_i|}|^2$

Gaussian Mixture — $\arg\max_{\{\mu_1,\ldots,\mu_d\}} \sum_{i=1}^{N} \log\left(\sum_{k=1}^{d} \mathbb{P}(k)\mathcal{N}(x_i \,|\, \mu_k, \sigma^2) \right)$

Figure 11   Alternative Bucketing Strategies Illustrated

Proof   Fix e, t, and ℓ. Let $\tilde{n}_k = n_k$ for $k \neq \ell$ and $\tilde{n}_\ell = n_\ell - 1$. Note that when a voter in bucket ℓ is considering whether or not to delegate, they see $\tilde{n}_k$ other voters of each type k. Hence, for a parameter setting $\vec{\varphi} = (\varphi^\ell(\eta_1), \ldots, \varphi^\ell(\eta_k))$, the probability they delegate to a specific voter of type k is $\frac{\varphi^\ell(\eta_k)}{\sum_{k'} \tilde{n}_{k'}\varphi^\ell(\eta_{k'})}$. In the data observed, a voter in bucket ℓ delegated to another in bucket k a total of z_k^ℓ times. Hence, the likelihood of $\vec{\varphi}$ can be written as

$$L(\vec{\varphi}) = \prod_k \left(\frac{\varphi^\ell(\eta_k)}{\sum_{k'} \tilde{n}_{k'}\varphi^\ell(\eta_{k'})} \right)^{z_k^\ell}$$

Notice that this probability does not change by changing all the φ values by the same constant factor, so when finding the most likely, it is without loss of generality to optimize over $\vec{\varphi}$ such that $\sum_{k'} \tilde{n}_{k'}\varphi^\ell(\eta_{k'}) = 1$. In such cases, the likelihood simplifies to

$$L(\vec{\varphi}) = \prod_k \left(\varphi^\ell(\eta_k) \right)^{z_k^\ell}.$$

Additionally, notice that optimizing over a constant times the likelihood results in the same optimization problem, so the MLE $\vec{\varphi}$ must also maximize, multiplying by $\prod_k (\tilde{n}_k)^{z_k^\ell}$,

$$L'(\vec{\varphi}) = \prod_k \left(\tilde{n}_k \varphi^\ell(\eta_k) \right)^{z_k^\ell}.$$

Finally, note that choosing $\vec{\varphi}$ such that $\sum_{k'} \tilde{n}_{k'}\varphi^\ell(\eta_{k'}) = 1$ so as to maximize L' is equivalent to finding the MLE of a multinomial distribution with B outcomes and probability $\tilde{n}_k \varphi^\ell(\eta_k)$ of ending up in outcome k. This is well known to be maximized with $\tilde{n}_k \varphi^\ell(\eta_k) \propto z_k$ Therefore, up to arbitrary normalization factors, the MLE $\vec{\varphi}$ satisfy $\varphi^\ell(\eta_k) = \frac{z_k}{\tilde{n}_k}$.   □

J.3.   Kendall tau rank correlation coefficient

Let $R(\varphi^\ell_{e,t}(\eta_k))$ be the rank of $\varphi^\ell_{e,t}(\eta_k)$ amongst all the estimated weights and $R(\eta_k)$ be the rank of η_k among all competenceS. We then compute

$$\tau = \frac{\sum_{k \in [N]} \mathbb{I}[R(\varphi^\ell_{e,t}(\eta_k)) = R(\eta_k)] - \sum_{k \in [N]} \mathbb{I}[R(\varphi^\ell_{e,t}(\eta_k)) \neq R(\eta_k)]}{N - N_\varphi - N_\eta}.$$

where N_φ, N_η represent the number of ties in $R(\varphi^\ell_{e,t}(\eta_k))$ and $R(\eta_k)$ respectively. Next, we compare the z-score for the Kendall tau rank, $z = \frac{3\tau\sqrt{N(N-1)^2(N-2)}}{\sqrt{8(N(N-1)+5)}}$ the values of a Gaussian distribution.[24] We further check the Kendall tau rank correlation coefficient between $\varphi^\ell_{e,t}(\eta_k)$ and η_k for a fixed ℓ.

J.4.   Task-specific Effects

In Table 8, we report the the Kendall tau correlation coefficients for each task to understand whether the observed behavior is constant across tasks. Note that for the first three tasks, there are 96 data points per regression. For those, we observe a clear trend that φ is increasing in its second coordinate across all three tasks. The following ones, however, rely on only 16 data points.

Table 8   Correlation and p-values Across Tasks

Task	Correlation	p-value
T_1	0.3647***	0.0011
T_2	0.2621**	0.0079
T_3	0.3830****	< 0.0001
T_4	-0.1540	0.6233
T_5	-0.2055	0.1004
T_6	-0.1372	0.5666
T_7	0.4768*	0.0201
T_8	-0.1925	0.5392
T_9	-0.5231**	0.0127
T_{10}	-0.2490	0.2300
T_{11}	0.5427**	0.0108
T_{12}	0.0321	0.9115
T_{13}	0.5644**	0.0059
T_{14}	0.5768**	0.0051
T_{15}	-0.2689	0.1948
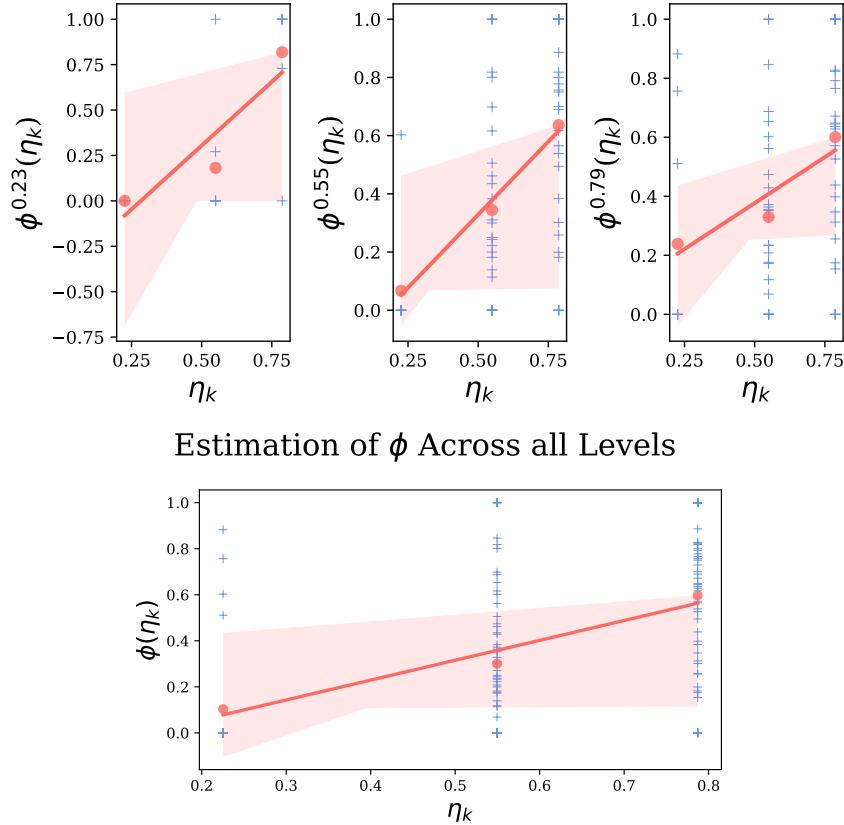
*p<0.1; **p<0.05; ***p<0.01; ****p<0.0001

J.5.   Alternative Bucketing Results

We run the same experiments as in Section 6.4 with $B \in \{3, 5, 7, 10\}$. Figures 12 to 15 and Tables 9 to 12 show the equivalent of Figure 2 and Table 2 for these various values of B. Similarly, Figure 16 and Table 13 show these results for $B = 7$ using quantile cut bucketing, but all different variations lead to qualitatively similar results.

---

[24] (Or, more precisely, we let the function stats.kendalltau in python's library scipy handle all of this for us.)

ϕ function for various expertise levels

Estimation of ϕ Across all Levels

Figure 12   **Estimation of φ with $B = 3$**

Table 9   **Estimation of φ with $B = 3$**
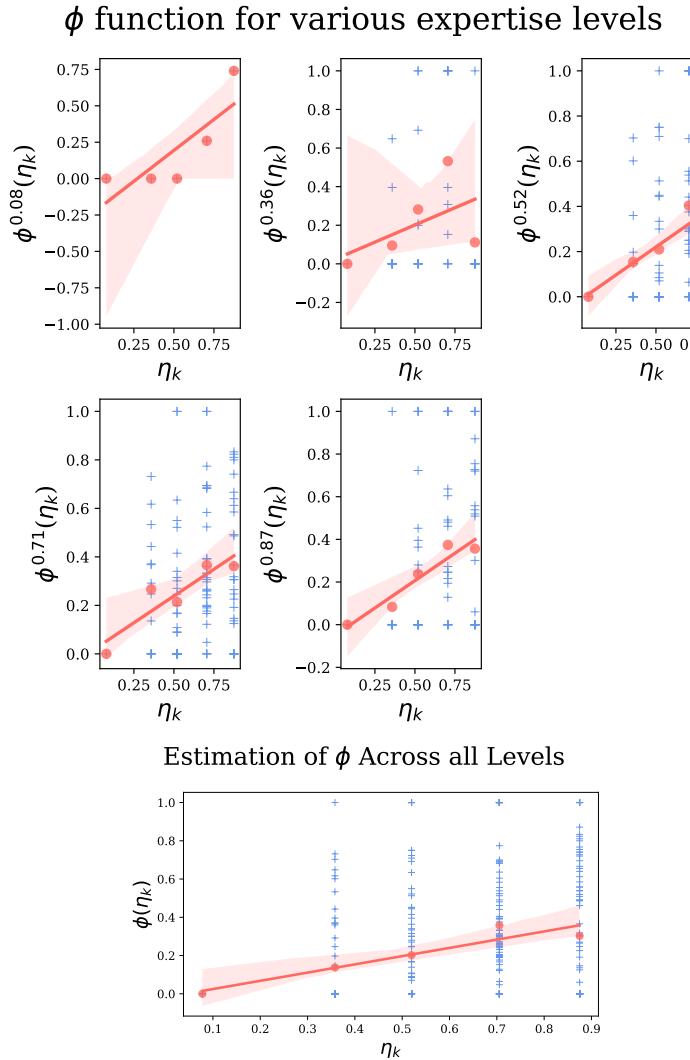
	Overall	For fixed ℓ		
		c_1	c_2	c_3
Correlation	0.44^{****}	0.60^{*88*}	0.43^{*888}	0.34^{****}
P-value	2×10^{-11}	2×10^{-3}	2×10^{-5}	7×10^{-4}

Note: *p<0.1; **p<0.05; ***p<0.01; ****p<0.0001

Table 10   **Estimation of φ with $B = 5$**

	Overall	For fixed ℓ				
		c_1	c_2	c_3	c_4	c_5
Correlation	0.13^{****}	0.14^{**}	0.06	0.15^{**}	0.10	0.21^{***}
P-value	6×10^{-5}	6×10^{-2}	3×10^{-1}	4×10^{-2}	10^{-1}	5×10^{-3}

Note: *p<0.1; **p<0.05; ***p<0.01; ****p<0.0001

ϕ function for various expertise levels

Estimation of ϕ Across all Levels
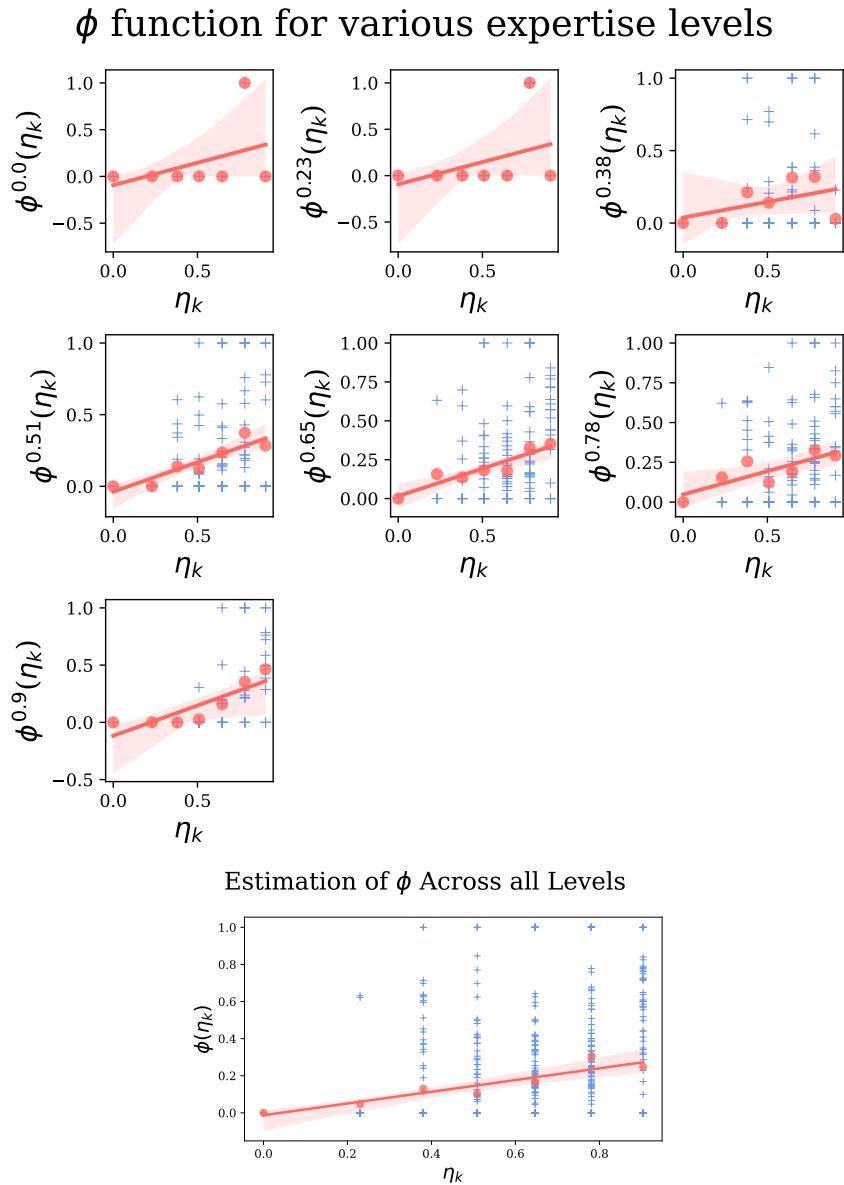
Figure 13   **Estimation of** φ **with** $B = 5$

Table 11   **Estimation of** φ **with** $B = 7$
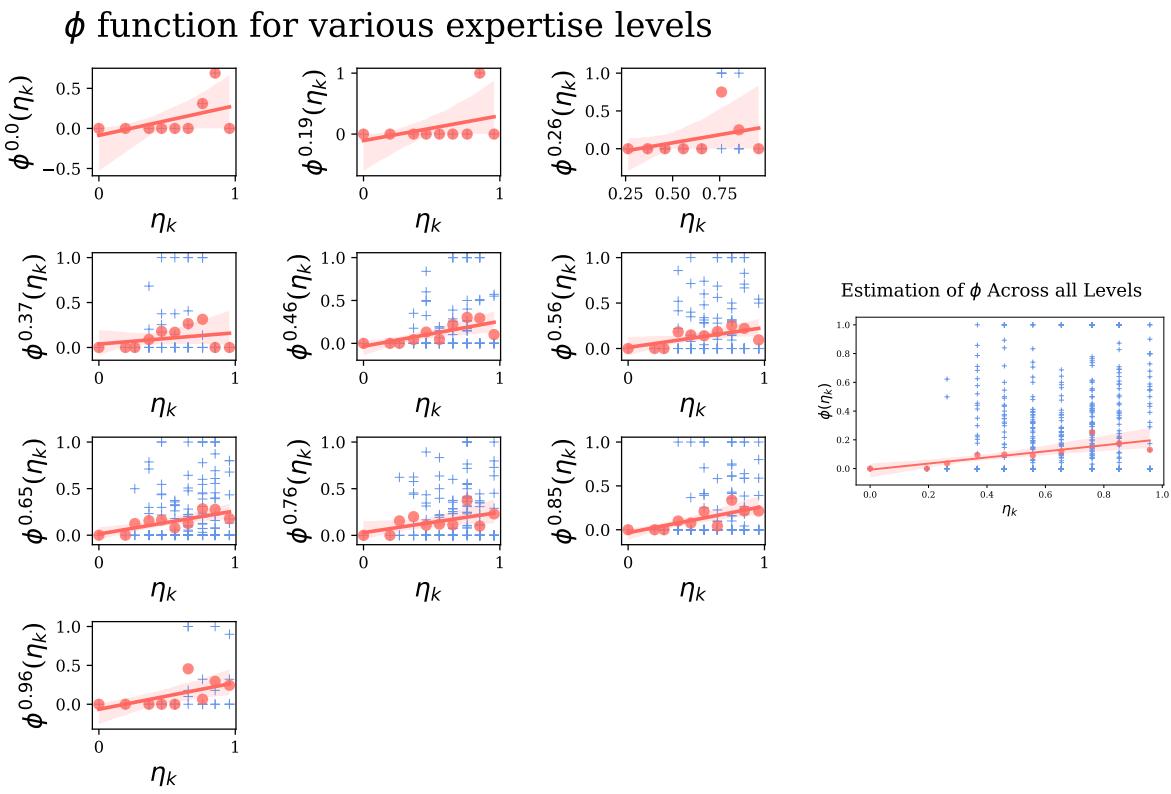
	Overall	For fixed ℓ						
		c_1	c_2	c_3	c_4	c_5	c_6	c_7
Correlation	0.23^{****}	0.37	0.36^{**}	0.05	0.18^{***}	0.25^{****}	0.16	0.45^{****}
P-value	10^{-11}	3×10^{-1}	10^{-2}	6×10^{-1}	10^{-1}	2×10^{-4}	2×10^{-2}	3×10^{-6}
Note:						*p<0.1; **p<0.05; ***p<0.01; ****p<0.0001		

ϕ function for various expertise levels

Estimation of ϕ Across all Levels

Figure 14 **Estimation of** φ **with** $B = 7$

ϕ function for various expertise levels

Figure 15   **Estimation of** φ **with** $B = 10$

Table 12   **Estimation of φ with $B = 10$**

	Overall	For fixed ℓ									
		c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
Correlation	0.14^{****}	0.47	0.35	0.35^{**}	0.00	0.16^{**}	0.03	0.15^{**}	0.10^{*}	0.21^{***}	0.30^{**}
P-value	1.2073×10^{-7}	0.1111	0.2453	0.0388	0.9763	0.0291	0.6754	0.0135	0.0957	0.0026	0.0418
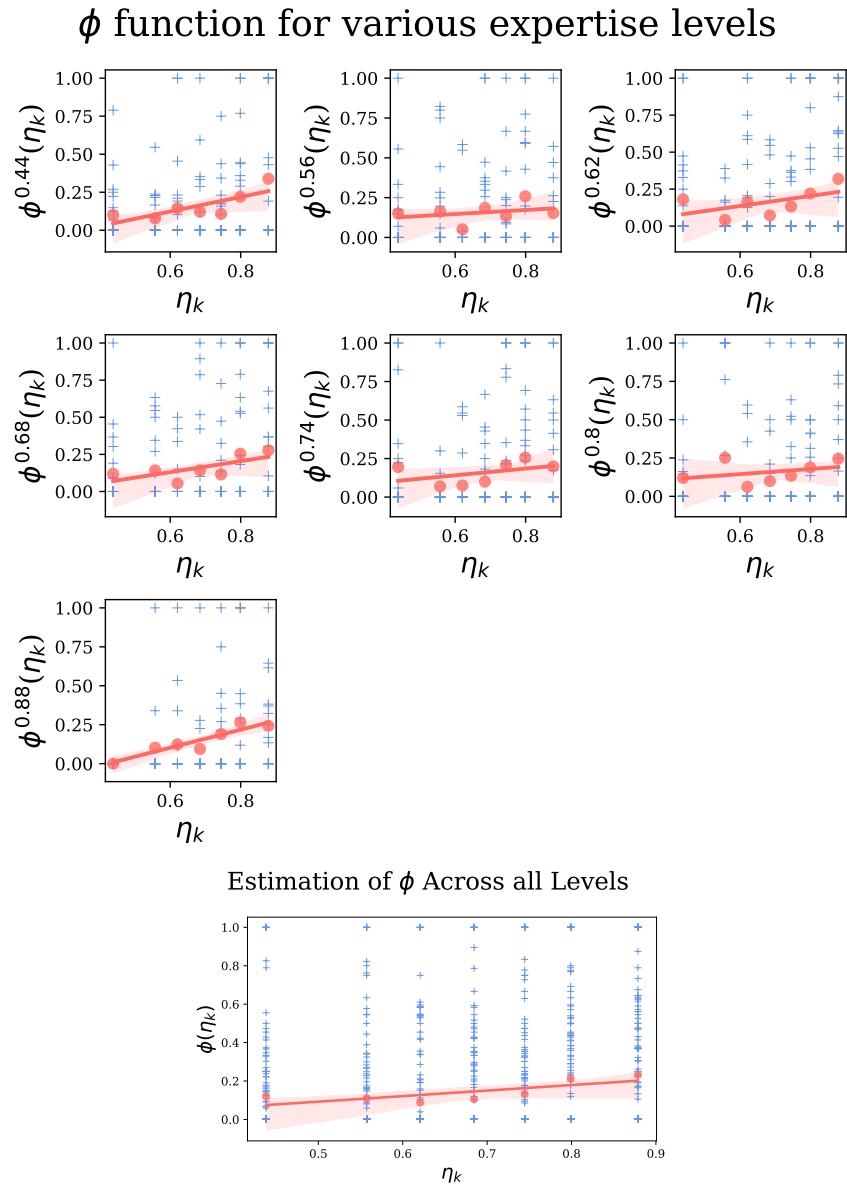
Note: *p$<$0.1; **p$<$0.05; ***p$<$0.01; ****p$<$0.0001

Table 13   **Estimation of φ with $B = 7$ using quantile bucketting.**

	Overall	For fixed ℓ						
		c_1	c_2	c_3	c_4	c_5	c_6	c_7
Correlation	0.10^{****}	0.07	0.03	0.08	0.10^{*}	0.11^{*}	0.09	0.29^{***}
P-value	3×10^{-5}	0.25	0.62	0.16	0.099	0.063	0.17	0.0003

Note: *p$<$0.1; **p$<$0.05; ***p$<$0.01; ****p$<$0.0001

ϕ function for various expertise levels

Estimation of ϕ Across all Levels

Figure 16   **Estimation of** φ **with** $B = 7$ **using quantile bucketting.**

K. Additional Tests and Statistics

K.1. Robustness Check on Attention

In order to check whether asking respondents to answer the questions associated with the tasks they delegated biased their answers, we run a robustness check on the time spent on the survey as a function of the delegation frequency.

First, note that the design's key components are that (i) respondents delegate in response to a prompt (e.g., historical landmark) without seeing the questions so that (ii) we collect the liquid vote along with how they would have voted had they not delegated (this information is also used to assess competence). Subsequently, participants are told "you may be asked to answer a few additional questions." Importantly, we do not say why, and participants answer questions they have not seen before. We deliberately designed the survey that way in order to minimize any biasing effects that might make a participant less motivated to answer questions they delegated.

Second, we computed the time participants spent on the experiment (standardized per experiment to account for the fact that different experiments had different questions) as a function of the proportion of time that they delegate, to see whether they are significantly likely to spend less time on the test if they delegate more (as one could hypothesize that they would be less motivated and then answer the second phase questions more rapidly). We found no relation between time spent on the survey and the likelihood delegation (with a confidence interval for the coefficient of [-0.505, 0.155] and a coefficient of -0.175).

K.2. Frequency of Correctness
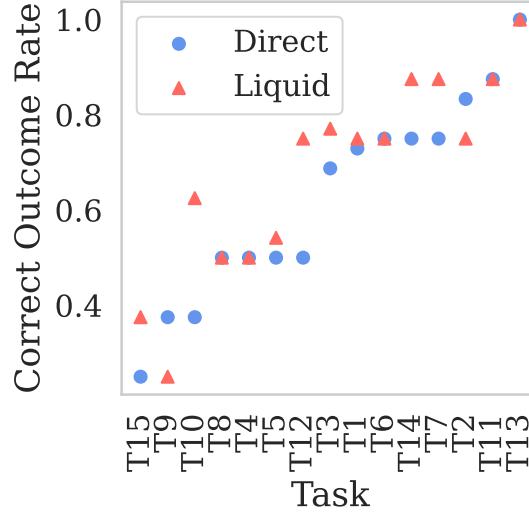
We begin by directly comparing how frequently liquid and direct democracy were correct. The corresponding data (separated by task) can be found in Figure 17. Although our experiments find evidence for sufficient conditions for DNH and PG, we do not observe statistically significant gains for LD over DD (but we do observe a positive difference in the point estimates). This is still consistent with DNH and PG holding (which both only kick in the limit, and even on these finite samples, LD is doing no worse). However, even if we believe such a gain could exist, it could not appear for a variety of reasons: it's plausible we simply didn't have enough statistical power to detect the improvement (we have 32 comparisons) or that the improvement appears in the large limit and there were not enough group sizes. We leave the study of this to future work.

On average, we see an overall improvement, with a few exceptions for some tasks. To understand the statistical significance of this improvement, we fit the following generalized linear model. We let $o_{e,t}^L$ and $o_{e,t}^D$ be the proportion of questions answered correctly using liquid and direct democracy in experiment e and task t. More formally, $o_{e,t}^L = \sum_r \mathbb{I}[\sum_{i \in [N_e]} v_{i,r}^\ell > N_e/2]/8$ and $o_{e,t}^D = \sum_r \mathbb{I}[\sum_{i \in [N_e]} v_{i,r}^D > N_e/2]/8$. We then fit the following generalized linear model

$$o_{e,t}^j = \alpha_0 + \alpha_e + \alpha_t + \beta^{LvD}\gamma_j + \varepsilon_{e,t}$$

where $\gamma_j = \mathbb{I}[j = L]$, i.e., it is an indicator that this proportion came from liquid democracy. Once fitted, β^{LvD} here represents the impact of using liquid democracy over direct democracy, while taking into account fixed effects from tasks and experiments (i.e., some tasks may be easier or harder, and some groups may overall perform better than others). We find $\beta^{\mathrm{LvD}} = 0.0313$ with $s.e. = 0.025$, $t = 1.229$ and $p = 0.23$. In other words, liquid democracy's average proportion of correct answers is 3 points above that of direct democracy, but this increase is not significant.

## Frequency at Which Liquid and Direct Democracies Are Correct

Figure 17   **Frequency of Correctness for Liquid and Direct Democracies**

K.3.   Increase in Competence

Next, we test liquid democracy's ability to increase the average competence. We let $\eta_{e,t}^L = \frac{\sum_{i \in [N_e]} w_{i,t} \eta_{i,t}}{N_e}$ (resp. $\eta_{e,t}^D = \frac{\sum_{i \in [N_e]} \eta_{i,t}}{N_e}$) to denote the average competence after (resp., before) delegation. We fit an analogous generalized linear model to the proportion of correctness, namely

$$\eta_{e,t}^j = \alpha_0 + \alpha_e + \alpha_t + \beta^{increase}\gamma_j + \varepsilon_{e,t}.$$

We find $\beta^{\text{increase}} = 0.031$ with $s.e. = 0.006$, $t = 4.78$ and $p = 0.000004$. In other words, across all tasks and experiments, the mean average competence post delegation is 3% higher than the mean average competence without delegation, and this result is quite statistically significant.
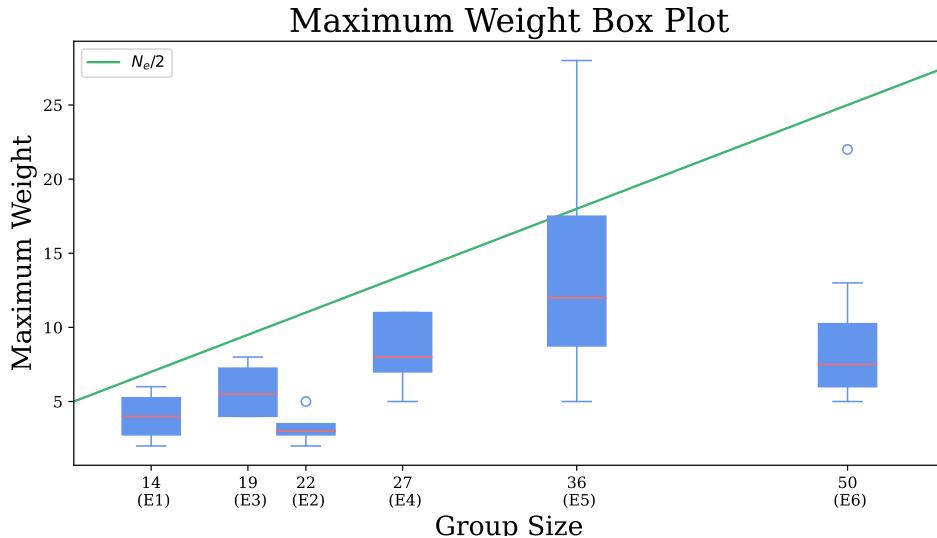
K.4.   Maximum Weight

Next, we consider the maximum weight received by any individual. Note that this weight should always be considered with respect to the size of the group considered. Boxplots of the observed maximum weights can be found in Figure 18. Along with this, we plot a line for $N_e/2$, the "dictator" line. Note that any voter with weight more than $N_e/2$ essentially has complete control over the outcome.

competence after delegation. For each experiment e and task t, we use $m_{e,t} = \max_{i \in [N_e]} w_{i,t}$ to denote the maximum weight, and $\eta_{e,t}^L = \frac{\sum_{i \in [N_e]} w_{i,t} \eta_{i,t}}{N_e}$ (resp. $\eta_{e,t}^D = \frac{\sum_{i \in [N_e]} \eta_{i,t}}{N_e}$) to denote the average competence post (resp. pre) delegation.

K.5.   Coalition Analysis

For this section, we will use powers of coalitions as a different way to measure concentration of power. We use the standard terminology of calling voters that do not delegate, *gurus*. Suppose the gurus have weights
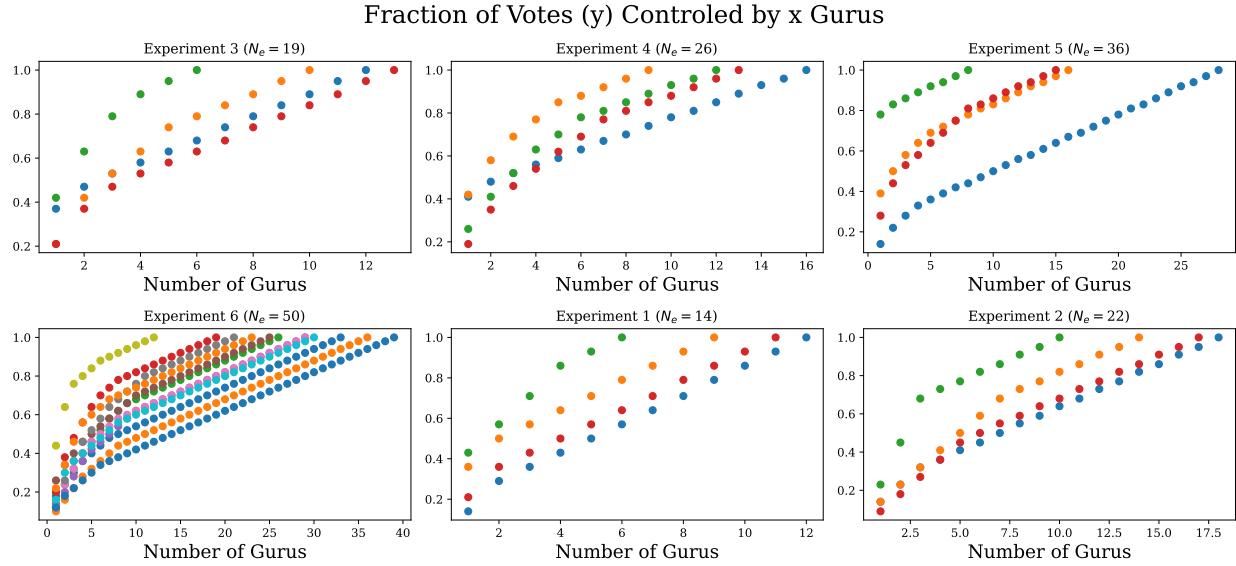
Figure 18 **Distribution of maximum weights received in tasks across experiments.**
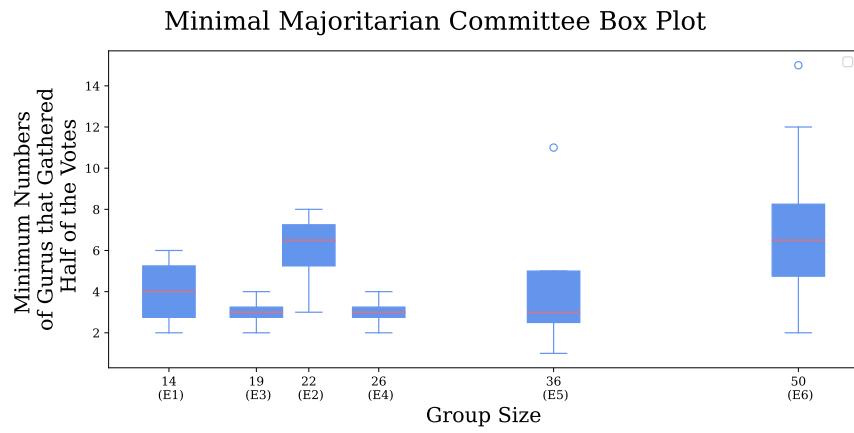
$w_1 \geq w_2 \geq \cdots$. We will say the *power* of a coalition of gurus S is the number of votes they control, i.e., $\sum_{i \in S} w_i$. Note that the maximum weight is exactly the most power a single guru holds. However, we can also consider other measures such as the largest amount of power that k gurus hold. This could potentially be as large as k times the max weight, but in certain instances could be much smaller. Of paticular interest is the minimal coalition size that constitutes a majority, i.e., the minimal k such that $\sum_{i=1}^{k} w_i \geq n/2$. Note that in the case of a dictatorship, this size is just 1, while in the case of direct democracy, it is $n/2$. Values among this spectrum give another measure of how concentrated power is.

Figure 19 shows this first property, i.e., in each experiment and each task, for any fixed k, what proportion of the votes do the top k gurus control. Figure 20 shows the second, that is, for each group, what the fewest number of gurus that collectively control half the votes. Note that, anecdotally, the size of the smallest coalition that controls half the votes is about the square root of the group size. This is based only on the six mean numbers measured in Figure 20. No rigorous conclusions can be drawn for this, but further empirical inquiry on this matter would be an interesting direction for future work.

Finally, as a way to understand the unequalness of this power, we plot the smallest majority coalition size against the total number of gurus. This can be found Figure 21. Overall, while this data is interesting, it is difficult to draw conclusions about how harmful this level of concentration of power is. Our theoretical results hold asymptotically, and either more data or larger group sizes would be needed to truly understand the implications.

Fraction of Votes (y) Controled by x Gurus

Figure 19 **The fraction of votes gathered by the smallest coalition that maximizes total weight. The x axis represents the number k of voters with the k largest weights. The y axis represents the fraction of votes they gathered altogether. Each color represents a task (the first five experiments comprised 4 tasks, whil the sixth comprised 13 tasks.**

Minimal Majoritarian Committee Box Plot

Figure 20 **Smallest coalition of gurus that control half the votes. These are box-plots with medians, quantiles and outliers.**

Number of Gurus as a Function
of Minimal Number of Gurus
that Gathered Half of the Votes
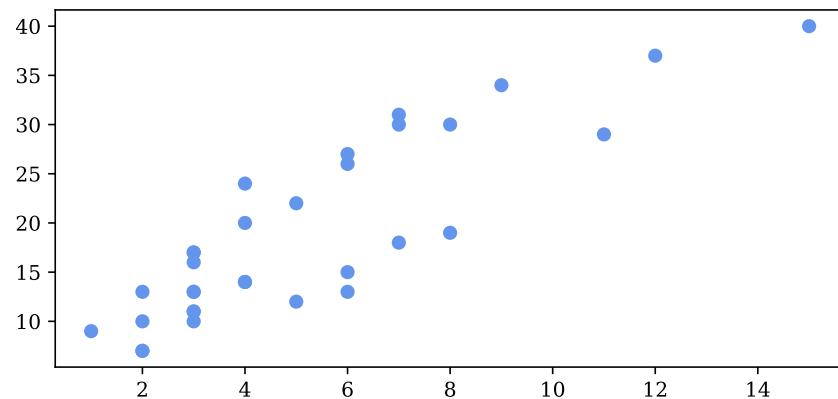
**Figure 21** **Number of Gurus as a Function of the Size of the Smallest Potentially Majority Coalition (SPMC)**

L. Pre-study

We provide in this section the results of the same analyses in the pre-study, that comprised 6 experiments and a similar design that was used to inform the survey flow and material of the main study.

L.1. Update in the Design for the Main Study

In the main study, the spatial reasoning task was dropped as participants were almost always correct. Ambiguous questions (for which multiple answers were true) taken from Simoiu et al. (2019) were rephrased to enforce a clear objective answer. Questions that were mislabeled in Simoiu et al. (2019) were removed. The survey flow was also changed so that more questions could be answered in the same amount of time, and all tasks were set to have the same number of questions $|R_t| = 8$.

L.2. Recruitment

The six groups with which the pre-study was run are described below. There were a total of 102 participants that participated in the surveys between March 21st and April 5th, 2022. Of the participants across all experiments, 29% were native English speakers, 16% were female, 4% non-binary, and 80% were male.

Table 14   **Qualitative group descriptions and sizes from the pre-study**

Group ID	Group Description	Group Size
1	Research Group	11
2	Graduate and Undergraduate Class	12
3	Graduate Class	32
4	Sports Team	14
5	Financial Association	18
6	Group of employees, students and faculty	15

L.3. Material

The tasks used in the pre-study are described below, Table 15 details the tasks and Table 16 some of the questions within the tasks.

L.4. Assessing competence

We run the IRT framework and find a correlation of 97% between the naive competence $\eta_{i,t}^{\text{naive}}$ and the competence computed accounting for task difficulty $\eta_{i,t}$. Note that, while the competence distribution for the main experiment is normal (Appendix H), the competence distribution in the pre-study is not, due to the spatial reasoning task for which almost all participants are correct. We show the distribution of competence computed with both the naive and IRT frameworks in Figure 22.

L.5. Delegation Statistics

We collected 505 delegation data points, one per participant per task. Of those, 28% were delegations and 57% are direct participants that did not receive any delegation besides their own. Among the delegates, 21% received only one delegation besides their own (hence had weight 2 in the decision), 11% received two delegations besides their own and just about 1% received five or more delegations besides their own, showing little sign of concentration of power.

Table 15   Task Descriptions

| ID	Number of Questions	Task Prompts	Corresponding Experiment(s)
Landmark	5	You will be shown images of architectural landmarks from around the world, and asked to select the country where the landmark is located.	$1, 2, 3, 4, 5, 6$
Movie	5	You will be provided with short audio files with theme songs from various movies, and asked to select the movie it was featured in.	$1, 2, 3, 4, 5, 6$
English	5	You will be given English idioms, and asked to identify their meaning. An idiom is a group of words that have a meaning not deducible from those of the individual words (e.g., rain cats and dogs, see the light).	$1, 2, 3, 4, 5, 6$
Spatial	3	You will be asked to watch a short video of the Cups and Balls magic trick, and identify the location of the ball at the end of the trick.	$1, 2, 3, 4, 5, 6$
Sports	7	You will be given US college basketball teams, and asked to predict which round they will make it to in the NCAA Tournament, taking place in March 2022?	$1, 2$
Sports	7	You will be given upcoming soccer games, and asked to predict the games' outcome?	$3, 4, 5$
Sports	7	You will be given upcoming sport events (soccer and tennis games), and asked to predict the games' outcome?	6
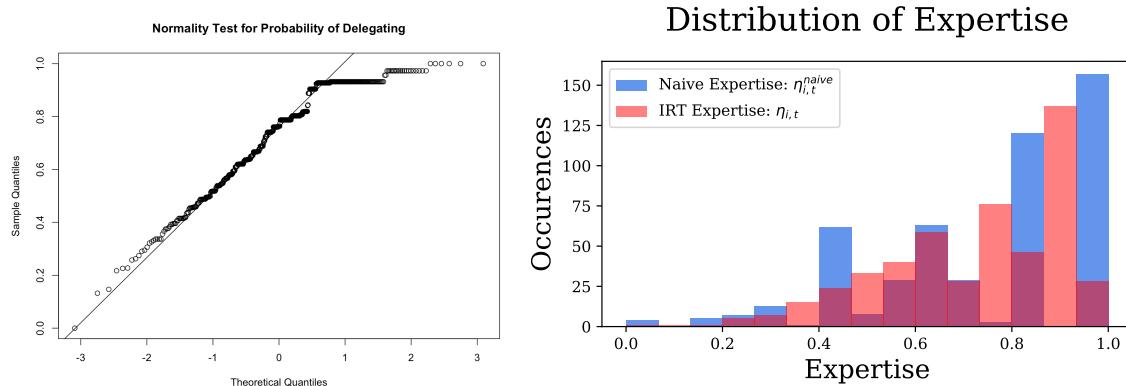
Figure 22   Normality Test for competence in pre-study (left) and Distribution of competence per delegation behavior (right)
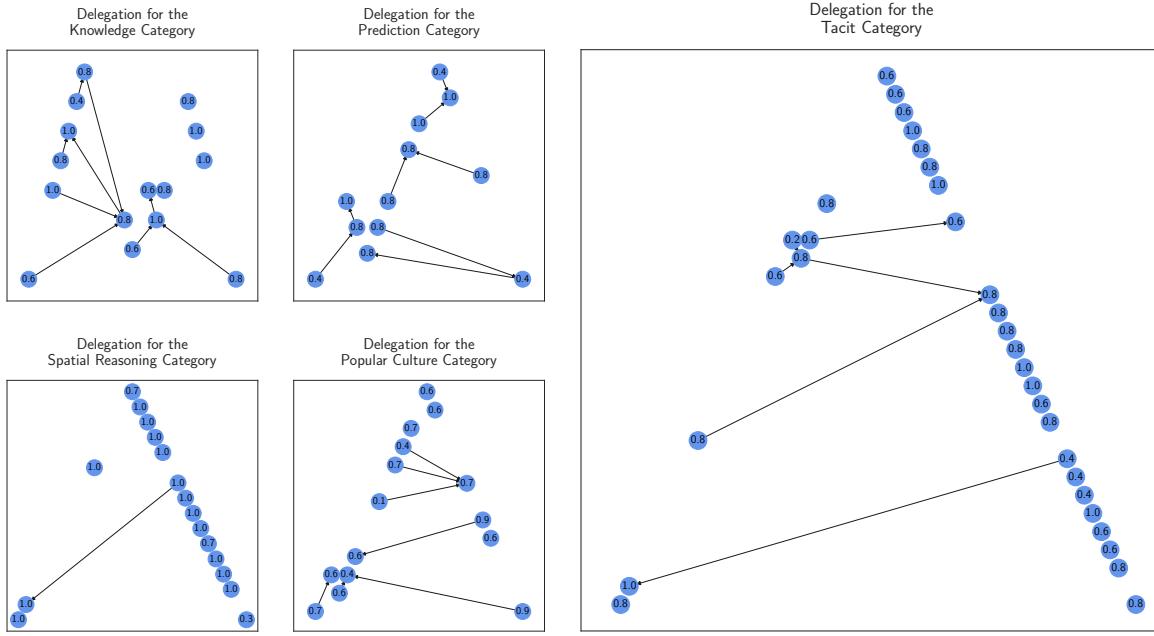
Figure 23 gives examples of some delegation graphs, where nodes labeled by their naive competence $\eta_{i,t}^{naive} = \sum_{i \in [N_e]} v_{i,e,t}^D / R_t$. The top left plot shows an example of a successful delegation chain to the right, where an expert from experiment 6 and task *landmarks*, with $\eta_{i,t} = 1$ was identified by six other participants either directly or transitively through a local expert j with $\eta_{j,t} = 0.8$. On the right, a smaller chain shows two participants delegating to a more competent expert, who in turn delegates to a non-expert.

Over the course of the experiments, we observed only two delegation cycles of size two (where A delegates to B, who delegates to A), both in Experiment 3 with $N_3 = 32$.

Table 16   Survey Material

Task	Prompt	Answer
Landmark	This landmark is located in Italy.	False
	This landmark is located in Turkey.	True
	This landmark is located in Myanmar.	False
	This landmark is located in France.	False
	This landmark is located in Brazil.	False
Movie	This music was featured as a theme song in the movie The Hobbit.	False
	This music was featured as a theme song in the movie The Empire of Sun.	False
	This music was featured as a theme song in the movie Gravity.	True
	This music was featured as a theme song in the movie Goodfellas.	False
	This music was featured as a theme song in the movie The Pianist.	False
	This music was featured as a theme song in the movie A Passage through India.	False
	This music was featured as a theme song in the movie The Schindler's List.	True
English	"A man of straw" means "A very active person".	False
	"To drive home" means "To emphasize".	True
	"To smell a rat" means "To suspect foul dealings".	True
	"To end in smoke" means "To excite great applause".	False
	"To catch a tartar" means "To deal with a person who is more than one's match".	False
Prediction for Experiments 1-2	The US college basketball team West Virginia Mountaineers will make it to the Elite Eight in the 2022 NCAA Tournament.	False
	The US college basketball team Michigan State Spartans will make it to the First Round in the 2022 NCAA Tournament.	True
	The US college basketball team Syracuse Orange will win the 2022 NCAA Tournament.	False
	The US college basketball team Purdue Boilermakers will make it to the 2nd round in the 2022 NCAA Tournament.	True
	The US college basketball team Arizona Wildcats will make it to the Elite Eight in the 2022 NCAA Tournament.	False
Prediction for Experiments 3-5	Galatasaray SK will beat FC Barcelona during the Europa League game on March 17th.	False
	Olympic de Marseille and OGC Nice will tie during the French League game on March 20th.	False
	VFL Wolfsburg will beat Bayer 04 Leverkusen during the German League game on March 20th.	False
	Salernitana will lose against Juventus during the Italian League game on March 20th.	True
	FC Barcelona and Real Madrid CF will tie during the Spanish League game on March 20th.	False
Prediction for Experiment 6	Eintracht Frankfurt will beat FC Barcelona during the Europa League game on April 7th.	False
	Olympic de Lyon and West Ham United will tie during the Europa League game on April 7th.	True
	Brazil will lose to Spain during the Women's International Friendly game on April 7th.	False
	Neither Rafael Nadal nor Novak Djokovic will qualify for the ATP Masters 1000 Monte Carlo Final on April 17th.	NA
	Stefanos Tsitsipas will win the ATP Masters 1000 Monte Carlo Tournament on April 17th.	NA
Spatial Reasoning	The object is located in the middle cup at the end of the trick.	False
	The object is located in the middle cup at the end of the trick.	False
	The object is located in the right cup at the end of the trick.	True

Note that in the pre-study, only 27% of the tasks were delegated, which is much less than the 47% delegation rate of the main study.

Figure 23   Examples of delegatoin graphs for each of the task categories

L.6.   Estimating q and φ

Next, we estimate the q and φ functions using the same methods applied on the main study. Table 17 shows the equivalent results of Tables 5 and 6 for the pre-study. We again, using the Kneedle algorithm, find the optimal number of clusters to be 4. Statistics about the clusters can be found in Table 18. Figure 24 and Table 19 show the equivalent information to Figure 2 and Table 2. We overall find qualitatively very similar results. We can conclude that both q is decreasing and φ is increasing in its second coordinate in statistically significant ways.

Table 17   Estimates of β_q
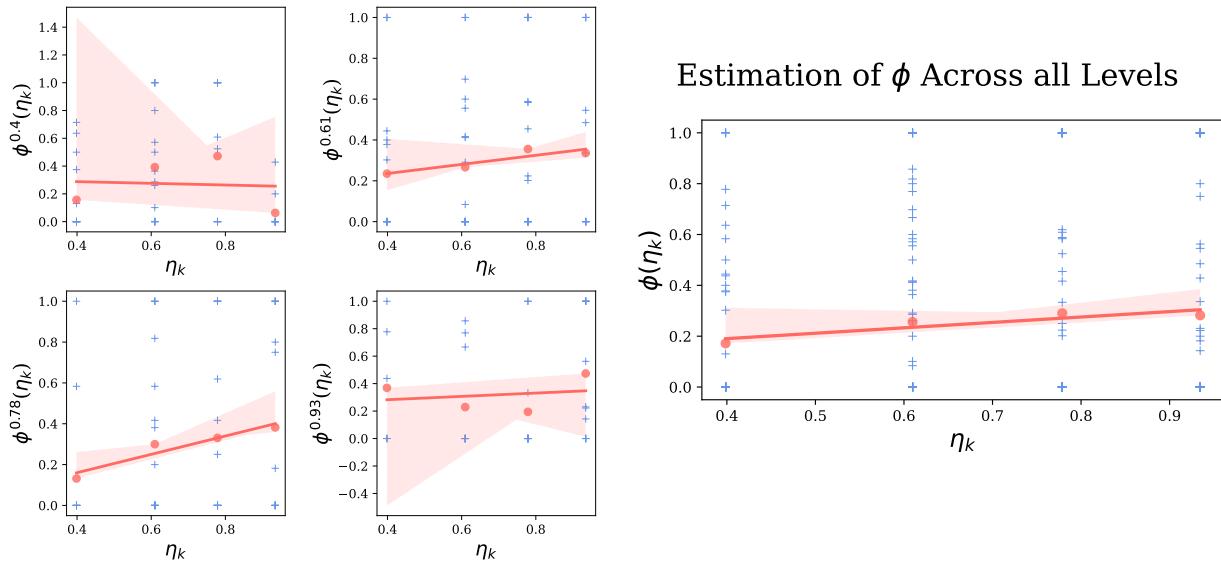
	Overall Model	Tasks Models						
		(Tasks)						
		(T_1)	(T_3)	(T_2)	(T_4)	(T_5)	(T_6)	(T_7)
Effect Size β^q	-2.45^{****}	-4.44^{**}	0.02	-5.11^{***}	0.48	-3.52	-3.10^{**}	-3.11
	(0.50)	(2.24)	(1.69)	(1.50)	(2.28)	(2.19)	(1.34)	(3.81)
Fixed Effects	NA	NA	NA	NA	NA	NA	NA	NA
Clustered S.E.	i	NA	NA	NA	NA	NA	NA	NA
Note:				*p<0.1; **p<0.05; ***p<0.01; ****p<0.0001				

Table 18   Bucket Descriptions

| Bucket	Interval	Mean competence	Proportion of participants
c_1	$[0.00, 0.50]$	0.40	15%
c_2	$[0.51, 0.69]$	0.61	26%
c_3	$[0.70, 0.84]$	0.78	27%
c_4	$[0.89, 1.00]$	0.94	33%

ϕ function for various expertise levels

Estimation of ϕ Across all Levels

Figure 24   **Pooled estimates of $\varphi_{e,t}^{\ell}$, both for each bucket individually, and grouped together. The blue crosses show the values computed for $\varphi_{e,t}^{\ell}(\eta_k)$. The pink dots show the average across all values for that η_k, and the pink lines correspond to a linear regression over the mean values. We observe increasing trends across the board, with slope (coefficient of determination) being $-0.05(0.005), 0.18(0.81), 0.39(0.91)$ and $0.10(0.05)$ respectively, for individual buckets, and $0.21(0.81)$ for the pooled test.**

Table 19   Summary of correlation statistics

	Overall	For fixed ℓ			
		c_1	c_2	c_3	c_4
Correlation	0.099^*	-0.009	0.079	0.20^*	0.19
P-value	8×10^{-1}	0.93	0.46	9×10^{-1}	0.11
Note:		*p<0.1; **p<0.05; ***p<0.01; ****p<0.0001			