Delegation under Liquid Democracy Two Experiments.*

Joseph Campbell, Alessandra Casella, Lucas de Lara,
Victoria Mooers, and Dilip Ravindran

May 6, 2022

Abstract

Under Liquid Democracy (LD), decisions are taken by referendum, but voters are allowed to delegate their votes to other voters. Theory shows that in common interest problems where the experts are correctly identified, the outcome can be superior to simple majority voting. However, delegation reduces the variety of independent information sources and can be advantageous only if used sparely. We report the results of two experiments, each studying two treatments: in one treatment, participants have the option of delegating to better informed individuals; in the second, participants can choose to abtain. The first experiment follows a tightly controlled design planned for the lab; the second is a perceptual task run online where information about signals' precision is ambiguous. Although the experiments are very different, they reach the same result: in both, LD underperforms relative to both universal voting and abstention. In the lab experiment, we observe systematic overdelegation relative to the best theoretical eqiulibrium; in the perceptual task, we lack a precise theoretical benchmark, but delegation rates remain significantly higher than abstention rates. Faced with better informed experts, voters neglect the value of their own independent information.

JEL codes: **C92, D70, D72, D91**

Keywords: voting rules, majority voting, information aggregation, laboratory experiments, Condorcet Jury theorem

*We thank audiences at Columbia, the Paris School of Economics, Princeton, the University of Bologna, and the Political Economy UK network. We are particualrly grateful to Mael Lebreton, Jonathan Nicholas, Nahuel Salem, and Camilla van Geen for their help with the second experiment. We thank the Program for Economic Research at Columbia and the Columbia Experimental Lab for the Social Sciences for their financial support. The experiments have been reviewed and approved by Columbia University's IRB.

1 Introduction

> I believe that some sort of computerized participation by large numbers of the public in opinion formation and direct policy-making is in the cards in the next ten to twenty years. It may be that we will be able to turn this new technology to the improvement and defense of democratic institutions. I hope so. However, it is by no means evident that this will be the result. (Martin Shubik, 1970, commenting on Miller, 1969)

In Western societies, the sense of living a crisis of traditional political institutions is bringing calls for different, more participatory forms of democracy. Among these, Liquid Democracy has caught the imagination of the young and the tech-savvy. It advocates a voting system where all decisions are submitted to referendum, but voters can delegate their votes freely. Beyond its intellectual roots in the writings of Charles Dogson (in particular, Dogson, 1884), Liquid Democracy was proposed more recently by James Miller in 1969. It has been adopted occasionally for internal decisions by European protest parties (the Swedish and the German Pirate parties being the most famous examples), and now finds vocal support in the tech community, where it aligns both with the emphasis on a non-hierarchical order and with the use of cryptographic tools to maintain confidentiality and reliability.[1] Although the details varies, the common point of different implementations is the ease and specificity of delegation. Supporters herald it as the golden medium between representative and direct democracy: better than the former because representatives can be chosen according to their specific competence on each decision, better than the latter because uninformed or uninterested voters can delegate their votes.

However, even if the difficult problem of identifying genuine experts were overcome, Condorcet has taught us that delegation reduces the aggregate information available in the system. Unless voters are able to modulate the extent of delegation correctly, a smaller number of independent voters may well lead to worse decision-making. This very basic trade-off is the necessary point of departure of Liquid Democracy and is the focus of this paper.

We study a canonical common interest model where voters receive independent signals, conditional on an unknown state ex ante equally likely to take one of two values. The

[1] See for example LiquidFeedback (https://liquidfeedback.com/en/), the Association for Interactive Democracy (https://interaktive-demokratie.org/association.en.html), or Democracy.Earth (https://democracy.earth/). Google ran a 3-year experiment on its internal network, implementing Liquid Democracy for decisions like food menu choices, tee-shirt designs, or logos for charitable events (Hardt and Lopes, 2015). Liquid Democracy is becoming the governance choice for cryptoworld DAO's (Decentralized Autonomous Organizations)–see for example Element Finance (https://medium.com/element-finance).

common objective is to identify the state correctly, aggregating information via majority voting. Signals vary across individuals in the probability of being correct—a variable we denote as *precision*. Experts are publicly identified and the precision of their signals is known; for all other voters, signals' precisions are private information but known to be weakly lower than the experts'. If a voter chooses to delegate, the vote is randomly assigned to one of the experts. We begin by showing theoretically that for any size of the group and any number of experts, there is an equilibrium with positive delegation such that the outcome is superior to majority voting without delegation. However, in such an equilibrium delegation must not be too frequent, given its informational cost. The finding is not surprising, but the equilibrium frequency seems counter-intuitively low. For example, consider one of the parametrizations we study: a group of 15 voters of which 3 are experts; the experts' information is correct with probability 70%, while the precision of non-experts' signals can take any value between 50 and 70%, with equal probability. Then a non-expert with information that she knows is only 55% likely to be correct should not delegate to experts whom she knows to be correct with probability 70%. And mistakes are costly: small errors towards over-delegation lead to expected losses that soon become severe. In actual implementations, other factors, for example overconfidence and overweighing of own information, could introduce countervailing forces. It is with these concerns in mind that we test Liquid Democracy with two very different experimental designs.

Before describing the experiments in detail, note that the informational benefit from overweighing voters with more precise signals can be achieved via abstention as well, as long as abstention correlates with more inaccurate signals, as it will in the equilibria of a pure common interest model. Abstention is a familiar option and does not require any transfer of votes, reducing the appearance of suspicious deals. It differs from delegation because the increase in voting weight concerns *all* individuals who choose to vote, not only those targeted as delegates. Yet, we know from McMurray (2013) that, under common interest and in the absence of voting costs, abstention too can lead to improvements over simple majority voting, and for reasons very similar to those favoring delegation. The relative performance of the two voting system–delegation and abstention–is thus an interesting question per se, and our experiments study both. We can the second system Majority Voting with Abstention (MVA)

The first experiment was designed for the lab and follows the theory very closely. We study groups of either 5 voters (of which 1 is an expert) or 15 voters (of which 3 are experts). We observe the frequency of delegation and the fraction of group decisions that yield the correct outcome. We then compare these results to a second treatment, where the option of abstention takes the place of delegation, and we evaluate both relative to simple majority voting with voting by all. We find systematic over-delegation: delegation rates that are

between 2 and 3 times the rate in the unique strict equilibriumt, given the realized experimental precisions. As a result, Liquid Democracy underperforms, relative to simple majority without delegation. Abstention rates, on the other hand, are very close to the theory, and the fraction of correct decisions is comparable to what majority voting without abstention (or delegation) would deliver. Interestingly, MVA suffers from its own sub-optimal behavior: the environment is symmetric, and voting in line with one's own signal is optimal, but experimental subjects occasionally deviate, and deviate more when abstention is allowed. The reason is that, in our data, voting according to signal correlates positively with the signal's precision, and since more subjects vote under abstention, at lower precisions, we also observe more votes against signal. Although the frequency of such deviations remains low, the result is a decline in correct group decisions that prevents MVA from reaping the gains over universal voting that theory predicts. The conclusion of our first experiment, then, is that even when experts are correctly identified and both Liquid Democracy and MVA have the potential to dominate universal majority voting, both systems in fact fail to do so. Liquid Democracy in particular shows a more clearly detectable negative effect.

We designed the experiment to replicate the simplest theoretical model of voting under common interest. But could the design itself be biasing results against Liquid Democracy and MVA? There are three reasons to consider the question. First, as mentioned, the theoretical thresholds for equilibrium delegation seem counter-intuitive. But we would like to argue that what makes them counter-intuitive is not so much the expected low frequency of delegation, but the detailed mathematical manner in which information is conveyed, and thus thresholds are chosen. Each participant is told a number for her own precision and is naturally induced to compare such number to the known precision of the experts–in the example we made earlier, escaping the obvious point that 55% is lower than 70% seems very difficult. In reality, voting decisions take place in an ambiguous world, where individuals do not have explicit numerical knowledge of the reliability of their and others' information. Evaluations are fuzzier and leave more room for bias. The large over-delegation we find in our first experiment under Liquid Democracy could be less pronounced in a less precise world. A similar argument may also affect voting against signal under MVA. The detailed mathematical design gives us clean theoretical predictions, but could in fact be confusing participants, no matter how much we clarify the instructions. At 55% precision, thinking that one should vote against signal about half the time is a reasonable thought. In an actual voting situation, however, lacking an explicit mathematical value for the probability that one's information is correct, it is unlikely that individuals would vote against their best estimate of the right decision.

Finally, there is a third reason for considering a less controlled environment. Could our detailed mathematical design be favoring universal majority voting? In a still current analy-

sis of the Condorcet Jury theorem, suggestively titled "A Note on Incompetence", Margolis (1976) commented on what he saw as the contradiction between the Condorcet theorem and political reality: in common interest problems, "Why is not decision-making trivially easy?". The co-existence of private interests, correlated signals, asymmetric scenarios, all may lead direct democracy to function less well, but Margolis proposed a different explanation, typically unexplored in the formal literature. What if, over some questions and for some voters, information is actually correct with probability *lower* than 1/2? For any individual voter this will not be true when averaging over many decisions, but may well be true over some. And it will affect the aggregate outcome. By opening some space for improvement over universal majority voting, note also that the possibility of worse than random information allows testing Liquid Democracy in larger electorates, where it is meant to be applied, while maintaining pure common interest and independent signals. With binary choices, conveying information with less than random accuracy requires noise.[2] And although the noise could be introduced in precise mathematical formulation, together with the arguments above, we see the need for noise as another reason to accept the realistic ambiguity of collective decision-making.

Our second experiment then is meant to capture a voting environment such that voters have " some sense" of how well-informed they are and of how likely they are to be correct, and similarly of how likely experts are to be correct, but such sense is vague and instinctive, and may in fact be incorrect. There is of course a cost: we lose the precise control granted by Experiment 1. But we can exploit a very rich literature that studies problems with exactly these features: the large literature in psychology and neuroscience that studies perceptual tasks. Our focus is not on measuring accuracy of perception, but on designing the task as a group decision problem.

The Random Dot Kinematogram (RDK) is a classic perceptual task amply used in vision and cognitive research.[3] A number of moving dots are displayed for a very short interval; some move in a coherent direction, either Left or Right in our binary implementation, others move at random; subjects report in which direction they think coherent dots are moving. We can label experts ex post as the individuals with performance in the highest quintile, and generate a collective decision by aggregating individual responses, with the additional option of delegation to the experts (in the Liquid Democracy treatments) or abstention (In the MVA treatments). We ran the experiment on Amazon Mechanical Turk with three electorate sizes: $N = 5$, and $N = 15$, as in Experiment 1, and a larger electorate of $N = 125$.

[2]In the absence of noise, if a signal is more likely to be wrong, its negative is more likely to be correct.

[3]It was originally developed to study the perception of motion under noisy conditions in humans and non-humans primates (e.g. van de Grind et al., 1983). In neuroscience, it has been used to study the neuronal correlates of motion perception (Newsome et al., 1989, Britten et al., 1992 and Roitman and Shadlen, 2002)

We reach three results. First, given our experimental parameters, it is not rare for individuals' accuracy to be worse than random. And this even over a large number of decisions: aggregating at the subject level over all 120 tasks, between 9 and 12% of subjects have ex post accuracy strictly below randomness; between 11 and 17% do no better than randomness. If we want to study voting and information aggregation when information may be faulty, perceptual tasks can provide a very useful tool.

Second, we do not know individuals' beliefs about their accuracy and the accuracy of others, and cannot compare the extent of delegation or abstention we see in the data to precise theoretical benchmarks. Yet, remarkably, we find the same patterns we saw in Experiment 1. The distributions of voters' accuracy we observe in the two samples–Liquid Democracy and MVA–are effectively identical, but delegation is twice as frequent than abstention when $N = 5$, and more than 50% more frequent when $N = 15$[4], if anything accentuating the disparity observed in the first experiment. Between one fourth and one third of subjects choose to abstain, but about half, in all treatments, choose to delegate.

Third, the high frequency of delegation exacts its expected informational costs. Even with a relatively high fraction of random, or below random subjects, universal majority voting remains the best information aggregator, delivering the highest frequency of correct group decisions in all treatments; MVA is only slightly less efficient, while Liquid Democracy is dominated by both in all treatments.

The robustness of the conclusions across two very different experimental designs is the main contribution of this study. The second design sacrifices experimental control in exchange for a less mathematical formulation; it conveys less information and leaves more space for idiosyncratic responses; it is run on MTurkers rather than students, and it is much shorter. And yet, and contrary to our expectations, treatment effects in the second design closely replicate the effects we observe in the first. In a pure common interest setting, where experts are correctly identified, individuals over-delegate. Increasing the voting weight of the experts does not lead to an increase in efficiency because the extent of delegation is too high, and thus the net informational effect is negative. Experimental subjects are less prone to abstaining, and thus the simpler, existing option of allowing abstention leads to better outcomes than allowing delegation.

The second contribution of this study is methodological. While the controlled design of our first experiment in the end delivers robust conclusions, we think that it is important to add to our experimental tool-kit designs that recognise the ambiguity in group decision-making. Experiments on ambiguity at the individual level are common; to our knowledge

[4]It is 75% more frequenct when $N = 125$.

they are much less so for collective decision-making.[5] In voting problems, in particular, the complexity of many questions, and the asymmetry between the cost of acquiring information and the small marginal impact of a single vote make the lack of precise information very likely. Perceptual tasks, with the large and sophisticated literature that accompanies them, can be a particularly usable tool. In this study, it is the combination of a strictly controlled design in the lab with the freer design of the perceptual task that we think teaches us the most. And it is this combination that we recommend.

Our work is related to three separate literatures. The informational costs and benefits of delegating to better informed individuals in pure common interest voting problems were the subject of early studies on the Condorcet Jury theorem (Margolis, 1976, Grofman et al., 1983 and 1983, Shapley and Grofman, 1984), highlighting the trade-off between the loss in aggregate information and the more precise information of the individuals actually casting votes. More recent work (Austen-Smith and Banks, 1996**;** Feddersen and Pesendorfer 1997**;** McLennan, 1998; Wit, 1998**)** put the analysis of the Condorcet Jury theorem on solid equilibrium grounds, but abstracted from the focus on delegation. We did not find in the literature our starting theoretical result, on the existence of an equilibrium with delegation that dominates universal majority voting, but it builds on the work of McLennan**.** The trade-off identified in the case of delegation exists also in the case of abstention. Here the best-known work includes partisan voters (Feddersen and Pesendorfer, 1996), but the analysis can also be profitable and rich in a pure common interest setting, as shown by Morton and Tyran (2011) in the case of three voters, and more generally by McMurray (2013). Battaglini et al (2010) and Morton and Tyran (2011) test the theoretical predictions in the lab. The latter model is closer to ours, and, contrary to Battaglini et al., allows for a range of information types and does not rely on the existence of perfectly informed voters. Interestingly, it finds that experimental subjects appear predisposed towards abstaining, doing so also even when abstention is dominated. In Morton and Tyran's words, subjects "follow a norm of "letting the experts decide"". According to our results, this tendency is strengthened further when the choice is explicitly phrased as delegation. We are not aware of existing experimental studies in this literature that expressly study delegation in voting. [**To be checked further**].

The second strand of related works are studies of Liquid Democracy. Most theoretical works belong either in normative political theory or in computer science. Green-Armytage (2015) and Blum and Zuber (2016) discuss what they see as normative advantages of Liq-

[5]There is an increasing focus on strategic uncertainty. But the question is different from the lack of basic information about the distributions of relevant parameters in the population, and even own parameters (accuracies, for us)

uid Democracy, on both epistemic and equalitarian reasons: decisions are taken by better informed voters, but Liquid Democracy avoids the creation of a detached class of semi-permanent professional representatives. Because the focus is on ideal properties, these studies do not analyze strategic incentives. The computer science literature is instead largely concerned with understanding how Liquid Democracy would work in practice. It models behavior via a priori algorithms and studies rich interactions where delegation takes place on networks (Christoff and Grossi, 2017; Kahng, Mackenzie and Procaccia, 2018; Bloembergen, Grossi and Lackner, 2019; Caragiannis and Michas, 2019). These authors connect Liquid Democracy to the social choice tradition, but here too strategic considerations are absent. An exception is Armstrong and Larson (2021) which discusses the informational trade-off involved in delegation and identifies a Nash equilibrium that always strictly improves over universal majority voting. The paper retains the algorithmic flavor of this literature by modelling the delegation choice as sequential; the common interest nature of the problem, together with the added assumptions of complete information and costly voting, then delivers the equilibrium superiority of delegation. Finally, strategic concerns are at the heart of two recent paper in economics, Ravindran (2021) and Dhillon et al. (2021). In Ravindran's model, voters' types are binary and known, with either higher or low information accuracy, and the goal is the characterization of the efficient equilibrium. With a single expert, optimal delegation is defined precisely; with multiple experts, complications can arise although, as in Armstrong and Larson, they can be solved if delegation decisions are sequential. Dhillon et al. study delegation in a model à la Feddersen and Pesendorfer (1996), with partisan voters and perfectly informed experts. As in the papers just discussed, they show that under complete information delegation has desirable properties: the game is dominance-solvable and delegation allows voters to coordinate on the best equilibrium. With incomplete information, multiple equilibria are more difficult to avoid and results are weaker. None of these works is experimental.

Finally, a literature in psychology studies a question that is closely related to our second experiment: if a group of individuals face, individually, a perceptual task but can then aggregate their reactions into a group decision, under what conditions will the group decision be more accurate? In Sorkin et al.(1998), a small group of subjects are faced with a signal detection task and asked whether the display reflects noise only or signal plus noise. Aggregation improves performance and simple majority rule leads to the highest accuracy. Individual behavior is modeled as reflecting two main parameters, detection sensitivity and confidence, and the emphasis on confidence shapes the direction this research has since taken. In small groups, decision typically follows discussion, and during discussion individual confidence translates into influence. Communication thus threatens group accuracy, unless confidence

correlates positively with individual sensitivity (Sorkin et al., 2001, Bahrami et al. 2012, Silver et al. 2021).[6] Although it seems a natural next step, we are not aware of similar studies that include the possibility of delegation.

In what follows, we begin by describing the theoretical model (Section 2) and its equilibrium properties (Section 3). We then discuss our first experiment: its parameterization and treatments (Section 4); its implementation (Section 5), and its results (Section 6). Section 7 describes the motivation and the design of our second experiment; Section 8 [Incomplete] reports its results.

2 The Model

We study the canonical problem of information aggregation through voting in a pure common interest problem. N odd voters face an uncertain state of the world ω and must take a decision d. There are two possible states of the world, $\omega \in \{\omega_1, \omega_2\}$, and two alternative decisions $d \in \{d_1, d_2\}$. Every voter's payoff equals 1 if the decision matches the state of the world ($d = d_s$when $\omega = \omega_s, s = 1, 2$), and 0 otherwise. Every voter's expected payoff, which we denote by EU, thus equals the probability that the group reaches the correct decision**.** Voters share a common prior $\pi = Pr(\omega_1)$ and receive conditionally independent signals $\sigma_i \in \{1, 2\}$. We call q_i the *precision* of individual i's signal, or the probability that i's signal is correct. Precision varies across individuals but is symmetric over the two possible states of the world: $q_i = Pr(\sigma_i = 1|\omega_1) = Pr(\sigma_i = 2|\omega_2)$.

The group of N voters is composed of K (odd) experts and M (even) non-experts. All experts receive signals of known precision $q_e = p$. The precisions of non-expert signals are instead private information; they are independent draws from a commonly known distribution $F(q)$ with support $[\underline{q}, \overline{q}]$, with $\underline{q} = 1/2$ and $\overline{q} = p$. The signals themselves are also private information, for both experts and non-experts, but the identity of each voter, as either non-expert or expert, is commonly known.

Each individual holds a single, non-divisible vote. After having learnt her precision q_i but before having received the signal, the voter decides whether to keep her vote or to delegate it. If the voter chooses delegation, she specifies whether it is to an expert or to a non expert; the vote is then assigned randomly, with equal probability, to any individual in the indicated category.[7] We denote the probability of delegation by δ_e, for delegation to

[6]The earlier experiments in this tradition studied a very large number of tasks, in the hundreds for each subject, but a very small group of subject, as small as 8 or 12.

[7]The random assignment of delegated votes is the natural assumption in the absence of distinguishing characteristics across experts. It also leads to some desirable spreading of delegated votes, as advocated for example by Gőlz et al. (2018).

an expert, and δ_{ne} for delegation to a non-expert. Once delegation is decided, signals are realized and individuals who have not delegated their vote indicate their preferred decision. At the voting stage, the strategic decision is the probability of voting according to one's signal, or against it, a decision we denote by $\{v_\sigma, v_{-\sigma}\}$. When counting votes, each voter who has chosen not to delegate receives a weight equals to the number of votes delegated to her, plus 1. The decision receiving more votes is chosen. Note two implications. First, because delegation is decided before the signal is received, the delegation decision itself cannot convey information about the signal. Second, because voters holding multiple votes communicate only their preferred decision, they are constrained to cast all votes in the same direction. Both implications will simplify the equilibrium characterization.

The game is thus a two-stage game, with strategies $\{\delta_e(q_i), \delta_{ne}(q_i)\}$ in the first stage, and, conditional on no delegation, $\{v_\sigma(q_i), v_{-\sigma}(q_i)\}$ in the second stage. With an eye to the experimental implementation, we will select equilibria that require little coordination, and in particular such that experts never delegate, and non-experts only delegate to experts. Hence, multi-step delegation (i delegates to j who delegates to z) will not be observed in equilibrium, and thus neither will circular delegation flows (i delegates to j who delegates to z who delegates to i). The model nevertheless needs to specify what would happen in such cases. We allow for multi-step delegation: if delegation targets a voter who has herself chosen delegation, the full packet of votes is delegated according to her instructions. However, if a set of delegation decisions results in a circular delegation flow, we specify that all delegations involved are void.[8]

3 Equilibrium

We study an environment that matches the experimental set-up, and where, specifically, $\pi = Pr(\omega_1) = 1/2$. In this symmetric environment, voting according to signal ($v_\sigma = 1$ for all q_i)is an undominated strategy, a result that holds whether delegation is allowed, as in our model, or is not, as in traditional majority voting. We focus on semi-symmetric Perfect Bayesian equilibria in undominated strategies where, when voting, voters follow their signal, and delegation strategies are symmetric for all voters of given type (expert or non-expert). In what follows, "equilibrium" refers to such a notion. We are interested in the welfare properties of delegation, and say that an equilibrium "strictly improves over majority voting" if in equilibrium the ex ante probability of reaching the decision that matches the state of the world is strictly higher than under (sincere) majority voting (MV), or $EU_{LD} > EU_{MV}$.

[8]We include all such delegation decisions: if i and j delegate to z, and z delegates to i, the original votes are returned to all three voters (all three are left with a single vote).

"Ex ante" indicates that the expectation is taken before the realization of the precisions, and of the signals. Our most general theoretical result is summarized in the following theorem:[9]

Theorem. *Suppose $\pi = Pr(\omega_1) = 1/2$. Then for any F, and any M and K finite there exists an equilibrium with delegation that strictly improves over MV.*

The result is of interest because the environment we are studying is particularly favorable to MV. Given the symmetric prior and information structure, the Condorcet Jury theorem applies to rational voting, and thus we know that MV both converges to the correct decision with probability 1 asymptotically, as the size of the electorate becomes unbounded. In addition, we are restricting our attention to semi-symmetric equilibria, and thus excluding asymmetric profiles of strategies that require demanding coordination.[10] And yet, the theorem states that there always exists an equilibrium where the possibility of delegation strictly improves over MV.

We prove the theorem in the appendix, but the intuition is both straightforward and interesting. The essence of the proof is that, when delegation is possible and some voters' information may be barely better than random, there cannot be an equilibrium where delegation is excluded with probability 1: every voter casting their vote with probability 1 (and thus replicating MV) is not an equilibrium. But in this common interest problem, we know from McLennan (1998) that the set of strategies that maximize expected utility must be an equilibrium. Together, the two observations imply that an equilibrium exists, that it must include a positive probability of delegation, and that it must dominate MV. In addition, again from McLennan (Theorem 2), because the environment is fully symmetric for all voters of given type, the conclusion continues to apply when we restrict attention to semi-symmetric strategies.

To see intuitively why there cannot be an equilibrium where every voter cast their vote with probability 1, consider, for ease of exposition only, the simple case of a single expert: $K = 1$. Suppose everyone else votes (sincerely) and consider the choice for a voter with precision q_i near the lower bound of the precision distribution $F(q)$. The voter must decide whether to cast her vote and not delegate (an action yielding expected utility $EUND(q_i)$), or delegate the vote to others, and in this case the best choice is to delegate to the expert.

[9] We grant the result the name of "theorem" because it refers to the default canonical model of information aggregation under common interest, and yet we have not found it in the literature. However, many results with similar flavor do exist. See for example Grofman, Owen and Feld (1982) and (1983).

[10] Building on Nitzan and Parousch (1982), and Shapley and Grofman (1984), Ravindran (2021) characterizes the highest welfare equilibrium in the case of a single expert. It is an asymmetric equilibrium where the total number of votes delegated to the expert must mirror the expert's precision (or, more precisely, must be proportional to $ln(p/(1-p))$.

Delegation to the expert yields expected utility $EUD(q_i)$. It is not difficult to verify that:

$$lim_{q_i \to \underline{q}}(EUND(q_i) - EUD(q_i)) =$$
$$\underline{q}(1-p)\left[\binom{M-1}{M/2}\mu^{M/2}(1-\mu)^{M/2-1}\right] - p(1-\underline{q})\left[\binom{M-1}{M/2-1}\mu^{M/2-1}(1-\mu)^{M/2}\right] \quad (1)$$

where $\mu = \int_{\underline{q}}^{\overline{q}} q dF(q)$.[11] The expression corresponds to the difference in expected utility from the two actions, conditional on voter i being pivotal and disagreeing with the expert, since delegation matters only in such a case. Not delegating is advantageous if the number of correct votes cast by non-experts falls short of majority by a single vote, and the expert's signal is not correct, while i's signal is–the first term on the right-hand side of equation (1). Delegating is advantageous if the number of correct votes cast by non-experts falls short of majority by 2 votes, and the expert's signal is correct, while i's signal is not–the second term on the right-hand side of equation (1). Since $\binom{M-1}{M/2} = \binom{M-1}{M/2-1}$,it follows that:

$$Sign[lim_{q_i \to \underline{q}}(EUND(q_i) - EUD(q_i))] = Sign\left[\underline{q}(1-p)\mu - p(1-\underline{q})(1-\mu)\right]$$

or:

$$lim_{q_i \to \underline{q}}(EUND(q_i) - EUD(q_i)) < 0 \Longleftrightarrow [\underline{q}(1-p)\mu - p(1-\underline{q})(1-\mu)] < 0.$$

With $\underline{q} = 1/2$, the condition simplifies to $\mu < p$, which is always satisfied. Hence there exists no equilibrium such that all non-experts prefer voting to delegation for any precision of their signal. Two observations conclude the argument. First, note that assuming $\underline{q} = 1/2$, while important for the proof, is a natural assumption: $\underline{q}$ cannot be inferior to 1/2 (because, with binary states, a signal correct with probability inferior to 1/2 is equivalent to the opposite signal being correct with probability higher than 1/2), and thus 1/2 is the natural lower boundary of the support of the precision distribution. Second, as we show in the appendix, generalizing the model to K experts does not modify the argument, and indeed the same simple condition identified here remains the critical condition when there are multiple experts.

The theorem does not characterize the equilibria with delegation. In the next section, we will do so when we specialize the model to the parameter values we use in the experiment. Here, to make sure the mechanisms that drive the model are intuitively clear, we describe in detail the unique semi-symmetric equilibrium in undominated strategies that exists in the

[11] Because signals are independent, when all non-experts but i are supposed to vote, the probability that x of them receive a correct signal simplifies to $\binom{M-1}{x}\mu^x(1-\mu)^{M-1-x}$

case of a single expert.

3.1 A Single Expert ($K = 1$)

Proposition. *Suppose $\pi = Pr(\omega_1) = 1/2$ and $K = 1$. Then for any M, there exists an equilibrium such that: (i) the expert never delegates her vote and always votes according to signal; (ii) there exists a unique $\widetilde{q}(M) \in (\underline{q}, \overline{q})$ such that non-expert i delegates his vote to the expert if $q_i < \widetilde{q}$ and votes according to signal otherwise. (iii) Such an equilibrium strictly improves over MV and is ex ante maximal among sincere semi-symmetric equilibria where the expert never delegates and nonexperts delegate to the expert only.*

The proposition is proved in the Appendix. The structure of the equilibrium, however, is intuitive. Note that in all equilibria (and in fact, for any number of experts), non-experts must adopt monotone threshold strategies–there must exist a precision threshold $\widetilde{q}$ such that voters with lower precision delegate, and voters with higher precision do not. The reason is immediate: if the voter delegates, expected utility does not depend on the voter's precision. But if the voter does not delegate, there is a non-zero probability the voter is pivotal, in which case expected utility increases with the voter's precision. The conclusion then follows. Given monotone threshold strategies, the Appendix then shows that the delegation directions in the proposition–the expert never delegating and non-experts delegating to the expert only–are indeed best responses when all others adopt them too. And since we know, from the earlier theorem, that an equilibrium with partial delegation must exist, it follows that an equilibrium with the strategies characterized in the proposition must exist. Finally, we also find that the condition identifying the equilibrium threshold corresponds to the first order condition from the maximization of ex ante expected utility, over all profiles of semi-symmetric monotone threshold strategies with sincere voting and the specified directions of delegation. Hence, again invoking the theorem, the equilibrium is maximal over such profiles and improves strictly over majority voting.

It is important to note that the equilibrium threshold $\widetilde{q}$ that supports the improvement over MV is strictly interior to the range $(\underline{q}, \overline{q})$. Because $\overline{q} = p$, that means that in equilibrium there are voters who know that their precision is strictly lower than the expert's precision, and yet cast their vote, rather than delegating. Since delegation decreases the aggregate information of the system, and yet the equilibrium with delegation is superior to MV, we expect the threshold $\widetilde{q}$ to be low–only voters with very imprecise information delegate in equilibrium. Indeed, this is what the numerical examples will show. Studying in detail a particularly simple example makes clear why.

3.1.1 The simplest example: $N = 3$, and $K = 1$

Suppose $N = 3$ and $K = 1$, and consider non-expert i's choice of whether or not to delegate to the expert. Note that i conditions on the other non-expert not delegating, because only in such a case does i's decision matter. If i delegates, the expert has 2 votes and chooses unilaterally. Hence $EUD_i = p$. If i does not delegate, the correct decision is reached if i's signal is correct and at least one of the other two signals is correct as well, or if i's signal is incorrect but both of the others are correct. Or:

$$EUND_i = q_i[1 - (1 - \mu_v)(1 - p)] + (1 - q_i)p\mu_v$$

where $\mu_v = E(q|q > \widetilde{q})$.

Equilibrium $\widetilde{q}$ solves $EUD_{q_i=\widetilde{q}} = EUND_{q_i=\widetilde{q}}$:

$$p = \widetilde{q}[1 - (1 - \mu_v)(1 - p)] + (1 - \widetilde{q})p\mu_v$$

Or

$$p(1 - \widetilde{q})(1 - \mu_v) = \widetilde{q}(1 - p)\mu_v$$

Note that the equilibrium condition thus equalizes the probability that the expert is correct and *both the non-experts* are not, with the probability that the expert is incorrect and *both the non-experts* are correct. Intuitively, the delegation decision only matters: (i) in case of disagreement with the expert (since the choice is otherwise moot), and (ii) if the vote is pivotal, and the other non-expert not only does not delegate (has precision above $\widetilde{q}$),but disagrees with the expert (and hence agrees with the non-expert voter with precision $\widetilde{q}$). The non-expert signal with precision $\widetilde{q}$ receives implicit validation from a second independent and more accurate signal. It is this implicit validation that pushes equilibrium behavior away from delegation, and the equilibrium $\widetilde{q}$ towards low values. For example, with $N = 3$ and $K = 1$, if $p = 0.7$ and $F(q)$ is Uniform over [0.5,0.7],$\widetilde{q} = 0.572$,below the mean non-expert precision of 0.6. A non-expert voter's precision is always lower than the expert's, but the ex ante individual probability of delegation is only 36 percent.

3.1.2 Arbitrary N, $K = 1$

The intuition we just discussed extends to more general cases. The delegation decision must be conditioned on pivotality, and the effect of pivotality is to validate the non-expert's signal by conjecturing independent realizations of other non-experts' signals that agree with one's own. Maintaining $K = 1$, consider for example an arbitrary (even) number $M$of non-experts.

The threshold $\widetilde{q}$ solves[12]:

$$\widetilde{q}(1-p)\sum_{s=0}^{M/2-1}\binom{M-1}{s}F(\widetilde{q})^s(1-F(\widetilde{q}))^{M-1-s}\binom{M-1-s}{M/2}\mu_v^{M/2}(1-\mu_v)^{M/2-1-s}=$$
$$=p(1-\widetilde{q})\sum_{s=0}^{M/2-1}\binom{M-1}{s}F(\widetilde{q})^s(1-F(\widetilde{q}))^{M-1-s}\binom{M-1-s}{M/2-1-s}\mu_v^{M/2-1-s}(1-\mu_v)^{M/2}$$

where the symbol s indicates the number of other non-experts delegating to the expert. Voter $\widetilde{q}$ conditions on $M/2,..M-1$ other non experts voting, and always on $M/2$ of them, a majority,, agreeing with his signal.

The good properties of the equilibrium with $\widetilde{q}\in(0,1)$ depend strongly on the optimal, spare use of delegation. But internalizing such reasoning is difficult. Our first experiment tests participants' behavior in an environment that mirrors the model. In the next section, we describe the treatments we study in the lab and their specific parametrizations.

4 Experiment 1: Treatments and Parametrizations

In the lab, we make two simplifications relative to the model. First, we constrain the direction of delegation: experts cannot delegate, and non-experts can only delegate to the experts. Second, we combine the delegation and voting decisions: after being told her own precision and receiving her signal, every participant faces a single ballot with three possible alternatives: delegate, or vote for one of the two options. In the theoretical model, the delegation decision is modeled as preceeding the signal to ensure that delegation itself does not convey information about the signal. In the lab, with random group formation and partial feed-back, the coordination required to use delegation informatively with respect to the signal is in practice impossible to achieve. Without such a concern, unifying the decisions in a single ballot streamlines the experiment.

We are interested in three main questions: (1) We consider first the simplest setting, when decisions are taken by a small group and a single expert. How well does LD perform, relative to MV? (2) According to the theorem, LD's potential to improve over MV persists with larger group sizes and multiple experts. In the lab, do results change qualitatively when the size of the group and the number of experts increase? (3) LD makes it possible to shift voting weight away from less informed voters and towards more informed ones. But reducing the weight of less informed voters can also be achieved, more simply, by allowing abstention.

[12]See the Appendix for the explicit derivation

Table 1: $p = 0.7$, $F(q)$ Uniform over $[0.5, 0.7]$

$LD1 : N = 5; K = 1$

$\widetilde{q}$	$F(\widetilde{q})$	EU_{LD}	EU_{MV}
0.7	1	0.7	0.717
0.543	0.215	0.731	

$LD3 : N = 15; K = 3$

$\widetilde{q}^3$	$F(\widetilde{q}^3)$	EU^3_{LD}	EU^3_{MV}
0.532	0.162	0.843	0.832

How does LD compare to MV with abstention? We denote such rule by MVA, and as in the case of LD, study it both in a small group with a single expert, and in a larger group with multiple experts.

In all experiments, we set $p = 0.7$, and $F(q)$ Uniform over $[0.5, 0.7]$. We study four treatments. Two treatments concern LD. In LD1, groups consist of 5 voters with a single expert: $N = 5$, $K = 1$. In LD3, each group has 15 voters in all, of which 3 are experts: $N = 15$, $K = 3$. Hence in both treatments one fifth of the group are experts: $K/N = 1/5$. The two treatments with abstention, MVA1 and MVA3, substitute abstention for the possibility of delegation, again either with $N = 5$, $K = 1$ (MVA1), or with $N = 15$, $K = 3$ (MVA3).

4.1 Liquid Democracy

Table 1 reports the theoretical predictions when delegation is possible.[13]

In treatment LD1, we find two semi-symmetric equilibria. For any realization of non-expert precisions, there always exists an equilibrium where the expert alone controls the outcome: no individual non-expert is ever pivotal, and delegating one's vote is a (weak) best response. With semi-symmetric strategies, such an equilibrium corresponds to $\widetilde{q} = \overline{q}$ and yields ex ante utility $EU_{LD}(\widetilde{q} = \overline{q}) = p = 0.7$. In addition, there is a unique equilibrium where $\widetilde{q}$ is strictly interior. As argued earlier, the $\widetilde{q}$ threshold is low, and the ex ante probability of delegation is only just above 20 percent. The ex ante probability of reaching the correct decision, equivalent to the expected utility measures, is lowest when the expert decides alone ($\widetilde{q} = \overline{q}$), intermediate under MV, and highest in the equilibrium with delegation and interior $\widetilde{q}$. The proportional increase in the probability that the group selects the correct option is very small, about 2 percent at each step.[14]

In LD3, full delegation is not an equilibrium any longer. Intuitively, when there are

[13]The details of the derivations are in the Appendix.

[14]With a single expert, the uniqueness of the semi-symmetric equilibrium with interior $\widetilde{q}$ can be proven analytically and holds for arbitrary N. Absent either communication or repetition, asymmetric equilibria are implausible in the lab if they require coordination. But when the expert is dictator and no other voter is pivotal, the choice to delegate or not becomes irrelevant and trivial asymmetric equilibria becomes more realistic. For example, for any realization of non-expert precisions, there are equilibria with 3 non-experts delegating and 1 voting.

multiple experts and all other non-experts delegate, voter i can be pivotal only if the experts disagree among themselves. The disagreement reduces the attraction of delegation and for sufficiently high q (still smaller than $\overline{q}$) casting a vote is preferable. As we know, the equilibrium with interior $\widetilde{q}$ continues to exist. Equilibrium delegation is rare: the threshold $\widetilde{q}$ is smaller than with a single expert and the expected frequency of individual delegation falls to 16 percent. As theory teaches, the interior equilibrium yields a higher probability of a correct decision than MV. However, with the increase in the size of the group, the Condorcet Jury Theorem effect becomes very pronounced: majority voting works very well and the scope for improvement is small. The percentage gain is only 1.3 percent.

The table conveys two main messages. First, we see concretely what the interior equilibrium entails for the experimental parametrizations. In particular, as expected, equilibrium delegation is rare and concerns only voters with precisions not far from randomness. Second, the improvement in the probability of taking the correct decision is small, almost certainly too small to be detectable in the lab. Setting a higher p, and/or setting $\overline{q} < p$ would increase the scope and expected gain from delegation. Increasing $\underline{q}$ or N would have the opposite effect. We have chosen a parametrization that delivers similar efficiencies for the two rules to leave the data free to favor either. The realistic challenge for the experiment will be to see whether indeed in this environment LD and MV's welfare properties are comparable.

4.2 Abstention

Like delegation, abstention can lower the voting weights of less informed voters, with the major advantage of being a simpler and familiar option. However, the two mechanisms are not equivalent: under abstention, voting weight is redistributed towards *all* voters who choose to vote; under LD delegated votes target the experts only.

We implement the MVA treatments in the identical environment we study under LD. After non-expert voters learn, privately, the precision and the content of their personal signal, they decide, simultaneously and independently, whether to vote or to abstain. Experts are not given the option of abstaining. Everything else remains unchanged. The model of abstention is closely related to McMurray (2013), and its main results– the existence of an equilibrium in monotone cutpoint strategies, and its superiority to MV–carry over to our setting.[15] As in the case of delegation, and for very similar reasons, abstention too is limited

[15]McMurray's model and ours differ under two main aspects. First, for comparison to LD, we assume the existence of a known group of experts with higher, known, but not perfect precision. McMurray does not distinguish experts, but widens the support of the distribution of precisions $F(q)$ to cover the full interval $[1/2, 1]$. Second, because of our experimental aim, we assume that the size of the electorate is known and need not be large, deviating from McMurray's large Poisson game set-up. The logic of the two models is otherwise identical. The central intuition comes from the fact that best response strategies are monotone in individual

Table 2: $p = 0.7$, $F(q)$ Uniform over $[0.5, 0.7]$

$MVA1 : N = 5; K = 1$

$\widetilde{\alpha}$	$F(\widetilde{\alpha})$	EU_{MVA}	EU_{MV}
0.7	1	0.7	0.717
0.580	0.40	0.724	
0.5	0	0.717	

$MVA3 : N = 15; K = 3$

$\widetilde{\alpha}^3$	$F(\widetilde{\alpha}^3)$	EU^3_{MVA}	EU^3_{MV}
0.7	1	0.784	0.832
0.580	0.40	0.849	
0.5	0	0.832	

to voters with weak information.

Table 2 shows the equilibria with abstention, for the experimental parametrizations. We denote by $\widetilde{\alpha}$ the precision threshold below which in equilibrium a non-expert abstains, and above which a non-expert votes.

For both group sizes, there are three semi-symmetric equilibria. Two are boundary equilibria, with either zero ($\widetilde{\alpha} = 0.5$) or full ($\widetilde{\alpha} = 0.7$) abstention; one is an interior equilibrium where, for both group sizes, a non-expert abstains if precision is below 0.58, i.e. with ex ante probability of 40 percent. The boundary equilibrium with zero abstention corresponds to MV; the one with full abstention, where the decision is delegated to the experts, is inferior to MV. As in McMurray's analysis, the interior equilibrium does deliver expected gains over MV, but these remain quantitatively small.[16] The interior equilibrium threshold for abstention is higher than the threshold for delegation, and remains constant in the two group sizes. It implies a larger expected number of abstentions than delegations: for example, and rounding up to integers, when $N = 15$, in equilibrium we expect 2 non-experts to delegate under LD, but 5 non-experts to abstain under MVA.

Under both LD and MVA, the expected improvements over MV are minor. It is natural to ask how sensitive such potential improvements are to strategic mistakes.

4.3 Robustness

We consider here a particularly simple parametrization of strategic mistakes: we suppose that behavior remains symmetric, but the precision threshold for delegation or abstention is chosen incorrectly. In Figure 1, the horizontal axis is the common threshold, and the vertical axis reports expected utilities, or the probability of making the correct choice under each of the three voting rules, LD (in Orange), MVA (in Green) and MV (in Blue). The highest points on the Orange and on the Green curves coincide with the respective equilibrium

precision and thus abstaining in equilibrium shifts voting weight towards better informed individuals.

[16]The existence of the boundary equilibria depends on N and Kodd. Note that because MV without abstention is an equilibrium when abstention is allowed, the simple proof used to establish the earlier theorem cannot be extended from delegation to abstention.

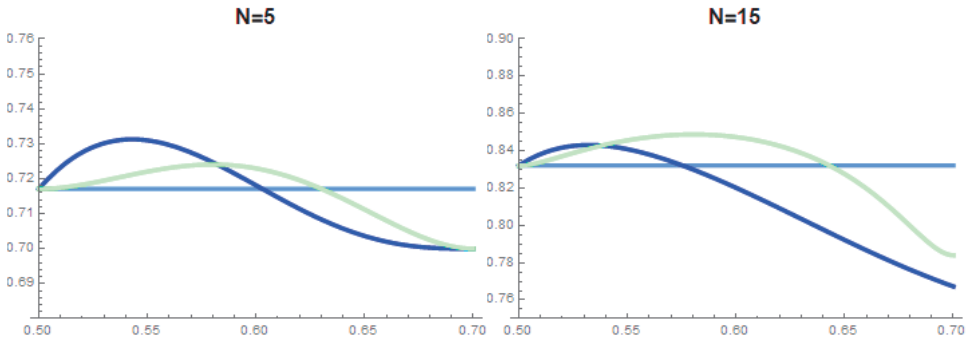


Figure 1: Robustness to strategic errors. The horizontal axis is $\widetilde{q}$ ($\widetilde{\alpha}$), the vertical axis is EU. Blue is LD, Green is MVA, and Grey is MV.

thresholds. The first panel corresponds to $N = 5$, $K = 1$; the second to $N = 15$, $K = 3$.

At $\widetilde{q}/\widetilde{\alpha} = 1/2$, no-one delegates their vote or abstains, and all curves equal MV and coincide. At $\widetilde{q}/\widetilde{\alpha} = 0.7$,all non-experts delegate or abstain, and only the expert/s decide(s).[17] In the first panel, with a small group, the maximum potential improvement over MV from delegation (from LD) is higher than from abstention (MVA). However, this is not true in the second panel, with the larger group. Both results were already shown in Tables 1 and 2. More interesting is the range of thresholds for which each voting rule dominates MV. Here the message is consistent across the two group sizes: in both cases, the range of thresholds that deliver improvements over MV is limited, and particularly limited for LD. When the group is larger, LD's potential for losses is evident in the figure, as is its increased fragility, relative to MVA: the range of thresholds that improve over MV is half as large under LD3 than under MVA3. With both voting schemes, but with LD in particular, while potential gains are small, there is the real danger of reaching worse decisions: under LD3, maximal potential losses are more than six times maximal potential gains.

5 Experiment 1: Implementation

We ran the experiment online over the Summer of 2021, using the Zoom videoconferencing software. Participants were recruited from the Columbia Experimental Laboratory for the Social Sciences (CELSS)' ORSEE website.[18] They received instructions and communicated with the experimenters via Zoom, and accessed the experiment interface on their personal computer's web browser. The experiment was programmed in oTree and, with the exception

[17] When $K = 3$, the orange and green curve do not coincide at $\widetilde{q} = 0.7$ because under MVA3 each expert has the same weight; while under LD3 the number of votes each of them commands is stochastic.

[18] Greiner (2015). CELSS' ORSEE subjects are primarily undergraduate students at Columbia University or Barnard College.

Table 3: Experiment 1: Experimental Design

Sessions	Treatments	Rounds	Subjects	Groups
1a	LD1, LD3	20,20	15	3, 1
1b	LD3, LD1	20,20	15	1, 3
2a	MVA1, MVA3	20,20	15	3, 1
2b	MVA3, MVA1	20,20	15	1, 3
3a, 3a'	LD3, MVA3	20,20	15	1, 1
3b, 3b'	MVA3, LD3	20,20	15	1, 1
4a	LD1, MVA1	20,20	15	3, 3
4b	MVA1, LD1	20,20	15	3, 3

of a more visual style for the instructions, developed very similarly to in-person experiments. Each session lasted about 90 minutes with average earnings of $26, including a show-up fee of $5.

Participants were asked to vote on the correct selection of a box containing a prize, out of two possible choices, a Green box and a Blue box. The computer selected the winning box putting equal probability on both; conditionally on the computer's random choice, participants then received a message suggesting a color, and were told the probability that the message was accurate.[19] The same screen also informed them of whether or not they were an expert (for that round). Participants were then asked to vote for one of the two boxes, or, if non-experts, to either delegate their vote to an expert (in the LD treatments), or abstain (in the MVA treatments). Across rounds, groups were re-formed randomly, if the session involved more than one group, and expert/non-expert identities were re-assigned randomly, under the constraint that groups of 5 voters had a single expert, and group of 15 had three. A copy of the instructions is reproduced in the Appendix.

We ran 10 sessions, each involving 15 subjects (150 subjects total). Participants played 20 rounds each of two treatments (40 rounds in total), according to the experimental design reproduced in the following table. Hence in total we have data for 240 rounds for LD1 and for MVA1, and 120 rounds for LD3 and for MVA3.

[19] We used examples to instruct participants on what the probability of the message being accurate meant. To limit decimal digits, the precision of the signal was drawn uniformly from a discrete distribution with bins of size 0.01. When comparing the experimental results to the theory, below, we compute equilibria using the corresponding discrete distribution of precisions. The differences are minute.

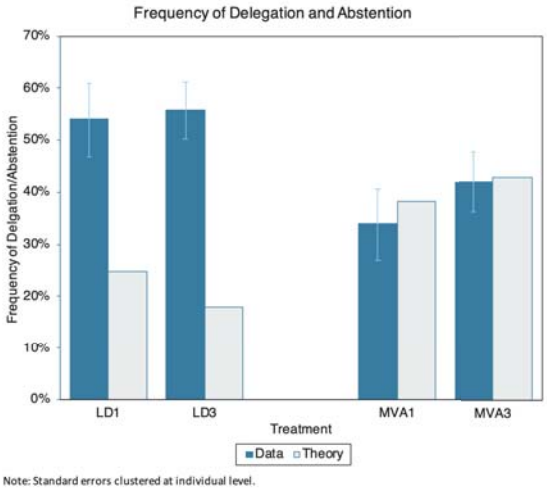


Figure 2: Aggregate frequency of delegation and abstention.

6 Experiment 1: Results

6.1 Frequency of delegation and abstention

Figure 2 reports the aggregate frequencies of delegation and abstention in the data (the dark columns), and according to the predictions of the interior equilibrium, given realized signal precisions in the experiment (the light columns). Columns on the left refer to LD treatments; columns on the right to MV. The light blue lines are 95% confidence intervals, calculated with standard errors clustered at the individual level.

The result is unambiguous: delegation rates in the experiment are between two and three times what theory predicts for the equilibrium that improves over MV. Abstention rates on the other hand are comparable. With such high propensity to delegate, the conclusion is robust to all plausible ways of cutting the data: disaggregating by session; considering only 10 final rounds; clustering standard errors at the session level. Under LD1, there is a second symmetric equilibrium with universal delegation. We do not see it in the data, but asymmetric equilibria with very high delegation also exist.[20] The specific claim here then is not that subjects were necessarily playing non-equilibrium strategies, but rather that they were not playing the symmetric interior equilibrium that dominates MV.

Regressions on individual behavior that control for signal quality and round or treatment order effects lead to the same conclusion. Table 2 report linear probability and probit regressions, with Panel A referring to LD1 and MVA1 ($N = 5$); Panel B to LD3 and MVA3

[20]Out of the 240 LD1 rounds, there are 87 rounds (36.25%) where at least 3 non-experts delegated. In such rounds, the expert is dictator and no individual delegation decision is pivotal. Thus the strategies are an equilibrium.

($N = 15$). In both panels, the excluded case is MVA, played as first treatment in MVA-only sessions. Standard errors are clustered at the individual level. The results are very similar for both group sizes. As expected, the propensity to abstain or delegate is affected negatively by higher precision of the signal, with no significant difference in impact over the two treatments. Order of treatments and number of rounds have no sustained effects, although there is some weak evidence for increasing abstention/delegation in later rounds in the larger groups (when $N = 15$), or when the treatment is played second in the smaller groups. The strongly significant and surprising result is the higher propensity to delegate, relative to abstaining. If we focus on interior equilibria, the theoretical prediction is the opposite.

Participants' choices appear coherent, if not optimal. Delegation and abstention decisions are not only negatively correlated to signals precision; they are also monotonic in signal precisions (if non-expert i votes at precision $q(i)$, then i votes at all $q'(i) > q(i)$). We report histograms of monotonicity violations for all four treatments in the Appendix. There is weak evidence of fewer violations under MVA, but the two treatments are effectively comparable. Just below 60 percent of subjects have no violations under LD; just above 60 percent under MVA, and the fraction becomes 80 percent if we limit attention to the last 10 rounds of each treatment. The results are invariant to the size of the group. Considering all rounds, in all cases, it is possible to generate perfect monotonicity for at least 80 percent of participants by changing at most 2 of their non-expert choices.[21]

We use monotonicity to estimate individual precision thresholds for delegation and abstention– the thresholds below which each participant delegates or abstains. Figure 3 reports, for each participant, the mean of the range of thresholds consistent with minimal monotonicity violations; the size of the dots is proportional to the number of participants at that threshold. The darker diamonds correspond to the average empirical thresholds, and the lighter ones to the theory. The figure confirms the over-delegation that characterizes LD, while again average values for abstention are close to the theoretical predictions. The dispersion in estimated thresholds is typical of similar experiments (for example, Levine and Palfrey, 2007; Morton and Tyran, 2011), but is in clear tension with the focus on symmetric equilibria.

The observation that thresholds tend to be lower for delegation rather than abstention is confirmed in Figure 4, where we plot the cumulative distribution functions of the estimated thresholds.

For both group sizes, the LD distribution first order stochastically dominates the MVA distribution: at any precision, including at the lower boundary of the support, the fraction of subjects estimated to delegate is above the corresponding fraction of abstainers (the fraction

[21] With 15 rounds, the maximum possible number of monotonicity violations is 7.

Frequency of Delegation or Abstention

	(1) Linear Probability	(2) Probit
LD	0.328***	0.938***
	(0.0971)	(0.333)
	[0.00107]	[0.00491]
Round	0.0309	0.0851
	(0.0442)	(0.158)
	[0.486]	[0.590]
Signal Quality	-0.777***	-2.624***
	(0.0759)	(0.283)
	[0.000]	[0.000]
Second	0.154*	0.534*
	(0.0868)	(0.307)
	[0.0801]	[0.0815]
Second * Mixed	-0.129	-0.451
	(0.0967)	(0.337)
	[0.186]	[0.182]
LD * Second	-0.0896	-0.332
	(0.118)	(0.413)
	[0.449]	[0.421]
LD * Second * Mixed	-0.0253	-0.0327
	(0.144)	(0.492)
	[0.861]	[0.947]
LD * Round	-0.112*	-0.345
	(0.0665)	(0.233)
	[0.0956]	[0.138]
LD * Signal	-0.0776	0.0653
	(0.0893)	(0.348)
	[0.387]	[0.851]
Constant	0.675***	0.582***
	(0.0759)	(0.216)
	[0.000]	[0.00705]
Observations	1,920	1,920
R-squared	0.309	

*** p<0.01, ** p<0.05, * p<0.1

Notes: Standard errors in parentheses, clustered at the individual subject level. P-values in brackets. Delegation/abstention is measured as a binary 0-1 subject decision. Experts are dropped because they cannot delegate or abstain. Only LD1 and MVA1 rounds are included. "Second" indicates that the treatment appeared second in the session. "Mixed" indicates that both an LD treatment and an MVA treatment appeared in the session.

Table 4: Determinants of delegation and abstention; N=5

Frequency of Delegation or Abstention

	(1) Linear Probability	(2) Probit
LD	0.208***	0.677**
	(0.0746)	(0.283)
	[0.00609]	[0.0167]
Round	0.0783*	0.274**
	(0.0401)	(0.136)
	[0.0532]	[0.0444]
Signal Quality	-0.861***	-2.691***
	(0.0553)	(0.220)
	[0.000]	[0.000]
Second	-0.0963	-0.341
	(0.0797)	(0.275)
	[0.229]	[0.214]
Second * Mixed	0.0784	0.295
	(0.0865)	(0.298)
	[0.367]	[0.322]
LD * Second	0.0367	0.125
	(0.113)	(0.384)
	[0.745]	[0.745]
LD * Second * Mixed	-0.166	-0.577
	(0.118)	(0.402)
	[0.160]	[0.151]
LD * Round	-0.0506	-0.174
	(0.0600)	(0.205)
	[0.401]	[0.395]
LD * Signal	0.0107	0.102
	(0.0690)	(0.287)
	[0.877]	[0.721]
Constant	0.832***	0.992***
	(0.0614)	(0.199)
	[0.000]	[0.000]
Observations	2,880	2,880
R-squared	0.309	

*** p<0.01, ** p<0.05, * p<0.1

Notes: Standard errors in parentheses, clustered at the individual subject level. P-values in brackets. Delegation/abstention is measured as a binary 0-1 subject decision. Experts are dropped because they cannot delegate or abstain. Only LD3 and MVA3 rounds are included. "Second" indicates that the treatment appeared second in the session. "Mixed" indicates that both an LD treatment and an MVA treatment appeared in the session.

Table 5: Determinants of delegation and abstention; N=15

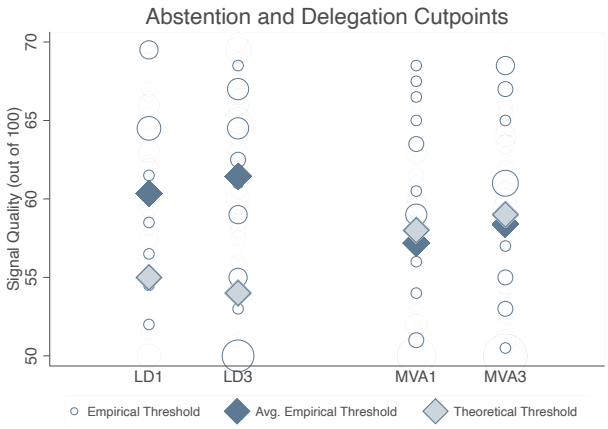


Figure 3: Delegation and Abstention Thresholds; all rounds

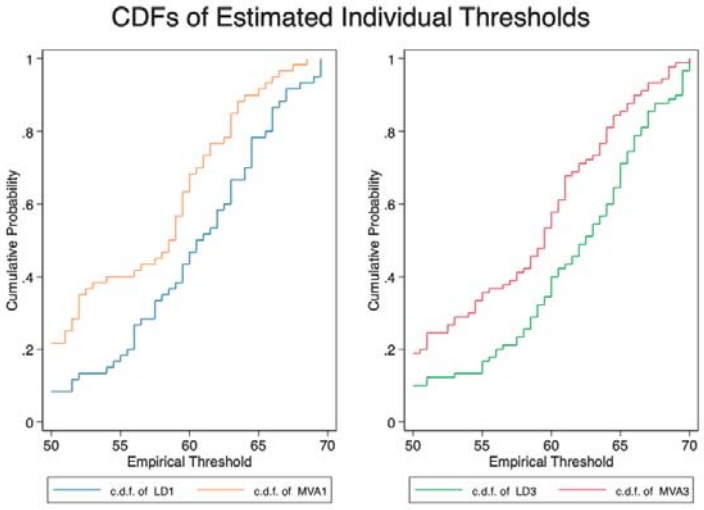


Figure 4: Individual delegation/ abstention thresholds. CDF's.

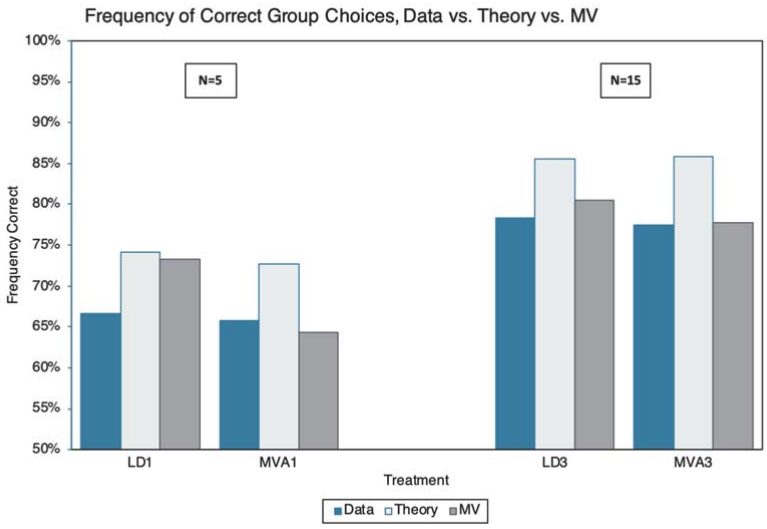


Figure 5: Frequency of correct outcomes

of subjects whose estimated threshold is below the threshold, and thus are voting, is lower). The data do not show substantive differences between the two group sizes.

6.2 Frequency of correct choices

Beyond regularities of delegation and abstention, the real variable of interest is the frequency with which the voting system leads to the correct choice. Figure 5 reports the experimental data.[22] We are comparing very similar systems, and a large majority of outcomes are identical (more than 70% of all outcomes under LD coincide with MV; and more than 80% for MVA). None of the differences in the figure are statistically significant. For now, the figure describes the data and clarifies the forces at play; later, we will test the relative performance of the different systems.[23]

We report results grouped by N. The vertical axis is the frequency of correct outcomes over the full data set for the corresponding treatment. The figure holds three main lessons. First, for both group sizes, LD and MVA yield very similar frequencies of correct decisions:

[22] All results are calculated given the experimental realizations of the state and of signals.

[23] The results under MV in Figure 5 were constructed by imputing votes for those subjects who delegated or abstained in the experiment. Because some of the votes we do observe were cast contrary to signal, we allowed for such a possibility. We calculate the outcome under MV by supposing that unobserved votes follow signals with probability equal to the average frequency observed in the treatment. The randomization is repeated 10,000 times; each time yields a share of correct decisions under MV for the relevant data set, and the figure reports the mean over the 10,000 randomizations. We have verified that calculating the probability of voting sincerely by using the subject's own frequency over the observed votes does not affect the results.

given N, the dark blue columns have very similar heights. Second, for both group sizes, both systems fall short of their possible best performance: the grey columns are consistently higher than the blue columns. Third, MVA outcomes are closely comparable to MV for both group sizes; LD outcomes fall short of MV, especially for small groups.

Two main deviations could be responsible for the systems' underperformance relative to the theory.[24] The first is the erroneous choice of delegation/abstention thresholds. Figures 2 and 3 support this interpretation for LD, with its consistent over-delegation, but not for MVA. The second is random voting in the form of voting against signal. As we show in the Appendix, the frequency of voting against signal correlates negatively and significantly with signal precision, and experts vote against signal more rarely than non-experts. Both because experts cast multiple votes under LD and because subjects choose to vote at lower precision under MVA, the share of votes cast against signal is lower in the LD treatments (at about 6%) than in the MVA ones (at about 10%), with little difference across group sizes.[25] MVA suffers from more random voting. Both systems thus fail to realize their potential gains over MV, but for different reasons.

The comparison to MV shows that the penalty is higher for LD. The better performance of MV in the LD samples reflects the random superiority of the signal draws in those samples: although signal realizations were drawn from the same probability distribution, the frequency of correct signals was higher in the LD treatments. Thus, although LD and MVA have similar shares of correct decision in our data, LD treatments could have performed better, given the superiority of the realized signals.[26]

6.3 Comparing LD and MVA to MV

Evaluating the statistical significance of the disparities observed between LD or MVA on one side, and MV on the other is complicated by several factors. Correct indices of statistical significance are not straightforward: we are interested in group results, with complex correlations across data points: the same individuals are observed over multiple rounds; the frequency with which they are assigned the role of experts is random and variable; the voting group is assembled randomly every round when $N = 5$, but is fixed when $N = 15$; finally, the imputation of missing votes under MV creates randomness in the MV outcomes. But the

[24] A priori, a third possibility would be non-monotonicity in delegation and abstention decisions. But as we described earlier, violations of monotonicity are rare.

[25] The numbers are comparable to those reported, for example, by Guarnaschelli et al. (2000) and Goeree and Yariv (2011) for juries voting under simple majority and pure common interest, in the absence of communication (6-9%).

[26] If votes were not cast against signals, all three voting systems would be more efficient, but we have verified that, as expected, the difference between LD and MV, and the lack of a difference between MVA and MV, would not be affected.

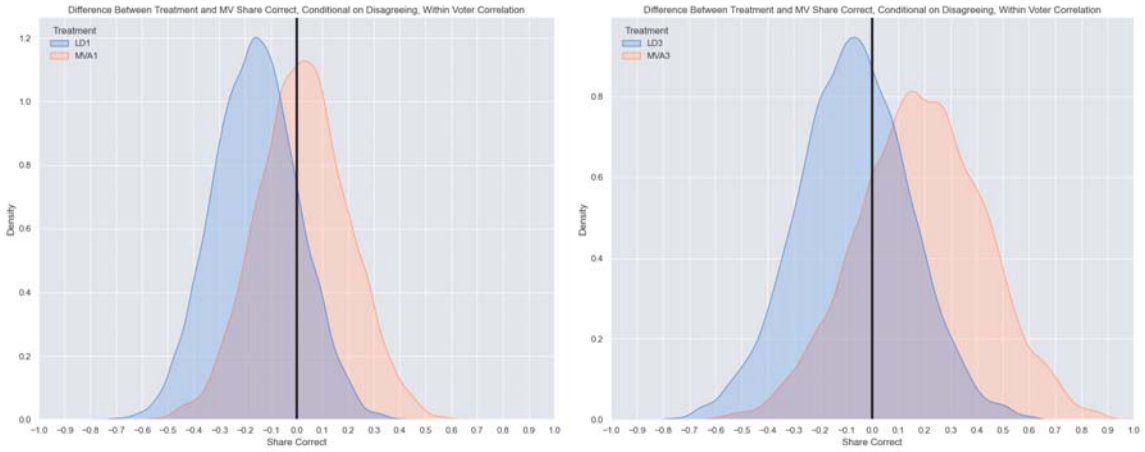


Figure 6: Differential frequency of correct decisions, relative to MV, conditional on different outcomes. Distributions over 100,000 bootstrap replications.

fundamental difficulty is simpler: as mentioned above, outcomes coincide in a large majority of cases. Bootstrapping methods allow us to generate a large sample of elections in which outcomes under delegation, or abstention, differ from outcomes under MV.

The procedure we implement allows for correlation across an individual's multiple decisions, and uses randomization to generate the correct balance of experts and non-experts. We assume that our experimental subjects are representative of an underlying population; for each voting system and group size, we generate outcomes by drawing subjects, with replacement, each with their full set of 20-round decisions, and matching them randomly into groups, and finally we study the outcomes corresponding to 100,000 replications of the experiment for each treatment, using the population of subjects for that voting system and group size. We describe the procedure in more detail in the Appendix. Figure 6 shows the distributions of the differential frequency of correct decisions between the voting systems we are studying and MV, for each group size, conditioning on the decisions being different. Consider for example LD1. For each of the 100,000 simulations, we focus on the subset of decisions $\widetilde{D}_{LD1}$such that LD1 and MV reach a different outcome. Call $\gamma_{LD1}(\widetilde{D}_{LD1})$ $(\gamma_{MV}(\widetilde{D}_{LD1}))$ the frequency with which LD1 (MV) is correct over subset$\widetilde{D}_{LD1}$. We are interested in $\gamma_{LD1}(\widetilde{D}_{LD1}) - \gamma_{MV}(\widetilde{D}_{LD1})$. The first panel of Figure 6 plots, in blue, the distribution of such variable over the 100,000 replications. The equivalent distribution for MVA is plotted in the same panel in red. The second panel reports the results for groups of size 15.[27] **[Check the FN]**

For both group sizes, the blue distribution is shifted to the left, relative to the zero point

[27]Averaging over all replications, the share of elections in which the outcome differs from MV is 23.5% for LD1, 15.4% for MVA1, 20.4% for LD3, and 14.9% for MVA3

indicated by the vertical black line: when LD and MV differ, the correct decision is more likely to be the one reached by MV. The asymmetry is more pronounced for $N = 5$, when the probability that MV is superior to LD1, conditional on disagreement, is 85%, versus 67% for LD3. MVA on the other hand, when disagreeing with MV, is more likely to be right than wrong: only barely when $N = 5$, and the probability that MV is superior to MVA1 is just below 50% (48%), but more substantially when $N = 15$ and the probability that MV is correct, conditional on disagreement with MVA3, falls to 26%. The distributions as a whole are informative of the quantitative gap in the probability of being correct, relative to MV. In the upper panel, for example, the mode of the blue distribution at -16% tells us that over the 100,000 replications, conditional on disagreement, the highest probability mass is around a frequency of correct decisions of around 42% for LD1, versus 58% for MV. [28]

In our first experiment then, LD falls short of the hopes if its supporters, even in a streamlined environment where experts are correctly identified. The simpler option of abstention, rather than delegation, captures some of the motivation for LD and allows voters with particularly weak information not to influence the final choice. In our data, its efficiency is either comparable or somewhat superior to universal majority voting. But do these results reflect some core feature of the systems we are studying, or are they artifacts of the lab? We analyze this question in our second experiment.

7 Experiment 2: The Random Dot Kinematogram

As we discussed in the Introduction, the goal of our second experiment is to evaluate whether the deviations from optimal behavior we see in the lab may stem from the over-mathematization of the environment. With this in mind, we chose a perceptual task–the Random Dot Kinematogram (RDK)–where individual signals correspond to the accuracy of individual perceptions, and neither own or others' accuracy is described or known in precise probabilistic terms. Because the task may be unfamiliar, we describe it in some detail. More information as well as the instructions are reproduced in the Appendix.

We ran the experiment on Amazon Mechanical Turk (with prescreening of subjects by CloudResearch) with three electorate sizes: $N = 5$, and $N = 15$, as in Experiment 1, but also $N = 125$, i.e. with a larger size than we could run in the lab or conveniently on Zoom. In our implementation, 300 moving dots appear in each subject's screen for 1 second; a small fraction of them (dependent on treatment) moves in a coherent direction, either Left

[28] We have verified that conclusions remain qualitatively similar if we ignore the possibility of correlation in individual behavior across rounds, and thus construct the bootstrapped data set by drawing individual decisions (as opposed to the full set of a subject's choices over 20 rounds).

or Right, with equal ex ante probability. After 1 second, the image disappears and each participant reports whether the perceived coherent direction was Left or Right. We divide the experiment into two parts; each preceded by a few practice tasks, but with the first part effectively playing the role of extended training. Both parts are divided into six blocs, with each bloc consisting of 20 tasks of equal coherence. We report in the Appendix the precise parameters we used for the task (the size and color of the dots, the movements per frame, the random process for the dots moving randomly, etc.), but it should be clear that our experiment does not aim at measuring perception per se–for example, we cannot control the ambient light, screen size, or contrast of the monitors our subjects use. Rather, our focus remains on collective decision-making.[29]

In Part 1, subjects are rewarded on the basis of their individual accuracy. Coherence ranges from 20% in bloc 1 (one fifth of the 300 dots move synchronously) to 10% in bloc 2, 8% in bloc 3, 6% in bloc 4, and finally the same coherence used in Part 2, and dependent on N, for the final two blocs. At the end of Part 1, each subject is informed of her fraction of correct answers in each bloc. In part 2, each task has both an individual component ("Choose the coherent direction"), and a group decision with the possibility of delegation (under LD), or abstention (under MVA). ("You said Left. Do you want to Vote or to Delegate/Abstain?"). When delegation is chosen, the vote is assigned randomly to an "expert", that is, one participants whose accuracy is in the top 20% of the group over the last 2 blocs (40 tasks); experts are not allowed to delegate (under LD) or to abstain (under MVA). Thus in line with Experiment 1, groups of 5 have 1 expert, and groups of 15 have 3; the group of 125 has 25 experts, and in line with our standard notation we denote the two treatments on the larger group by LD25 and MVA25. The group decision corresponds to the majority of votes cast, and individuals are rewarded both for their individual accuracy and for the accuracy of the group. As in Part 1, feedback about average individual accuracy in each bloc is provided at the end of Part 2.[30] In Part 2, coherence is kept constant across all blocs. We chose its value according to two main criteria: the task should not be so difficult that subjects are discouraged and act randomly, and should not be so easy that MV accuracy, especially in the large group, leaves effectively no room for possible improvement. Based on the results of two preliminary pilots, we fixed coherence in Part 2 at 5% for electorates of sizes 5 and 15, and at 3% for the electorate of size 125. The task is objectively hard, as the reader can verify at the following link: `https://blogs.cuit.columbia.edu/ac186/files/2022/05/rdk-video.gif`

The experiment used the RDK plugin n jsPsych (Rajananda et al., 2018) and was hosted

[29]Heer and Bostock (2010) and Woods et al. (2015) report on the replication successes and challenges of conducting research on perceptual stimuli online.

[30]Feedback over group accuracy cannot be provided because it depends on choices made by others and is calculated ex post.

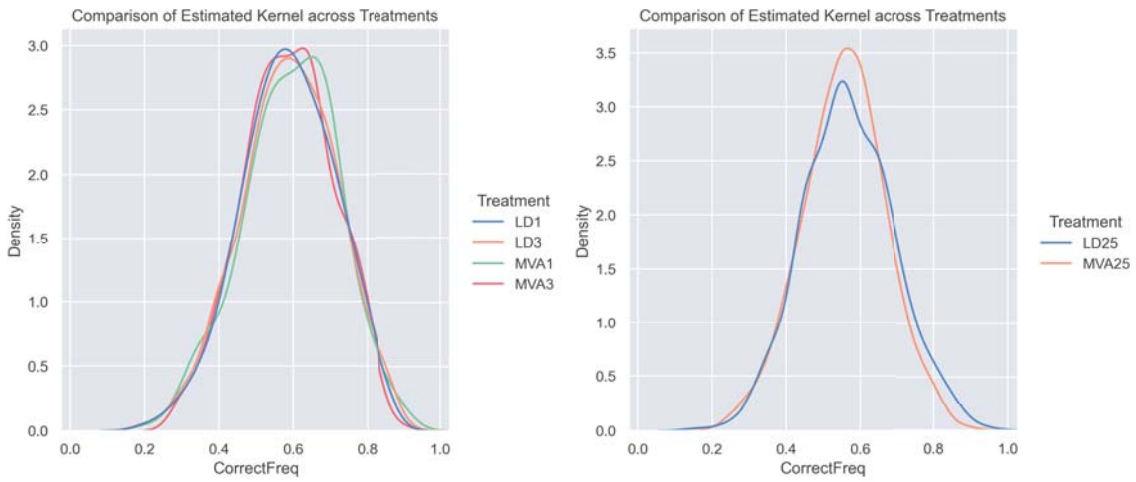


Figure 7: Accuracies per bloc per subject. Distributions. The panel on the left corresponds to coherence of 0.05; the panel on the right to coherence of 0.03.

on cognition.run. For each of LD and MVA, we recruited 60 subjects divided into 12 groups for the $N = 5$ treatment and 90 subjects divided into 6 groups for $N = 15$, (thus replicating the corresponding number of subjects and groups in Experiment 1), and an additional 125 subjects for the largest group. There were then 275 subjects for each voting system, or 550 in total. The group size and the relevant number of experts were always made public. The experiment lasted about 20 minutes. Subjects earned \$1 for participation and a bonus proportional to the number of correct responses, for a total average compensation of \$4.92, or just below \$15 an hour.

8 Experiment 2: Results

8.1 Accuracy

One motivation for considering a perceptual experiment is the possibility of observing signals, here accuracies, of less than random precision. Figure 8 reports the distributions of correct responses in Part 2, aggregating by level of coherence (0.05 on the left; 0.03 on the right). Accuracy is the fraction of correct responses calculated over each of the 6 blocs for each subject, that is, over 20 tasks.

For given coherence, the distributions are very similar. In all cases, the spread in the distribution of accuracies is large, ranging from 25% all the way to 95%. The frequency of blocs with accuracy below 50% is non-negligible and, surprisingly, persists when we aggregate over a larger number of tasks. Aggregating at the subject level, over all 120 tasks, 9% of subjects have accuracy below randomness with coherence 0.05, and 12% with coherence 0.03;

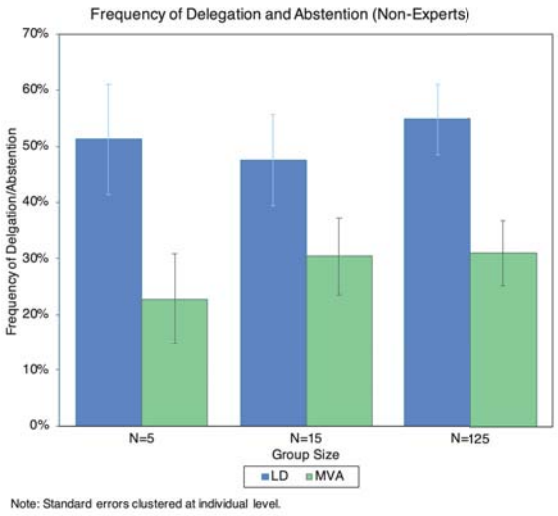


Figure 8: Frequency of delegation and abstention (non-experts only)

11% and 17% respectively do no better than randomness.[31] If we want to study voting and information aggregation when information may be faulty, perceptual tasks can provide a very useful tool.

8.2 Frequency of delegation and abstention

Absent information about the distributions of subjects' beliefs, we do not have theoretical reference point to evaluate the extent of delegation and abstention we see in the data. We can however compare the two, under the assumption, which seems plausible and is supported by Figure 7, that accuracies and beliefs about accuracies are comparable across the LD and MVA samples. Figure 8 plots the frequencies of delegation and abstention for each group size, calculated over non-experts only for possible comparison to Experiment 1.[32] The 95% confidence intervals are calculated form s.e's clustered at the individual level.

We find that in Experiment 2 delegation is much more common than abstention, for all group sizes. In groups of 5, where the disparity is largest, delegation is more than twice as frequent; in groups of 15, where we see the least disparity, delegation is still 60% more common. Interestingly, the decline in coherence, from $N = 15$ to $N = 125$, is not obviously noticeable in the data, with some small increase in delegation but no effect on abstention. Unexpectedly, considering the rather radical change in experimental design, Figure 8 looks quite similar to Figure 2, for the group sizes for which we have data from both experiments.

The higher frequency of delegation is confirmed in the regressions reported in Table 6.

[31] Note however that individual subjects' accuracies show high variability across blocs, evidence of random noise in perceiving and recording the stimulus, as formalized in psychophysics research.

[32] The figure is effectively identical if the frequencies are calculated over the full sample.

The unit of analysis is the bloc at the individual subject level (hence 6 blocs per subject), with data grouped by coherence level.[33] The regressions reported below confirm the results of the figure: in all treatments, delegation is significantly more frequent than abstention. In Experiment 2, accuracy is at best a very weak predictor of participation in voting, never significant at conventional levels and with the wrong sign in the probit estimation, confirming the high uncertainty in subjects' evaluation of their own accuracy. The probit regressions detect a decline in abstention as blocs proceed, which would be consistent with increased familiarity with the task.

8.3 Frequency of correct outcomes

How well did the three voting systems do in Experiment 2? Figure 9 reports the frequency of correct group decisions, aggregated over all groups and tasks for given treatment. As in the case of Figure 5, a large fraction of decisions are identical, and none of the differences are statistically significant.

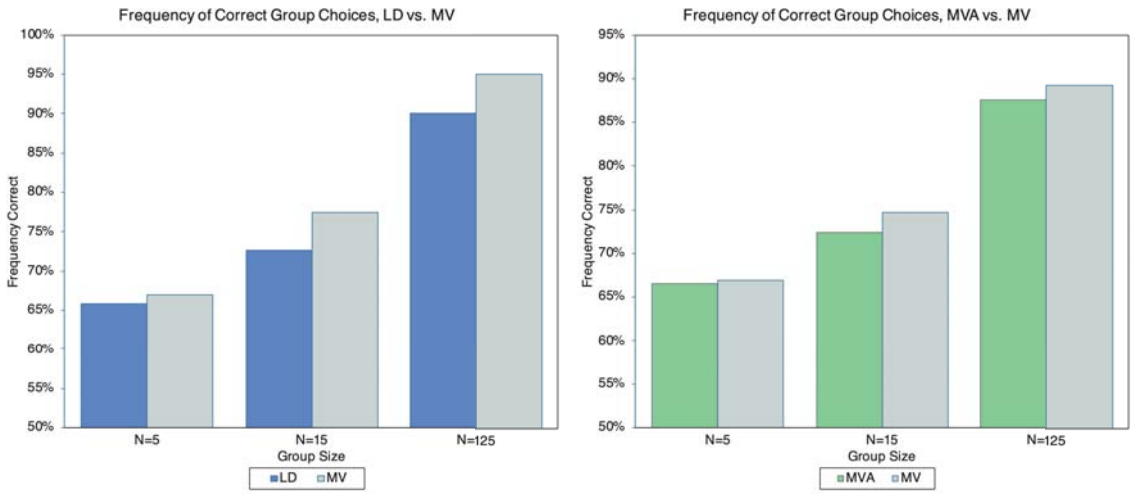


Figure 9: Frequency of correct outcomes

As expected, for all three systems, the fraction of correct decisions increases with the size of the group, ranging from about 65% at $N = 5$ to 90-95% at $N = 125$. In a pure common interest environment, the aggregation of independent signals remains very powerful, even in the presence of weak accuracies. In line with this logic, and with Condorcet, universal voting consistently achieves the best results. Here too, for the group sizes that are common to both experiments, the numbers are surprisingly similar to the numbers in Figure 5. Calibrating

[33]The results are unchanged if the data are separated by group size

Frequency of Delegation/Abstention, N=5 and N=15

	(1) Linear Probability	(2) Linear Probability	(3) Probit	(4) Probit
Accuracy	-0.122	-0.122	0.356	0.354
	(0.0839)	(0.0839)	(0.309)	(0.308)
	[0.146]	[0.146]	[0.249]	[0.251]
LD	0.226***	0.226***	0.547***	0.548***
	(0.0378)	(0.0378)	(0.139)	(0.140)
	[0.000]	[0.000]	[0.000]	[0.000]
N=15	0.00451	0.00446	0.0506	0.0503
	(0.0384)	(0.0384)	(0.139)	(0.140)
	[0.907]	[0.908]	[0.717]	[0.719]
Keys: [E][Y]		-0.00884		-0.0401
		(0.0379)		(0.139)
		[0.816]		[0.772]
Bloc		0.000511		-0.256***
		(0.0115)		(0.0654)
		[0.965]		[0.000]
Constant	0.339***	0.343***	0.163	0.313
	(0.0627)	(0.0656)	(0.222)	(0.234)
	[0.000]	[0.000]	[0.463]	[0.181]
Observations	1,800	1,800	1,800	1,800
R-squared	0.100	0.100		

*** p<0.01, ** p<0.05, * p<0.1

Notes: Standard errors in parentheses, clustered at the individual subject level. P-values in brackets. Delegation/abstention is measured as the share of rounds in a given bloc that a subject chose to delegate/abstain (with a range from 0 to 1). Accuracy is the share of rounds in the bloc that subject answered correctly. Subjects randomly use either keys [E] and [Y] or [V] and [N] to decide whether to delegate; a dummy for being assigned [E][Y] is included. The values for bloc have been scaled to be between 0 and 1; the coeffient for "bloc" thus indicates the effect of going from the first to last bloc.

Frequency of Delegation/Abstention, N=125

	(1) Linear Probability	(2) Linear Probability	(3) Probit	(4) Probit
Accuracy	0.000304	0.00380	0.479	0.455
	(0.101)	(0.102)	(0.357)	(0.359)
	[0.998]	[0.970]	[0.180]	[0.205]
LD	0.224***	0.223***	0.661***	0.663***
	(0.0415)	(0.0415)	(0.164)	(0.164)
	[0.000]	[0.000]	[0.000]	[0.000]
Keys: [E][Y]		0.0291		-0.0430
		(0.0414)		(0.161)
		[0.484]		[0.790]
Bloc		0.00192		-0.0363**
		(0.00258)		(0.0147)
		[0.459]		[0.0138]
Constant	0.316***	0.282***	0.313	0.694**
	(0.0623)	(0.0721)	(0.225)	(0.288)
	[0.000]	[0.000119]	[0.163]	[0.0162]
Observations	1,500	1,500	1,500	1,500
R-squared	0.095	0.097		

*** p<0.01, ** p<0.05, * p<0.1

Notes: Standard errors in parentheses, clustered at the individual subject level. P-values in brackets. Delegation/abstention is measured as the share of rounds in a given bloc that a subject chose to delegate/abstain (with a range from 0 to 1). Accuracy is the share of rounds in the bloc that subject answered correctly. Subjects randomly use either keys [E] and [Y] or [V] and [N] to decide whether to delegate; a dummy for being assigned [E][Y] is included. The values for bloc have been scaled to be between 0 and 1; the coeffient for "bloc" thus indicates the effect of going from the first to last bloc.

Table 6: Frequency of delegation/abstention.

performance relative to MV, hence controlling for the specific visual stimuli, LD is once again dominated by MVA, although barely so in the smallest groups (Figure 10).

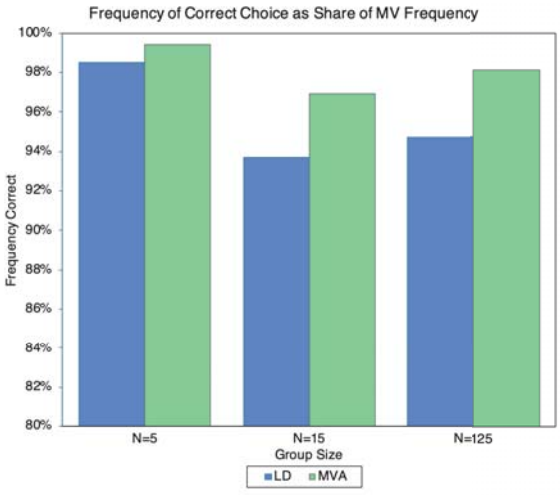


Figure 10: Calibrating LD and MVA performance relative to MV

The different performance arises from the different frequencies of delegation and abstention shown in Figure 8: with abstention being less common, the fraction of decisions in which MVA and MV disagree is about the half fraction under LD (between 8 and 12% across the different treatments under MVA, v/s 12 to 24% under LD). And if MVA and MV agree more, their performance overall is more similar. To achieve better estimates of the differences across the three voting systems, we need to bootstrap the data and replicate the group decisions a large number of times, generating a large sample of decisions over which the voting results would differ, as we did with the data from Experiment 1.

[**Results from the bootstrap**]

References

[1] Armstrong, Ben and Kate Larson, 2021, "On the Limited Applicability of Liquid Democracy". *Appears at the 3rd Games, Agents, and Incentives Workshop (GAIW2021). Held as part of the Workshops at the 20th International Conference on Autonomous Agents and Multiagent Systems*, London, UK, May 2021, IFAAMAS, 7 pages.

[2] Austen-Smith, David, and Jeffrey Banks, 1996, "Information aggregation, rationality, and the Condorcet Jury theorem", *American Political Science Review*, 90:34-45.

[3] Battaglini, Marco, Rebecca Morton, and Thomas Palfrey, 2010, "The swing voter's curse in the laboratory", *Review of Economic Studies*, 77:61-89.

[4] Bahrami, B., K. Olsen, D. Bang, A Roepstorff, G Rees and C. Frith, 2012, "What failure in collective decision-making tells us about metacognition", *Philosophical Transactions of the Royal Society* 367: 1350-1365.

[5] Blum, Christian and Christina Zuber, 2016, "Liquid Democracy: Potentials, Problems, and Perspectives", *Journal of Political Philosophy* 24: 162–182.

[6] Britten K.H., Michael Shadlen, William Newsome, and J. Anthony Movshon, 1992, "The analysis of visual motion: a comparison of neuronal and psychophysical performance", *Journal of Neuroscience* 12: 4745–4765.

[7] Caragiannis, Ioannis and Evi Michas, 2019, "A Contribution to the Critique of Liquid Democracy", *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI-19),* 116-122.

[8] van de Grind, W.A., A. J. van Doorn, and J. J. Koenderink, 1983, "Detection of coherent movement in peripherally viewed random-dot patterns," *Journal of the Optical Society of America* 73: 1674-1683.

[9] Dhillon, Amrita, Grammateia Kotsialou, and Dimitrios Xefteris, 2021, "Information Aggregation with Delegation of Votes", Working Paper 2021-22, Quantitative Political Economy Group, King's College, London.

[10] Dodgson, Charles, 1884, *The Principles of Parliamentary Representation*, London: Harrison and Sons.

[11] Feddersen, Timothy and Wolfgang Pesendorfer, 1996, "The swing voter's curse", *American Economic Review* 86: 408–424.

[12] Feddersen, Timothy and Wolfgang Pesendorfer, 1997, "Voting Behavior and Information Aggregation in Elections With Private Information", *Econometrica* 65: 1029-1058.

[13] Goeree Jacob and Leeat Yariv, 2011, "An Experimental Study of Collective Deliberation", *Econometrica*, 79: 893–921.

[14] Gölz, Paul, Anson Kahng, Simon Mackenzie, and Ariel D. Procaccia, 2018, "The Fluid Mechanics of Liquid Democracy", *arXiv:1808.01906v1* [cs.GT].

[15] Green-Armytage, James, 2015, "Direct voting and proxy voting", *Constitutional Political Economy* 26:190–220.

[16] Grofman, Bernard, Guillermo Owen and Scott Feld. 1983. “Thirteen theorems in search of the truth”, *Theory and Decision*, 15:261-278.

[17] Grofman, Bernard, Guillermo Owen and Scott Feld. 1982. “Evaluating the competence of experts, pooling individual judgments into a collective choice, and delegating decision responsibility to subgroups”. In Felix Geyer and Hans van der Zouwen (Eds.), *Dependence and Inequality*. NY: Pergamon Press, 221-238.

[18] Guarnaschelli, Serena, Richard McKelvey and Thomas Palfrey, 2000, “An Experimental Study of Jury Decision Rules”, *The American Political Science Review*, 94: 407-423.

[19] Hardt, Steve and Lia Lopes, 2015, "Google Votes: A Liquid Democracy Experiment on a Corporate Social Network", *Technical Disclosure Commons*, http://www.tdcommons.org/dpubs_series/79

[20] Heer, Jeffrey and Michael Bostock, 2010, “Crowdsourcing Graphical Perception: Using Mechanical Turk to Assess Visualization Design”, *CHI 2010*: April 10–15, 2010, Atlanta, Georgia, USA: 203-212.

[21] Kahng, Anson, Simon MacKenzie and Ariel Procaccia, 2018, “Liquid Democracy: An Algorithmic Perspective”, *Proceedings of the 32nd AAAI Conference in Artificial Intelligence (AAAI)*, 1095-1102.

[22] Margolis, Howard, 1976, “A Note on Incompetence”, *Public Choice* 26: 119-127.

[23] McMurray, Joseph, 2013, “Aggregating Information by Voting: The Wisdom of the Experts versus the Wisdom of the Masses*”, Review of Economic Studies* 80: 277-312.

[24] McLennan, Andrew, 1998, “Consequences of the Condorcet Jury Theorem for Beneficial Information Aggregation by Rational Agents”, *American Political Science Review* 92: 413-418.

[25] Morton, Rebecca and Jean-Robert Tyran, 2011, “Let the experts decide? Asymmetric information, abstention, and coordination in standing committees”, *Games and Economic Behavior*, 72: 485–509.

[26] Nitzan, Shmuel and Jacob Paroush, 1982, “Optimal decision rules in uncertain dichotomous choice situations”, *International Economic Review* 23: 289-297.

[27] Newsome William, K.H. Britten, and J. Anthony Movshon, 1989, “Neuronal correlates of a perceptual decision”, *Nature* 341: 52–54.

[28] Pilly, Praveen and Aaron R. Seitz, 2009, "What a difference a parameter makes: A psychophysical comparison of random dot motion algorithms", *Vision Research*, 49: 1599-1612.

[29] Rajananda, S., Lau, H. & Odegaard, B., 2018, "A Random-Dot Kinematogram for Web-Based Vision Research", *Journal of Open Research Software*, 6(1), p.6. doi:10.5334/jors.194.

[30] Ravindran, Dilip, 2021, "Liquid Democracy", in *Essays in information and behavioral economics*, PhD Dissertation, Columbia University, New York, NY.

[31] Roitman, Jamie and Michael Shadlen, 2002, "Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task", *Journal of Neuroscience* 22: 9475–9489.

[32] Schütz, Alexander, Doris Braun, J. Anthony Movshon, and Karl Gegenfurtner, 2010, "Does the noise matter? Effects of different kinematogram types on smooth pursuit eye movements and perception", *Journal of Vision*, 10(13):26. doi: https://doi.org/10.1167/10.13.26.

[33] Shapley, Lloyd and Bernard Grofman, 1984, "Optimizing group judgmental accuracy in the presence of interdependencies", *Public Choice* 43: 329-343.

[34] Shubik, Martin, 1970, "On Homo Politicus and the Instant Referendum", *Public Choice* , 9: 79-84.

[35] Silver, Ike, Barbara Mellers and Philip Tetlock, 2021, "Wise teamwork: collective confidence calibration predicts the effectiveness of group discussion", *Journal of Experimental Social Psychology*, 96: Case Report 104157.

[36] Sorkin, Robert, Ryan West and Donald Robinson, 1998, "Group Performance Depends on the Majority Rule", *Psychological Science* 9: 456-463.

[37] Sorkin, Robert, Christopher Hays, and Ryan West, 2001, "Signal-Detection Analysis of Group Decision Making", *Psychological Review* 108: 183-203.

[38] Wit, Jőrgen, 1998, "Rational Choice and the Condorcet Jury Theorem", *Games and Economic Behavior*, 22: 364-376.

[39] Woods Andy et al., 2015, "Conducting perception research over the internet: a tutorial review", *PeerJ 3:e1058*; DOI 10.7717/peerj.1058

A Appendix [Preliminary]

A.1 Theoretical results

Theorem. *Suppose $\pi = Pr(\omega_1) = 1/2$. Then for any M and K, there exists an equilibrium with delegation that strictly improves over MV.*

Proof. The proof proceeds by steps, but the important steps are well-known. First, as we noted, with $\pi = 1/2$, and $Pr(\sigma_i = 1|\omega_1) = Pr(\sigma_i = 2|\omega_2)$ for all i, voting according to signal is an undominated strategy, and we can focus on sincere equilibria. Second, because this is a voting game of pure common interest, we can apply McLennan (1998) first result: a set of strategies that maximizes each voter's expected utility is an equilibrium. Hence an equilibrium exists. Third, the environment is fully symmetric for voters of each given type. All voters have identical preferences and are endowed with a single vote. All experts have equal precision and equal probability of receiving any delegated vote, and thus, for any delegation strategy by non-experts, each expert's vote has equal expected weight on the final decision. Non-experts will generally have heterogeneous realized precisions, and the equilibrium strategy will depend on individual precision, but each precision $q(i)$ is an independent draw from the same distribution F. Hence any permutation of realized precisions to different non-expert voters is assigned equal probability, and each voter holds equal beliefs about the others' precisions. For each type of voters, then, these symmetry conditions satisfy the requirement of McLennan's Theorem 2: if a semi-symmetric strategy profile maximizes each voters' expected utility in the set of semi-symmetric profiles, then it is an equilibrium. Thus we can restrict attention to semi-symmetric profiles.

The only new step is the fourth and last: can there be an equilibrium where no-one delegates? If the answer is negative, then the equilibrium that is maximal for expected utility in the profile of semi-symmetric strategies, and that we know exists, must involve a strictly positive probability of delegation for some realized precisions. But note that if no-one delegates, LD is exactly equivalent to sincere MV. Thus if it is not an equilibrium, then there must exist an equilibrium of the LD voting game that both involves a strictly positive probability of delegation and strictly dominates MV.

Consider the perspective of non-expert voter i, with q_i in the neighborhood of $\underline{q}$. Conjecture an equilibrium with no delegation. We show in what follows that i's best response is to delegate his vote. Note first that if no-one delegates, all non-i voters cast a single vote and have equal weight on the group decision. Hence if i delegates, he delegates the vote to the expert, with precision $p \geq q_j$ for all j.

Calculating i's expected utilities from non delegating or delegating is somewhat cumber-

some but conceptually straightforward. We find:

$$lim_{q_i \to \underline{q}} EUND(q_i) = \sum_{c_n=0}^{M-1} \binom{M-1}{c_n} \mu^{c_n} (1-\mu)^{M-1-c_n} \times \\ \times \left[\sum_{c_e=0}^{K} \binom{K}{c_e} p^{c_e} (1-p)^{K-c_e} \left(\underline{q} I_{c_n+c_e+1>\frac{(M+K)}{2}} + (1-\underline{q}) I_{c_n+c_e>\frac{(M+K)}{2}} \right) \right]$$

$$lim_{q_i \to \underline{q}} EUD = \sum_{c_n=0}^{M-1} \binom{M-1}{c_n} \mu^{c_n} (1-\mu)^{M-1-c_n} \times \\ \times \left[\sum_{c_e=0}^{K} \binom{K}{c_e} p^{c_e} (1-p)^{K-c_e} \left(\left(\frac{c_e}{K} \right) I_{c_n+c_e+1>\frac{(M+K)}{2}} + \left(\frac{K-c_e}{K} \right) I_{c_n+c_e>\frac{(M+K)}{2}} \right) \right]$$

where $c_n(c_e)$ indexes the number of non-experts (experts) whose signals are truthful, and I_C is an indicator function that takes value 1 if condition C is satisfied and 0 otherwise. As described in the text, μ is the expected precision of non-experts who choose to vote, and thus in this conjectured scenario, $\mu = \int_{\underline{q}}^{\overline{q}} q dF(q)$.

Voter i strictly prefers delegation if it yields higher expected utility, or:

$$lim_{q_i \to \underline{q}} \left(EUND(q_i) - EUD \right) < 0$$

Denoting $r \equiv (M+K+1)/2 - c_n$, (the number of additional correct votes required to reach a majority, given the votes of the non-experts, excluding i), after some simplifications, we can write:

$$lim_{q_i \to \underline{q}} \left(EUND(q_i) - EUD \right) = \\ = \sum_{r=1}^{K+1} \binom{M-1}{\frac{M+K+1}{2} - r} \mu^{\frac{M+K+1}{2} - r} (1-\mu)^{\frac{M-(K+3)}{2} + r} \binom{K}{r-1} p^{r-1} (1-p)^{K-(r-1)} \left(\underline{q} - \frac{r-1}{K} \right) \tag{2}$$

Signing this expression is not immediate because the sign depends on the last term. However, the problem is greatly simplified by noticing that:

$$\binom{M-1}{\frac{M+K+1}{2} - r} = \binom{M-1}{\frac{M+K+1}{2} - (K+2-r)}.$$

Equation (2) can then be written as:

$$lim_{q_i \to \underline{q}} \left(EUND(q_i) - EUD\right) = \sum_{x=1}^{(K+1)/2} \binom{M-1}{\frac{M+K+1}{2} - x} (\mu(1-\mu))^{\frac{M+K+1}{2} - (K+2-x)} \binom{K}{x} p^x (1-p)^x \times$$
$$\times \left\{ (\mu(1-p))^{K+2-2x} \left[\underline{q} - \frac{x-1}{K} \right] + ((1-\mu)p)^{K+2-2x} \left[\frac{x-1}{K} - (1-\underline{q}) \right] \right\}$$

or, with $\underline{q} = 1 - \underline{q} = 1/2$:

$$lim_{q_i \to \underline{q}} \left(EUND(q_i) - EUD\right) = \sum_{x=1}^{(K+1)/2} \binom{M-1}{\frac{M+K+1}{2} - x} (\mu(1-\mu))^{\frac{M+K+1}{2} - (K+2-x)} \binom{K}{x} p^x (1-p)^x \times$$
$$\times \left[\underline{q} - \frac{x-1}{K} \right] \left\{ (\mu(1-p))^{K+2-2x} - ((1-\mu)p)^{K+2-2x} \right\}$$

With $\underline{q} = 1/2$, $\left[\underline{q} - (x-1)/K\right] > 0$ for all $x < (K+2)/2$, and thus for all relevant x values. It follows that:

$$(\mu(1-p))^{K+2-2x} < ((1-\mu)p)^{K+2-2x} \Rightarrow lim_{q_i \to \underline{q}} \left(EUND(q_i) - EUD\right) < 0$$

With all $x < (K+2)/2$, the exponent on both sides is positive, and we can compare the roots:

$$\mu(1-p) < (1-\mu)p \Rightarrow lim_{q_i \to \underline{q}} \left(EUND(q_i) - EUD\right) < 0,$$

a condition that reduces to:

$$\mu < p$$

and is always satisfied. Hence $lim_{q_i \to \underline{q}} \left(EUND(q_i) - EUD\right) < 0$: delegation is preferred, and $h_e(\underline{q}) = 1$. A profile of strategies such that all non-voters cast their vote with probability 1 cannot be an equilibrium. □

Proposition. *Suppose $\pi = Pr(\omega_1) = 1/2$ and $K = 1$. Then for any M, there exists an equilibrium such that: (i) the expert never delegates her vote and always votes according to signal; (ii) there exists a unique $\widetilde{q}(M) \in (\underline{q}, \overline{q})$ such that non-expert i delegates his vote to the expert if $q_i < \widetilde{q}$ and votes according to signal otherwise. Such an equilibrium strictly improves over MV and is maximal among sincere semi-symmetric equilibria where the expert never delegates and non experts delegate to the expert only.*

We prove the proposition in three steps. Following the argument in the text, we focus on equilibria with sincere voting. We begin by taking as given the direction of delegation: the expert never delegates, and non-experts only delegate to the expert. The first lemma

shows that, if the assumed pattern of delegation holds, then there must exist a $\widetilde{q} \in (\underline{q}, \overline{q})$ such that non-expert i delegates her vote to the expert if $q_i < \widetilde{q}$ and votes according to signal otherwise. Because $\widetilde{q}$ is strictly interior to the interval $[\underline{q}, \overline{q}]$, the conjectured equilibrium must have strictly interior probability of delegation: $h_{i,e} \in (0, 1)$. Lemma 2 shows that given Lemma 1, the assumed pattern of delegation is indeed a best response when others follow it. Hence the profile of strategies described in Proposition 1 is indeed an equilibrium. Finally, Lemma 3 proves that such interior equilibrium maximizes ex ante expected utility among sincere, semi-symmetric equilibria with the specified direction of delegation. Hence it strictly improves over MV.

Lemma 1. *Suppose $\pi = Pr(\omega_1) = 1/2$ and $K = 1$. Suppose in addition that the expert never delegates her votes, and non-experts only delegate to the expert. Then there exists $\widetilde{q} \in (\underline{q}, \overline{q})$ such that non-expert i delegates her vote to the expert if $q_i < \widetilde{q}$ and votes according to signal otherwise.*

Proof. Under the conjectured pattern of delegation, the expert never delegates, and cycles of delegation are not possible. The only question is the delegation decision by non-experts. Keeping in mind that all q_i's are independent draws from $F(q)$, non-expert i's expected utility conditional on delegation, EUD_i, does not depend on q_i (because i would not be voting). Hence $EUD_i = EUD$ for all i. On the other hand, i's expected utility when not delegating, $EUND_i$, must be weakly increasing in q_i (because, given sincere voting, the probability of reaching the correct outcome must be weakly increasing in each voter's precision), and strictly increasing if i's probability of being pivotal is positive (because, given sincere voting, the probability of reaching the correct outcome must be strictly increasing in a pivotal voter's precision). Suppose first $\widetilde{q} = \overline{q}$: all non-experts delegate to the expert. In such a case, the expert is dictator, and no individual non-expert can make herself pivotal by deviating: $EUND_i = EUND = EUD$, and delegating is indeed a (weak) best response for any q_i. Hence $\widetilde{q} = \overline{q}$ can be supported as a system of mutual best responses by all non-experts. However, we know from the Condorcet Jury Theorem that in the environment we are studying, MV dominates dictatorship by any voter. When the expert never delegates, MV corresponds to $\widetilde{q} = \underline{q}$: all voters cast their vote ($h_{i,e} = 0$ for all i). And yet we know from our earlier theorem that MV cannot be an equilibrium when delegation is allowed. By the logic of the theorem, there must be a profile of mutual best responses with positive but not full expected delegation.[34] But if delegation is not full (or $\widetilde{q} \neq \overline{q}$), then any voter who casts a vote has strictly positive probability of being pivotal. Thus for any non-expert i,

[34] We proved the theorem for arbitrary directions of delegation. But notice that the theorem continues to hold if we fix delegation to the pattern conjectured here: the expert never delegates and non-experts only delegate to the expert.

$(EUND_i(q_i) - EUD)$ is strictly increasing in q_i. It then follows that there must exist a $\widetilde{q} \in (q, \overline{q})$ such that $EUND_i(q_i) = EUD$ at $q_i = \widetilde{q}$, $EUND_i(q_i) > EUD$ for all $q_i > \widetilde{q}$, and $EUND_i(q_i) < EUD$ for all $q_i < \widetilde{q}$: all non-experts i with $q_i > \widetilde{q}$ prefer to vote, and all non-experts i with $q_i < \widetilde{q}$ prefer to delegate to the expert. □

Lemma 2. *Suppose $\pi = Pr(\omega_1) = 1/2$ and $K = 1$. Then: (i) If non-experts adopt the threshold strategies described in Lemma 1 and only delegate to the expert, it is optimal for the expert never to delegate. (ii) If the expert never delegates, and all other non-experts delegate to the expert if and only if $q_i < \widetilde{q}$, with $\widetilde{q} \in (\underline{q}, \overline{q})$, and never to each other, then it is optimal for each non-expert to delegate to the expert if and only if $q_i < \widetilde{q}$, and never to each other. The strategies are mutual best responses and constitute an equilibrium.*

Proof. We begin by proving claim (i): delegation from the expert to a non-expert cannot be optimal. First consider expert e delegating to some non-expert j when j does not delegate to e (and thus there is no cycle). At equilibrium with cutoff $\widetilde{q}$, expected payoff from all M non-experts using cutoff $\widetilde{q}$ is higher than the expected payoff from $M-1$ non-experts using cutoff $\widetilde{q}$ and 1 non-expert i always delegating to the expert regardless of i's precision (because i delegating when $q_i > \widetilde{q}$ strictly decreases expected payoff). This is in turn higher than the expected payoff from the same actions if the expert's precision were drawn from $[\widetilde{q}, \overline{q}]$ according to distribution F, rather than being p. But the expected payoff from this last scenario is identical to the expected payoff from e delegating to j: $M-1$ non-experts delegate using cutoff $\widetilde{q}$, but all their delegated votes are turned over to non-expert j, with precision randomly drawn from $[\widetilde{q}, \overline{q}]$, and e, who always delegates regardless of precision, is the analogue of voter i in the constructed scenario. Now suppose that when e delegates to j, j also delegates to e. This creates a cycle, and all delegations are cancelled. Note that having one's delegation cancelled conveys no information about signal realizations. Hence the voting game after delegations are returned is simply majority voting. But we know that majority voting is strictly dominated by the cutoff equilibrium with positive probability of delegation. In particular, cancelling all delegations must strictly decrease expected payoff since j only delegates to e when $q_j < \widetilde{q}$, and thus equilibrium delegation in this scenario is strictly positive.

Consider now claim (ii). At equilibrium $\widetilde{q}$, consider the difference in expected utility for a non-expert i between delegating to expert e or instead delegating to some non-expert j who has not herself delegated her vote to e (the choice would otherwise be irrelevant). Non-expert j has unconditional expected precision $\mu \equiv \mathbb{E}_F[q_j | q_j > \widetilde{q}] < p$, the unconditional precision of the expert. Voter i's expected utility from the two forms of delegation can differ only if i's vote is pivotal and $s_e \neq s_j$: for any number of delegated votes d, the expert agrees with

$\frac{M}{2} - (d+1)$ nonexperts and disagrees with $\frac{M}{2}$ (j included, i not included). The delegation choice is not trivial because, although $\mu < p$, when i is pivotal and $s_e \neq s_j$, there must be fewer independent signals agreeing with e than with j. Let$piv_i(d)$ be the set of signal realizations at which i's vote is pivotal, conditional on the expert being delegated d votes, and $pivj_i(d)$ be the same event additionally conditioning on $s_e \neq s_j$. Note that in both events the expert agrees with $\frac{M}{2} - (d+1)$ nonexperts and disagrees with $\frac{M}{2}$: the events contain the same information content and $Pr(s_e = \omega | piv_i(d)) = Pr(s_e = \omega | pivj_i(d))$ for all d. Then, noting that the summation below is to $M-2$ to exclude i and j:

$$EU(i \text{ delegate to } e) - EU(i \text{ delegate to } j) =$$
$$= \sum_{d=0}^{M-2} \binom{M-2}{d} F(\widetilde{q})^d (1-F(\widetilde{q}))^{M-2-d} Pr(pivj_i(d)|d)[Pr(s_e = \omega | pivj_i(d)) - Pr(s_j = \omega | pivj_i(d))]$$

Define $r(d) \equiv Pr(s_e = \omega | pivj_i(d)) = Pr(s_e = \omega | piv_i(d))$. Then:[35]

$$r(d) = \frac{Pr(pivj_i(d)|s_e = \omega) Pr(s_e = \omega)}{Pr(pivj_i(d))} =$$
$$= \frac{p\mu(\widetilde{q})^{\frac{M}{2}-(d+1)}(1-\mu(\widetilde{q}))^{\frac{M}{2}}}{p\mu(\widetilde{q})^{\frac{M}{2}-(d+1)}(1-\mu(\widetilde{q}))^{\frac{M}{2}} + (1-p)(1-\mu(\widetilde{q}))^{\frac{M}{2}-(d+1)}\mu(\widetilde{q})^{\frac{M}{2}}}$$

As $Pr(s_j = \omega | pivj_i(d)) = 1 - Pr(s_e = \omega | pivj_i(d))$, we can rewrite:

$$EU(i \text{ delegate to } e) - EU(i \text{ delegate to } j) = \sum_{d=0}^{M-2} \binom{M-2}{d} F(\widetilde{q})^d (1-F(\widetilde{q}))^{M-2-d} Pr(pivj_i(d)|d)[2r(d)-1].$$

Consider the difference in EU between i delegating to e and i voting when her precision is $q_i = \widetilde{q}$:

$$EU(i \text{ delegate to } e) - EU(i \text{ voting}) =$$
$$\sum_{d=0}^{M-2} \binom{M-2}{d} F(\widetilde{q})^d (1-F(\widetilde{q}))^{M-2-d} Pr(piv_i(d)|d)[Pr(s_e = \omega | piv_i(d)) - \widetilde{q}] =$$
$$= \sum_{d=0}^{M-2} \binom{M-2}{d} F(\widetilde{q})^d (1-F(\widetilde{q}))^{M-2-d} Pr(piv_i(d)|d)[r(d) - \widetilde{q}]$$

[35] Note that $r(d)$ is strictly decreasing in d for all d small enough that e is not dictator.

Note that $Pr(pivj_i(d)|d) = \frac{\frac{N}{2}}{N-1} Pr(piv_i(d)|d)$. For equilibrium $\widetilde{q}$, $EU(i \text{ delegate to } e) - EU(i \text{ voting}) = 0$ which implies:

$$\frac{M-1}{\frac{M}{2}} \sum_{d=0}^{M-2} \binom{M-2}{d} F(\widetilde{q})^d (1-F(\widetilde{q})^{M-2-d} Pr(pivj_i(d)|d) r(d) =$$

$$= \frac{M-1}{\frac{M}{2}} \sum_{d=0}^{M-2} \binom{M-2}{d} F(\widetilde{q})^d (1-F(\widetilde{q}))^{M-2-d} Pr(pivj_i(d)|d) \widetilde{q}$$

or:

$$r(d) = \widetilde{q}.$$

Hence:

$$EU(i \text{ delegate to } e) - EU(i \text{ delegate to } j) = \sum_{d=0}^{M-2} \binom{M-2}{d} F(\widetilde{q})^d (1-F(\widetilde{q}))^{M-2-d} Pr(pivj_i(d)|d)[2\widetilde{q}-1]$$

But $(2\widetilde{q}-1) > 0$ for all $\widetilde{q} \in (\underline{q}, \overline{q})$, and thus $EU(i \text{ delegate to } e) - EU(i \text{ delegate to } j) > 0$. □

Lemma 3. *Suppose $\pi = Pr(\omega_1) = 1/2$ and $K = 1$. Then the equilibrium strategies characterized in Lemma 2 maximize ex ante expected utility over sincere, semi-symmetric equilibria where the expert never delegates and nonexperts delegate to the expert only.*

Proof. We begin by defining interim expected utilities–expected utility for a voter with precision realization q_i, given cutpoint $\widetilde{q}$. A voter who delegates her vote to the expert e has interim expected utility EUD. Recall that upon delegation, expected utility does not depend on the voter's precision, and that d denotes the number of delegated votes:

$$EUD(\widetilde{q}) = \sum_{d=1}^{M} \binom{M-1}{d-1} F(\widetilde{q})^{d-1} (1-F(\widetilde{q}))^{M-d} EU(d, \widetilde{q})$$

where:

$$EU(d, \widetilde{q}) = \int_{\widetilde{q}}^{\overline{q}} ... \int_{\widetilde{q}}^{\overline{q}} U(d, q^1, ..., q^{M-d}) f(q^1 | q_{i_1} > q) ... f(q^{M-d} | q_{i_{M-d}} > q) dq^1 ... dq^{M-d}$$

The function $U(d, q^1, ..., q^{M-d})$ corresponds to the probability of the group choosing the correct alternative when d votes are delegated to e and $i_1, ..., i_{M-d}$ nonexperts, with precisions $q_{i_1}, ..., q_{i_{M-d}}$, have not delegated their votes.

If voter i does not delegate her vote, interim expected utility $EUND_i$is given by:

$$EUND_i(\widetilde{q}, q_i) = \sum_{d=0}^{M} \binom{M-1}{d} F(\widetilde{q})^d (1 - F(\widetilde{q}))^{M-d-1} EU_i(d, \widetilde{q}, q_i)$$

where:

$$EU_i(d, \widetilde{q}, q_i) = \int_{\widetilde{q}}^{\overline{q}} ... \int_{\widetilde{q}}^{\overline{q}} U_i(d, q_i, q^1, ..., q^{M-d}) f(q^1 | q_{i_1} > q) ... f(q^{M-d} | q_{i_{M-d}} > q) dq^1 ... dq^{M-d}$$

The function $(U_i(d, q_i, q^1, ..., q^{M-d})$ denotes the probability of choosing the correct alternative from the perspective of voter i with precision q_i who has not delegated). An equilibrium in stricty interior threshold strategies exists if there exists $\widetilde{q} \in (\underline{q}, \overline{q})$ such that $EUD(\widetilde{q}) = EUND_{q_i=\widetilde{q}}(\widetilde{q})$.

Consider now ex ante expected utility, before the realization of any non-expert precision, given cutpoint $\widetilde{q}$:[36]

$$EU(\widetilde{q}) = \sum_{d=0}^{M} \binom{M}{d} F(\widetilde{q})^d (1 - F(\widetilde{q}))^{M-d} EU(d, \widetilde{q})$$

We prove the lemma by showing that if there exists $\widetilde{q}' \in (\underline{q}, \overline{q})$ such that $EUD(\widetilde{q}') = EUND_{q_i=\widetilde{q}'}(\widetilde{q}')$, then $\widetilde{q}' = argmax EU(\widetilde{q})$: if there is an equilibrium in strictly interior threshold strategies, then the equilibrium threshold maximizes $EU(\widetilde{q})$. But we know from the Theorem that an interior threshold equilibrium must exist, and thus such equilibrium threshold must be maximal, in the class of equilibrium strategies embodied in $EU(\widetilde{q})$ (sincere, semi-symmetric strategies where the expert never delegates and nonexperts delegate to the expert only).

Consider the first order condition corresponding to the maximization of $EU(\widetilde{q})$ with respect to $\widetilde{q}$:

$$\frac{dEU(\widetilde{q})}{d\widetilde{q}} = 0$$

where

[36]Expectations are taken for given realization of precisions, and thus do not account for all possible payoff-irrelevant permutations of such realizations across different individuals.

$$\frac{dEU(\widetilde{q})}{d\widetilde{q}} = \sum_{d=0}^{M} \binom{M}{d} f(\widetilde{q})(d)F(\widetilde{q})^{d-1}(1-F(\widetilde{q}))^{M-d}EU(d,\widetilde{q}) -$$
$$-\sum_{d=0}^{M} \binom{M}{d} f(\widetilde{q})(M-d)F(\widetilde{q})^{d}(1-F(\widetilde{q}))^{M-d-1}EU(d,\widetilde{q}) + \sum_{d=0}^{M} \binom{M}{d} F(\widetilde{q})^{d}(1-F(\widetilde{q}))^{M-d}\frac{d}{d\widetilde{q}}EU(d,\widetilde{q}). \tag{3}$$

The first term can be simplified to:

$$Mf(\widetilde{q})\sum_{d=1}^{M} \binom{M-1}{d-1} F(\widetilde{q})^{d-1}(1 - F(\widetilde{q}))^{M-d}EU(d,\widetilde{q}) = Mf(\widetilde{q})EUD(\widetilde{q}).$$

For the third term, first note that $\frac{d}{d\widetilde{q}}EU(d,\widetilde{q}) = \frac{d}{d\widetilde{q}}\int_{\widetilde{q}}^{\overline{q}} ... \int_{\widetilde{q}}^{\overline{q}} U(d,q^1,...,q^{M-d})f(q^1|q_{i_1} > \widetilde{q})...f(q^{M-d}|q_{i_{M-d}} > \widetilde{q})dq^1...dq^{M-d}$. We have:

$$\frac{d}{d\widetilde{q}}EU(d,\widetilde{q}) = -\int_{\widetilde{q}}^{\overline{q}} ... \int_{\widetilde{q}}^{\overline{q}} U(d,q^1=\widetilde{q},q^2,...,q^{M-d})\frac{f(\widetilde{q})}{(1-F(\widetilde{q}))}f(q^2|q_{i_2} > \widetilde{q})...f(q^{M-d}|q_{i_{M-d}} > \widetilde{q})dq^2...dq^{M-}$$
$$+\int_{\widetilde{q}}^{\overline{q}} \frac{d}{d\widetilde{q}}\left(\int_{\widetilde{q}}^{\overline{q}} ... \int_{\widetilde{q}}^{\overline{q}} U(d,q^1,...,q^{M-d})f(q^1|q_{i_1} > \widetilde{q})...f(q^{M-d}|q_{i_{M-d}} > \widetilde{q})dq^2...dq^{M-d}\right)dq^1 =$$
$$= -(M-d)\int_{\widetilde{q}}^{\overline{q}} ... \int_{\widetilde{q}}^{\overline{q}} U(d,q^1=\widetilde{q},q^2,...,q^{M-d})\frac{f(\widetilde{q})}{(1-F(\widetilde{q}))}f(q^2|q_{i_2} > \widetilde{q})...f(q^{M-d}|q_{i_{M-d}} > \widetilde{q})dq^2...dq^{M-d} +$$
$$+\int_{\widetilde{q}}^{\overline{q}} ... \int_{\widetilde{q}}^{\overline{q}} \frac{d}{d\widetilde{q}}U(d,q^1,...,q^{M-d})f(q^1|q_{i_1} > \widetilde{q})...f(q^{M-d}|q_{i_{M-d}} > \widetilde{q})dq^1dq^2...dq^{M-d} =$$
$$= .-(M-d)\frac{f(\widetilde{q})}{(1-F(\widetilde{q}))}EU_i(d,q_i=\widetilde{q},\widetilde{q}) + \int_{\widetilde{q}}^{\overline{q}} ... \int_{\widetilde{q}}^{\overline{q}} \frac{d}{d\widetilde{q}}U(d,q^1,...,q^{M-d})\frac{f(q^1)...f(q^{M-d})}{(1-F(\widetilde{q}))^{M-d}}dq^1...dq^{M-d} =$$
$$= -(M-d)\frac{f(\widetilde{q})}{(1-F(\widetilde{q}))}EU_i(d,q_i=\widetilde{q},\widetilde{q}) + (M-d)\frac{f(\widetilde{q})}{(1-F(\widetilde{q}))}EU(d,\widetilde{q}) \tag{4}$$

where the second equality is by Leibniz's integral rule and the third comes from applying the rule recursively to the inner integrals.

Thus, simplifying the third term in 3, we have:

$$\sum_{d=0}^{M}\binom{M}{d}F(\widetilde{q})^d(1-F(\widetilde{q}))^{M-d}\frac{d}{d\widetilde{q}}EU(d,\widetilde{q})=\sum_{d=0}^{M}\binom{M}{d}f(\widetilde{q})(M-d)F(\widetilde{q})^d(1-F(\widetilde{q}))^{M-d-1}EU(d,\widetilde{q})-$$
$$-Mf(\widetilde{q})\sum_{d=0}^{M-1}\binom{M-1}{d}F(\widetilde{q})^d(1-F(\widetilde{q}))^{M-d-1}EU_i(d,q_i=\widetilde{q},\widetilde{q})=$$
$$\sum_{d=0}^{M}\binom{M}{d}f(\widetilde{q})(M-d)F(\widetilde{q})^d(1-F(\widetilde{q}))^{M-d-1}EU(d,\widetilde{q})-Mf(\widetilde{q})EUND_i(\widetilde{q},q_i)$$

Substituting 4 in 3, we thus obtain:

$$\frac{dEU(\widetilde{q})}{d\widetilde{q}}=Mf(\widetilde{q})EUD(\widetilde{q})-\sum_{d=0}^{M}\binom{M}{d}f(\widetilde{q})(M-d)F(\widetilde{q})^d(1-F(\widetilde{q}))^{M-d-1}EU(d,\widetilde{q})+$$
$$+\sum_{d=0}^{M}\binom{M}{d}f(\widetilde{q})(M-d)F(\widetilde{q})^d(1-F(\widetilde{q}))^{M-d-1}EU(d,\widetilde{q})-Mf(\widetilde{q})EUND_i(\widetilde{q},q_i)=$$
$$Mf(\widetilde{q})\left(EUD(\widetilde{q})-EUND_i(\widetilde{q},q_i)\right)$$

or: if there exists $\widetilde{q}'$ such that $EUD(\widetilde{q}')-EUND_{q=\widetilde{q}'}(\widetilde{q}')$, then $\frac{dEU(\widetilde{q})}{d\widetilde{q}}\Big|_{\widetilde{q}=\widetilde{q}'}=0$.

As mentioned above, this result, together with the Theorem, establishes the Lemma. We know from the Theorem that an equilibrium with $\widetilde{q}\in(\underline{q},\overline{q})$ exists and must be such that $EU(\widetilde{q})>EU(\underline{q})>EU(\overline{q})$,(where the last inequality follows from the Condorcet Jury Theorem). Hence there exists $\widetilde{q}'=argmaxEU(\widetilde{q})$, with $\widetilde{q}'\in(\underline{q},\overline{q})$ and by the result above, such $\widetilde{q}'$ must be an equilibrium. □

A.1.1 A Single Expert - Expected Utilities

$$EUD(\widetilde{q})=\sum_{z=0}^{M-1}\binom{M-1}{z}(1-F(\widetilde{q}))^z\,F(\widetilde{q})^{M-1-z}\sum_{c_n=0}^{z}\binom{z}{c_n}\mu_v(\widetilde{q})^{c_n}(1-\mu_v(\widetilde{q}))^{z-c_n}\times$$
$$\times\left[pI_{c_n+M-z+1>\frac{(M+1)}{2}}+(1-p)I_{c_n>\frac{(M+1)}{2}}\right]$$

$$EUND(q_i) = \sum_{z=0}^{M-1} \binom{M-1}{z} (1-F(\widetilde{q}))^z F(\widetilde{q})^{M-1-z} \sum_{c_n=0}^{z} \binom{z}{c_n} \mu_v(\widetilde{q})^{c_n} (1-\mu_v(\widetilde{q}))^{z-c_n} \times$$
$$\times \left[q_i \left(pI_{c_n+M-z+1>\frac{(M+1)}{2}} + (1-p) I_{c_n+1>\frac{(M+1)}{2}} \right) + (1-q_i) \left(pI_{c_n+M-z>\frac{(M+1)}{2}} + (1-p) I_{c_n+>\frac{(M+1)}{2}} \right) \right],$$

A.2 Monotonicity Violations

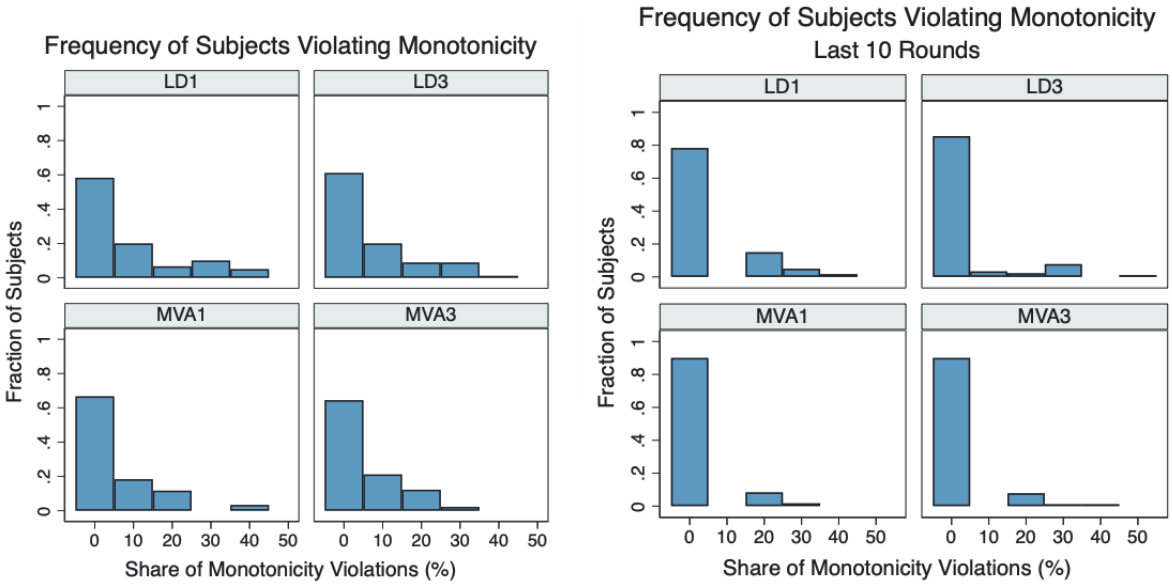


Figure 11: Monotonicity violations - histograms

A.3 Voting against signal and signal quality

Note that the only significant effect is from signal quality (precision). Neither the voting system nor the size of the group matter, but a decrease in signal quality is strongly correlated with an increase in voting against signal. Although voting against signal is an inferior action, the loss from doing so is indeed increasing in signal quality. There could be a small learning effect (voting against signal is less frequent in the treatments run second), but it is very variable and counterbalanced by an equally small and insignificant increase in voting against signal as rounds proceed, possibly from fatigue.

Frequency of Voting Against Signal

	(1) Linear Probability	(2) Probit
Signal Quality	-0.472***	-2.244***
	(0.0560)	(0.235)
	[0.000]	[0.000]
Round	0.00264	0.00454
	(0.0158)	(0.0832)
	[0.868]	[0.956]
LD	-0.00539	-0.0447
	(0.0270)	(0.150)
	[0.842]	[0.765]
N = 15	-0.00637	-0.0728
	(0.0362)	(0.197)
	[0.860]	[0.712]
Second	-0.0115	-0.0700
	(0.0383)	(0.200)
	[0.764]	[0.726]
Second * Mixed	-0.0372	-0.208
	(0.0332)	(0.195)
	[0.264]	[0.286]
N = 15 * Second	0.0288	0.164
	(0.0614)	(0.338)
	[0.640]	[0.626]
N = 15 * Mixed	-0.00469	0.00404
	(0.0398)	(0.242)
	[0.906]	[0.987]
Constant	0.452***	0.221
	(0.0692)	(0.239)
	[0.000]	[0.356]
Observations	2,552	2,552
R-squared	0.154	

*** p<0.01, ** p<0.05, * p<0.1

Notes: Standard errors in parentheses, clustered at the individual subject level. P-values in brackets. The dependent variable is a 0-1 indicator of whether the subject voted against their signal. Only non-experts are included, and only the instances in which they did not delegate or abstain are included. "Second" indicates that a treatment appeared second in the session. "Mixed" inidcates that both an LD treatment and MVA treatment appeared in the session.

A.4 The bootstrapping procedure: allowing for individual correlation across rounds

Replicating what happens in an individual session, we draw with replacement 15 subjects from the relevant treatment, each with all choices made over the 20 rounds. Among these 15 subjects, we draw, with replacement, 3 subjects, assigning to each of them one choice they made as expert,[37] and 12 subjects, assigning to each one choice made as non-expert. In $N = 15$, that constitute the group and yields one group decision; in $N = 5$ treatments, we divide the 12 subjects randomly into three groups of 4 and assign to each group one of the experts drawn earlier, generating three group decisions. We repeat new draws of 3 experts and 12 experts as above 20 times, generating 20 decisions from the same sample of 15 subjects if the treatment has $N = 15$, and 60 if the treatment has $N = 5$, thus simulating one experimental session. We then draw, always with replacement, a new group of 15 subjects, and repeat the procedure, each time generating 20 (60) decisions from the same group of 15 subjects, depending on the size of the group in the treatment. We repeat the whole procedure 4 times ($N = 5$), or 6 times ($N = 15$) generating 240 (120) decisions, as in our data from each of the $N = 5$ ($N = 15$) treatments. We then calculate the frequency of correct decisions, and consider that one data point for that treatment. We repeat the whole process 10,000 times and generate a distribution of the frequency with which the correct decision was reached.

A.5 The Random Dot Kinematogram

In a Random Dot Kinematogram (RDK), the perceptual stimulus consists of a number of dots being displayed on a screen. A proportion of these dots are determined to be signal dots, while the remaining are noise dots. Signal dots all move in a determined direction, while noise dots move at random according to an algorithm. The tasks consists in reporting the direction in which the signal dots are moving. This direction is called the coherent direction and the proportion of signal dots, the coherence, is the main factor in determining the difficulty of the task.

The task can be programmed in various ways, using a variety of parameters (e.g., color, duration, algorithm, number of dots). Research has been done to study the various effects of using different combinations of these parameters (Pilly and Seitz, 2009, Schütz et al., 2010). We take advantage of the recent development of a customizable version of the RDK (Rajananda et al., 2018) which can be implemented as a plugin in jsPsych. This version allows for the configuration of various parameters in order to adjust the task as desired by

[37] If the subject was never an expert, the subject is dropped and another one is drawn.

Duration:	1 second
Directions:	Left/Right
Number of dots:	300
Background color:	Black
Color of dots:	White
Dot radius:	2 pixels
Dot movement per frame:	1 pixel
Aperture width:	600 pixels
Aperture height:	400 pixels
Signal selection:	Same
Noise type:	Random direction
Aperture shape:	Ellipse
Reinsertion:	Dots reappear randomly when hitting edge
Fixation cross:	No
Aperture border:	No
Coherence:	20% to 3% (according to treatment)

Table 8: Experiment 2: RDK parameters

the researchers. We report in the table below the parameters that we used in our experiment. The reader can find details about how they affect the task in Rajananda et al., 2018 and in the following link: https://www.jspsych.org/6.3/plugins/jspsych-rdk/. It is important to emphasize again that our objective in using the RDK was not to study perception in itself, but rather to create a common task that is reasonably well controlled and calibrated and where, nevertheless, the information about the accuracy of the signals remains ambiguous.