The Concreteness of Social Knowledge and the Quality of Democratic Choice*

Kai Ou[†] Scott A. Tyson[‡]

Abstract

Democracy relies on citizens who are politically knowledgeable and engaged. However, when a voter gains political knowledge regarding important issues, through television, town halls, or social media, she also learns that there are many other politically knowledgeable voters, highlighting the importance of social knowledge in political participation. Will a voter with concrete—as opposed to hypothetical—knowledge about other voters' political knowledge have an increased incentive to participate? Or instead, will concrete social knowledge about other voters actually inhibit participation? In this article, we develop a novel experimental design that focuses on whether concrete knowledge about other voters' political knowledge influences political participation. Our main result shows that concrete social knowledge decreases individual voters' willingness to vote, and thereby reduces the probability democracy chooses the majority preferred alternative, i.e. the quality of democratic choice.

*The financial support from New York University and Florida State University is gratefully acknowledged. The study is pre-registered at EGAP. We are indebted to comments by James Bisbee, Chris Dawes, Chris Fariss, Guillaume Fréchette, Jana Freundt, Jens Großer, Dimitri Landa, Arthur Lupia, Max Lykins, Rebecca Morton, Melissa Sands, Nicholas Valentino, Jim Vreeland, members at the Behavioral Models of Politics Conference, ESA North America Meeting, APSA Annual Meeting, SPSA Annual Meeting, SJTU Experimental Economics Workshop, and seminar audiences at New York University and Florida State University. All errors remain the authors' own responsibility.

[†]Assistant Professor of Political Science, and faculty of the xs/fs Experimental Social Science Research Group at Florida State University, 554 Bellamy Hall, Tallahassee, FL 32306-2230. kou@fsu.edu.

[‡]Assistant Professor of Political Science, University of Rochester, and Research Associate, W. Allen Wallis Institute of Political Economy, University of Rochester. styson2@ur.rochester.edu.

1 Introduction

Perhaps the most important role of democratic institutions is to represent the interests of the majority, and ensuring this key function in modern democracies requires having an engaged and politically knowledgeable electorate. But democracy is undermined by at least two problems. First is the democratic dilemma (Lupia and McCubbins 1998), where voters lack the expertise to make decisions that protect their interests.[1] Second, and more problematic, is when voters who are politically knowledgeable lack the willingness or desire to engage in democratic decision-making. These two issues have a similar impact on *the quality of democratic choice*: the probability that the majority preferred alternative is achieved through democratic decision-making (Tyson 2016).

In response to recent declines in political engagement, scholars and policymakers have suggested that better achieving the democratic ideal of a knowledgeable and engaged electorate requires improving voters' overall interest and knowledge of politics. The context where political knowledge is acquired, for example, through town halls and most recently social media, typically bundles together a number of different characteristics, each of which influences the quality of democratic choice differently (Kenny 1998). Social media and polls have been shown to significantly increase voters' interest in politics which can increase voter turnout (e.g., Bernhardt, Krasa and Polborn 2008), but this impact may be context-dependent (Großer and Schram 2010).

In the interest of improving political participation, scholars have focused on the importance of enriching community, with a particular emphasis on the political role of social connections (Putnam 1966, 2001; Huckfeldt 1979; Gerber, Green and Larimer 2008), and appeals based on group membership (Herrera and Martinelli 2006; Schnakenberg 2014). Giles and Dantico (1982) find that voters are more likely to participate in politics when surrounded by similar peers (see also Großer and Schram 2006). Enriching social ties, however, entails a number of distinct mechanisms, each of whose all-else-equal influence is not well understood. In this study we address one component entailed by stronger community ties, which is the extent to which increased social ties leads to a more con-

[1] Druckman (2019) gives an overview of how the term "democratic dilemma" is used in the literature.

crete understanding about other voters' political knowledge. In particular, are voters with concrete social knowledge more (or less) motivated to participate in the democratic process? Are voters who know how many other voters are politically knowledgeable more engaged than those for whom such knowledge is hypothetical?

Consider, for example, a voting scenario where voters must choose between a policy that is tough on crime, using a large portion of the public budget on policing, and another that uses it for other things (perhaps mental health infrastructure). There may be partisan voters who prefer heavy police funding regardless of the current crime level, and partisan voters who want to defund the police in any state of the world. A majority of voters, however, are independent and want more police funding only when crime is high, and not when crime is low. Now suppose that independent voters learn about the level of crime through town halls (instead of, for instance, local news), where they also gain a concrete understanding of how many other independent voters are politically knowledgeable as well as the relative strength of partisan support. Does this kind of concrete social knowledge, compared to a situation where the presence of other politically knowledgeable voters is hypothetical, influence voters' propensity to participate in democratic decision-making?

To understand the importance of concretely knowing who is politically knowledgeable, we need to study the concreteness of social knowledge in isolation. To achieve this goal, we conduct experiments to isolate how the way political knowledge is obtained affects voter participation and the quality of democratic choice. Our experiments consist of an election where some voters are exogenously endowed with *political expertise*, telling them which alternative is best. Other voters, who do not receive expertise, do not know which alternative best serves the common good, and thus, cannot beneficially contribute to democratic decision-making.[2] Our main experimental treatment varies whether voters have concrete (as opposed to hypothetical) knowledge about the expertise of other voters. Specifically, in our *Treatment Group*, we reveal the number of expert voters among the electorate, whereas in our *Control Group*, this information remains hidden.

[2]See Brennan (2009) or Saunders (2020) for normative arguments against uninformed voting. In our experiment, nonexpert voters lack instrumental incentives to vote.

Our main findings address the quality of democratic choice, which is the probability that the majority preferred alternative is collectively chosen by voters. Our first main result is that concretely revealing the number of expert voters causes a 12 percentage point drop in the quality of democratic choice, suggesting that having concrete knowledge that there are other expert voters, perhaps because political knowledge is obtained in social settings, causes a decline in the quality of democratic choice.

We identify our main result using a novel experimental design that measures voters' desire to politically engage by directly eliciting each voter's *willingness to vote*, i.e. the largest cost a voter is willing to incur to ensure that her vote counts. We show that revealing the number of expert voters among the electorate causes a significant decline in the average willingness to vote by expert voters, establishing that the 12 percentage point decline in the quality of democratic choice is in fact because of reduced engagement by expert voters at the individual level, and not due to outliers, ecological differences between treatment groups, or incorrect 2^{nd}-order beliefs about nonexpert voting.

From the perspective of policy making, our results suggest that increased levels of political knowledge are not necessarily beneficial toward improving democratic outcomes. This is relevant to understanding the role of new media technologies, such as various social media platforms (e.g., Lupia 2016) and how social media affects the interplay between independent voters and partisan voters (e.g., Druckman, Levendusky and McLain 2018). Our results imply that distributing political information in public settings will be less effective than providing that same information in a decentralized manner. However, decentralized dissemination of information can be quite costly, thus suggesting key tradeoffs between several factors—material, normative, and strategic—affecting collective decision-making and the design of democratic institutions.

Our main experimental treatment distinguishes between voters who have concrete social knowledge and voters for whom this knowledge is hypothetical. In a related article, Esponda and Vespa (2014) show that, even when political information is perfectly reliable (as in our experiment), voters fail to accurately make inferences from hypothetical events (see also Esponda and Vespa (2016)). Since the acquisition and circulation of political knowledge is a social activity, it is important to understand how various aspects of political knowledge acquisition influence political participation, and consequently the quality of democratic choice. Existing empirical studies have found mixed results (Huckfeldt 1979; Huckfeldt and Sprague 1987). Specifically, McClurg (2006) finds that the societal supply of political expertise can encourage higher levels of political participation, whereas Mutz (2002) finds that people in social interactions that cross lines of political difference are less likely to participate in politics. It is important to stress that we are not studying whether a better informed electorate is good for democracy (see, e.g., Schnakenberg 2016; Prato 2018; Schnakenberg and Turner 2019). Additionally, our study is not about the tension between minority and majority interests, as we study a common value setting.

In our study, expert voters essentially face a kind of threshold public good environment, potentially under population uncertainty (see, e.g., Myerson 1998; Bouton and Castanheira 2012). McBride (2006) studies a threshold public good game where the threshold is uncertain, and shows that increased threshold uncertainty increases equilibrium contributions, provided the value of the public good is sufficiently high. In an experimental study, McBride (2010) evaluates behavior in a context similar to a public goods game, where subject beliefs are elicited to measure uncertainty. He finds that contributions increase in subjects' subjective assessment of the likelihood their contribution will be pivotal in public good provision.

The quality of democratic choice, which is the probability the majority preferred alternative is adopted, is also typically used in studies (experimental or otherwise) on Condorcet's Jury Theorem (e.g., Austen-Smith and Banks 1996; Feddersen and Pesendorfer 1997, 1998; Guarnaschelli, McKelvey and Palfrey 2000; Gerling, Grüner, Kiel and Schulte 2005). Specifically, studies that focus on information aggregation in elections measure the probability that a policy choice is matched with an unknown state of the world and refer to this quantity as "informational efficiency," see, in particular, Battaglini (2017). Each expert voter in our experiment is truthfully told the state of the world and the votes of partisan voters with certainty, thus the benefits of voting do not come from aggregating dispersed signals, but instead, come from an increased likelihood of defeating partisan voters.[3]

Finally, our study is connected to the literature on the swing voter's curse (Feddersen and Pesendorfer 1996), which is the phenomenon whereby uninformed voters cast votes (often according to a mixed strategy) to offset a known partisan bias, thus increasing the likelihood that informed voters are pivotal (see also Feddersen and Pesendorfer (1999) and Fey and Kim (2002)). Battaglini, Morton and Palfrey (2010) evaluate the swing voter's curse in a laboratory setting, where human subjects play the role of independent voters and partisan votes are played by the computer (i.e. partisan bias is exogenously assigned). In our experiment each level of the partisan bias is equally likely, making our setting similar to Battaglini, Morton and Palfrey (2010)'s treatment where each state of the world is equally likely and the partisan bias equals zero.[4] Because Battaglini, Morton and Palfrey (2010)'s study was designed to evaluate the swing voter's curse, both informed and uninformed voters are told the realization of the partisan bias.

There are three substantial differences between our study and those focusing on the swing voter's curse. First, in our experiment subjects do not receive noisy signals but are told the state of the world truthfully.[5] Second, our outcome measure is subjects' willingness to vote, which is the maximum cost a subject is willing to pay to ensure that her vote is counted; there are no vote costs in Battaglini, Morton and Palfrey (2010)'s experiment. Third, nonexpert voters in our experiment are *not* told the partisan bias. This is a critical difference between our experiment and Battaglini, Morton and Palfrey (2010), since this implies that the swing voter's curse cannot arise in our experiment. It is important to stress that avoiding swing voter's curse style behavior in our experiment was intentional, as our experiment was designed to isolate the role of the concreteness of

[3]Such concerns are also different from tactical coordination in multi-candidate elections (Myatt and Fisher 2002; Myatt 2007).

[4]In their study, subjects are randomly assigned to an electorate of seven voters and rematched throughout the session. Participants are asked to make discrete—and costless—vote choices under three different partisan biases ($0, 2$ and 4), and two states of the world. Participants receive noisy signals of the state of the world under two different prior distributions over the state of the world ($\frac{1}{2}$ and $\frac{5}{9}$).

[5]The swing voter's curse, as it is originally articulated in Feddersen and Pesendorfer (1996), does not rely on noisy information or information aggregation. Thus the experiments conducted by Battaglini, Morton and Palfrey (2010) do not measure a pure swing voter's curse. Instead, a modification of our design where realizations of the partisan bias are told to nonexpert voters, and the state of the world is told to expert voters, would isolate the swing voter's curse experimentally.

social knowledge.

2 The Experiment

Our experiment corresponds to an election where the best alternative for every independent voter (or subject) is the same (i.e. our setup is one of common values). Each election consists of 5 subjects and two additional "partisan" votes that we refer to as the *partisan bias*, which in our experiment is randomly determined by the computer. Each value of the partisan bias, corresponding to either two votes for A, two votes for B, or one vote for each, A and B, is equally likely, and thus, each occurs with probability $\frac{1}{3}$. Voters are tasked with collectively choosing an alternative, A or B, and collective decisions are made by simple plurality, i.e. whichever alternative receives more votes is the decision that is applied to everyone. Ties are broken by a fair coin toss.

To determine the best alternative for voters, a state of the world, which is either A or B, is drawn from a uniform distribution, so $P(A) = P(B) = \frac{1}{2}$. Each voter receives a "high" payoff whenever the collective decision matches the state of the world and a "low" payoff whenever the collective decision differs from the state of the world. The partisan bias can represent the behavior of a set of voters whose preferences are independent of the state of the world, and thus, may (or may not) be in agreement with the interests of experimental subjects.

The alternative which best serves voters' interest is not known to every voter. Instead, some voters are *expert voters*, who are told the state of the world and are thus perfectly informed of the alternative that best serves the common interest. Both expert voters and nonexpert voters know the distribution of expertise, i.e. how expertise is exogenously assigned among voters, as well as the distribution of the partisan bias. However, only expert voters are told the realized state of the world and the realized partisan bias.

The partisan bias in our experiment varies expert voters' instrumental "importance" in the election outcome by altering the hurdle that expert voters must cross in order to achieve their most preferred alternative. There are three different scenarios. First, the partisan bias is *supportive* when two partisan votes favor independent voters' preferred alternative. Second, when one partisan votes for A and the other partisan votes for B, we say the partisan bias is *neutral*. Third, the partisan bias is *against* when the computer casts two votes for the alternative that does not serve independent voters' interests.

We do not inform nonexpert voters of the realized partisan bias in the experiment to avoid the strategic incentives such information might create (e.g., the swing voter's curse) which would pollute the treatment effect we consider. In our setting, expert voters know perfectly the state of the world and partisan bias and so there is no benefit from aggregating information through voting. This feature of our design distinguishes our study from those that focus on the Condorcet jury theorem where aggregating information is crucial.

2.1 Procedures and Treatments

In each round (election), before voters make any decisions, it is *exogenously* determined which subjects are given expertise. An expert voter is told the state of the world, and the value of the partisan bias. In each round, to determine whether a subject privately receives expertise, the computer generates a random integer for each subject from a uniform distribution between 1 and 100. If an individual's number is below or equal to the exogenous (symmetric) cutoff of 70, then she is told the state of the world and the partisan bias. When a subject's number is higher than 70, she is not given this information.[6] This design implies that on average about 70% of subjects have expertise in each election. We provide expertise in this way because using a cutoff on a discrete uniform distribution is relatively simple for subjects to understand, and hence, manage calculations.

After the level of expertise is exogenously determined, each subject faces a voting decision, and our goal is to elicit each subject's *desire to engage*. Specifically, a subject's decision in our experiment consists of two parts. First, a subject can either vote for A, vote for B, or *Abstain*. If a subject abstains, her willingness to vote is recorded as 0. Second, each subject who did not abstain is asked how much she is willing to pay (in experimental points) to ensure that her vote is counted. To elicit a subject's *willingness to*

[6] It is important to stress that subjects do not incur a cost for expertise.

vote, we use a modification of the Becker-Degroot-Marshak (BDM) mechanism (Becker, DeGroot and Marschak 1964). That is, if a subject wants to vote, either for A or B, she must enter her willingness to vote, which is an integer in the interval $[1, 100]$. At this point, subjects do not know their own voting cost, nor do they know the voting costs of other subjects (although they know the distribution from which voting costs are drawn). Once a subject has reported her willingness to vote, the computer independently draws a vote cost from a uniform distribution between 1 and 100 experimental points, and implements the subject's decision rule. Namely, if a subject's vote cost is higher than her reported willingness to vote, then she does not pay her vote cost and her recorded vote is not counted toward the collective decision. Instead, if a subject's randomly generated vote cost is equal to or lower than her reported willingness to vote, then her vote is counted and she pays her voting cost.

When the democratically chosen alternative matches the state of the world, all subjects (expert and nonexpert) receive 110 experimental points. If the collective decision does not match the state of the world, then all subjects receive 10 experimental points. Each subject is given an endowment of 100 experimental points from which her vote cost is deducted should her vote be counted. This endowment ensures that every voter's experimental points are always positive. In every session, the exogenously assigned cutoff determining whether a participant receives expertise is common knowledge. The state of the world, the level of the partisan bias, and the realized number of expert voters are *ex ante* unknown. To eliminate possible cross-election influences like session effects (Fréchette 2012), subjects do not know who they are paired with in an electorate, and between election rounds subjects receive *no* feedback regarding the outcome of previous elections, previous realized voting costs, or the voting decisions of other subjects.

To investigate the role of the concreteness of social knowledge, we compare the *Treatment Group*, where voters are told the number of expert voters, with the *Control Group*, where voters are not given this information. If the number of expert voters is not revealed, subjects must infer how many voters have expertise (which is stochastic) from the distribution determining expertise among the population (described above). In both groups, the level of expertise (measured by the ex ante probability an individual voter is an expert) is the same, and hence, our results are not driven by differences in the level of political expertise.

2.2 Subjects

Our experiment was conducted at the XS/FS Experimental Social Science Laboratory at Florida State University where subjects were registered university students. No one participated in more than one session. The experiment was programmed using z-Tree (Fischbacher 2007) and conducted via computerized network. Subjects were seated at individual computer terminals and not allowed to see other subjects' choices. A total of 120 subjects participated in our study, where 60 were randomly assigned to the Control Group and 60 to the Treatment Group. We ran a total of 8 sessions. In each session, subjects were randomly assigned to groups of 5 participants, who interact throughout a session, i.e. there was not rematching between participants regarding their group assignment. Critically, subjects did not know the identity of members of their group, and we did not provide subjects with feedback regarding the outcome or distribution of votes in any previous election. We conducted 4 sessions with 12 independent electorates for each treatment group. They participated in a total of 30 elections in each session.

In order to help subjects understand the procedure of the experiment, we used numerical examples in the instructions as well as a comprehension quiz at the end of the instructions to illustrate the setup and process of the experiment. Subjects were given time to read the instructions at their own pace, provided an explanation of any incorrect answers, and allowed to ask questions about the quiz. A whole session lasted for about 50 minutes and one of the 30 elections was randomly selected by the computer as the election to be paid. The average earning for subjects was about $22 in which a $7 show-up payment was included. Figure 1 summarizes the progression of activities within the experiment and the instructions received by subjects are reported in Supplemental Appendix E online.

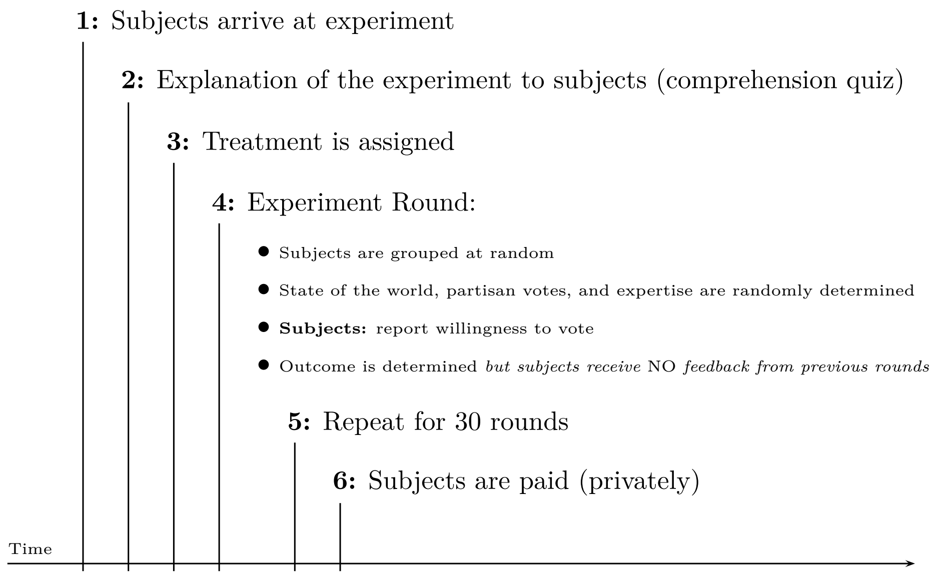

Figure 1: Structure of Experiment

3 Experimental Results

In this section, we present our main experimental findings regarding the quality of democratic choice and the willingness to vote among expert and nonexpert voters. We use electorate averages as the unit of observation, leaving us with 12 independent electorates for each treatment group, and perform non-parametric tests on electorate averages. All statistical tests reported in this section are two-sided and non-parametric. We use both Mann-Whitney tests (whose p-values are denoted by p_{MW}), and the less conservative Fisher-Pitman permutation test (whose p-values are denoted by p_{FP}). Recalling that subjects receive *no* feedback throughout our experiment, the statistical analysis reported can be viewed as an exceedingly conservative approach, since we take the same approach to our statistical analysis as other experiments where participants are provided with feedback and information of previous election periods.

3.1 Expert Voters

We are interested in how the concreteness of social knowledge affects the extent to which democratic decision-making represents the interests of the majority (absent other frictions of collective decision-making). First, we explore our main treatment effect regarding how knowing the number of voters with political expertise affects the *quality of democratic* *choice*—which is the probability the majority preferred alternative is selected through democratic decisionmaking. To see that the decline in the quality of democratic choice results from decreased engagement of expert voters, we look directly at the propensity of expert voters to vote, by comparing the average willingness to vote, when voters are told the number of expert voters, to the willingness to vote when this information remains hidden.

In the Control Group, the quality of democratic choice is 76 %, whereas in the Treatment Group the quality of democratic choice is 64 %. Comparing the Control Group with the Treatment Group shows that concretely knowing the number of expert voters causes the quality of democratic choice to decline by 12 percentage points, a statistically significant decline of about 16% ($p_{MW} = 0.021, p_{FP} = 0.022$).

Having established that knowing how many voters have expertise reduces the quality of democratic choice, we need to show that this decline is indeed the result of declining engagement among expert voters. In our experiment, expert voters' level of engagement is measured by computing the average willingness to vote among expert voters in the treatment and control groups. Figure 2 reports the average willingness to vote in the control and treatment groups along with the associated confidence intervals for each treatment, averaging over realizations of the number of expert voters, states of the world, and levels of the partisan bias. In the Control Group the average willingness to vote among expert voters is 40.3 (measured in experimental points) and in the Treatment Group the average willingness to vote among expert voters is 24.9. Thus, concretely informing voters of the number of expert voters reduces their willingness to vote by 15.4 experimental points, which is a statistically significant decline of about 38% ($p_{MW} < 0.001, p_{FP} < 0.001$).

While expert voters respond differently to the concreteness of social knowledge, this feature would be concerning if expert voters responded qualitatively differently to treatment at different levels of the partisan bias. To show that our main results are not because of averaging across different levels of the partisan bias, we break down our analysis of the willingness to vote by different levels of the partisan bias. When the partisan

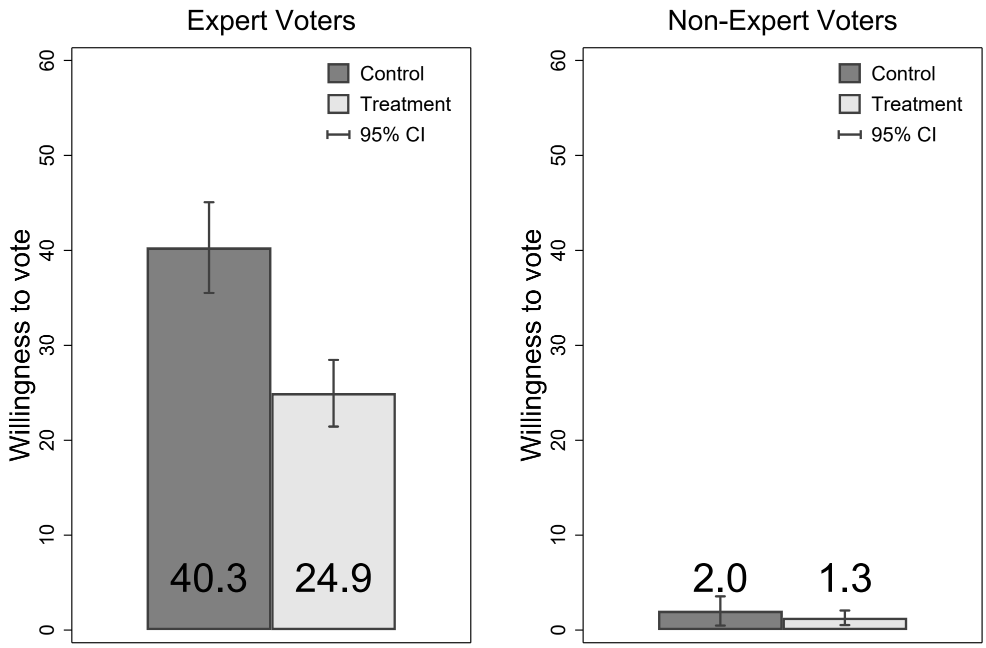

Note: We average (non)expert voters' willingness to vote and use electorate level averages as the unit of independent observation in the statistical analysis.

Figure 2: Average Willingness to vote by Treatment

bias is against, the average willingness to vote is 51.5 in the Control Group and 32.4 in the Treatment Group, which is significantly different ($p_{MW} < 0.001, p_{FP} < 0.001$) and corresponds to a decline in political engagement of about 37% due to revealing the number of expert voters. When the partisan bias is neutral (i.e. equal to zero), the average willingness to vote is 46.4 in the Control Group and 28.5 in the Treatment Group, which is statistically significant ($p_{MW} < 0.001, p_{FP} < 0.001$) and implies about a 39% decline in the willingness to vote of expert voters. When the partisan bias is supportive, the average willingness to vote in the Control Group is 21.3 and it is 13.6 in the Treatment Group, which is also marginally significant ($p_{MW} = 0.095, p_{FP} = 0.026$) and implies about a 36% drop in the willingness to vote among expert voters. These results show that when broken down by the partisan bias, although the size of our treatment effect changes, our main finding remains: revealing the number of expert voters consistently reduces political participation by expert voters.

3.1.1 Realization of Expert Voters

We now break down the analysis by investigating individual voters' willingness to vote as a function of the realized number of expert voters. We note that if expert voters were fully rational and purely instrumental, then because expert voters know the number of realized expert voters only in the Treatment Group, we would expect a flat average willingness to vote in the Control Group, but in the Treatment Group, expert voters should change their willingness to vote with the realized number of expert voters.

Figure 3 shows the relationship between willingness to vote and the realized number of expert voters. These results suggest that the concreteness of political knowledge changes voters' desire to engage, but voters' willingness to vote is not (completely) driven by strategic considerations. Indeed, in a regression using willingness to vote as the dependent variable, and the number of expert voters as the independent variable, we find little evidence that individuals' willingness to vote changes significantly over the realized number of expert voters, whether we pool the data together or separately by partisan bias. We analyze the behavior of expert voters in the Treatment Group with the same method and find that the average willingness to vote does not change with the realized number of expert voters.

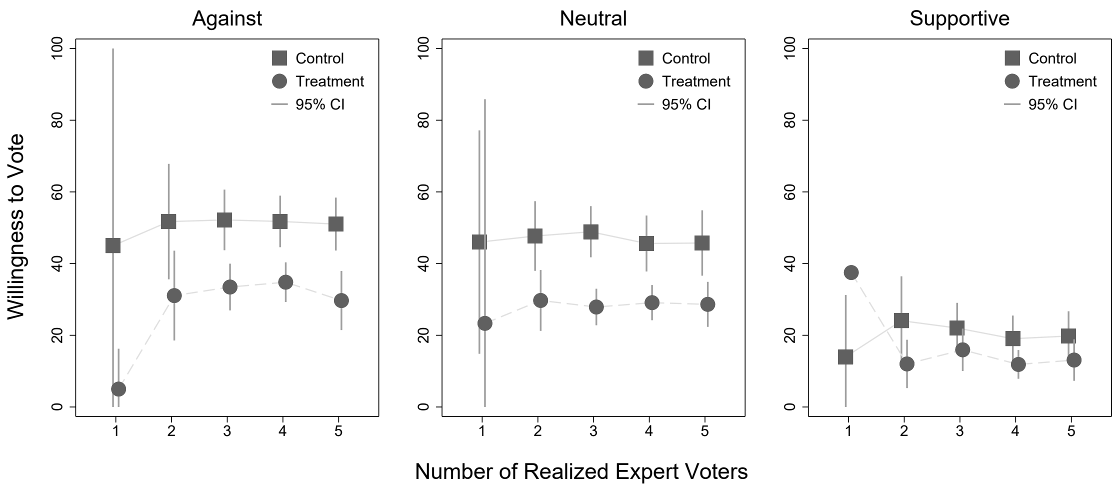

Note: In both the Control Group and the Treatment Group, on average about 70% of subjects have expertise in each election. There are few observations in which there are only one expert voter. For the cases in which there are only one expert voter, the lower bound of 95% CI is restrict at 0 and the upper bound of 95% CI is restrict at 100.

Figure 3: Willingness to Vote by the Realized Number of Expert Voters

When the partisan bias is against or neutral, the potential role of nonexpert voters is confounded with other potential incentives for expert voters. However, we can explore the potential strategic influence of nonexpert voting by focusing on the case of a supportive partisan bias. If expert voters anticipate the votes of nonexpert voters, and respond accordingly, then when the partisan bias is supportive, expert voters' willingness to vote should depend on the realized number of expert voters only through the corresponding difference in the number of nonexpert voters. By examining the average willingness to vote by the number of expert voters when the partisan bias is supportive in the Treatment Group, we find no evidence that expert voters' average willingness to vote changes across different realizations of the number of expert voters. This finding suggests that expert voters do not strategically change their willingness to vote to defeat potential nonexpert voting (which does not generally occur, see Section 3.2).

3.1.2 The Potential Influence of Outliers

To show that our main results are not driven by outliers, but instead, reflect a systematic change in expert voters' willingness to vote, we present the distributions of expert voters' willingness to vote across treatments. Figure 4 reports cumulative density distributions of the willingness to vote by treatment. Generally, the distribution of subjects' willingness to vote is shifted across treatments, suggesting that the average treatment effect is consistent and systematic. Putting everything together shows that our results hold at each level of the partisan bias and are not unduly influenced by outliers.

In sum, our results provide strong evidence that a particular kind of social knowledge—being told how many expert voters are present at election time—decreases the quality of democratic choice. Moreover, our results also provide strong evidence that the decrease in the quality of democratic choice is the direct result of decreased engagement by expert voters. Taken together, these results suggest that when political knowledge is obtained in a social context that also concretely conveys that other people are also gaining political knowledge, the quality of democratic choice decreases.

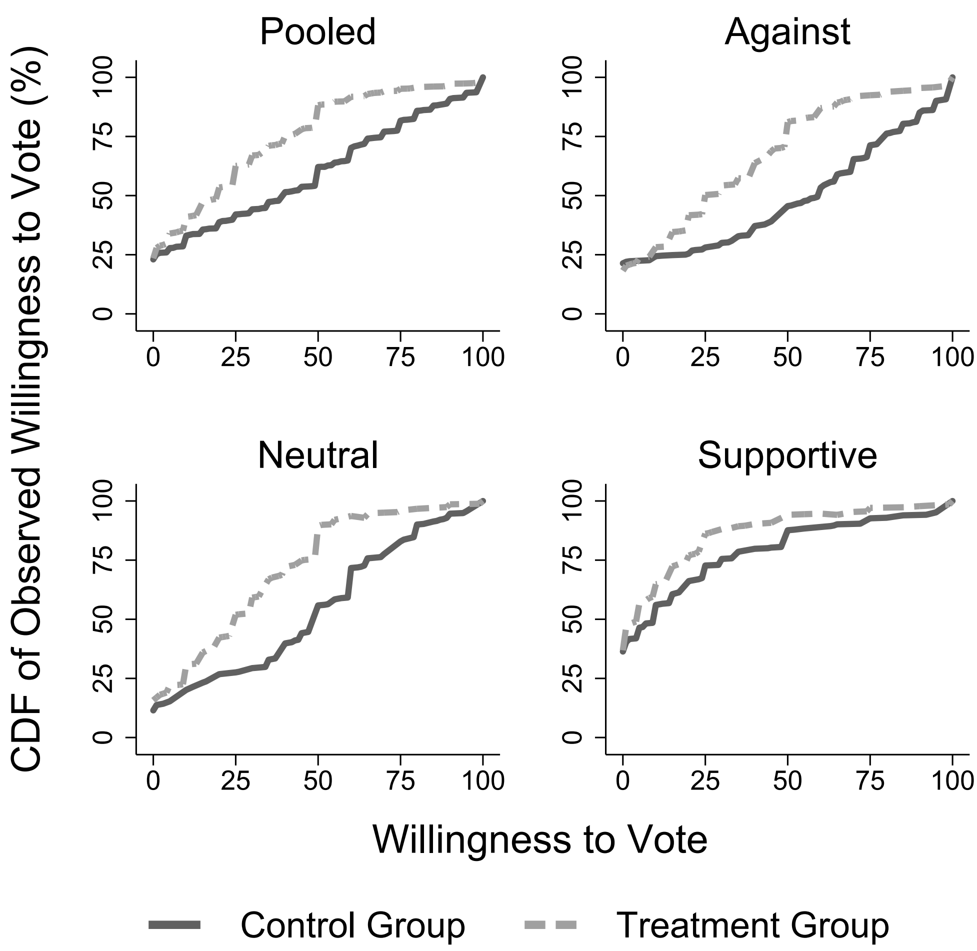

Figure 4: Cumulative Density Distribution of Expert Voters' Willingness to Vote

3.2 Nonexpert Voters

Our experiment is explicitly designed to focus on expert voters, and the influence of concrete social knowledge on their political participation. Since a nonexpert voter in our experiment has not been told the state of the world or the realized partisan bias, she does not know which alternative best serves her interests. If a nonexpert voter were to cast a vote, she may undermine the efforts of expert voters, and as a consequence, we expect that nonexpert voters should not be willing to incur much cost to ensure that an uninformed vote counts. In this section we explore two things. First, do nonexpert voters vote in our experiment? Second, are expert voters' choices influenced by the hypothetical possibility that nonexpert voters vote differently based on their presence in the Treatment

Group vs the Control Group?

In our experiment, nonexpert voters abstain about 93% of the time, meaning they do not express a vote choice or a positive willingness to vote. While a small minority of nonexpert voters express a vote choice, we find that both in the control and treatment groups, nonexpert voters' average willingness to vote is extremely small and significantly lower than that of expert voters. In the Control Group, the average willingness to vote of nonexpert voters is about 2.0 (out of 100), while in the Treatment Group, the average willingness to vote of nonexpert voters is about 1.3 (out of 100); which are not statistically distinguishable ($p_{MW} = 0.641, p_{FP} = 0.456$). Since nonexpert voters' willingness to vote is quite small, in our experiment their realized voting cost is very often prohibitively high, leading non-expert voters' choices to *not* count more than 98% of the time.

To further examine the treatment effect on willingness to vote, we compare the effect of becoming experts for each individual relative to being non-expert voters across the treatment and control groups, and perform a kind of difference-in-differences analysis. Expert voters' willingness to vote is significantly higher than nonexpert voters, whether in the Control Group (40.3 vs 2.0, $p_{MW} < 0.001, p_{FP} < 0.001$) or Treatment Group (24.9 vs 1.3, $p_{MW} < 0.001, p_{FP} < 0.001$). Figure 5 reports the difference of willingness to vote between expert and nonexpert voters in the control and treatment groups along with the associated confidence intervals for each treatment. The difference is 38.3 in the Control Group and 23.7 in the Treatment Group, which is statistically significant at the 0.001 level. These results suggest that individuals' voting choices respond to whether they are expert voters.

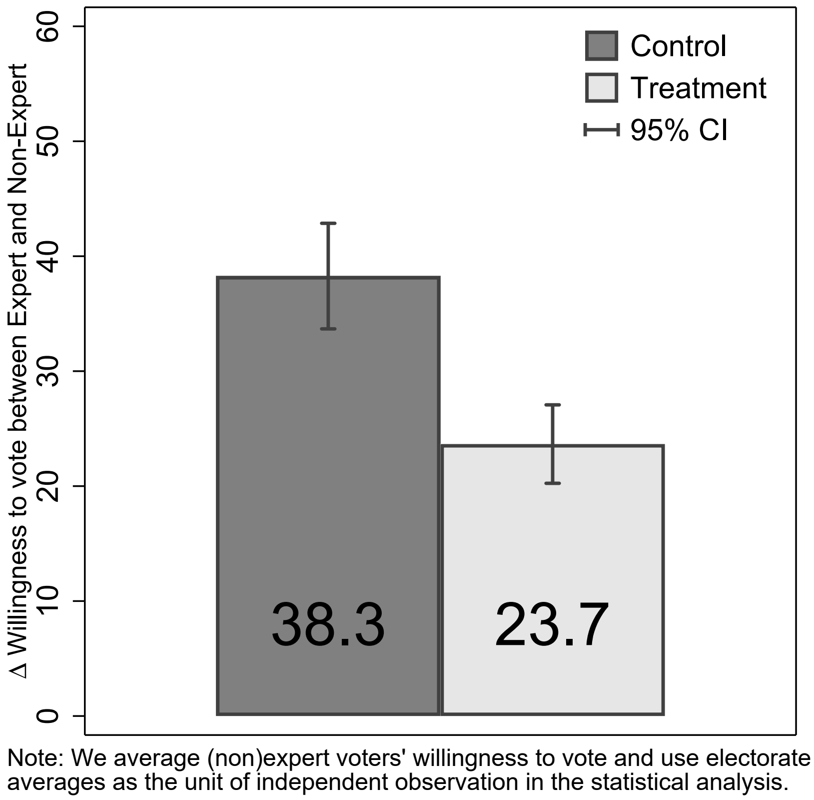

Figure 5: Marginal Effect between Expert and Non-expert Voting

3.3 Pivotal Voting Models

To what extent can a simple pivotal voting model, where a voter's incentive to vote is motivated purely by the likelihood that her vote is pivotal, explain our experimental results? In Supplemental Appendix G we present two pivotal voting models, one resembling our Control Group and the other resembling our Treatment Group, and characterize the vector of vote cost cutoffs that are symmetric Bayesian Nash equilibria of these pivotal voting models.[7] To assess the ability of such a pivotal voting model to provide insights into the behavior in our experiment, we consider what cutoff willingness to vote would be "predicted" by such a model, which are presented in Table 1.[8]

Taking the predicted values from Table 1, we find the average willingness to vote is computed to be 17.2 in the Control Group and 7.8 in the Treatment Group, which

[7]In Online Appendix H, we present a symmetric quantal response equilibrium, attributing deviations from the pivotal voting model to logistic decision errors.

[8]Relative to the model in the appendix we set the solidarity and duty terms to 0, and using the parameters of our experiment.

suggests that the average willingness to vote should be higher in the Control Group than in the Treatment Group. Our experimental observation is that the average willingness to vote is 40.3 in the Control Group and 24.9 in the Treatment Group. Thus, the results derived from pivotal voting models are qualitatively consistent with our experimental results. So, on the surface, a pivotal voting model seems to capture our main treatment effect well.

Table 1: Symmetric Bayesian Nash equilibrium Predictions

		Partisan Bias			
		Against	Neutral	Supportive	Average
Control Group	Obs.[Pred.]	51.5 [0 or 27.6]	46.4 [24.0]	21.3 [0]	40.3 [17.2]
Treatment Group	Aggregate	32.4 [0 or 1.6]	28.5 [21.8]	13.6 [0]	24.9 [7.8]
	M= 1	5 [0]	23.3 [50.0]	37.5 [0]	
	M= 2	31.1 [0]	29.7 [33.3]	11.8 [0]	
	M= 3	31.2 [0]	29.4 [26.8]	15.3 [0]	
	M= 4	35.0 [0 or 33.3]	27.8 [22.9]	12.0 [0]	
	M= 5	27.6 [0 or 34.7]	28.4 [20.2]	13.8 [0]	

Note: We report the experimental observations (Bayesian Nash equilibria computed based on pivotal voting models) outside (inside) of brackets. In the Control Group, given voters do not know concretely about how many expert voters are there, only average willingness to vote following the distribution of expertise in our experiment is estimated. In the Treatment Group, besides reporting the average willingness to vote of all realizations of the number of expert voters, since voters know how many expert voters are there in each election period, the results are also reported by $M = x$ that represents the number of realized expert voters. Note that when the partisan bias is against, there are multiple equilibria and zero is always an equilibrium willingness to vote, and when the partisan bias is supportive, the predicted willingness to vote is always zero.

Does the qualitative congruence of our treatment effect with what is produced by comparing pivotal voting models suggest that such models can give insights into our observed treatment effect? When breaking down the analysis by treatment and partisan bias, we find little to no evidence that voters' willingness to vote is primarily motivated by reasoning consistent with pivotality concerns. In Table 1, we present the observed average willingness to vote for each scenario of our experiment with the Bayesian Nash equilibrium vote cost cutoff from a pivotal voting model in brackets. Aside from the results that are inconsistent with pivotality concerns presented in previous sections, Table 1 provides a number more "predictions" that are not borne out in our data. First, when the partisan bias is neutral and there is only one expert voter, the voter's willingness to vote should be the highest, yet, in our experiment we find several scenarios where the average willingness to vote is higher. Second, a pivotal voting model suggests that expert voters' willingness to vote should decrease as the number of realized expert voters increases in the Treatment Group. Looking at the left column of Figure 1 and Table 1 in which the partisan bias is against, voters in our experiment consistently report a higher willingness to vote than would be warranted purely from pivotality concerns. A natural explanation for this "over-voting" might be the presence of a duty term, a la Riker and Ordeshook (1968), which could easily be incorporated into a pivotal voting model.[9] However, when we move to the middle column in the Treatment Group where the partisan bias is neutral, the pattern reverses for some values of the realized number of expert voters, namely, voters in our experiment are "under-voting" relative to the Bayesian Nash equilibrium of a pivotal voting model. This cannot be easily reconciled by simply appealing to a duty term. These patterns suggest that a pivotal voting model does not capture salient concerns motivating voters' willingness to vote at the level where decisions take place, despite the relatively good match at the aggregate level, where the factors that cause variation in voters' decisions are averaged out.

4 Normative Implications

In this section, we investigate the normative implications of concretizing social knowledge using two different consequentialist welfare criteria. First, we consider a maximum engagement criterion where all expert voters engage in voting and examine the quality of democratic choice. Second, we consider a utilitarian criterion, where full engagement is not endorsed as it is too costly, and explore the extent to which voters fall short of economic efficiency.

4.1 Maximum Engagement

We are interested in how the concreteness of social knowledge affects the ability of democracy to represent the interests of the majority, measured by the quality of democratic choice, which is the probability that the alternative preferred by the majority of voters is

[9]Indeed, in Supplemental Appendix G, we include such a term.

the alternative that is collectively chosen through democratic decision-making.[10] In our experiment the majority preferred option always corresponds to the alternative favored by subjects in our experiment. The partisan bias may, however, complicate voters' ability to achieve their most preferred alternative by acting as a hurdle voters must defeat.

In our experiment, a nonexpert voter is not told the state of the world, and consequently, does not know whether a vote for A or a vote for B best serves her interest. Additionally, were a nonexpert voter to caste a vote, she may undermine the efforts of expert voters, who know which alternative best serves her interests. Thus, an uninformed vote inhibits the ability of democratic choice to represent the interests of the majority.

To represent the quality of democratic choice, let the parameter $\rho \in [0,1]$ be the *level of political engagement by expert voters*, which is the probability that an expert vote is cast. In the context of our experiment, and assuming that nonexpert voters do not vote, we can represent the quality of democratic choice as $Q(\rho) =:$

$$\underbrace{\frac{1}{2}\left(\frac{1}{3}\left[1-\frac{1}{2}\binom{5}{0}\rho^0(1-\rho)^5\right]+\frac{1}{3}\left[1-\binom{5}{0}\rho^0(1-\rho)^5-\binom{5}{1}\rho^1(1-\rho)^4-\frac{1}{2}\binom{5}{2}\rho^2(1-\rho)^3\right]\right)}_{\text{the state of the world is } A \text{ and the partisan bias favors } B}$$
$$+\underbrace{\frac{1}{2}\left(\frac{1}{3}\left[1-\binom{5}{0}\rho^0(1-\rho)^5-\binom{5}{1}\rho^1(1-\rho)^4-\frac{1}{2}\binom{5}{2}\rho^2(1-\rho)^3\right]+\frac{1}{3}\left[1-\frac{1}{2}\binom{5}{0}\rho^0(1-\rho)^5\right]\right)}_{\text{the state of the world is } B \text{ and the partisan bias favors } A}$$
$$+\underbrace{\frac{1}{2}\left(\frac{1}{3}+\frac{1}{3}\right)}_{\text{the partisan bias favors voters' preferred alternative}}.$$

The first line represents the probability voters achieve their most preferred alternative when the political engagement of expert voters is ρ, the state of the world is A, and the partisan bias favors B. The second line represents a similar probability but when the state of the world is B and the partisan bias favors A. The last line is the probability that the partisan bias favors voters' preferred alternative.

Because of the symmetry in our experiment (between states of the world and levels of the partisan bias), whether the state of the world is A or B is immaterial for the quality

[10]It is important to stress that lower participation can improve welfare when alternatives are endogenous (Prato and Wolton 2020).

of democratic choice, and we can simplify to

$$Q(\rho) = \frac{1}{3}\left(2 - \frac{3}{2}\binom{5}{0}\rho^0(1-\rho)^5 - \binom{5}{1}\rho^1(1-\rho)^4 - \frac{1}{2}\binom{5}{2}\rho^2(1-\rho)^3\right) + \frac{1}{3}. \qquad (1)$$

This expression gives the connection between the level of political engagement of expert voters, ρ, and the quality of democratic choice, which is the probability that voters achieve their most preferred alternative through democratic decision-making. If expert voters are fully engaged, the quality of democratic choice is determined solely by the level of political expertise, which in our experiment is 70%. Thus, by substituting $\rho = 0.7$ in expression (1), a straightforward calculation establishes that the *maximum* probability voters collectively choose the alternative that is best for all of them, when expert voters are fully engaged, is $Q(0.7) = 0.967$. This corresponds to the majority-preferred alternative being adopted as often as possible.

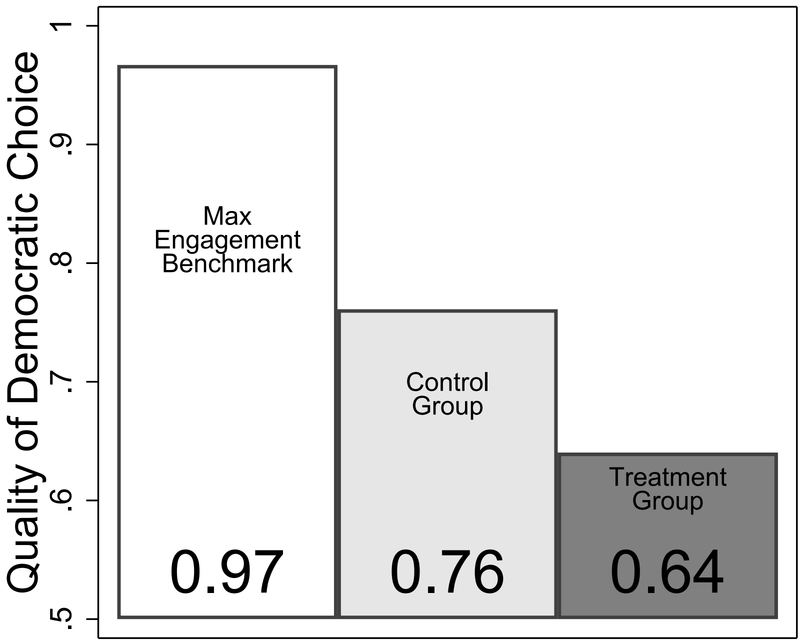

Note: In both the control and treatment groups, the estimates are based on actual votes of expert voters whose willingness to vote ≥ voting cost. We use electorate as the unit of independent observation. We use electorate level averages in the analysis and perform Mann-Whitney and exact Fisher-Pitman permutation tests. The achieved quality of democratic choice in each group is smaller than the max. eng. benchmark (p<0.003).

Figure 6: The Quality of Democratic Choice by Treatment

To interpret our experimental results with respect to this maximum engagement benchmark, in Figure 6 we compare the quality of democratic choice achieved in our experiment with the normative benchmark in which all expert voters engage in voting. We find that relative to when expert voters are maximally engaged, the actual quality of democratic choice is significantly smaller in both the Control Group ($p_{MW} = 0.002, p_{FP} < 0.001$) and Treatment Group ($p_{MW} = 0.002, p_{FP} < 0.001$), representing the (significant) drop resulting from reduced engagement by expert voters.

4.2 Utilitarian Planner

In this section we consider utilitarian welfare, under the implicit assumption that the only aspect relevant to voter welfare is their net payments in experimental points. To facilitate the analysis we focus on the problem of a utilitarian planner, who prescribes a symmetric cost cutoff for expert voters, which in principle, can depend on the partisan bias, denoted by β, and in the Treatment Group, the number of expert voters, $M \leq N$ (where N is the total number of voters). Suppose that an individual subject i's vote cost is c_i and that she votes if her vote cost is below a threshold $\overline{c}$.

In our analysis, we assume that the utilitarian planner, although benevolent, is not omniscient, and does not know individual voters' realized vote costs, realizations of the partisan bias, or the realized number of expert voters. It is important to emphasize that our utilitarian calculations are focused only on our experimental subjects, i.e. we do not consider partisan votes as being cast by people (which indeed they are not in our experiment). This is a reasonable way to proceed when considering the ethical implications of our particular experiment, but when applying our results to a context where partisan voters are people, some modifications may be necessary.

Focusing first on the Treatment Group, denote the utilitarian planner's symmetric cutoff prescription by $\overline{c}^*(M;\beta)$. For a supportive partisan bias, i.e. when $\beta = -2$, since no expert voter is pivotal, any level of voting is economically inefficient, and hence $\overline{c}^*(M;\beta = -2) = 0$. When the partisan bias is not supportive, the utilitarian planner must consider the (average) cost of endorsing voting among expert voters. Specifically, the total expected cost is

$$\sum_{j=1}^{M}\binom{M}{j}\overline{c}^{j}(1-\overline{c})^{M-j}j\cdot\frac{1}{\overline{c}}\int_{0}^{\overline{c}}cdc,$$

and, since $\int_0^{\overline{c}} cdc = \frac{1}{2}\overline{c}^2$, we can write the total expected cost as

$$\frac{1}{2}\overline{c}\sum_{j=1}^{M}\binom{M}{j}\overline{c}^{j}(1-\overline{c})^{M-j}j=\frac{M}{2}\overline{c}^{2}.$$

Thus, for a fixed partisan bias, $\beta \in \{0, 2\}$, and a fixed number of expert voters, M, the utilitarian planner's problem in the Treatment Group is

$$\max_{\overline{c}\in[0,1]} N\left(\sum_{k=\beta+1}^{M}\binom{M}{k}\overline{c}^{k}(1-\overline{c})^{M-k}+\frac{1}{2}\binom{M}{\beta}\overline{c}^{\beta}(1-\overline{c})^{M-\beta}\right)-\frac{M}{2}\overline{c}^{2}.$$

Using standard results (e.g., Hartley and Fitch 1951), and the Leibniz integral rule, the first-order condition can be reduced to (details in the Appendix):

$$\frac{\overline{c}^{\beta}(1-\overline{c})^{M-\beta-1}}{\int_0^1 t^{\beta}(1-t)^{M-\beta-1}dt}+\frac{1}{2}\binom{M}{\beta}\overline{c}^{\beta-1}(1-\overline{c})^{M-\beta-1}[\beta-M\overline{c}]=\frac{M}{N}\overline{c}.$$

Then, since $N = 5$, and using standard properties of the beta function, an interior $\overline{c}^*(M;\beta)$ solves:

$$5(M-1)(M-2)\overline{c}^{2}(1-\overline{c})^{M-3}+5\binom{M}{2}\overline{c}(1-\overline{c})^{M-3}[2-M\overline{c}]=\overline{c}.$$

Moving on to the Control Group, the utilitarian planner chooses a symmetric vote cost cutoff depending only on the partisan bias, β, which we denote by $\overline{c}^*(\beta)$.[11] As above, for a supportive partisan bias, $\overline{c}^*(\beta = -2) = 0$. For a partisan bias of $\beta \in \{0, 2\}$, and a probability of having political expertise of γ, the utilitarian planner's problem in the

[11] If the planner's prescriptions depended on the realized number of expert voters, then this information would be revealed to voters through the planner's prescription, and the Control Group would be the same as the Treatment Group.

control group is

$$\max_{\bar{c}\in[0,1]} N\left[\sum_{k=\beta+1}^{N}\binom{N}{k}(\gamma\bar{c})^k(1-\gamma\bar{c})^{N-k}+\frac{1}{2}\binom{N}{\beta}(\gamma\bar{c})^{\beta}(1-\gamma\bar{c})^{N-\beta}\right]-\frac{N}{2}\gamma\bar{c}^2.$$

Using that $\gamma = 0.7$ and $N = 5$ in our experiment, similar techniques as above allow us to reduce the planner's first-order condition in the Control Group to

$$10.5(0.7\bar{c})^2(1-0.7\bar{c})^2+0.35\binom{5}{2}\bar{c}(1-\bar{c})^2[2-3.5\bar{c}] = 3.5\bar{c}.$$

We report the numerical values of the utilitarian planner's choice problems, both in terms of willingness to vote and quality of democratic choice, in Figure 7 and compare them to our experimental results.[12] From these results, it is straightforward to see that subjects in our experiment dramatically under-participate when the partisan bias is against (relative to what is economically efficient). When the partisan bias is against, the observed willingness to vote is always significantly different from the planner's prescription at the one percent level except in the Treatment Group when there is one expert voter. We see similar under-participation when the partisan bias is neutral by members of the Treatment Group. When the partisan bias is neutral, the willingness to vote in the experiment is always significantly different from the planner's prescription at the one percent level except when there is only one expert voter. This type of under-participation relative to what is economically efficient is consistent with free-ridership, and thus, it is not surprising that we see similar behavior in our experiment. Interestingly, in the Control Group, and when the partisan bias is neutral, the average willingness to vote are statistically identical to the utilitarian planner's prescription. And perhaps more surprisingly, in contrast to the other cases, when the partisan bias is supportive, voters *over-participate* relative to what is economically efficient, in both the Treatment and Control groups (and the difference is always statistically significant at one percent level). Ultimately, our experiment is not designed to adjudicate what causes voters to under-participate or over-participate relative to what is economically efficient. Yet, the results

[12]Detailed information is reported in Online Appendix I.3.

of our experiment suggest that even in a laboratory environment, voters are motivated by non-material concerns, suggesting also that voter behavior may not be fully captured by pivotality concerns.

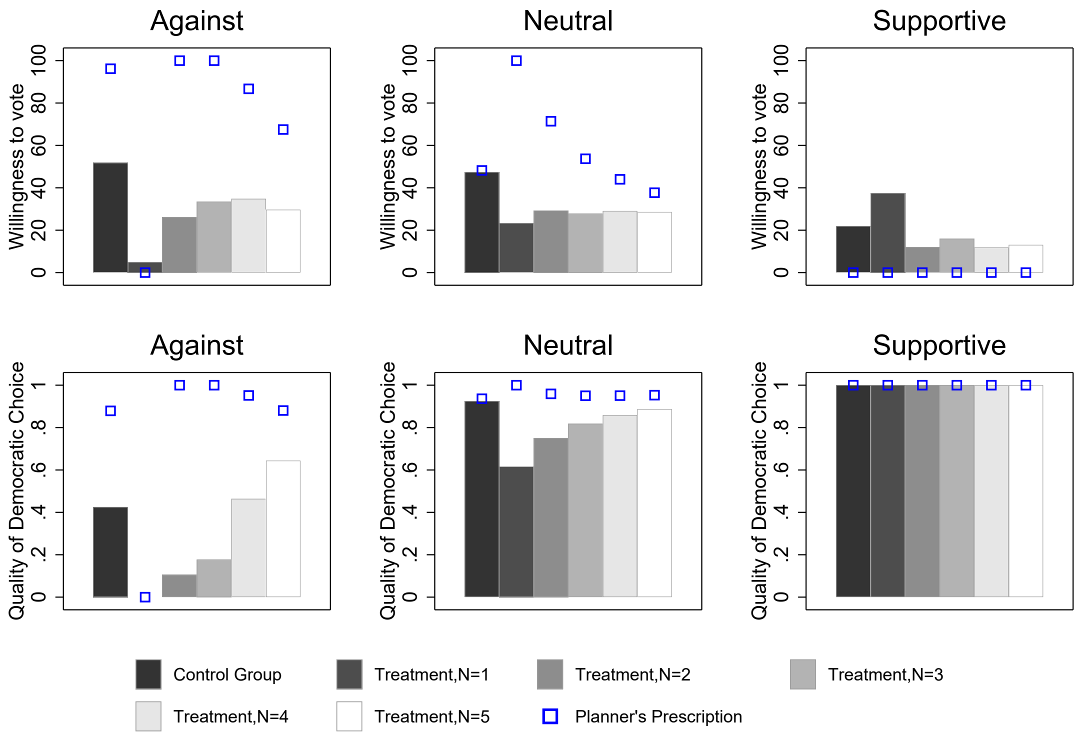

Note: The vertical axis of figures in the first row shows the willingness to vote. The vertical axis of figures in the second row shows the quality of democratic choice.

Figure 7: Utilitarian Optimal and Experimental Results by Partisan Bias

5 Conclusion

We show that even if the general distribution of political expertise is the same, when voters know concretely that there are other voters with expertise, they are less willing to politically engage. We report evidence from a laboratory experiment, where our main treatment effect isolates the relationship between voters' (possibly degenerate) beliefs regarding the number of expert voters and expert voters' willingness to vote, and consequently, the quality of democratic choice, which is the probability that the majority preferred alternative is reached through voting.

From a methodological perspective, our study departs from canonical designs where voters are only asked to make discrete vote choices. Instead, we directly elicit expert voters' willingness to vote, meaning the largest cost they are willing to incur to ensure that their vote is counted, which has at least two advantages over canonical designs. By directly eliciting a voter's desire to participate, our design contains more identifying variation than canonical designs. Second, canonical designs are forced to "back out" voters' incentives, often positing some structural model, which typically undermines causal interpretations of the experimental evidence. We are thus able to precisely investigate the causal effect of the concreteness of social knowledge, and moreover, show that our main results are the consequence of decreased engagement at the individual level.

References

Austen-Smith, David and Jeffrey S Banks. 1996. "Information aggregation, rationality, and the Condorcet jury theorem." *American Political Science Review* 90(1):34–45.

Battaglini, Marco. 2017. "Public protests and policy making." *The Quarterly Journal of Economics* 132(1):485–549.

Battaglini, Marco, Rebecca B Morton and Thomas R Palfrey. 2010. "The swing voter's curse in the laboratory." *Review of Economic Studies* 77(1):61–89.

Becker, Gordon M, Morris H DeGroot and Jacob Marschak. 1964. "Measuring utility by a single-response sequential method." *Behavioral Science* 9(3):226–232.

Bernhardt, Dan, Stefan Krasa and Mattias Polborn. 2008. "Political polarization and the electoral effects of media bias." *Journal of Public Economics* 92(5):1092–1104.

Bouton, Laurent and Micael Castanheira. 2012. "One person, many votes: Divided majority and information aggregation." *Econometrica* 80(1):43–87.

Brennan, Jason. 2009. "Polluting the polls: When citizens should not vote." *Australasian Journal of Philosophy* 87(4):535–549.

Druckman, James N. 2019. "The Democratic Dilemma Then and Now." *PS: Political Science & Politics* 52(4):625–629.

Druckman, James N, Matthew S Levendusky and Audrey McLain. 2018. "No need to watch: How the effects of partisan media can spread via interpersonal discussions." *American Journal of Political Science* 62(1):99–112.

Esponda, Ignacio and Emanuel Vespa. 2014. "Hypothetical thinking and information extraction in the laboratory." *American Economic Journal: Microeconomics* 6(4):180–202.

Esponda, Ignacio and Emanuel Vespa. 2016. Hypothetical thinking: Revisiting classic anomalies in the laboratory. Technical report Working Paper.

Feddersen, Timothy J and Wolfgang Pesendorfer. 1996. "The swing voter's curse." *American Economic Review* pp. 408–424.

Feddersen, Timothy J and Wolfgang Pesendorfer. 1999. "Abstention in elections with asymmetric information and diverse preferences." *American Political science review* 93(2):381–398.

Feddersen, Timothy and Wolfgang Pesendorfer. 1997. "Voting behavior and information aggregation in elections with private information." *Econometrica* pp. 1029–1058.

Feddersen, Timothy and Wolfgang Pesendorfer. 1998. "Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting." *American Political science review* 92(1):23–35.

Fey, Mark and Jaehoon Kim. 2002. "The swing voter's curse: comment." *American Economic Review* 92(4):1264–1268.

Fischbacher, Urs. 2007. "z-Tree: Zurich toolbox for ready-made economic experiments." *Experimental Economics* 10(2):171–178.

Fréchette, Guillaume R. 2012. "Session-effects in the laboratory." *Experimental Economics* 15(3):485–498.

Gerber, Alan S, Donald P Green and Christopher W Larimer. 2008. "Social pressure and voter turnout: Evidence from a large-scale field experiment." *American Political Science Review* 102(1):33–48.

Gerling, Kerstin, Hans Peter Grüner, Alexandra Kiel and Elisabeth Schulte. 2005. "Information acquisition and decision making in committees: A survey." *European Journal of Political Economy* 21(3):563–597.

Giles, Michael W and Marilyn K Dantico. 1982. "Political participation and neighborhood social context revisited." *American Journal of Political Science* pp. 144–150.

Großer, Jens and Arthur Schram. 2006. "Neighborhood information exchange and voter participation: An experimental study." *American Political Science Review* 100(02):235–248.

Großer, Jens and Arthur Schram. 2010. "Public opinion polls, voter turnout, and welfare: An experimental study." *American Journal of Political Science* 54(3):700–717.

Guarnaschelli, Serena, Richard D McKelvey and Thomas R Palfrey. 2000. "An experimental study of jury decision rules." *American Political Science Review* 94(02):407–423.

Hartley, HO and ER Fitch. 1951. "A chart for the incomplete beta-function and the cumulative binomial distribution." *Biometrika* 38(3/4):423–426.

Herrera, Helios and Cesar Martinelli. 2006. "Group formation and voter participation." *Theoretical Economics* 1(4):461–487.

Huckfeldt, R Robert. 1979. "Political participation and the neighborhood social context." *American Journal of Political Science* pp. 579–592.

Huckfeldt, Robert and John Sprague. 1987. "Networks in context: The social flow of political information." *American Political Science Review* 81(4):1197–1216.

Kenny, Christopher. 1998. "The behavioral consequences of political discussion: Another look at discussant effects on vote choice." *Journal of Politics* 60(1):231–244.

Lupia, Arthur. 2016. *Uninformed: Why people know so little about politics and what we can do about it.* Oxford University Press.

Lupia, Arthur and Mathew D McCubbins. 1998. *The democratic dilemma: Can citizens learn what they need to know?* Cambridge University Press.

McBride, Michael. 2006. "Discrete public goods under threshold uncertainty." *Journal of Public Economics* 90(6-7):1181–1199.

McBride, Michael. 2010. "Threshold uncertainty in discrete public good games: an experimental study." *Economics of Governance* 11(1):77–99.

McClurg, Scott D. 2006. "The electoral relevance of political talk: Examining disagreement and expertise effects in social networks on political participation." *American Journal of Political Science* 50(3):737–754.

Mutz, Diana C. 2002. "The consequences of cross-cutting networks for political participation." *American Journal of Political Science* pp. 838–855.

Myatt, David P. 2007. "On the theory of strategic voting." *Review of Economic Studies* 74(1):255–281.

Myatt, David P and Stephen D Fisher. 2002. "Tactical coordination in plurality electoral systems." *Oxford Review of Economic Policy* 18(4):504–522.

Myerson, Roger B. 1998. "Population uncertainty and Poisson games." *International Journal of Game Theory* 27(3):375–392.

Prato, Carlo. 2018. "Electoral competition and policy feedback effects." *Journal of Politics* 80(1):195–210.

Prato, Carlo and Stephane Wolton. 2020. "Wisdom of the Crowd? Information Aggregation in Representative Democracy.".

Putnam, Robert D. 1966. "Political attitudes and the local community." *American Political Science Review* 60(3):640–654.

Putnam, Robert D. 2001. *Bowling alone: The collapse and revival of American community*. Simon and Schuster.

Riker, William H and Peter C Ordeshook. 1968. "A Theory of the Calculus of Voting." *The American political science review* 62(1):25–42.

Saunders, Ben. 2020. "Against Detaching the Duty to Vote." *The Journal of Politics* 82(2):753–756.

Schnakenberg, Keith E. 2014. "Group Identity and Symbolic Political Behavior." *Quarterly Journal of Political Science* 9(2):137–167.

Schnakenberg, Keith E. 2016. "Directional Cheap Talk in Electoral Campaigns." *Journal of Politics* 78(2):527–541.

Schnakenberg, Keith E and Ian R Turner. 2019. "Signaling with Reform: How the Threat of Corruption Prevents Informed Policy-making." *American Political Science Review* pp. 1–16.

Tyson, Scott A. 2016. "Information acquisition, strategic voting, and improving the quality of democratic choice." *Journal of Politics* 78(4):1016–1031.

Supplemental Appendices for: The Concreteness of Social Knowledge and the Quality of Democratic Choice*

Kai Ou† Scott A. Tyson‡

*The financial support from New York University and Florida State University is gratefully acknowledged. The study is pre-registered at EGAP. All errors remain the authors' own responsibility.

†Assistant Professor of Political Science, and faculty of the xs/fs Experimental Social Science Research Group at Florida State University, 554 Bellamy Hall, Tallahassee, FL 32306-2230. kou@fsu.edu.

‡Assistant Professor of Political Science, University of Rochester, and Research Associate, W. Allen Wallis Institute of Political Economy, University of Rochester. styson2@ur.rochester.edu.

A Parameter Realizations

The experiment was programmed using z-Tree (Fischbacher 2007) and conducted via computerized network. We used the built-in function from z-Tree to generate a uniform distribution of the state of the world and partisan bias. Realizations of the state of the world in each treatment are: A 51% of the time and B 49% of the time. As reported in Figure A1, on average about 70% of voters are expert voters both in the control and treatment groups. Each level of the partisan bias occurred about $\frac{1}{3}$ of the time. In the Control Group, the partisan bias is supportive 31.94% of the time, neutral 32.22% of the time, and against 35.83% of the time. In the Treatment Group, the partisan bias is supportive 32.78% of the time, neutral 33.33% of the time, and against 33.89% of the time. The distribution of partisan bias is statistically identical between the Control Group and Treatment Group.

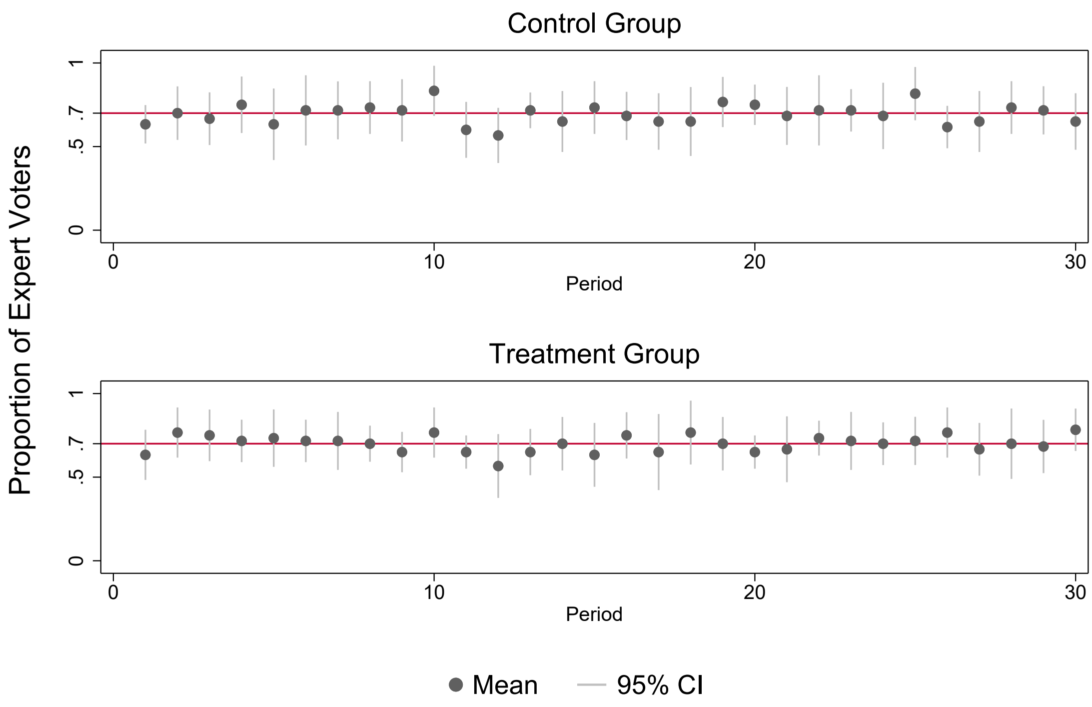

Figure A1: Distribution of Realized Expertise by Election-Period

B Regression Analysis of Quality of Democratic Choice

As a supplement to the nonparametric statistical analysis reported in the main text, we conduct linear regressions to investigate the extent to which the quality of democratic choice is different between the Control Group and Treatment Group. The results are reported in Table A1. We find that the quality of democratic choice is significantly smaller in the Treatment Group. These results are consistent with results that we report in the main text.

Table A1: Analyses of Treatment Effects on the Quality of Democratic Choice

Dependent Variable: Quality of Democratic Choice		
	(1)	(2)
Treatment	-0.121**	-0.121**
	(0.047)	(0.047)
Period		0.000
		(0.002)
Constant	0.761	0.755
	(0.031)	(0.046)
Observations	720	720

Note: OLS specification. Standard errors clustered at the electorate level. Clustered standard errors are reported in the parentheses. The signs $*$, $**$, $***$ indicate significance at 10%, 5%, and 1% level, respectively.

C Regression Analysis of Willingness to Vote

We conduct linear regressions to investigate the extent to which individual voters' willingness to vote is different between the Control Group and Treatment Group. The results are reported in Table A2. *Expert* is a dummy variable that we use to investigate the difference of willingness to vote when a subject is assigned expertise as compared to when the subject is a nonexpert voter. *Treatment* is a dummy variable that we use to indicate the Treatment Group. *ElectionPeriod* is the number of rounds which we use to control for learning effects. *Against* and *Neutral* are dummy variables that are used to identify what a specific level of partisan bias a subject plays. The inclusion of the interaction of the Treatment dummy and the Partisan Bias dummies explore treatment effects conditional on a specific level of the partisan bias. In Model 7 of Table A2, we further control for the demographic variables including gender, age, and subjects' performance on a Cognitive Reflection Task (CRT) developed by Frederick (2005).[1]

The results of Table A2 can be summarized as follows. First, on average, a subject's willingness to vote is significantly larger when she has expertise as compared to when she is a nonexpert voter. Second, expert voters in the Treatment Group always have a significantly lower cutoff willingness to vote, whether we only regress the voting data on the *Treatment* dummy or include additional control variables such as election periods, the level of partisan bias, and the demographic variables. Third, relative to the voting decisions in the scenario in which the level of the partisan bias is supportive, a voter's willingness to vote is significantly higher when the partisan bias is against and neutral. All these results yield conclusions that are the same as what we report in the main text.

[1] The CRT is conducted after subjects have finished all voting decisions. It is not incentivized; subjects are not paid for the CRT task.

Table A2: Analysis of Treatment Effects on Expert Voters' Willingness to Vote

	Dependent Variable: Cutoff Willingness to Vote						
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treatment		-15.421***			-15.049***	-7.607**	-8.121**
		(2.897)			(2.876)	(3.623)	(3.653)
Election Period			-0.045		0.017	0.031	0.031
			(0.078)		(0.074)	(0.073)	(0.073)
Against				24.991***	24.516***	30.323***	30.164***
				(2.481)	(2.456)	(3.878)	(3.883)
Neutral				19.647***	19.813***	25.089***	25.234***
				(2.383)	(2.330)	(4.224)	(4.255)
Treatment×Against						-11.623**	-11.141**
						(4.805)	(4.812)
Treatment×Neutral						-10.281**	-10.289**
						(4.666)	(4.705)
Age							-0.521
							(0.324)
Female= 1							-1.852
							(2.892)
CRT Performance							-0.472
							(1.206)
Expert= 1	30.966***						
	(1.549)						
Constant	1.635***	40.339***	33.300***	17.392***	24.790***	20.770***	32.972***
	(0.398)	(2.360)	(1.935)	(1.826)	(2.768)	(3.252)	(8.201)
Observations	3,600	2,511	2,511	2,511	2,511	2,511	2,511
Only Expert Voters	No	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.227	0.062	0.000	0.120	0.179	0.186	0.191

Note: OLS specification. Standard errors are clustered at the individual level. reported in the parentheses. The signs ∗, ∗∗, ∗ ∗ ∗ indicate significance at 10%, 5%, and 1% level, respectively.

D Learning Effects

To address the role of subject learning over the course of the experiment, we look at how the percentage of times voters succeed changes over time in our experiment. We consider the estimated quality of democratic choices as a function of the *Period* in a session. The results of the OLS regressions are summarized in Table A3.

We find that, although the quality of democratic choice generally increases with the number of periods in the Control Group, it decreases with the number of periods in the Treatment Group. We test whether the coefficients are significantly different by treatment, and find that the coefficient of each treatment is statistically indistinguishable by treatment. We conducted similar regressions to investigate whether individuals' willingness to vote changes over time, and we find the same qualitative results. In sum, we find no evidence that differences in the quality of democratic choice can be attributed to subjects learning over the course of the experiment.

Table A3: Quality of Democratic Choice as a Function of Period

Treatment	Coefficient	t	$Pr > \|z\|$
Control	0.002	0.71	0.495
Treatment	-0.001	0.34	0.739

Note: OLS specification. Standard errors clustered at electorate level. Dependent variable is the estimate of the quality of democratic choice. Independent variable is the number of election-period. The signs $*$, $**$, $***$ indicate significance at 10%, 5%, and 1% level, respectively.

E Experiment Instructions

You are participating in group decision-making experiment, where you will be making decisions as a member of a committee. We will start with a brief instruction period. During the instruction period, you will be given a complete description of the experiment and will be shown how to use the computers. If you have any questions during the instruction period, please raise your hand and your question will be answered out loud so everyone can hear. If you have any questions after the experiment has begun, raise your hand, and an experimenter will come and assist you. The instruction period will be followed by the paid session. The experiment consists of 30 rounds, and at the end one of the 30 rounds will be randomly selected by the computer as the round to be paid. You will also receive an additional show-up fee of $7. Everyone will be paid in private and you are under no obligation to tell others how much you earned. Your earnings during the experiment are denominated in Points (experimental currency). Everyone will be given an initial starting budget that is 100 Points. The experimental currency will be converted to US dollars and you will be paid by check. The exchange rate is 10 Points = $1.

Procedure

We begin the first round by dividing you into committees of five members each. Each of you is assigned to exactly one of these committees and you are not told the identity of the other members of your committee. Your committee is tasked with making a group decision. The decision is simply a choice between one of two alternatives: A and B. Committees make decisions by voting, where whichever alternative receives more votes is the committee's decision, and ties are broken by a fair coin toss.

At the beginning of each round, the computer will randomly select a *state of the world* for your committee. The state of the world will be either A or B. Each state of the world will be selected with a 50% chance. For each committee, and for each round, the computer will choose the state of the world separately, which means the selected state of the world could be different for different committees, and the selected state of the world would vary from one round to the other. Within one committee, the assigned state of the world will be the same for every member. Please note that you do not know the randomly assigned state of the world at first.

The computer will also randomly cast two votes. There is no connection between the assigned state of the world for your committee and how the computer cast the 2 votes. This means that:

- the state of the world and computerized votes operate independently and separately.
- The computer may either cast two votes for A, two votes for B, or split their votes evenly between A and B. Please note that each of these events is equally likely. For example, on average, 33% of the time the computer will cast two votes for A, 33% of the time the computer will cast two votes for B, and 33% of the time the computer will cast one vote for A and one vote for B.

Decision Task

Before voting, you may be given *expertise*, in which case you are told which state of the world was assigned to your committee, and how the computer has voted. [IN THE TREATMENT ONLY: You are also told exactly how many people in your committee are given expertise.] Whether you are given expertise is randomly decided. Specifically, the computer will randomly generate a number for each of the 5 committee members that is an independent random number between 1 and 100. The computer may generate a different number for different participants. If the number generated for you is equal to or lower than 70, you will be told the state of the world and computerized votes. If the number generated for you is greater than 70, you will not be given expertise. Each number between 1 and 100 is equally likely, and based on this mechanism, on average 70% of your committee members will be given expertise. Note that there is no cost for expertise.

Then, once you find out whether you have been given expertise, you and your other committee members will be asked to make a vote choice. You can vote for A, B, or you can abstain. If your decision is voting for A or B, you will be asked to input your highest willingness to pay in Points to make your ballot count. If your decision is to abstain, you will not be asked to input such a willingness to pay.

If you decide to vote for A or B, note that whether your ballot will be used in determining your committee's decision depends on your reported willingness to pay and a randomly drawn voting cost. Specifically, the computer will independently generate a random voting cost for you, which will be a number between 1 and 100, and where each number is equally likely.

- If the randomly generated voting cost is equal to or lower than your reported willingness to pay, then your vote will be counted as part of your committee's decision, and the *voting cost* will be deducted from your income.

- If the randomly generated voting cost is greater than your reported willingness to pay, then your vote choice will not be used in determining your committee's decision, and you do not need to pay the voting cost.

- The higher (lower) your willingness to pay for the ballot to make it count, the more (less) likely your ballot will be a valid vote that is used to determine the committee's decision.

- Your personal voting cost is randomly drawn by the computer independently, which means different committee members may have different voting costs.

- If you abstain, you will not be charged a voting cost and your ballot will not be cast.

Please note that when you are voting, the computer has cast the 2 votes. If you get expertise, you will know how the computer has voted; otherwise, you do not know this information.

For your committee, **only the ballots of those who decide to vote meanwhile whose willingness to pay are equal to or higher than their individual cost of voting will be counted.** The 2 votes cast by the computer will always be counted. The valid ballots of your committee and the 2 votes cast by the computer will jointly make a decision for your committee. Your committee's decision will be applied to everyone in your committee, and it is determined by simple plurality, i.e. whichever alternative receives more votes is the committee's decision. For example, **including the 2 votes cast by the computer**, if there are 4 votes for A, 3 votes for B, then your committee's decision is A. If there are 0 votes for A and 2 votes for B, then your committee's decision is B.

Ties (1-1, 2-2, 3-3) are broken randomly by a fair coin toss. So, for example, including the 2 votes cast by the computer, if 2 votes are for A, 2 votes are for B, and other participants abstain, then the total vote would be 2-2, which is a tie, and the tie is broken randomly, meaning that with a 50% chance A will be the committee's decision, and with a 50% chance B will be the committee's decision.

The other committees in the room face a similar task, but the correct decision may be different for different committees. Remember that committees are completely independent, and they act independently.

Payment

Payoffs are determined as follows. If your committee's decision matches the assigned state of the world, then all participants in your committee will receive a High payoff equal to 110 points.

- For example, if your committee decision is A and the assigned state of the world is A, then all participants in your committee will receive a High payoff.
- Whether you vote or abstain, and whether your vote is A or B, if your committee's decision is correct, you will receive a High payoff.

If your committee's decision does not match the assigned state of the world, then all participants in the committee receive a Low payoff equal to 10 points.

- For example, if your committee's decision is A but the assigned state of the world is B, then all participants in your committee will receive a Low payoff.
- Whether you vote or abstain, and whether your vote is A or B, if your committee's decision is wrong, you will receive a Low payoff.

Thus, your earning will be the realized payoff (High or Low) minus the cost of voting (if any).

After the first round is completed, we will proceed to the second round and the process repeats itself. The actual experiment will consist of 30 rounds. In each round, a new state of the world will be chosen and new votes by the computer will be cast. Each round is operated independently. One of the 30 rounds will be randomly chosen as the round to be paid, and you will be paid by check. Please treat every round seriously as each round is equally likely to be selected as the round to be paid.

[Experimenter: Do you have any questions?]

[Experimenter: We will first go through a practice round. During the practice round, please do not hit any keys until you are asked to, and when you are instructed to enter information, please do exactly as asked. You will not be paid for this practice round.]

The key information of the experimental procedure and decision task is summarized below.

1

- A *state of the world* is randomly assigned for your committee.
- At the same time, the computer randomly casts two votes.
- These two things are operated independently and separately.

2

- The 5 committee members do not know the assigned state of the world *ex ante*.
- The computer determines whether you get expertise.
- On average, about 70% of voters get expertise.

3

- Every voter will be asked to decide whether to vote for A, vote for B or to abstain.
- You will be asked to input your willingness to pay to make the ballot count if you decide to vote for A or B.
- Only the ballots of those whose willingness to pay is equal to or higher than their voting cost will be counted to determine the committee's decision.
- The 2 votes cast by the computer will always be counted in the election.

4

- Collective decisions are made by simple plurality; i.e. whichever alternative receives more votes is the committee's decision.
- Ties are broken by a fair coin toss.
- If your committee's decision matches the assigned state of the world, then all participants in your committee receive a High payoff equal to 110 points.
- If your committee's decision does not match the assigned state of the world, then all participants in the committee receive a Low payoff equal to 10 points.

F Comprehension Quiz

To make sure subjects understand how the experiment works, we ask them to answer the following comprehension quiz before they start making voting decisions in the experiment. The comprehension quiz is programmed using z-Tree (Fischbacher 2007) and conducted via computerized network. If a subject provides a wrong answer to a question, detailed explanations pop up on the computer screen. A subject can also raise her hand and an experimenter would answer her questions in private. Subjects cannot skip questions. They can make voting decisions only after they have answered all the questions correctly.

1. If in your committee, there are 2 votes for A, 0 votes for B, but in the other committee, there are 2 votes for A, 5 votes for B, what would be your committee's decision? [Answer: A]

2. Because each state of the world will be chosen with a 50% chance, if A was chosen for your committee in the last three rounds, does it mean that B will be chosen for your committee in this round (with a somewhat higher probability)? [Answer: No]

3. Regarding the 2 votes cast by the computer, on average 33% of the time the 2 votes are for A, 33% of the time the 2 votes are for B, and 33% of the time 1 vote is for A and 1 vote is for B. If you saw that in the first round the 2 votes are for A, in the second round there is 1 vote for A and 1 vote for B, does it mean that for the current round, it is somewhat more likely that the 2 votes are for B? [Answer: No]

4. Assume that in the round to be paid, A was the state of the world, and including the computerized votes your committee's decision was A. Your willingness to pay was 38 Points, in the end of the session, you find that your voting cost was 38 Points, what are your earnings for this round? [Answer: 72]

5. Assume that in the round to be paid, B was the state of the world, and including the computerized votes your committee's decision was A. Assume that your highest willingness to pay to make your vote count was 61, but the voting cost for your was 83, what are your earnings for this round? [Answer: 10]

6. Assume that in the round to be paid, including the computerized votes your committee's decision was the state of the world. You decided to abstain. Your voting cost was 48 Points. What are your earnings in this round? [Answer: 110]

7. Assume that the state of the world was A, the computer cast 2 votes for B. In your committee, 4 voters decided to vote for A, 1 voter decided to abstain. For those who decided to vote, their willingness to pay to make their vote count were 30, 50, 70, 90, respectively. If the computer randomly generated a voting cost 50 for everyone, how many votes were used to decide the committee's decision? [Answer: 3]

8. True or False: Based on the mechanism of how a voter is exogenously given expertise, on average about 70% of voters will have expertise. But in a particular round, it is possible that all the 5 voters get expertise or only 1 of the 5 voters get expertise. [Answer: True]

9. If A was the state of the world, the computer cast 1 vote for A and 1 vote for B. There was 1 committee member who got expertise and whose willingness to pay was 100 Points but he mistakenly voted for B, and the other committee members' votes are not counted (either because they abstain or their willingness to vote is less than their cost of voting), what was your committee's decision? [Answer: B]

10. You decide to vote for A, but your group decision is B. If the state of the world is B, and your willingness to pay is equal to or higher than your voting cost (which is 25 Points), how many points do you earn? [Answer: 85]

11. If A was the state of the world, and the computer cast 2 votes for A. Assume everyone in your committee gets expertise and no one will vote for the wrong alternative, how many votes are needed for your committee to guarantee a High payoff? [Answer: 0]

12. If B was the state of the world, and the computer cast 2 votes for A. Assume everyone in your committee gets expertise and no one will vote for the wrong alter-

native, how many votes are needed for your committee to guarantee a High payoff? [Answer: 3]

13. Including the votes cast by the computer, 6 votes are used to determine the committee's decision. Among the 6 votes, there are 3 votes for A and 3 votes for B. What is the probability that everyone in your committee will get a High payoff (110 Points)? [Answer: 50%]

14. In the selected round to be paid, if you decided to vote for A, and you would like to pay 75 points to make your vote count. The computer randomly generated a voting cost for you, which is 11 points. Regardless of the election outcome, how much voting cost will actually incur to your payoff? [Answer: 11]

15. In the selected round to be paid, if you decided to vote for B, and you would like to pay 15 points to make your vote count. The computer randomly generated a voting cost for you, which is 88 points. Regardless of the election outcome, how much voting cost will actually incur to your payoff? [Answer: 0]

G Pivotal Voting Models

There are N individuals deciding between two alternatives, $w = A$ and $w = B$, who make a collective decision by holding a vote. There are two equally likely states of the world, $\omega = A$ or $\omega = B$, and voters receive a utility normalized to 1 if alternative $w = \omega$ is adopted in state of the world ω, and 0 otherwise. In addition to this uncertainty, voters also face a hurdle represented by a partisan bias, denoted by β, which determines how many votes need to cast against a given alternative. This might represent partisan voters (as in the case of the motivating example) whose preferences are independent of the state of the world, in which case the total electorate would actually be larger. When $\beta > 0$, then alternative B receives β votes and voters must cast more than β votes in favor of alternative A in order to overturn alternative B. Similarly, when $\beta < 0$, then alternative A receives partisan support of $|\beta|$, and voters need to place at least $|\beta|$ votes in favor of alternative B to achieve the desired alternative. The partisan bias is drawn from a uniform distribution on $\{-(N-1), \ldots, 0, \ldots, N-1\}$. Ties are broken by a fair coin toss between the two alternatives A or B.

At the beginning of the game, voter i gains expertise with probability $\gamma \in [0, 1]$. Expertise informs a voter of the state of the world and the partisan bias: (ω, β). A voter's private ballot cost is independently drawn from a uniform distribution on $[0, 1]$. Following the realization of her private ballot cost, c_i, a voter must choose either to abstain, vote for alternative A, or vote for alternative B.

Denote by $d \geq 0$ a voter's private benefit to casting a ballot independent of the election's result, i.e. her duty term a la Riker and Ordeshook (1968). Additionally, let $s \geq 0$ represent a voter's sense of solidarity that she might receive by voting with other voters, which linearly scales the number of other voters who vote. Let V represent a voter's conjecture of the number of *other* informed voters who cast a vote, voter i's linear utility if she votes is

$$\mathbb{1}_{\{\omega=w\}} - c_i + d + s \cdot V,$$

and

$$\mathbb{1}_{\{\omega=w\}},$$

when i does not vote.

To summarize, the timing of the game is as follows:

1. The state ω, the partisan bias β, and private ballot costs are independently drawn.

2. Voters must decide whether to vote, and place their ballots.

3. The collective decision is reached and payoffs are received.

Instrumental Considerations. In a pivotal voter model, each voter is motivated by her assessment of the likelihood that she will be pivotal. In general, a voter must consider how many other voters have expertise. Let M be the random variable representing an individual voter's conjecture regarding the *total* number of expert voters. Notice that when a voter learns this quantity, then her belief is simply a point mass on the correct value, but when a voter is not informed of the number of other expert voters, then she must form a belief regarding this quantity. For a given partisan bias β, the probability a voter's preferred alternative wins if she votes is

$$P(M \geq \beta) \cdot P(V \geq \beta \mid M) + \frac{1}{2} P(M \geq \beta - 1) \cdot P(V = \beta - 1 \mid M).$$

The first term is the probability that an expert voter turns a tie into a win and the second term is the probability that an expert voter turns a loss into a tie. If an expert voter instead abstains, then the probability her preferred alternative is chosen is

$$P(M \geq \beta + 1) \cdot P(V \geq \beta + 1 \mid M) + \frac{1}{2} P(M \geq \beta) \cdot P(V = \beta \mid M).$$

From these expressions, an expert voter is pivotal with probability

$$\frac{1}{2}\Big[P(M \geq \beta) \cdot P(V = \beta \mid M) + P(M \geq \beta - 1) \cdot P(V = \beta - 1 \mid M)\Big]. \tag{1}$$

In the proceeding sections we analyze two different cases of the model, which differ in that the pivot probability, Equation (1), manifests differently in each case. First, we consider an expert voter's strategic decision to vote when she does not know how many other voters also have expertise, which resembles the Control Group in our experiment. Second, we consider the same strategic scenario, but where an expert voter knows the number of other expert voters, which resembles the Treatment Group from our experiment. We focus on the non-zero willingness to vote equilibria when they exist.

G.1 Unknown Number of Expert Voters

In this section we consider an expert voter's decision to vote, based on her assessment of the likelihood that she will affect the outcome of the election. We are interested in finding a symmetric vote cost cutoff where an expert voter with cost c votes if and only if $c \leq \overline{c}(\gamma; \beta)$, and abstains if $c > \overline{c}(\gamma; \beta)$. Denote by $\hat{c}^*(\gamma) = (\overline{c}(\gamma; 0), \overline{c}(\gamma; 1), \ldots, \overline{c}(\gamma; N-1))$. An expert voter does not know the number of other expert voters, so she instead uses the distribution determining expertise to form a belief about the number of other expert voters. Additionally, an expert voter, who knows other voters' symmetric cost cutoff,

$\overline{c}(\gamma;\beta)$, calculates the likelihood that other voters will vote. At a given partisan bias, β, suppose that an expert voter expects other voters to use the symmetric cutoff rule $\overline{c}$, then from (1), the pivot probability for an individual expert voter is

$$
\begin{aligned}
Piv(\gamma,\overline{c};\beta) \equiv & \sum_{m=\beta}^{N-1} \binom{N-1}{m} \gamma^m (1-\gamma)^{N-1-m} \binom{m}{\beta} (\overline{c})^{\beta} (1-\overline{c})^{m-\beta} \\
& + \sum_{m=\beta-1}^{N-1} \binom{N-1}{m} \gamma^m (1-\gamma)^{N-1-m} \binom{m}{\beta-1} (\overline{c})^{\beta-1} (1-\overline{c})^{m-\beta-1} \\
= & \binom{N-1}{\beta} (\gamma\overline{c})^{\beta} (1-\gamma\overline{c})^{N-1-\beta} + \frac{1}{2}\binom{N-1}{\beta-1} (\gamma\overline{c})^{\beta-1} (1-\gamma\overline{c})^{N-1-(\beta-1)} .
\end{aligned}
$$

For a given γ, the symmetric best-response for an expert voter at a given state β, is characterized by the balloting cost c^* that solves

$$
Piv(\gamma, c^*;\beta) + d + s((N-1)\gamma c^*) = c^*. \tag{2}
$$

The left-hand side is the probability an individual committee member is pivotal, the voter's duty benefit, and her solidarity payoff, and the right-hand side is the private cost of voting. Denote a solution to equation (2) by $c^*(\gamma, d, s;\beta)$, which gives the vector of symmetric pure-strategy Bayesian Nash equilibrium balloting cutoffs. Straightforward calculations of the binomial density allows us to write (2) as:

$$
\begin{aligned}
\binom{N-1}{\beta} \left(\gamma c^*_{\beta}\right)^{\beta} \left(1-\gamma c^*_{\beta}\right)^{N-1-\beta} + \frac{1}{2}\binom{N-1}{\beta-1} \left(\gamma c^*_{\beta}\right)^{\beta-1} \left(1-\gamma c^*_{\beta}\right)^{N-1-(\beta-1)} \\
+ d + s((N-1)\gamma c^*_{\beta}) = c^*_{\beta}.
\end{aligned}
$$

To generate point predictions, we set $d = s = 0$, thus focusing on pure pivotality considerations, and compute c^*_{β} using the parameters of our experiment, $N = 5$ and $\gamma = 0.7$. When the partisan bias is against ($\beta = -2$), by solving

$$
\frac{1}{2}\binom{4}{2}(0.7c^*)^2(1-0.7c^*)^2 + \frac{1}{2}\binom{4}{1}(0.7c^*)(1-0.7c^*)^3 = c^*,
$$

we get $c^*_{-2} = 0$ and $c^*_{-2} = 0.2755$.

Similarly, when the partisan bias is neutral, by solving

$$\frac{1}{2}\binom{4}{0}(0.7c^*)^0(1-0.7c^*)^4 = c^*,$$

we get $c_0^* = 0$ and $c_0^* = 0.2398$. Finally, since when the partisan bias is supportive, there is not a pivotal independent voter, the symmetric Bayesian Nash equilibrium cutoff willingness to vote is $c_2^* = 0$.

The average symmetric Bayesian Nash equilibrium willingness to vote, averaged over equally likely realizations of the partisan bias, is given by

$$\frac{1}{3}\left(c_{-2}^* + c_0^* + c_2^*\right). \tag{3}$$

From the symmetric Bayesian Nash equilibrium calculated when $d = s = 0$, and using the parameters from our experiment, the average willingness to vote when the number of expert voters remains unknown is

$$\frac{1}{3}(c_{-2}^* + c_0^* + c_2^*) = \frac{1}{3}(27.55 + 23.98 + 0) = 17.18.$$

G.2 Known Number of Expert Voters

In this section we consider the case where expert voters are also informed about the number of other expert voters. Suppose the realized number of *other* expert voters is M, then for a given level of the partisan bias β, the pivot probability, (1), is

$$Piv(\gamma, \overline{c}, M; \beta) \equiv \frac{1}{2}\left[\binom{M-1}{\beta}(\overline{c})^{\beta}(1-\overline{c})^{M-1-\beta} + \binom{M-1}{\beta-1}(\overline{c})^{\beta-1}(1-\overline{c})^{M-1-(\beta-1)}\right].$$

The best-response cutoff for expert voters, $c^*(\beta)$, must satisfy the equality

$$Piv(\gamma, c_\beta^\dagger, M; \beta) + d + s(Mc_\beta^\dagger) = c_\beta^\dagger.$$

Similar to above, the left-hand side is the probability an individual committee member is pivotal, the voter's duty benefit, and her solidarity payoff, and the right-hand side is the private cost of voting. Denote a solution to equation (G.2) by $c^*(\gamma, d, s; M, \beta)$, which gives the vector of symmetric pure-strategy Bayesian Nash equilibrium balloting cutoffs. We can write (G.2) as:

$$\frac{1}{2}\Big[\binom{M-1}{\beta}\left(c^{\dagger}_{\beta}(M)\right)^{\beta}\left(1-c^{\dagger}_{\beta}(M)\right)^{M-1-\beta}+\binom{M-1}{\beta-1}\left(c^{\dagger}_{\beta}(M)\right)^{\beta-1}\left(1-c^{\dagger}_{\beta}(M)\right)^{M-1-(\beta-1)}\Big]$$
$$+\,d+s(Mc^{\dagger}_{\beta}(M))=c^{\dagger}_{\beta}(M).$$

To generate point predictions, we set $d = s = 0$, and compute $c^*_{\beta}(M)$. Based on the setting of our experiment, when the partisan bias is against and there is only one expert voter, which means $M = 1$, $c^{\dagger}_{-2}(1) = 0$. When there are two expert voters, by solving

$$\frac{1}{2}c^{\dagger}(2) = c^{\dagger}(2),$$

we get $c^{\dagger}_{-2}(2) = 0$. When there are three expert voters, by solving

$$\frac{1}{2}\binom{2}{2}c^{\dagger}(3)^2 + \frac{1}{2}\binom{2}{1}c^{\dagger}(3)(1-c^{\dagger}(3)) = c^{\dagger}(3),$$

we get $c^{\dagger}_{-2}(3) = 0$. When there are four expert voters, by solving

$$\frac{1}{2}\binom{3}{2}c^{\dagger}(4)^2(1-c^{\dagger}(4)) + \frac{1}{2}\binom{3}{1}c^{\dagger}(4)(1-c^{\dagger}(4))^2 = c^{\dagger}(4),$$

we get $c^{\dagger}_{-2}(4) = 0$ and $c^{\dagger}_{-2}(4) = 0.3333$. When there are five expert voters, by solving

$$\frac{1}{2}\binom{4}{2}c^{\dagger}(5)^2(1-c^{\dagger}(5))^2 + \frac{1}{2}\binom{4}{1}c^{\dagger}(5)(1-c^{\dagger}(5))^3 = c^{\dagger}(5),$$

we get $c^{\dagger}_{-2}(5) = 0$ and $c^{\dagger}_{-2}(5) = 0.3473$.

Next, when the partisan bias is neutral and there is only one expert voter ($M = 1$),

$c_0^\dagger(1) = 0.5$. When there are two expert voters, by solving

$$\frac{1}{2}\binom{1}{0}(1 - c^\dagger(2)) = c^\dagger(2)$$

we get $c_0^\dagger(2) = 0.3333$. When there are three expert voters, by solving

$$\frac{1}{2}\binom{2}{0}(1 - c^\dagger(3))^2 = c^\dagger(3)$$

we get $c_0^\dagger(3) = 0.2679$. When there are four expert voters, by solving

$$\frac{1}{2}\binom{3}{0}(1 - c^\dagger(4))^3 = c^\dagger(4)$$

we get $c_0^\dagger(4) = 0.2291$. When there are five expert voters, by solving

$$\frac{1}{2}\binom{4}{0}(1 - c^\dagger(5))^4 = c^\dagger(5)$$

we get $c_0^\dagger(5) = 0.2024$. Finally, when the partisan bias is supportive, there is not a pivotal independent voter, so $c_2^\dagger(M) = 0$ for all M.

The average symmetric Bayesian Nash equilibrium, averaged over equally likely realizations of the partisan bias and number of expert voters, is given by

$$\frac{1}{3}c^*_{-2}+\frac{1}{3}\sum_{j=1}^{N}\binom{N}{j}(c_0^\dagger(j))^j(1-c_0^\dagger(j))^{N-j}c_0^\dagger(j)+\frac{1}{3}\sum_{j=1}^{N}\binom{N}{j}(c_2^\dagger(j))^j(1-c_2^\dagger(j))^{N-j}c_2^\dagger(j). \quad (4)$$

From the symmetric Bayesian Nash equilibrium calculated when $d = s = 0$, and using the parameters from our experiment, the average willingness to vote when the number of expert voters is known is

$$
\begin{aligned}
&\frac{1}{3}c_2^\dagger + \frac{1}{3}\sum_{j=1}^{N}\binom{N}{j}(c_0^\dagger(j))^j(1-c_0^\dagger(j))^{N-j}c_0^\dagger(j) + \frac{1}{3}\sum_{j=1}^{N}\binom{N}{j}(c_2^\dagger(j))^j(1-c_2^\dagger(j))^{N-j}c_2^\dagger(j)\\
=&\frac{1}{3}c_2^\dagger + \frac{1}{3}\sum_{j=1}^{N}\binom{N}{j}(c_0^\dagger(j))^{j+1}(1-c_0^\dagger(j))^{N-j} + \frac{1}{3}\sum_{j=1}^{N}\binom{N}{j}(c_2^\dagger(j))^{j+1}(1-c_2^\dagger(j))^{N-j}\\
=&\frac{1}{3}(0+0.2180+0.0155) = 0.0778.
\end{aligned}
$$

G.3 Combining the Models

We now compare the two cases of the pivotal voter model, one resembling our Control Group and the other resembling our Treatment Group. Subtracting (4) from (3) we obtain the difference

$$
\begin{aligned}
&\frac{1}{3}\left[\sum_{j=1}^{N}\binom{N}{j}(c_0^\dagger(j))^{j+1}(1-c_0^\dagger(j))^{N-j} + \sum_{j=1}^{N}\binom{N}{j}(c_2^\dagger(j))^{j+1}(1-c_2^\dagger(j))^{N-j} - c_0^* - c_2^*\right]\\
=&17.18 - 7.78 = 9.4,
\end{aligned}
$$

which corresponds to the same calculation we use to obtain our main treatment effect, where here the calculation is conducted on the symmetric Bayesian Nash equilibria of the pivotal voting model that best resembles our experimental conditions.

To summarize, the symmetric Bayesian Nash equilibrium for the various cases covered in our experiments are reported in Table A4. The point predictions are reported in brackets. Based on Mann-Whitney sign-rank tests and exact Fisher-Pitman permutation tests, the purely instrumental symmetric Bayesian Nash equiibrium predictions are significantly different from the experimental observations in every case at 0.01 level of significance. Moreover, the numerical value for the average difference in willingness to vote between pivotal voting models is 9.4, and the treatment effect from our experiment is 15.4. The difference, $15.4 - 9.4 = 6$ is statistically significant at 0.05 level.

Table A4: Symmetric Bayesian Nash equilibrium Predictions

		Partisan Bias			
		Supportive	Neutral	Against	Average
Unknown Expertise	Obs.[Pred.]	21.3 [0]	46.4 [24.0]	51.5 [0 or 27.6]	40.3 [17.2]
Known Expertise	Average	13.6 [0]	28.5 [21.8]	32.4 [1.6]	24.9 [7.8]
	M= 1	37.5 [0]	23.3 [50.0]	5 [0]	
	M= 2	11.8 [0]	29.7 [33.3]	31.1 [0]	
	M= 3	15.3 [0]	29.4 [26.8]	31.2 [0]	
	M= 4	12.0 [0]	27.8 [22.9]	35.0 [0 or 33.3]	
	M= 5	13.8 [0]	28.4 [20.2]	27.6 [0 or 34.7]	

Note: In the estimation of average willingness to vote for each treatment group, we use the non-zero cutoff equilibrium when it exists.

G.4 Existence Argument

We conclude our analysis of pivotal voting models by sketching the argument that establishes the existence of at least one symmetric Bayesian Nash equilibrium in each of the two cases we analyze above. It is important to note that there are generally multiple equilibria in each of the models presented.

Let $z \in \{0, 1\}$ be an indicator, and define the following smooth mapping:

$$\Psi^z_\beta(c; s, d) = z(Piv(\gamma, c; \beta) + s(N-1)\gamma c) + (1-z)(Pv(\gamma, c, M; \beta) + sMc) + d,$$

and note that a symmetric Bayesian Nash equilibrium is characterized by fixed points of $\Psi^z_\beta(c; s, d)$ in c at each β. Note first that $\Psi^z_\beta(0; s, d) > 0$ when $\beta \geq 0$. Thus, when $\Psi^z_\beta(1; s, d) < 1$, then, when $\beta \geq 0$, existence of a symmetric Bayesian Nash equilibrium, c^*_β, follows by the Intermediate Value Theorem. When the partisan bias favors voters, so that $\beta < 0$, then there is a unique cutoff, characterized by

$$d + zs(N-1)\gamma c) + (1-z)sMc = c,$$

which, after rearranging, is

$$c^*_\beta = \frac{d}{1 - zs(N-1)\gamma c) - (1-z)sMc}.$$

Finally, when $\Psi_{\beta}^{z}(1; s, d) > 1$, there is a symmetric Bayesian Nash equilibrium in which $c_{\beta}^{*} = 1$.

H Quantal Response Equilibrium Analysis

In this section, and building off the pivotal voting models presented in the previous sections, we introduce stochastic terms into the calculus of otherwise instrumentally minded voters, and focus on symmetric quantal-response equilibria. We follow the logit specification used in Goeree and Holt (2005) and discussed in Goeree, Holt and Palfrey (2016). Specifically, let the parameter $\mu \geq 0$ represent the degree of noise in voters' decisions so that individual voter i's expected utility includes a stochastic disturbance $\mu\varepsilon_i$, where ε_i is drawn independently across i from a logistic distribution. This is equivalent to idiosyncratic duty terms, i.e. where $d_i = \mu\varepsilon_i$ is the specification above. Denoting i's pivot probability by PIV_i, i votes if and only if

$$PIV_i \geq \mu\varepsilon_i \Rightarrow \frac{PIV_i}{\mu} \geq \varepsilon_i,$$

which occurs with probability

$$Prob(vote_i) = F\left[\frac{PIV_i}{\mu}\right]$$

Taking the inverse of $Prob(vote) = F(\cdot)$ and multiplying both sides by μ yields

$$\mu F^{-1}(Prob(vote)) = PIV.$$

Using the logistic distribution, $F(x) = 1/(1 + e^{-x})$, we derive the quantal response equilibrium condition:

$$\mu\left[-\ln\left(\frac{1 - Prob(vote)}{Prob(vote)}\right)\right] = PIV,$$

where the right-hand side is obtained from (1). In the Control Group, when $\mu = 0$, then the symmetric quantal-response equilibrium condition is same as that for a symmetric

Bayesian Nash equilibrium when $s = 0$, expressed above by (2). As a function of μ, Figure A2 illustrates how the symmetric quantal-response equilibrium provides a description of how the disturbance term, ε, can influence vote choices, which are otherwise dependent on pivotality concerns. Specifically, as μ increases, the symmetric quantal-response equilibrium cutoff willingness to vote also increases.

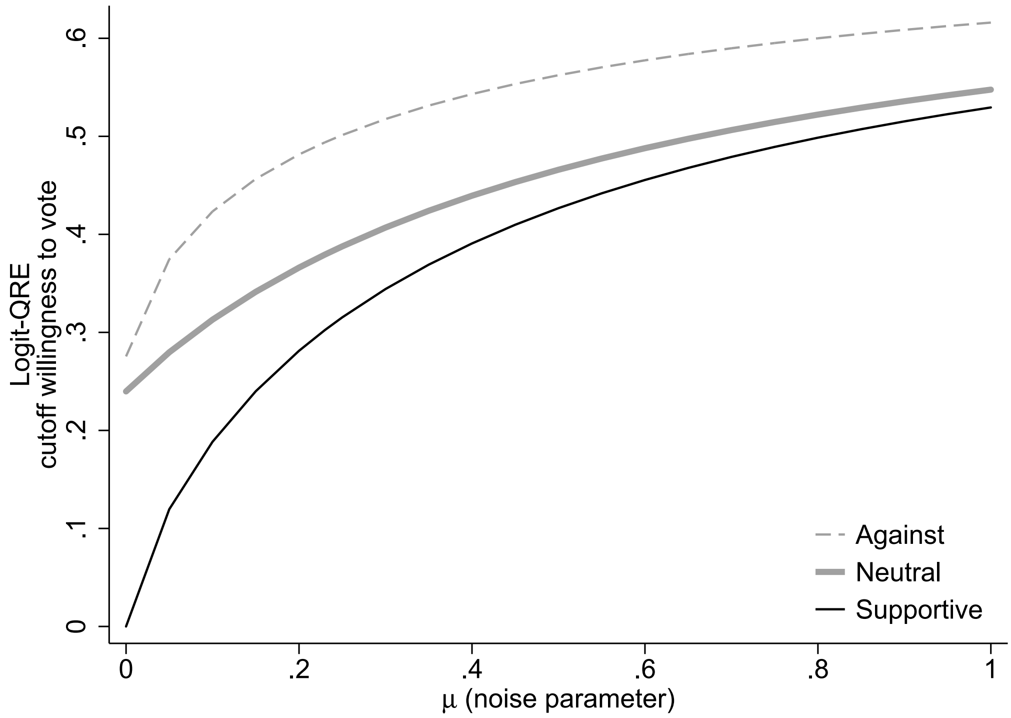

Figure A2: Symmetric Logit-QRE cutoff willingness to vote by μ (Control Group)

For the Treatment Group, when $\mu = 0$, then the symmetric quantal response equilibrium condition is equivalent to the symmetric Bayesian Nash equilibrium. As above, the symmetric quantal-response equilibrium cutoff willingness to vote increases as the logistic errors have a greater influence on vote choices. Figure A3 illustrates the relationship between the symmetric quantal-response equilibrium cutoff willingness to vote at each level of the partisan bias, as a function of μ.

Figure A4 plots the pointwise difference between symmetric vote cost cutoff quantal-response equilibrium willingness' to vote. It is straightforward to see that the symmetric vote cost cutoff quantal-response equilibrium willingness to vote is pointwise in μ higher in the Control Group than in the Treatment Group.

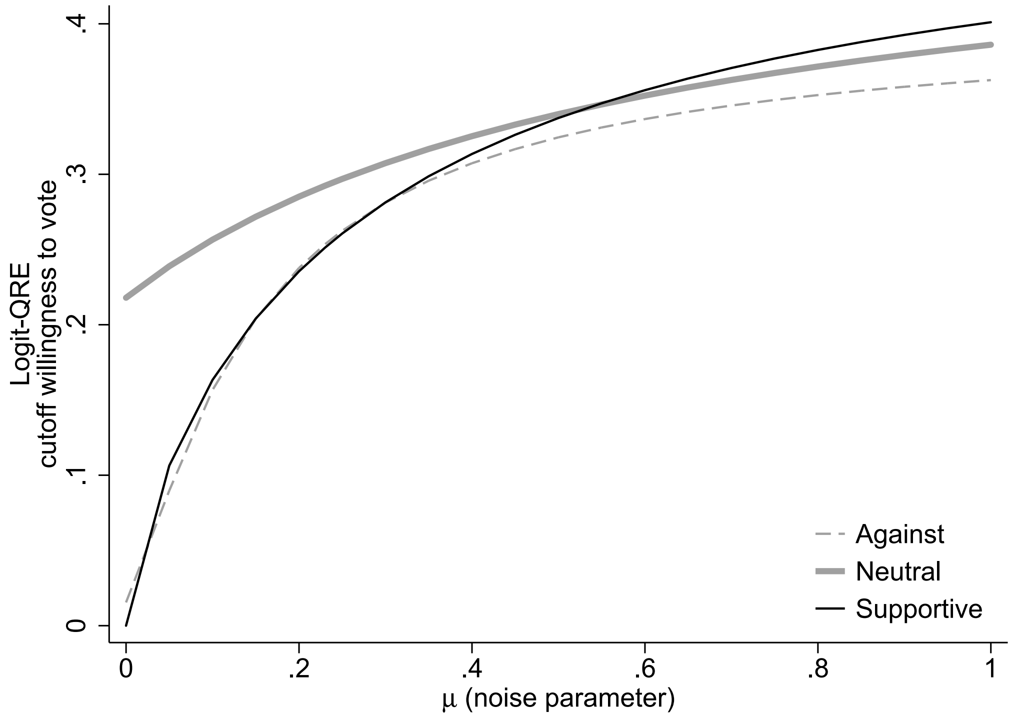

Figure A3: Symmetric Logit-QRE cutoff willingness to vote by μ (Treatment Group)

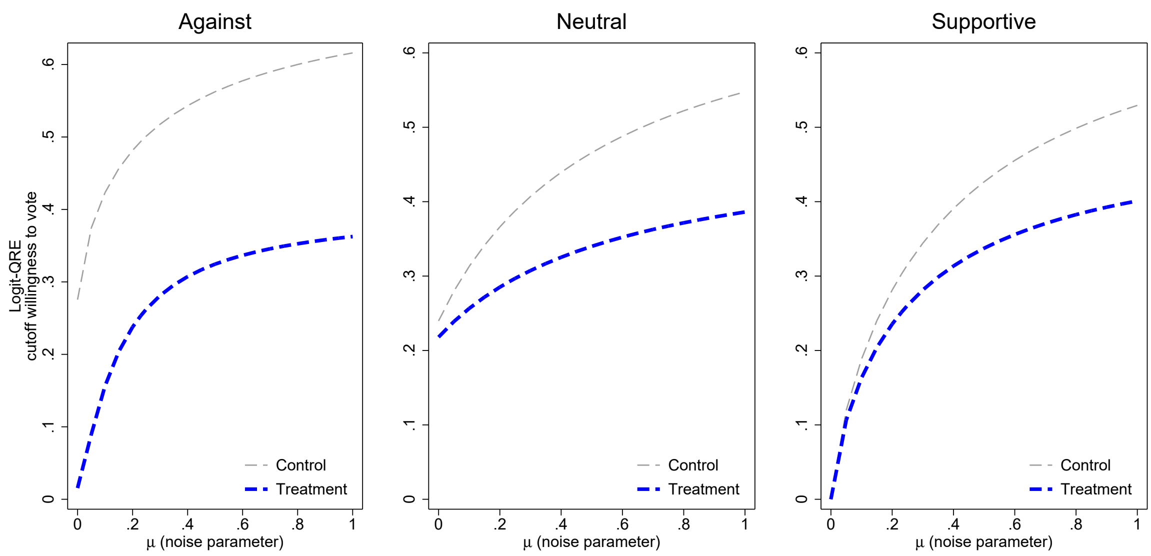

Figure A4: Symmetric Logit-QRE cutoff willingness to vote by μ and Partisan Bias

Finally, we compute the value of the weight on logistic noise, μ, that best fits our experimental data via maximum likelihood, and denote the estimated value by $\hat{\mu}$. To avoid issues related to over-fitting, we conduct the maximum likelihood estimation for the pooled data of all treatments and all levels of the partisan bias, each weighted by one-sixth, for all thirty periods of each session. We find that the value $\hat{\mu}^* = 0.23$ provides the best match between the symmetric quantal-response equilibrium and our experimental data. Then, following standard practice (e.g., Goeree, Holt and Palfrey 2016), we use the estimate of $\hat{\mu}^*$ we obtained from our data to derive point predictions of the symmetric quantal-response equilibrium vote cost cutoffs for that data. These results are reported in Table A5. Not surprisingly, the symmetric quantal-response equilibrium vote cost cutoff predictions, calibrated from our data, fit our data better than those produced by the symmetric Bayesian Nash equilibria of the pivotal voting model in the previous section, which were independent of our data. However, when we break down the analysis by the partisan bias, in 4 out of 6 cases, the predictions obtained by the quantal-response model are significantly different from the experimental observations. These results may suggest that a purely instrumental game-theoretical model, even with logistically distributed errors added to vote choice, does not fit the experimental results well.

Table A5: Logit-QRE Estimations ($\mu = 0.23$)

		Partisan Bias			
		Supportive	Neutral	Against	Average
Control Group	Obs.	21.3	46.4	51.5	40.3
	Est.	30.2	37.9	49.4	39.2
Treatment Group	Obs.	13.6	28.5	32.4	24.9
	Est.	25.1	29.2	25.3	26.6

I Utilitarian Welfare

Suppose that the utilitarian planner chooses a symmetric cost cutoff for voters to follow.

I.1 Treatment Group

In this case the utilitarian planner chooses a symmetric vote cost cutoff which depends on the number of expert voters, $M \leq N$ where N is the total number of voters, and the partisan bias, β. We denote the utilitarian's optimal cost cutoff by $\overline{c}_U^*(M;\beta)$.

For a supportive partisan bias, $\overline{c}_U^*(M;\beta=-2)=0$.

For a partisan bias of β, utilitarian welfare in the treatment group when the number of expert voters is M is

$$W_\beta^T(\overline{c};M) = N\left(\sum_{k=\beta+1}^{M}\binom{M}{k}\overline{c}^k(1-\overline{c})^{M-k} + \frac{1}{2}\binom{M}{\beta}\overline{c}^\beta(1-\overline{c})^{M-\beta}\right) - \frac{M}{2}\overline{c}^2.$$

When $\beta=0$, and since $N=5$, the problem is

$$\max_{\overline{c}\in(0,1]} 5(\sum_{k=1}^{M}\binom{M}{k}\overline{c}^k(1-\overline{c})^{M-k} + \frac{1}{2}\binom{M}{0}\overline{c}^0(1-\overline{c})^{M}) - \frac{M}{2}\overline{c}^2,$$

which can be simplified to

$$\max_{\overline{c}\in(0,1]} 5(1-\frac{1}{2}\binom{M}{0}\overline{c}^0(1-\overline{c})^{M} + \frac{1}{2}\binom{M}{0}\overline{c}^0(1-\overline{c})^{M}) - \frac{M}{2}\overline{c}^2,$$

which can be simplified to

$$\max_{\overline{c}\in(0,1]} 5(1-\frac{1}{2}\binom{M}{0}\overline{c}^0(1-\overline{c})^{M}) - \frac{M}{2}\overline{c}^2,$$

which can be simplified to

$$\max_{\overline{c}\in(0,1]} 5(1-\frac{1}{2}(1-\overline{c})^{M}) - \frac{M}{2}\overline{c}^2.$$

The first derivative of this function

$$\frac{5}{2}M(1-\overline{c})^{M-1} - M\overline{c},$$

so the first-order condition is

$$\frac{5}{2}(1-\overline{c})^{M-1} = \overline{c}. \tag{5}$$

If $M = 1$, this first-order condition reduces to

$$\frac{5}{2} = \overline{c},$$

and so

$$\overline{c}_U^*(1; \beta = 0) = \min\{\tfrac{5}{2}, 1\},$$

implying that $\overline{c}_U^*(1; \beta = 0) = 1$.

Case $M = 2$ then becomes

$$\frac{5}{2}(1-\overline{c}) = \overline{c},$$

which solving for $\overline{c}$ gives

$$\frac{7}{2}\overline{c} = \frac{5}{2},$$

which rearranges to

$$\overline{c} = \frac{5}{7}.$$

Cases $M = 3, 4, 5$ follow by substituting for M and solving (5) for $\overline{c}$ in each case.

When the partisan bias is against, it is obvious that $\overline{c}_U^*(1; \beta = -2) = 0$, so we can focus on $M = 2, 3, 4, 5$. For $M = 2$, because voters can only trigger a tie, the utilitarian planner's problem reduces to

$$\max_{\overline{c}\in[0,1]} N\overline{c}^2 - \overline{c}^2,$$

which is linear and increasing in c^2, so $\overline{c}_U^*(2; \beta = 2) = 1$. Finally, for $M = 3, 4, 5$, the utilitarian planner's problem is

$$\max_{\overline{c}\in[0,1]} N\left(\sum_{k=\beta+1}^{M} \binom{M}{k}\overline{c}^k(1-\overline{c})^{M-k} + \frac{1}{2}\binom{M}{\beta}\overline{c}^{\beta}(1-\overline{c})^{M-\beta}\right) - \frac{M}{2}\overline{c}^2.$$

Using standard results (e.g., Hartley and Fitch 1951), we can write

$$\sum_{k=\beta+1}^{M} \binom{M}{k} \overline{c}^k (1-\overline{c})^{M-k} = \frac{\int_0^{\overline{c}} t^{\beta}(1-t)^{M-\beta-1}dt}{\int_0^1 t^{\beta}(1-t)^{M-\beta-1}dt},$$

which using the Leibniz integral rule, the first-order condition is

$$\frac{\overline{c}^{\beta}(1-\overline{c})^{M-\beta-1}}{\int_0^1 t^{\beta}(1-t)^{M-\beta-1}dt} + \frac{1}{2}\binom{M}{\beta}\overline{c}^{\beta-1}(1-\overline{c})^{M-\beta-1}[\beta(1-\overline{c}) - (M-\beta)\overline{c}] = \frac{M}{N}\overline{c},$$

which reduces to

$$\frac{\overline{c}^{\beta}(1-\overline{c})^{M-\beta-1}}{\int_0^1 t^{\beta}(1-t)^{M-\beta-1}dt} + \frac{1}{2}\binom{M}{\beta}\overline{c}^{\beta-1}(1-\overline{c})^{M-\beta-1}[\beta - M\overline{c}] = \frac{M}{N}\overline{c}.$$

Using that the beta function is

$$B(x,y) = \int_0^1 t^{x-1}(1-t)^{y-1}dt,$$

and thus

$$\int_0^1 t^{\beta}(1-t)^{M-\beta-1}dt = B(\beta+1, N-\beta),$$

and using that $B(x+1,y) = B(x,y)\frac{x}{x+y}$, we can write

$$B(\beta+1, M-\beta) = B(\beta, M-\beta)\frac{\beta}{\beta+M-\beta} = B(\beta-1, M-\beta)\frac{\beta}{M}\cdot\frac{\beta-1}{\beta-1+M-\beta},$$

and since $\beta = 2$, we have that

$$B(3, M-\beta) = B(1, M-2)\cdot\frac{2}{M}\cdot\frac{2}{M-1},$$

and finally since $B(1,x) = \frac{1}{x}$,

$$B(3, M-\beta) = \frac{2}{M(M-1)(M-2)}.$$

Thus, since $N = 5$ we can write the planner's first-order condition as

$$5(M-1)(M-2)\overline{c}^2(1-\overline{c})^{M-3} + 5\binom{M}{2}\overline{c}(1-\overline{c})^{M-3}[2 - M\overline{c}] = \overline{c}. \tag{6}$$

Solving (6) yields the utilitarian planner's symmetric vote cost cutoffs when the partisan bias is against and $M = 3, 4, 5$.

I.2 Control Group

In this case the utilitarian planner chooses a symmetric vote cost cutoff which depends only on the partisan bias, β, which we denote by $\overline{c}_U^*(\beta)$.

For a supportive partisan bias, $\overline{c}_U^*(\beta = -2) = 0$.

For a partisan bias of β, utilitarian welfare in the control group is

$$W_\beta^C(\overline{c}) = N\left(\sum_{k=\beta+1}^{N}\binom{N}{k}(0.7\overline{c})^k(1-0.7\overline{c})^{N-k} + \frac{1}{2}\binom{N}{\beta}(0.7\overline{c})^\beta(1-0.7\overline{c})^{N-\beta}\right) - \frac{N}{2}0.7\overline{c}^2.$$

For a neutral partisan bias, we can write the utilitarian planner's problem as

$$\max_{\overline{c}\in[0,1]} N\left(1 - \frac{1}{2}\binom{N}{0}(0.7\overline{c})^0(1-0.7\overline{c})^N\right) - \frac{N}{2}0.7\overline{c}^2,$$

which simplifies to

$$\max_{\overline{c}\in[0,1]} N\left(1 - \frac{1}{2}(1-0.7\overline{c})^N\right) - \frac{N}{2}0.7\overline{c}^2.$$

The first-order condition is

$$N(1-0.7\overline{c})^{N-1} = 0.7\overline{c},$$

which since $N = 5$ is

$$5(1-0.7\overline{c})^4 = 0.7\overline{c}. \tag{7}$$

A solution to (7) gives $\overline{c}_U^*(\beta = 0)$.

When the partisan bias is against, and using the same techniques as above, we write the first-order condition as

$$\frac{0.7N(N-1)(N-2)(0.7\overline{c})^2(1-0.7\overline{c})^{N-3}}{2} + 0.35\binom{N}{2}\overline{c}(1-\overline{c})^{N-3}[2-0.7N\overline{c}] = 0.7N\overline{c},$$

which since $N = 5$, is

$$10.5(0.7\overline{c})^2(1-0.7\overline{c})^2 + 0.35\binom{5}{2}\overline{c}(1-\overline{c})^2(2-3.5\overline{c}) = 3.5\overline{c}. \tag{8}$$

A solution to (8) gives $\overline{c}_U^*(\beta = 2)$.

I.3 Comparisons

Table A6 summarizes the utilitarian planner's symmetric cost vote cutoff prescription in each case.

Table A6: Optimal Willingness to Vote by Treatment and Partisan Bias

		Partisan Bias		
		Supportive	Neutral	Against
Control Group	Exp. Obs.	21.3	46.4	51.5
	Pred.	0	48.2	96.2
Treatment Group	Exp. Obs.	13.6	28.5	32.4
	Pred. $M = 1$	0	100	0
	Pred. $M = 2$	0	71.4	100.0
	Pred. $M = 3$	0	53.7	100.0
	Pred. $M = 4$	0	44.0	86.7
	Pred. $M = 5$	0	37.7	67.5

We use both Mann-Whitney sign-rank test and the exact Fisher-Pitman permutation test to examine whether the experimental observations are different from the utilitarian planner's prescription. We use the electorate averages as the unit of independent observation in the statistical analysis to examine the difference on the quantities of interest. The numbers reported in Table A7 are p-values of statistical tests. All statistical tests reported in this table are two-sided and non-parametric.

Table A7: Comparisons between Experimental Results and Utilitarian Planner's Prescriptions

Partisan Bias	Statistical tests	Control Group	Treatment Group				
			M=1	M=2	M=3	M=4	M=5
Panel A: Willingness to Vote							
Against	Wilcoxon signed-rank	0.002	0.166	0.012	0.003	0.002	0.012
	Fisher-Pitman permutation	0.000	0.500	0.008	0.001	0.000	0.008
	# observations	12	3	8	11	11	7
Neutral	Wilcoxon signed-rank	0.938	0.109	0.008	0.002	0.003	0.037
	Fisher-Pitman permutation	0.825	0.250	0.004	0.000	0.001	0.033
	# observations	12	3	9	12	12	10
Supportive	Wilcoxon signed-rank	0.002	n/a	0.006	0.002	0.002	0.003
	Fisher-Pitman permutation	0.000	n/a	0.004	0.000	0.000	0.001
	# observations	12	1	10	12	12	11
Panel B: The Quality of Democratic Choice							
Against	Wilcoxon signed-rank	0.002	n/a	0.012	0.003	0.006	0.311
	Fisher-Pitman permutation	0.000	n/a	0.008	0.001	0.003	0.313
	# observations	12	3	8	11	11	7
Neutral	Wilcoxon signed-rank	0.695	0.109	0.008	0.002	0.003	0.017
	Fisher-Pitman permutation	0.492	0.250	0.004	0.000	0.001	0.018
	# observations	12	3	9	12	12	10
Supportive	Wilcoxon signed-rank	n/a	n/a	n/a	n/a	n/a	n/a
	Fisher-Pitman permutation	n/a	n/a	n/a	n/a	n/a	n/a
	# observations	12	1	10	12	12	11

Note: When the partisan bias is supportive, the achieved quality of democratic choice is always 100%, which is identical to the utilitarian planner's prescription, so statistical tests are not applicable in these cases. In the Treatment Group, when the partisan bias is against and there is only one expert voter, the quality of democratic choice is zero, which is identical to the utilitarian planner's prescription, so statistical tests are not applicable in this case. When the partisan bias is supportive and there is only one expert voter, we have only one observation, so statistical tests are omitted for this case.

References

Fischbacher, Urs. 2007. "z-Tree: Zurich toolbox for ready-made economic experiments." *Experimental Economics* 10(2):171–178.

Frederick, Shane. 2005. "Cognitive reflection and decision making." *Journal of Economic Perspectives* 19(4):25–42.

Goeree, Jacob K and Charles A Holt. 2005. "An explanation of anomalous behavior in models of political participation." *American Political Science Review* pp. 201–213.

Goeree, Jacob K, Charles A Holt and Thomas R Palfrey. 2016. *Quantal Response Equilibrium: A Stochastic Theory of Games*. Princeton University Press.

Hartley, HO and ER Fitch. 1951. “A chart for the incomplete beta-function and the cumulative binomial distribution.” *Biometrika* 38(3/4):423–426.

Riker, William H and Peter C Ordeshook. 1968. “A Theory of the Calculus of Voting.” *The American political science review* 62(1):25–42.