Journal of Economic Behavior & Organization 174 (2020) 332–359

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

Ignorance and bias in collective decisions☆

Alexander Elbittar[a,*], Andrei Gomberg[b], César Martinelli[c], Thomas R. Palfrey[d]

[a] CIDE, Mexico
[b] Instituto Tecnológico Autónomo de México, Mexico
[c] George Mason University, United States
[d] California Institute of Technology, United States

ARTICLE INFO

Article history:
Received 15 February 2016
Received in revised form 8 December 2016
Accepted 9 December 2016
Available online 27 December 2016

JEL classification:
D72
D83

Keywords:
Condorcet jury theorem
Rational ignorance
Noisy priors

ABSTRACT

We study theoretically and experimentally committee decision making with common interests. Committee members do not know which of two alternatives is optimal, but each member can acquire a private costly signal before casting a vote under either majority or unanimity rule. In the experiment, as predicted by Bayesian equilibrium, voters are more likely to acquire information under majority rule, and vote strategically under unanimity rule. As opposed to Bayesian equilibrium predictions, however, many committee members vote when uninformed. Moreover, uninformed voting is strongly associated with a lower propensity to acquire information. We show that an equilibrium model of subjective prior beliefs can account for both these phenomena, and provides a good overall fit to the observed patterns of behavior both in terms of rational ignorance and biases.

> If the probability that each voter is correct is larger than $\frac{1}{2}$, that is, if each voter is more likely to be right than wrong, then as the number of voters increases, the probability of the correct decision will increase: at the limit it will be certain [. . .]
>
> An assembly that is too large cannot be composed of enlightened men; it is likely that the members will be carried by ignorance and prejudices.
>
> Condorcet (1785) [1986, pp. 29–30]

☆ The authors thank CONACYT (grant 61501) and Asociación Mexicana de Cultura for financial support. Palfrey is grateful for financial support from the NSF (SES-1426560), the Gordon and Betty Moore Foundation (SES-1158), and the Russell Sage Foundation for a visiting scholar grant. The paper also benefited from comments and suggestions from audiences at the 2012 Latin American Meetings of the Econometric Society at University of the Pacific, Lima, the University of Technology, Sydney, the Department of Industrial Engineering at University of Chile, Cornell University, University of Cambridge, George Mason University, Columbia University, the 2014 Santa Barbara Conference on Experimental and Behavioral Economics, the 2014 Wallis Political Economy Conference at the University of Rochester, the 2015 Behavioral Political Economy Workshop at NYUAD, and the 2016 Behavioral Models of Politics Conference at the University of Pittsburgh. The authors especially thank Navin Kartik and Sourav Bhattacharya for detailed comments. Kirill Pogorelskiy provided excellent research assistance.

* Corresponding author.
E-mail addresses: alexander.elbittar@cide.edu (A. Elbittar), gomberg@itam.mx (A. Gomberg), cmarti33@gmu.edu (C. Martinelli), trp@hss.caltech.edu (T.R. Palfrey).

https://doi.org/10.1016/j.jebo.2016.12.011

1. Introduction

The idea that a committee or a jury can make better choices than a single individual, by aggregating the information dispersed among the group members, was first given a statistical foundation by Condorcet (1785), and has been very influential in social choice and in democratic theory, providing an epistemic foundation for the use of majority rule in a variety of contexts. During the last two decades, Condorcet's *jury theorem* has been studied from a game-theoretic viewpoint, starting with the pioneering work of Austen-Smith and Banks (1996). The game theoretic approach has led to some valuable insights about strategic voting behavior and the comparative performance of different voting rules in terms of information aggregation and efficiency. In particular, Feddersen and Pesendorfer (1996) have shown that whenever some voters are more informed than others, less informed voters have an incentive to abstain, effectively delegating the decision to informed voters, an effect dubbed *the swing voter's curse*. Moreover, when the voting rule is biased toward one of the alternatives, as is the case of unanimity rule, strategic voters may try to offset the built in bias by voting for the other alternative (Feddersen and Pesendorfer, 1998).

An additional layer of difficulty for information aggregation in committees occurs when heterogeneity in information is endogenous. If information is costly, group members may attempt to free ride on others. Moreover, realizing that their individual influence on the collective decision is small, group members may allow themselves to be carried away by prejudice. As a result, group members' opinions may actually contain little information about the alternatives, weakening the aggregation result. Indeed, as the quotation above makes clear, Condorcet was aware of the possibility of both ignorance and biased judgment clouding the opinion of jury members.

In this paper, we investigate theoretically and experimentally the problem of information aggregation in committees where information is costly, thus providing incentives for group members to attempt to free ride on others. In particular, we explore a version of the model developed by Martinelli (2006), in which each committee member is allowed to obtain costly private information about which of two alternatives is best for the group, with the individual cost of information being a privately observed random variable. After each committee member independently decides whether or not to acquire information, they vote in favor of either alternative, or abstain.[1] Because information costs are heterogeneous, only voters with low costs would acquire information, with the individual threshold for information acquisition depending on the voting rule and declining with the committee size. Thus the model generates clear predictions about the effects of group size and the voting rule on information aggregation and efficiency. We consider two voting rules here, *simple majority*, the classical setting for the analysis of information aggregation in committees, and *unanimity*, which is known to make strategic behavior more involved (Feddersen and Pesendorfer, 1998). Under majority, equilibrium behavior is characterized by a cost cutoff: if a committee member's cost of information acquisition falls below the cutoff, the individual acquires information, and votes according to the signal received. If instead a committee member's cost of information acquisition falls above the cutoff, the individual does not acquire information and abstains, as predicted by the swing voter's curse. Under unanimity, equilibrium behavior is also characterized by a cost cutoff, with individuals trying to offset the rule bias by either abstaining with some probability rather than casting a vote for the status quo when obtaining information favoring the status quo, or by voting against the status quo with some probability when uninformed.

We conducted laboratory experiments based on this model at Instituto Tecnológico Autónomo de México in Mexico City. The experiments involved four treatments, distinguished by committee size (three or seven subjects) and voting rule (majority or unanimity). In all treatments, the value of a correct decision, the informativeness of the signal, and the distribution of information costs were held constant, in order to deliver sharp comparative static predictions. Consistent with equilibrium, we find that there is more information acquisition under majority rule than under unanimity rule. Moreover, individuals attempt to counter the built-in bias in favor of the status quo under unanimity rule. In sharp contrast with equilibrium predictions, however, we find that uniformed individuals persistently cast votes – sometimes even in favor of the alternative favored by the voting rule under unanimity. There is, in fact, substantial heterogeneity in behavior, with some voters being very likely to acquire information, and preferring to abstain while uninformed, and others being very unlikely to acquire information, and usually casting an uninformed vote.

Our experimental results are surprising given the experimental support for equilibrium predictions when information is exogenous obtained by Guarnaschelli et al. (2000) and Battaglini et al. (2008, 2010). In addition, Goeree and Yariv (2010) have found that behavior in the lab under different voting rules tracks theoretical predictions in the jury setting without communication. Bhattacharya et al. (2014) provide further experimental support that voters adapt their behavior to institutions in accordance with game theoretic predictions.[2]

[1] There is a large and growing literature in political economy that explores a broad range of theoretical questions related to endogenous information acquisition in Condorcet jury environments. Besides Martinelli (2006), this includes contributions by Persico (2004), Mukhopadhaya (2005), Martinelli (2007), Gerardi and Yariv (2008), Gershkov and Szentes (2009), Koriyama and Szentes (2009), Oliveros (2013), Triossi (2013), and Tyson (2016).

[2] To the best of our knowledge, this paper provides the first experimental work on information acquisition in committees, together with the work of Grosser and Seebauer (2016), which originated independently from ours. Our work is different from theirs in that they focus on the difference between compulsory and voluntary voting, while we attempt to explain the patters of behavior under voluntary voting, and compare different voting rules. It is important to note that Grosser and Seebauer observe similar patterns of behavior in terms of uninformed voting when information is costly. Bhattacharya (personal communication) reports a similar pattern of uninformed voting when information is costly in a recent laboratory experiment conducted with John Duffy and SunTak Kim.

Current behavioral theories do not seem able to account for the puzzling behavior observed in the experiment. "Cursed" voters, that is voters who ignore the informational content of other voters' actions (as in Eyster and Rabin, 2005), would be indifferent between abstaining or voting in case of being uninformed, so they could account for uninformed voting. But, as we explain later on, cursed voters would be willing to acquire even more information, at higher costs, than voters who are not cursed. Intuitively, they would not understand that they could free-ride on information acquisition by other voters. This contradicts our finding that voters acquire less information than standard game theory predicts. "Loss aversion" (Kahneman and Tversky, 1983) also fails to account for our findings. While loss averse voters would be less willing than rational voters to acquire costly information, as we observe, such voters would not be willing to vote if uninformed. We discuss the predictions of these and other behavioral theories, including "level-k" models, expressive voting, and ethical voting, in a separate section.

Motivated by the experimental results, we propose an alternative behavioral theory, *subjective beliefs equilibrium*. The proposed theory (1) postulates that some individuals are *biased*, i.e. hold prior beliefs that are biased in favor of one or the other alternative, and (2) allows individuals to make random mistakes. Biased individuals can be interpreted as attributing informational content to hunches or guesses, even though – in the spirit of agreeing to disagree – they are aware that some other individuals may not. In the spirit of quantal response equilibrium (McKelvey and Palfrey, 1995, 1998), individuals are assumed to know that other voters make random errors. The subjective beliefs equilibrium model gives us enough behavioral freedom to perform a structural estimation using the experimental data with two parameters, the fraction of biased voters and the probability of random mistakes.

The notion of subjective beliefs equilibrium has a precedent in the experimental literature, introduced as "homemade priors" in Camerer and Weigelt (1988) to explain deviations from sequential equilibrium predictions in a reputation formation game. The random beliefs equilibrium concept, introduced by Friedman and Mezzetti (2005) in the context of finite normal form games is also related to our definition, though in their case random beliefs occur with respect to others' strategy choices. Heterogeneous priors have been employed before in the context of information aggregation by Che and Kartik (2009) and Sethi and Yildiz (2012, 2016). In our context, we believe they emerge as a form of overconfidence in the ability of the subjects to predict the state of the world, perhaps making inferences on the basis of previous observations which are unwarranted due to independence. There is a large literature in psychology and economics suggesting that people are overconfident about what they know. Nyarko and Schotter (2002), for instance, elicited extreme beliefs about the opponent's next move in an asymmetric matching pennies game in which in fact "true" frequencies of other subjects' play were close to 50/50. (See also Palfrey and Wang, 2009.)[3]

Subjective beliefs equilibrium delivers predictions that are consistent with the observed behavior at the lab. In particular, under either voting rule, sufficiently biased individuals will vote without acquiring information. Moreover, compared to the standard model without biased voters, the introduction of biased committee members makes unbiased committee members more willing to acquire information, but reduces the overall acquisition of information and the probability of making the correct decision under either voting rule, as observed in the data. Subjective beliefs equilibria also reconcile our empirical findings with the earlier successes of the swing voter's curse in the lab. In our experiment, as opposed to those conducted by Battaglini et al. (2008, 2010), signals acquired by voters are not fully revealing, and there is strategic uncertainty about whether other voters have acquired information. Under those circumstances, it is less apparent to uninformed voters, who may believe other voters are likely to be uninformed, that indulging in guessing behavior is detrimental to the group.[4]

The estimated subjective beliefs equilibrium fits the empirical distribution of strategies in the lab, yielding relatively similar parameters across treatments. The estimated probability that a voter is biased is about 40% in three of the four treatments, and the probability that a voter makes a random mistake is about 20–25% in all treatments. We then apply the results of our estimation to conduct a classification analysis of individual subject behavior. Using a 95% confidence interval, 96% of individuals are classified as either biased or unbiased. Again, we obtain that the probability that an individual is biased is 40%. The similarity in the results across treatments is quite remarkable, given that the variation in voting rules and committee sizes delivers very different equilibrium behavior.

The remainder of this paper is organized as follows. Section 2 presents the theoretical model and predictions. In this section, Bayesian equilibrium is treated as a special, extreme case of the subjective beliefs equilibrium in which the distribution of priors is degenerate and gives probability one to the unbiased priors. Section 3 explains the experimental design and hypotheses. Section 4 describes the experimental results and the structural estimation. Section 5 discusses alternative behavioral theories. Section 6 concludes. An appendix presents some of the proofs and provides a translation of experimental instructions.

---

[3] Idiosyncratic departures from a common symmetric priors can be replaced by small random payoff perturbations in the definition of the subjective beliefs equilibrium concept. Such payoff perturbations have the disadvantage of not having any obvious source (as opposed to belief perturbations being caused by overconfidence). Of course, there is a long tradition of considering payoff perturbations, under the guise of probabilistic voting, in explaining voting behavior. "Magical reasoning" a la Daley and Sadowski (2016) could also account for some of the anomalous features of the data; in particular, if some voters believe their votes can influence the state of the world with some small probability.

[4] In the context of a very different experiment involving sequential voting, Esponda and Vespa (2014) argue that it is particularly difficult for subjects to update beliefs conditioning on being pivotal.

2. The model

2.1. Basics

We consider a committee with $n \geq 2$ members which must choose between two alternatives, A and B. There are two possible states of the world, ω_A and ω_B. Each committee member receives a payoff of 1 if the committee reaches the decision A and the state of the world is ω_A, or if the committee reaches the decision B and the state of the world is ω_B, and a payoff of 0 otherwise.

Both states of the world are equally likely, and committee members do not know which state obtains. Each committee member, however, may choose to acquire some costly information. The cost at which information may be acquired is independently and identically distributed across voters according to a distribution function F, which is strictly increasing and continuously differentiable over the interval $[0, \bar{c}]$ for some $\bar{c} > 0$, with $F(0) = 0$, $F'(0) > 0$, and $F(\bar{c}) = 1$. After observing their idiosyncratic cost of information acquisition, each committee member decides whether to privately acquire information or not. Each committee member who acquires information receives a private signal $s \in \{s_A, s_B\}$. Conditional on the state of the world, private signals are independently and identically distributed across voters. The probability or receiving signal s_d in state ω_d is equal to $1/2 + q$ for $d \in \{A, B\}$, where $q \in (0, 1/2]$.

After the information acquisition stage, the committee votes over the two alternatives. A committee member may vote for A, vote for B, or abstain. A voting rule,

$$V : \{0, \ldots, n\} \times \{0, \ldots, n\} \rightarrow [0, 1],$$

specifies a probability that the committee selects alternative A for any feasible combination of votes for A and votes for B, with alternative B being selected by the committee with the complementary probability. We consider two possible voting rules: simple majority and unanimity.

Under *simple majority*, V_M, the alternative with most votes is chosen, with ties broken by a fair coin toss. That is:

$$V_M(\nu^A, \nu^B) = \begin{cases} 1 & \text{if } \nu^A > \nu^B \\ 1/2 & \text{if } \nu^A = \nu^B \\ 0 & \text{if } \nu^A < \nu^B \end{cases} .$$

where ν^d denotes the number of votes for decision d.

Under *unanimity*, V_U, in our specification, A is chosen unless every vote that is cast favors B, with A being chosen if every member abstains. That is:

$$V_U(\nu^A, \nu^B) = \begin{cases} 0 & \text{if } \nu^B > 0 = \nu^A \\ 1 & \text{otherwise} \end{cases} .$$

Given a voter's cost of information c_i, the utility, U_i, of voter i net of information acquisition costs is given by:

$$U_i = \begin{cases} 1 - c_i & \text{if the decision is } d \text{ and the state is } \omega_d, \text{ for } d \in \{A, B\} \\ -c_i & \text{otherwise} \end{cases}$$

if the voter acquires information. If voter i does not acquire information, then

$$U_i = \begin{cases} 1 & \text{if the decision is } d \text{ and the state is } \omega_d, \text{ for } d \in \{A, B\} \\ 0 & \text{otherwise} \end{cases} .$$

2.2. Subjective beliefs equilibrium

We allow voters to (1) hold private noisy prior beliefs that can deviate from the correct prior probability of each state, and (2) make random mistakes.

With respect to (1), we let the prior belief of voter i that the state of the world is ω_A be given by $1/2 + \epsilon_i$, where $\epsilon_i = 0$ with probability $1 - p$, $\epsilon_i = \beta$ with probability $p/2$, and $\epsilon_i = -\beta$ with probability $p/2$, where $p \in [0, 1)$ and $\beta \in (0, 1/2]$. That is, the voter is *biased* with probability p and *unbiased* with probability $1 - p$.

A voter's *type* is a triple $t = (\epsilon, c, s)$ specifying prior beliefs, the private cost of information acquisition, and a private signal, where we denote "no signal" by s_0. For a given voter, an *action* is a pair $a = (\iota, \nu)$, $\iota \in \{1, 0\}$, $\nu \in \{A, B, \phi\}$, indicating whether or not the voter acquires information in the first stage, and whether the voter casts a vote for alternative A, for alternative B, or abstains in the second stage.

A *strategy* for voter i is a pair of measurable mappings, $\sigma = (\sigma^\iota, \sigma^\nu)$, where σ^ι specifies the information acquisition decision as a function of the voter's type, and σ^ν specifies the (possibly mixed) voting decision as a function of the voter's type. With slight abuse of terminology, we denote by $\sigma(a|t) = (\sigma^\iota(\iota|t), \sigma^\nu(\nu|t))$ the probability of action $a = (\iota, \nu)$ if the voter's type is t.

We call a strategy *informative* if σ^ι puts positive probability on the set of actions such that $\iota = 1$, and *uninformative* otherwise. A *strategy profile* is a vector $(\sigma_1, \ldots, \sigma_n)$ that assigns to each voter $i = 1, \ldots, n$ a strategy σ_i.

With respect to (2), we assume that a voter chooses an action according to σ with probability $Q \in (0, 1]$, and randomizes uniformly over actions with probability $1 - Q$. That is, they become informed with probability 1/2, and they vote for A, for B, or abstain, with probability 1/3 each, regardless of whether they are informed or not, and regardless of the signal received. In the spirit of quantal response equilibrium, we assume that subjects are aware that other subjects, as well as themselves, make a mistake with probability $1 - Q$.

Given a strategy profile $(\sigma_1, \ldots, \sigma_n)$, let $EU_i(\sigma_1, \ldots, \sigma_n|\omega_d)$ be the expected utility of voter i in state ω_d.[5] Then the ϵ*-subjective expected utility* of voter i is equal to

$$\left(\frac{1}{2} + \epsilon\right) EU_i(\sigma_1, \ldots, \sigma_n|\omega_A) + \left(\frac{1}{2} - \epsilon\right) EU_i(\sigma_1, \ldots, \sigma_n|\omega_B).$$

We say that σ_i is a *subjective best response* to the strategies of other voters if for every realization ϵ of voter's i prior beliefs, σ_i maximizes the ϵ-subjective expected utility of voter i.

A voter playing a subjective best-response realizes that other voters behavior is influenced by their own noisy priors, but – in the spirit of agreeing to disagree – does not draw inferences from the priors held by other voters. In particular, a voter playing a subjective best-response is not "cursed," since the voter recognizes that the behavior of other voters depends on the state of the world.

A *subjective beliefs equilibrium* is a strategy profile such that for each voter i, σ_i is a subjective best response; that is, σ_i maximizes the subjective expected utility of voter i given the strategies of other voters and given voter i's prior beliefs about the states. We restrict attention to symmetric informative equilibrium, where a *symmetric* equilibrium is an equilibrium such that every voter uses the same strategy.

Note that if $p = 0$ and $Q = 1$, all voters have correct prior beliefs and make no mistakes, so the subjective equilibrium reduces to the standard Bayesian equilibrium for a common prior belief of 1/2. Looking ahead, in terms of the estimation we will assume that β is large enough so that biased voters will prefer not to acquire information and instead vote for the alternative selected by their priors. Intuitively, in terms of observed behavior, the presence of biased voters will be reflected by the actions of not acquiring information and voting for either alternative being more frequent than predicted by equilibrium augmented by uniform mistakes.

For any given voter, let $x = (x_A, x_B, x_\phi) \in \mathbb{N}^3$ represent the vote profile of other voters, that is the number of votes cast by other voters in favor of A, B, and abstention. From the perspective of each voter, this is the realization of a random vector with a probability distribution that depends on the strategy profile of other voters, the distribution of priors, and the state of the world. Given a voting rule, a voter is *decisive* at x if the committee decision may be different depending on whether the voter votes for A, B, or abstains. As it is well-understood, a best responding voter needs to be concerned only with vote profiles such that the voter is decisive. We next characterize symmetric informative equilibrium under simple majority and unanimity rules.

2.3. Simple majority

Under simple majority, a voter is decisive only if the difference between the number of votes cast by other voters in favor of each of the alternatives is zero or one. In particular, for a given voter i, let $D(z|\sigma_{-i}, \omega)$ be the probability that the difference between the number of votes for A and for B cast by other voters is equal to z when the strategy profile of other voters is σ_{-i} and the state of the world is ω. If the difference is zero, voting for one alternative rather than abstaining increases the probability of that alternative winning the election from 1/2 to 1. If the difference is one, voting for the alternative that is behind rather than abstaining increases the probability of that alternative winning the election from 0 to 1/2.

If the voter with prior ϵ acquires information, the difference in interim expected utility between voting for A and abstaining after observing signal s_A is:

$$\begin{aligned} G(s_A|\epsilon, \sigma_{-j}) \equiv \; & \frac{1}{2}\left(\frac{1}{2} + \epsilon\right)\left(\frac{1}{2} + q\right)(D(0|\sigma_{-j}, \omega_A) + D(-1|\sigma_{-j}, \omega_A)) \\ & - \frac{1}{2}\left(\frac{1}{2} - \epsilon\right)\left(\frac{1}{2} - q\right)(D(0|\sigma_{-j}, \omega_B) + D(-1|\sigma_{-j}, \omega_B)). \end{aligned}$$

[5] The dependence of EU_i on Q is understood. Except where necessary, the dependence on Q in these and similar expressions is suppressed to reduce notation.

Similarly, the difference in expected utility between voting for B and abstaining after observing signal s_B is:

$$G(s_B|\epsilon,\sigma_{-i}) \equiv -\frac{1}{2}\left(\frac{1}{2}+\epsilon\right)\left(\frac{1}{2}-q\right)(D(0|\sigma_{-i},\omega_A)+D(1|\sigma_{-i},\omega_A)) + \frac{1}{2}\left(\frac{1}{2}-\epsilon\right)\left(\frac{1}{2}+q\right)(D(0|\sigma_{-i},\omega_B)+D(1|\sigma_{-i},\omega_B)).$$

If the voter does not acquire information, the difference in expected utility between voting for A and abstaining is:

$$G(A|\epsilon,\sigma_{-i}) \equiv \frac{1}{2}\left(\frac{1}{2}+\epsilon\right)(D(0|\sigma_{-i},\omega_A)+D(-1|\sigma_{-i},\omega_A)) - \frac{1}{2}\left(\frac{1}{2}-\epsilon\right)(D(0|\sigma_{-i},\omega_B)+D(-1|\sigma_{-i},\omega_B)).$$

Similarly, the difference in expected utility between voting for B and abstaining is

$$G(B|\epsilon,\sigma_{-i}) \equiv -\frac{1}{2}\left(\frac{1}{2}+\epsilon\right)(D(0|\sigma_{-i},\omega_A)+D(1|\sigma_{-i},\omega_A)) + \frac{1}{2}\left(\frac{1}{2}-\epsilon\right)(D(0|\sigma_{-i},\omega_B)+D(1|\sigma_{-i},\omega_B)).$$

It is never optimal for a voter to become informed and then vote against their signal. That is, if a voter acquires information, the voter will either vote for A or abstain in case of receiving signal s_A, and either vote for B or abstain in case of receiving signal s_B. Thus, the difference in expected utility between acquiring information and not, net of the cost of information acquisition, is

$$c(\epsilon,\sigma_{-i}) \equiv \max\{G(s_A|\epsilon,\sigma_{-i}), G(s_B|\epsilon,\sigma_{-i}), G(s_A|\epsilon,\sigma_{-i})+G(s_B|\epsilon,\sigma_{-i})\} - \max\{0, G(A|\epsilon,\sigma_{-i}), G(B|\epsilon,\sigma_{-i})\}.$$

From the preceding argument it follows that best-reply behavior has the familiar cutoff property: given any strategy profile of other voters, a best-responding voter only acquires information if the cost is below a well-defined cutoff value.

Lemma 1. *Under majority rule, voter i with priors given by ϵ plays a best response to σ_{-i} if for almost every c,*

(1) *if $c \le c(\epsilon,\sigma_{-i})$ then the voter acquires information, and after signal s_d votes for d if $G(s_d|\epsilon,\sigma_{-i}) > 0$ and abstains if $G(s_d|\epsilon,\sigma_{-i}) < 0$,*
(2) *if $c > c(\epsilon,\sigma_{-i})$, then the voter does not acquire information, and votes for d only if $G(d|\epsilon,\sigma_{-i}) = \max\{0, G(A|\epsilon,\sigma_{-i}), G(B|\epsilon,\sigma_{-i})\}$ and abstains only if $G(A|\epsilon,\sigma_{-i}) \le 0$ and $G(B|\epsilon,\sigma_{-i}) \le 0$.*

We say that a strategy σ is *neutral* if

$$\sigma((0,A)|(\epsilon,c,s_d)) = \sigma((0,B)|(-\epsilon,c',s_{d'}))$$

for all d, d' and almost all ϵ, c, c', and

$$\sigma((1,A)|(\epsilon,c,s_A)) = \sigma((1,B)|(-\epsilon,c',s_B))$$

and

$$\sigma((1,A)|(\epsilon,c,s_B)) = \sigma((1,B)|(-\epsilon,c',s_A)) = 0$$

for all ϵ and almost all c, c'. A voter who plays a neutral strategy does not discriminate between the alternatives except on the basis of the private signal and prior beliefs, and does not vote for one alternative if receiving a signal in favor of the other alternative.

Neutrality is a very natural restriction under majority rule. Intuitively, the voting rule itself does not discriminate between the alternatives, and biased voters have equal probabilities of being biased in favor of one or the other alternative, so an unbiased voters has no reason to discriminate between the alternatives. Similarly, random mistakes do not discriminate between the alternatives. Thus, a natural conjecture for a voter is that other voters play neutral strategies.

If every voter other than i plays a neutral strategy, since mistakes are uniform, it is straightforward that

$$D(0|\sigma_{-i},\omega_A) = D(0|\sigma_{-i},\omega_B)$$

and

$$D(1|\sigma_{-i},\omega_A) = D(-1|\sigma_{-i},\omega_B) \ge D(-1|\sigma_{-i},\omega_A) = D(1|\sigma_{-i},\omega_B),$$

where the inequality is strict if at least one player other than i plays an informative strategy. With a slight abuse of notation, we now write

$$D(0|\sigma_{-i}) \equiv D(0|\sigma_{-i},\omega_A), \quad D(1|\sigma_{-i}) \equiv D(1|\sigma_{-i},\omega_A) \quad \text{and} \quad D(-1|\sigma_{-i}) \equiv D(-1|\sigma_{-i},\omega_A)$$

to indicate the probability that the correct alternative is tied, one vote ahead or one vote behind when all other voters are using neutral, informative strategies given by σ_{-i}.

We have:

Proposition 1. *In any symmetric, neutral equilibrium, unbiased voters intend to acquire information if the cost is below some positive threshold*

$$c^* = q\left(D(0|\sigma^*) + \frac{1}{2}D(-1|\sigma^*) + \frac{1}{2}D(1|\sigma^*)\right), \tag{1}$$

and intend to abstain when uninformed, and vote for the alternative favored by their signal when informed. If the bias β is large enough, biased voters do not intend to acquire information, and intend to vote for the alternative favored by their bias.

Proof. See Appendix A.□

From Proposition 1, we get that in a symmetric, neutral equilibrium, if the bias is large enough, the probabilities of an abstention, a vote for the correct alternative, and a vote for the incorrect alternatives are, respectively,

$$v_0 = (1-p)Q(1-F(c^*)) + \frac{1}{3}(1-Q),$$

$$v_r = (1-p)QF(c^*)\left(\frac{1}{2}+q\right) + \frac{1}{2}pQ + \frac{1}{3}(1-Q),$$

$$v_w = (1-p)QF(c^*)\left(\frac{1}{2}-q\right) + \frac{1}{2}pQ + \frac{1}{3}(1-Q).$$

Note that

$$D(0|\sigma^*) = \sum_{j=0}^{\lfloor (n-1)/2 \rfloor} \binom{n-1}{2j}\binom{2j}{j}(v_r)^j(v_w)^j(v_0)^{n-2j-1},$$

$$D(1|\sigma^*) = \sum_{j=0}^{\lfloor (n-3)/2 \rfloor} \binom{n-1}{2j+1}\binom{2j+1}{j}(v_r)^{j+1}(v_w)^j(v_0)^{n-2j-2},$$

$$D(-1|\sigma^*) = \sum_{j=0}^{\lfloor (n-3)/2 \rfloor} \binom{n-1}{2j+1}\binom{2j+1}{j}(v_r)^j(v_w)^{j+1}(v_0)^{n-2j-2}.$$

We can calculate equilibrium thresholds for different values of p, Q, n, q and the function F using v_0, v_r, v_w and Eq (1).

2.4. An example under majority

Suppose $q = 1/6$, c is distributed uniformly in [0, 0.1] and $n = 3$ or $n = 7$, as in our lab experiments. As a benchmark, suppose $p = 0$ and $Q = 1$ as in a Bayesian equilibrium. Eq. (1) has a unique solution for either committee size, given by $c^* \approx 0.056$ for $n = 3$, and by $c^* \approx 0.039$ for $n = 7$. It follows that there is a unique symmetric, neutral Bayesian equilibrium in either case. The probability the committee makes the correct decision, given the equilibrium with cutoff c^*, is given by:

$$\sum_{j=0}^{\lceil n/2-1 \rceil}\sum_{k=j+1}^{n-j}\frac{n!}{j!k!(n-j-k)!}\left(\frac{1}{2}+q\right)^k\left(\frac{1}{2}-q\right)^j F(c^*)^{j+k}(1-F(c^*))^{n-j-k}$$

$$+\left(\frac{1}{2}\right)\sum_{j=0}^{\lfloor n/2 \rfloor}\frac{n!}{j!j!(n-2j)!}\left(\frac{1}{4}-q^2\right)^j F(c^*)^{2j}(1-F(c^*))^{n-2j}.$$

This probability is approximately 0.67 for $n = 3$ and 0.71 for $n = 7$.

Suppose now $p > 0$ and $Q = 1$. We can calculate the cutpoint c^* by solving Eq. (1). We also need to check the inequality conditions guaranteeing that biased voters do not acquire information and vote for the alternative favored by their bias. That is, $G(A|, \beta, \sigma^c) \geq 0$ and $G(s_A|, \beta, \sigma^c) - G(A|, \beta, \sigma^c) \leq 0$:

$$\beta D(0|\sigma^c) + \frac{1}{2}\left(\frac{1}{2}+\beta\right)D(-1|\sigma^c) - \frac{1}{2}\left(\frac{1}{2}-\beta\right)D(1|\sigma^c) \geq 0,$$

$$(q-\beta)D(0|\sigma^c) - \left(\frac{1}{2}+\beta\right)\left(\frac{1}{2}-q\right)D(-1|\sigma^c) + \left(\frac{1}{2}-\beta\right)\left(\frac{1}{2}+q\right)D(1|\sigma^c) \leq 0.$$

As an example, suppose $p = 1/2$ and $\beta = 1/5$, so that the prior beliefs that the state of the world is ω_A are either 3/10, 1/2 or 7/10 with probability 1/4, 1/2 and 1/4, respectively. For $n = 3$, we have $c^* \approx 0.076$, so the probability of information acquisition is approximately 0.38. For $n = 7$ we get $c^* \approx 0.048$, so the probability of information acquisition is approximately 0.24.

Columns 3M and 7M of Table 1A summarize the standard ($p = 0$, $Q = 1$) Bayesian equilibrium when all voters share correct prior beliefs. As a comparison, the corresponding columns in Table 1B summarize the subjective equilibrium when with

Table 1
Comparison of SBE with $p = 0$ and $p = 0.5$ without random mistakes ($Q = 1$), for the experimental treatments.

Treatment (size, rule)	3M	7M	3U		7U	
(A) Equilibrium for $p = 0$, $Q = 1$						
Predicted frequencies of individual decisions						
Info acquisition	0.56	0.39	0.46	0.44	0.25	0.22
Vote *A* if uninformed	0	0	0	0	0	0
Vote *B* if uninformed	0	0	0	[0.07,1]	0	[0.08, 1]
Abstain if uninformed	1	1	1	[0, 0.93]	1	[0, 0.92]
Vote *A* if signal s_A	1	1	0.5	1	0.45	1
Abstain if signal s_A	0	0	0.5	0	0.55	0
Vote *B* if signal s_B	1	1	1	1	1	1
Abstain if signal s_B	0	0	0	0	0	0
Predicted frequency of group decision						
Correct decision	0.67	0.71	0.64	0.63	0.64	0.63

Treatment (size, rule)	3M	7M	3U	7U
(B) Equilibrium for $p = 0.5$, $Q = 1$				
Predicted frequencies of individual decisions				
Info acquisition	0.38	0.24	0.22	0.08
Vote *A* if uninformed	0.40	0.33	0.32	0.27
Vote *B* if uninformed	0.40	0.33	[0.32, 0.68]	[0.27, 0.73]
Abstain if uninformed	0.19	0.34	[0, 0.36]	[0, 0.46]
Vote *A* if signal s_A	1	1	1	1
Abstain if signal s_A	0	0	0	0
Vote *B* if signal s_B	1	1	1	1
Abstain if signal s_B	0	0	0	0
Predicted frequency of group decision				
Correct decision	0.60	0.51	0.55	0.51

probability 1/2 each voter has incorrect prior beliefs ($p = 1/2$). For either committee size, introducing deviations from correct priors reduces the unconditional probability of information acquisition, even though it increases the probability of information acquisition conditional on holding unbiased prior beliefs. Introducing deviations from correct priors also increases the probability of uninformed voting, and reduces the probability of the group reaching the correct decision.

The probability of information acquisition is decreasing in the committee size, both in the standard Bayesian equilibrium and after introducing deviations from correct priors. In the former case, however, the probability of reaching the correct decision increases in the size of the committee, while in the latter it is decreasing. In fact, with seven subjects, after introducing biased individuals, collective choice is little better than a coin toss.

Under majority rule, the probability of voting for either alternative conditional on being uninformed if $Q = 1$, as in Table 1 is given by $0.5p/(1 - pF(c^*))$, while the probability of voting against one's signal, or of abstaining after getting a signal, is zero. If $Q < 1$, the probability of voting for alternative A if uninformed is $Q0.5p/(1 - (1 - p)F(c^*)) + (1 - Q)/6$, and the probability of voting against one's signal for A, or of abstaining after getting a signal for A, is $(1 - Q)/6$, and similarly for alternative B. Intuitively, the difference between the probability of voting for either alternative if uninformed and the probability of voting against one's signal is indicative of behavior that cannot be explained by random errors.

2.5. Unanimity

Under unanimity, given our definition of this rule, a voter is decisive if and only if either every other voter has abstained, or at least one voter has voted for B and no voter has voted for A. In the former case, a vote for A or an abstention leads to a decision in favor of A, and a vote for B leads to a decision in favor of B. In the latter case, a vote for A leads to a decision in favor of A, and an abstention or a vote for B lead to a decision in favor of B.

Let $P(0|\sigma_{-i}, \omega)$ be the probability that all other voters abstain given the strategy profile σ_{-i} of other voters and the state of the world ω. Similarly, let $P(1|\sigma_{-i}, \omega)$ be the probability that all other voters abstain or vote for B and at least one other voter votes for B given the strategy profile σ_{-i} of other voters and the state of the world ω.

If the voter with prior ϵ acquires information, the difference in interim expected utility between voting for A and abstaining after observing signal s_A is:

$$H(s_A|\epsilon, \sigma_{-i}) = \left(\frac{1}{2} + \epsilon\right)\left(\frac{1}{2} + q\right)P(1|\sigma_{-i}, \omega_A) - \left(\frac{1}{2} - \epsilon\right)\left(\frac{1}{2} - q\right)P(1|\sigma_{-i}, \omega_B).$$

Similarly, the difference in expected utility between voting for B and abstaining after observing signal s_B is:

$$H(s_B|\epsilon, \sigma_{-i}) = -\left(\frac{1}{2} + \epsilon\right)\left(\frac{1}{2} - q\right)P(0|\sigma_{-i}, \omega_A) + \left(\frac{1}{2} - \epsilon\right)\left(\frac{1}{2} + q\right)P(0|\sigma_{-i}, \omega_B).$$

If the voter did not acquire information, the difference in expected utility between voting for A and abstaining is

$$H(A|\epsilon,\sigma_{-i}) = \left(\frac{1}{2}+\epsilon\right)P(1|\sigma_{-i},\omega_A) - \left(\frac{1}{2}-\epsilon\right)P(1|\sigma_{-i},\omega_B).$$

Finally, if the voter has not bought information, the difference in expected utility between voting for B and abstaining is

$$H(B|\epsilon,\sigma_{-i}) = -\left(\frac{1}{2}+\epsilon\right)P(0|\sigma_{-i},\omega_A) + \left(\frac{1}{2}-\epsilon\right)P(0|\sigma_{-i},\omega_B).$$

As was the case with majority rule, it is easy to show that it is never optimal for a voter to acquire information and then vote the opposite of the observed signal. Thus, the difference in expected utility between acquiring information and not, net of the cost of information acquisition, is

$$c(\epsilon,\sigma_{-i}) \equiv \max\{H(s_A|\epsilon,\sigma_{-i}), H(s_B|\epsilon,\sigma_{-i}), H(s_A|\epsilon,\sigma_{-i})+H(s_B|\epsilon,\sigma_{-i})\} - \max\{0, H(A|\epsilon,\sigma_{-i}), H(B|\epsilon,\sigma_{-i})\}.$$

From the preceding argument we obtain a result parallel to Lemma 1: given any strategy profile of other voters, a best-responding voter only acquires information if the cost is low enough.

Lemma 2. *Under unanimity rule, voter i with priors given by ϵ plays a best response to σ_{-i} if for almost every c,*

(1) *if $c \le c(\epsilon,\sigma_{-i})$ then the voter acquires information, and after signal s_d votes for d if $H(s_d|\epsilon,\sigma_{-i}) > 0$ and abstains if $H(s_d|\epsilon,\sigma_{-i}) < 0$,*
(2) *if $c > c(\epsilon,\sigma_{-i})$, then the voter does not acquire information, and votes for d only if $H(d|\epsilon,\sigma_{-i}) = \max\{0, H(A|\epsilon,\sigma_{-i}), H(B|\epsilon,\sigma_{-i})\}$ and abstains only if $H(A|\epsilon,\sigma_{-i}) \le 0$ and $H(B|\epsilon,\sigma_{-i}) \le 0$.*

We next characterize symmetric, informative Bayesian equilibria, corresponding to $p=0$, $Q=1$. Consider the strategy $\sigma^{c,y}$ of acquiring information if the cost is below some $c \ge 0$, voting for A with probability $1-y$ and abstaining with probability y after signal s_A, voting for B after signal s_B, and abstaining if uninformed. If every voter other than i follows this strategy we get

$$\begin{aligned}
P(0|\sigma^{c,y},\omega_A) &= \left(1-F(c)+F(c)\left(\frac{1}{2}+q\right)y\right)^{n-1},\\
P(0|\sigma^{c,y},\omega_B) &= \left(1-F(c)+F(c)\left(\frac{1}{2}-q\right)y\right)^{n-1},\\
P(1|\sigma^{c,y},\omega_A) &= \left(1-F(c)\left(\frac{1}{2}+q\right)(1-y)\right)^{n-1} - \left(1-F(c)+F(c)\left(\frac{1}{2}+q\right)y\right)^{n-1},\\
P(1|\sigma^{c,y},\omega_B) &= \left(1-F(c)\left(\frac{1}{2}-q\right)(1-y)\right)^{n-1} - \left(1-F(c)+F(c)\left(\frac{1}{2}-q\right)y\right)^{n-1}.
\end{aligned}$$

Note that $P(1|\sigma^{c,y},\omega_A) \le P(1|\sigma^{c,y},\omega_B)$ and $P(0|\sigma^{c,y},\omega_A) \ge P(0|\sigma^{c,y},\omega_B)$, implying $H(A|0,\sigma^{c,y}) \le 0$ and $H(B|0,\sigma^{c,y}) \le 0$, so that voters rather abstain than vote if uninformed and if other voters follow a strategy $\sigma^{c,y}$. For $\sigma^{c,y}$ to be a symmetric equilibrium strategy for $c>0$ and $0<y<1$, it is necessary and sufficient that

$$c = -\frac{1}{2}\left(\frac{1}{2}-q\right)P(0|\sigma^{c,y},\omega_A) + \frac{1}{2}\left(\frac{1}{2}+q\right)P(0|\sigma^{c,y},\omega_B) > 0 \tag{2}$$

and

$$\frac{P(1|\sigma^{c,y},\omega_A)}{P(1|\sigma^{c,y},\omega_B)} = \frac{(1/2)-q}{(1/2)+q}. \tag{3}$$

Eq. (3) implies that $H(s_A|0,\sigma^{c,y})=0$, so that voters are willing to randomize between abstention and voting for A after receiving signal s_A. Eq. (2) follows from $c(0,\sigma^{c,y}) = H(s_B|0,\sigma^{c,y})$, and the inequality implies that voters are willing to vote for B after receiving signal s_B. To verify that Eqs. (2) and (3) have a solution (not necessarily unique), one shows that for every $0 \le F(c) \le 1$ there is some $0<y<1$ such that the pair (c, y) solves Eq. (2). Similarly, for every $0 \le y \le 1$ there is some $0 < c < \bar{c}$ such that (c, y) solves Eq. (3). Existence of an interior solution to both equations follows from a standard fixed point argument.

There may be symmetric, informative Bayesian equilibria other than the one described above. In particular, consider the strategy $\tilde{\sigma}^{c,z}$ of acquiring information if the cost is below some $c \ge 0$, voting for A if receiving the signal s_A, voting for B if receiving the signal s_B, and abstaining with probability z and voting for B with probability $1-z$ when uninformed. If every voter other than i follows this strategy we get

$$\begin{aligned}
P(0|\tilde{\sigma}^{c,z},\omega_A) &= (1-F(c))^{n-1}z^{n-1},\\
P(0|\tilde{\sigma}^{c,z},\omega_B) &= (1-F(c))^{n-1}z^{n-1},\\
P(1|\tilde{\sigma}^{c,z},\omega_A) &= \left(1-F(c)\left(\frac{1}{2}+q\right)\right)^{n-1} - (1-F(c))^{n-1}z^{n-1},\\
P(1|\tilde{\sigma}^{c,z},\omega_B) &= \left(1-F(c)\left(\frac{1}{2}-q\right)\right)^{n-1} - (1-F(c))^{n-1}z^{n-1}.
\end{aligned}$$

Note $P(0|\tilde{\sigma}^{c,z},\omega_A)=P(0|\tilde{\sigma}^{c,z},\omega_B)$ and $P(1|\tilde{\sigma}^{c,z},\omega_A)\le P(1|\tilde{\sigma}^{c,z},\omega_B)$, implying $H(B|0,\tilde{\sigma}^{c,z})=0$, $H(s_B|0,\tilde{\sigma}^{c,z})\ge 0$ and $H(A|0,\tilde{\sigma}^{c,z})\le 0$. For $\tilde{\sigma}^{c,z}$ to be a symmetric equilibrium strategy for $c>0$ and $0<z\le 1$, it is necessary and sufficient that

$$c=\frac{1}{2}\left(\frac{1}{2}+q\right)\left[1-F(c)\left(\frac{1}{2}+q\right)\right]^{n-1}-\frac{1}{2}\left(\frac{1}{2}-q\right)\left[1-F(c)\left(\frac{1}{2}-q\right)\right]^{n-1} \tag{4}$$

and

$$0\le z\le\frac{(c/q)^{1/(n-1)}}{(1-F(c))}. \tag{5}$$

Eq. (4) implies $c=H(s_A|0,\tilde{\sigma}^{c,z})+H(s_B|0,\tilde{\sigma}^{c,z})$ (satisfying Lemma 3), and Eq. (5) implies

$$\frac{P(1|\tilde{\sigma}^{c,z},\omega_A)}{P(1|\tilde{\sigma}^{c,z},\omega_B)}\ge\frac{(1/2)-q}{(1/2)+q},$$

so that $H(s_A|0,\tilde{\sigma}^{c,z})\ge 0$. It is straightforward to check that Eq. (4) has a solution $c^*\in(0,\bar{c})$.

We have

Proposition 2. *Under unanimity rule, if $p=0$ and $Q=1$,*

(1) *For any solution (c,y) to Eqs. (2) and (3), there is a symmetric, informative equilibrium in which each voter acquires information if the voter's cost is below c, votes for B after receiving signal s_B, votes for A with probability y after receiving signal s_A, and abstains otherwise.*
(2) *For any solution (c,z) to Eqs. (4) and (5), there is a symmetric, informative equilibrium in which each voter acquires information if the voter's cost is below c, votes for A after receiving signal s_A, abstains with probability z if uninformed, and votes for B otherwise.*
(3) *There are no other symmetric, informative equilibria.*

Parts 1 and 2 of the proposition are proved above; the proof of part 3 is given in Appendix A. Proposition 2 shows that there are multiple symmetric, informative Bayesian equilibria under unanimity rule, involving either abstaining when receiving a signal favoring the status quo, or voting against the status quo when uninformed. For $p>0$ or $Q<1$, however, the behavior of biased voters or random mistakes may make it a best response for unbiased voters to vote for the alternative favored by the signal received and to abstain if uninformed. We illustrate this point below, continuing the example from the majority rule section.

2.6. An example under unanimity rule

As in the earlier example and the lab experiments, suppose $q=1/6$, c is distributed uniformly in [0, 0.1] and $n=3$ or $n=7$, and suppose the rule is unanimity.

Let $p=0$ and $Q=1$. A symmetric, informative Bayesian equilibrium strategy can be calculated solving Eqs. (2) and (3) or equivalently

$$c=-\frac{1}{6}\left(1-10c+\frac{20}{3}cy\right)^{n-1}+\frac{1}{3}\left(1-10c+\frac{10}{3}cy\right)^{n-1}$$

and

$$\frac{\left(1-\frac{20}{3}c(1-y)\right)^{n-1}-\left(1-10c+\frac{20}{3}cy\right)^{n-1}}{\left(1-\frac{10}{3}c(1-y)\right)^{n-1}-\left(1-10c+\frac{10}{3}cy\right)^{n-1}}=\frac{1}{2}.$$

The probability of reaching the correct decision is given by

$$\frac{1}{2}\left[1-\left(1-\frac{20}{3}c(1-y)\right)^{n}+\left(1-10c+\frac{20}{3}cy\right)^{n}\right]+\frac{10}{2}\left[\left(1-\frac{10}{3}c(1-y)\right)^{n}-\left(1-10c+\frac{10}{3}cy\right)^{n}\right];$$

solutions for $n=3$ and $n=7$ are given by the left column corresponding to the treatments 3U and 7U in Table 1A on experimental predictions.

Other symmetric, informative Bayesian equilibria can be calculated solving Eqs. (4) and (5), or equivalently

$$c=\frac{1}{3}\left(1-\frac{20}{3}c\right)^{n-1}-\frac{1}{6}\left(1-\frac{10}{3}c\right)^{n-1}$$

and

$$0\le z\le\frac{(6c)^{1/(n-1)}}{(1-10c)}.$$

The probability of reaching the correct decision is given by

$$\frac{1}{2}\left(1-\left(1-\frac{20}{3}c\right)^{n}\right)+\frac{1}{2}\left(1-\frac{10}{3}c\right)^{n};$$

solutions for $n=3$ and $n=7$ are given by the right column corresponding to the treatments 3U and 7U in Table 1A. The column to the left under each of the unanimity treatments corresponds to the equilibrium in which voters randomize after receiving a signal favoring the status quo, while the column to the right corresponds to the equilibria in which voters randomize when uninformed. In the latter case, there is an interval of equilibrium mixed strategies, which is indicated with square brackets.

Suppose now β is large enough for voters not to acquire information and vote for the alternative favored by their prior beliefs rather than abstaining when $\epsilon=\beta, -\beta$. In this case, in a symmetric strategy profile in which unbiased voters vote for the alternative favored by the signal received, and abstain with probability z and vote for B with probability $1-z$ if uninformed,

$$P(0|\sigma_{-i},\omega_A)=(1-p)^{n-1}(1-10c)^{n-1}z^{n-1},$$
$$P(0|\sigma_{-i},\omega_B)(1-p)^{n-1}(1-10c)^{n-1}z^{n-1},$$
$$P(1|\sigma_{-i},\omega_A)=\left(\frac{1}{2}p+(1-p)\left(1-\frac{20}{3}c\right)\right)^{n-1}-(1-p)^{n-1}(1-10c)^{n-1}z^{n-1},$$
$$P(1|\sigma_{-i},\omega_B)=\left(\frac{1}{2}p+(1-p)\left(1-\frac{10}{3}c\right)\right)^{n-1}-(1-p)^{n-1}(1-10c)^{n-1}z^{n-1}.$$

Thus, using $c(0,\sigma_{-i})=H(s_A|0,\sigma_{-i})+H(s_B|0,\sigma_{-i})$,

$$c=\frac{1}{3}\left(\frac{1}{2}p+(1-p)\left(1-\frac{20}{3}c\right)\right)^{n-1}-\frac{1}{6}\left(\frac{1}{2}p+(1-p)\left(1-\frac{10}{3}c\right)\right)^{n-1}.$$

We need to check

$$\frac{\left(\frac{1}{2}p+(1-p)\left(1-\frac{20}{3}c\right)\right)^{n-1}-(1-p)^{n-1}(1-10c)^{n-1}z^{n-1}}{\left(\frac{1}{2}p+(1-p)\left(1-\frac{10}{3}c\right)\right)^{n-1}-(1-p)^{n-1}(1-10c)^{n-1}z^{n-1}}>\frac{1}{2},$$

so that unbiased voters are willing to vote for A after signal s_A, and

$$\beta\geq\frac{1}{2}\max\left\{\frac{P(1|\sigma_{-i},\omega_B)-P(1|\sigma_{-i},\omega_A)}{P(1|\sigma_{-i},\omega_A)+P(1|\sigma_{-i},\omega_B)},\frac{2P(1|\sigma_{-i},\omega_A)-P(1|\sigma_{-i},\omega_B)}{2P(1|\sigma_{-i},\omega_A)+P(1|\sigma_{-i},\omega_B)}\right\},$$

so that voters with biased priors do not acquire information and vote for the alternative favored by their bias. The probability of reaching the correct decision is equal to

$$\frac{1}{2}\left[1-\left(\frac{1}{2}p+(1-p)\left(1-\frac{20}{3}c\right)\right)^{n}\right]+\frac{1}{2}\left[\left(\frac{1}{2}p+(1-p)\left(1-\frac{10}{3}c\right)\right)^{n}\right].$$

In particular, for $p=1/2$, we get that for $n=3$, $c\approx0.045$, so that the probability of information acquisition is near 0.22, and the various inequalities are satisfied for $\beta>0.14$. For $n=7$, $c\approx0.015$, so that the probability of information acquisition is near 0.075, and the various inequalities are satisfied for $0\leq z\leq1$ and $\beta>0.124$. Table 1B illustrates the solution for $p=1/2$. Under unanimity rule, there is a range of equilibrium mixtures between abstention and voting against the status quo if uninformed that are consistent with equilibrium behavior, as indicated by the square brackets, but there is no longer an equilibrium in which voters randomize after receiving a signal in favor of the status quo.

For either committee size, the standard Bayesian equilibrium requires either voters abstaining with positive probability when receiving a signal favoring s_A or voting for B with positive probability when uninformed. Introducing voters with subjective prior beliefs induces voters with correct priors to vote with their signals if informed and also reduces the probability of the group reaching the correct decision.

Increasing the size of the committee reduces information acquisition both in the standard Bayesian equilibrium and the subjective equilibrium with extreme prior voters; the theoretical effect of group size on the probability of the group reaching the correct decision is negligible under unanimity rule. As in the case of majority, with seven subjects, after introducing biased individuals, collective choice is barely better than a coin toss.

Under unanimity rule, the probability of voting for the status quo (A) conditional on being uninformed if $Q=1$, as in Table 1 is given by $0.5p/(1-pF(c^*))$, while the probability of voting against one's signal for B, or of abstaining after getting a signal for B, is zero. More generally, if $Q<1$, the probability of voting for A if uninformed is $Q0.5p/(1-(1-p)F(c^*))+(1-Q)/6$, and the probability of voting against one's signal for B, or of abstaining after getting a signal for B, is $(1-Q)/6$. The difference between the probability of voting for the status quo if uninformed and the probability of voting for the status quo against one's signal is $Q0.5p/(1-(1-p)F(c^*))$; similar to majority, this difference is indicative of behavior that cannot be explained by random errors.

3. Experimental design and hypotheses

3.1. Experimental design

The design of our experiment was guided by the comparative statics implications of the standard Bayesian equilibrium with $p = 0$ and $Q = 1$. While there are several dimensions of the model that yield clear comparative statics, we focus on two: the number of voters ($n = 3$ or $n = 7$) and the decision rule (V_M or V_U). In all treatments, the informativeness of the signal, q, was held constant, as was the distribution of signal costs. The design also allows us to explore the influence of subjective beliefs on behavior. In particular, we use the data from the experiment to estimate the parameters of our subjective beliefs equilibrium model in order to measure the extent of this phenomenon.

The procedures and framing of the experiment were based on the Condorcet jury "jar" interface introduced by Guarnaschelli et al. (2000) and adapted by Battaglini et al. (2008, 2010) in their initial laboratory studies of the swing voter's curse.[6] The two states of the world are represented as two jars, a red jar and a blue jar. The game proceeds as follows. First, either the master computer or a subject-monitor tosses a fair coin to determine the state of the world (i.e., selects the jar). The red jar contains 8 red balls and 4 blue balls, and the blue jar contains a 8 blue balls and 4 red balls, in order to induce a signal informativeness of $q = 1/6$. The red jar corresponds to state B and the blue jar corresponds to state A in the theoretical model. This labeling only matters for the unanimity committees, where decision A is the status quo.

Each committee member i was assigned an integer-valued signal cost, c_i, drawn from a commonly known uniform distribution over 0.001, 0.002, . . ., 0.100.[7] Then each committee member, acting independently of other committee members, could choose to pay their signal cost in order to privately observe the color of exactly one of the balls randomly drawn from the jar. The randomization was done as follows. A jar appears on the subject screen with 12 balls inside it, with 8 of them one color and 4 of them the other color. The locations of the 12 balls are randomly shuffled on each screen and the colors are greyed out. If a subject pays his or her signal cost, the computer prompts them to click on one of the greyed-out balls, which then reveals the color of that ball. In case they chose not to pay the cost, they do nothing at this point.

Once all subjects selected a ball or indicated their choice not to do so, each committee member is given three choices: vote for Red; vote for Blue; or Abstain. At no time was any communication between the subjects allowed, so both the information (or lack thereof) and vote decisions remained private until all the votes were cast, at which point only the votes were announced, and the committee decision was implemented according to the voting rule (either majority, or unanimity with red as the status quo).

If the committee choice was correct (i.e. the committee voted for the same color as the plurality of the balls in the jar) each committee member received a payoff of 1000 points, less whatever the private cost incurred for observing the color of a ball. If the decision was wrong (i.e. the color chosen by the committee and the color of the plurality of the balls in the jar did not coincide), each committee member received a payoff of 0 and still had to pay the private cost of acquiring information, if any had been incurred.

Each committee decision, as described above, constituted a single experimental round, upon completion of which committees were randomly re-matched and new jars and private observation costs were drawn independently from the previous rounds. Detailed instructions were read aloud before the experiment began. A translated copy of these instructions is provided in Appendix. Fig. 3 in Appendix presents a sample of the computer screen as it appeared to subjects after they observed the color of a ball. In case they chose not to observe, the screen would be identical, except that all balls would appear grey.

All experimental sessions (generally involving 21 subjects each, except for one 15-subject session with 3-member committees deciding by majority rule) consisted of 25 rounds of the same treatment with random re-matching between rounds, and were conducted at ITAM in Mexico City with student subjects recruited from introductory economics courses. At the end of each session each subject was paid the sum of their earnings across all rounds, in cash, using the exchange rate of 1000 points to 8 Mexican pesos (rounded to the nearest peso) plus 20 pesos as a show-up fee. Average earnings, including the show-up fee, were 133 pesos for M3, 141 pesos for M7, 125 pesos for U3, and 127 pesos for U7. (At the time of the experiment, 1 US dollar was worth around 12 pesos.) Each session lasted approximately 1 h.

3.2. Hypotheses

As described earlier, Table 1A summarizes equilibrium strategies and probability of a correct group decision for the standard ($p = 0$, $Q = 1$) Bayesian equilibria of the four treatments, 7M, 7M, 3U, and 7U. As explained in the theory section, there are multiple Bayesian equilibria under unanimity rule. The column to the left under each of the unanimity treatments corresponds to the equilibrium in which voters randomize after receiving a signal favoring the status quo, while the column to the right corresponds to the equilibrium in which voters randomize when uninformed. Based on Table 1A, we summarize the main hypotheses below:

---

[6] Sample instructions are provided in Appendix B, including a screenshot illustrating the jar interface.

[7] In the experiment, all payoffs were designated in points, using integer amounts. The possible costs, in points were drawn from 1, . . ., 100, and the value of a correct decision was 1000.

(H1) Under both voting rules, members of smaller committees acquire more information.
(H2) For both committee sizes, members of majority rule committees acquire more information than members of unanimity rule committees.
(H3) Under majority rule, committee members who do not acquire information abstain.
(H4) Under unanimity rule, committee members who do not acquire information abstain or vote for *B*.
(H5) Under both voting rules, committee members who acquire information never vote against their signal.
(H6) Under majority rule, committee members who acquire information vote their signal.
(H7) Under unanimity rule, committee members who acquire information and receive a *B* signal vote for *B*.
(H8) Under unanimity rule, committee members who receive an *A* signal vote for *A* or abstain.
(H9) With majority rule, larger committees make better decisions.
(H10) Majority committees make better decisions than unanimity committees.

Under the SBE, that is allowing for biases ($p > 0$), some of these hypotheses can reverse. Inspecting Table 1B, in particular, we can see that *H*1 and *H*2 hold after introducing biased voters, and *H*10 may hold at least for small sized committees. *H*3, *H*4 and *H*9, however, may fail after introducing biased voters. Introducing random mistakes ($Q < 1$) may lead to violations of *H*5 to *H*8, but we expect these violations to be small or negligible.

4. Experimental results

We first present some summary statistics that provide a simple test of the comparative static predictions of the baseline model with respect to treatment effects, as detailed in Table 1A of the previous section. We then present estimation results, using a structural approach to estimate the parameters of the subjective belief equilibrium model. In the last section, we take a close look at individual behavior, and use those estimates to classify individual subject behavior.

4.1. Treatment effects

Table 2 summarizes treatment effects.

4.1.1. Information acquisition (H1–H2)

With respect to the frequency of information acquisition, there are three notable observations. First, there is no significant effect of committee size on information acquisition. Because the baseline theory predicts a large effect of committee size on information acquisition (H1), this finding is surprising. Also note that the lack of a statistically significant effect of the committee size is not just an artifact of large standard errors. Quantitatively, the average effect of committee size is precisely zero (to two decimal places) for both voting rules. Second, consistent with H2, there is more information acquisition under majority rule than under unanimity rule. The size of this effect is the same for both committee sizes; in both cases, there is about 20% more information acquisition under majority rule than under unanimity. This percentage difference is somewhat higher than predicted by theory, although the raw difference in information acquisition (0.06) close to the theory. Third, we observe significantly less information acquisition than predicted, except for unanimity committees with seven members. The magnitude of this difference is large for committees with three members, where we observe 50% less information acquisition than predicted.

Table 2
Summary of experimental data. Standard errors in parenthesis treat each individual's 25 decisions as a single observation. The unit observation for a committee decision is one committee.

Treatment (size, rule)	3M	7M	3U	7U
Observed frequencies of individual decisions				
Info acquisition	0.33	0.33	0.27	0.27
	(0.05)	(0.06)	(0.06)	(0.06)
Vote *A* if uninformed	0.38	0.33	0.29	0.20
Vote *B* if uninformed	0.37	0.28	0.35	0.35
Abstain if uninformed	0.24	0.39	0.37	0.45
Vote *A* if signal s_A	0.96	0.96	0.80	0.82
Vote *B* if signal s_A	0.04	0.02	0.04	0.03
Abstain if signal s_A	0.00	0.02	0.16	0.15
Vote *A* if signal s_B	0.03	0.05	0.04	0.00
Vote *B* if signal s_B	0.97	0.93	0.89	0.89
Abstain if signal s_B	0.00	0.02	0.07	0.02
Observed frequency of group decision				
Correct decision	0.60	0.63	0.56	0.56
	(0.06)	(0.04)	(0.06)	(0.04)

4.1.2. Voting behavior (H3–H8)

With respect to the voting frequency, the most striking observation from Table 2 is the amount of voting by uninformed voters, which strongly contradicts H3. In the majority treatments, participation by uninformed voters exceeds 60%, while the baseline theory predicts zero turnout. This is strikingly different from the finding in the swing voter's curse experiments by Battaglini et al. (2008, 2010), where under majority rule, uninformed voters abstained nearly all the time when the two states were equally likely. We discuss this finding more in the concluding section. In the unanimity treatments, uninformed voting in favor of the status quo exceeds 20%, while the baseline theory predicts zero, contradicting H4. Uninformed voting in favor of either alternative under majority rule, as well as uninformed voting in favor of the status quo decline with committee size.

A second observation regarding voting frequency is that informed voters almost never vote against their signal, which is consistent with H5 and with past findings.

A third observation is that we observe significant levels of abstention among informed voters only in the case of voters who obtain a signal favoring the status quo under unanimity rule. The level of abstention is small, and informed voters do tend to vote for the alternative favored by the signal received. This supports hypotheses H6, H7, and H8.

4.1.3. Group decision accuracy (H9–H10)

With respect to the frequency of correct group decisions, the effect of group size is negligible for both voting rules, which is consistent with the theory for the unanimity committees, but contradicts H9. We do observe that the probability of correct decisions is higher under majority rule than under unanimity rule, which supports H10. (In fact, the difference in the probability of correct decision in majority versus unanimity is larger in the data than under the Bayesian equilibrium.)

To summarize, we find support for the qualitative hypotheses H2, H5, H6, H7, H8, and H10. The failure of H3 and H4 goes in line with the predictions of the subjective beliefs model with $p > 0$, to which we turn our attention now.

4.2. Structural estimation of the subjective beliefs equilibrium model

There are several key features of observed behavior that are consistent with the subjective beliefs equilibrium model with $p > 0$ but were not predicted by the standard model with $p = 0$. First, many subjects usually vote when uninformed. Second, subjects acquire information much less frequently than predicted. Third, informed voters usually vote their signal. These features are commonly shared across all four treatments. With this in mind, we perform a maximum likelihood estimation of the subjective beliefs equilibrium model for general (p, Q).

We assume that player types are *persistent*, so that unbiased subjects hold objective priors for all rounds of the experiment, while biased subjects draw a new subjective prior every round. [8] We assume that β is sufficiently high so that biased subjects prefer not to acquire information and simply vote their hunch, as explained in the theory section.

To simplify the estimation, we blind ourselves to the actual cost draws of subjects. Thus, there are nine possible sequences of actions for a given subject in a period that are relevant for the estimation, given by the signal observed by the voter and the vote cast.

(AA) Acquires information, draws signal s_A, and votes for A.
(AB) Acquires information, draws signal s_A, and votes for B.
($A\phi$) Acquires information, draws signal s_A, and abstains.
(BA) Acquires information, draws signal s_B, and votes for A.
(BB) Acquires information, draws signal s_B, and votes for B.
($B\phi$) Acquires information, draws signal s_B, and abstains.
(ϕA) Does not acquires information, and votes for A.
(ϕB) Does not acquires information, and votes for B.
($\phi\phi$) Does not acquires information, and abstains.

Under our assumption of uniform random mistakes, since the unconditional probability of receiving either signal is 1/2, each of the six action sequences, AA, AB, $A\phi$, BA, BB, and $B\phi$, will occur by mistake with probability $(1 - Q)/12$, and each of the three sequences, ϕA, ϕB and $\phi\phi$, will occur with probability $(1 - Q)/6$.

4.2.1. Likelihood function for majority rule.

The likelihood function is constructed as follows. Denote the number of times subject i took each of the nine action sequences as k^i_{sv}, where

$$sv \in \{AA, AB, A\phi, BA, BB, B\phi, \phi A, \phi B, \phi\phi\}.$$

[8] Treating subjects as persistent types in this way does not change the equilibrium strategies derived in the theory section.

Our data for subject i is simply i's profile of actions, $D^i = (k^i_{sv})$. For any pair of parameters, (p, Q), and for group size n, the majority rule likelihood of D^i is given by:

$$L_M(D^i|p,Q,n) = p\left\{\left[\left(\frac{1}{2}Q + \frac{1}{6}(1-Q)\right)\right]^{k^i_{\phi A}+k^i_{\phi B}}\left[\frac{1}{6}(1-Q)\right]^{k^i_{\phi\phi}}\left[\frac{1}{12}(1-Q)\right]^{k^i_{AA}+k^i_{BB}+k^i_{AB}+k^i_{A\phi}+k^i_{BA}+k^i_{B\phi}}\right\}$$
$$+(1-p)\left\{\left[\frac{1}{2}\iota^*_M(p,Q,n)Q + \frac{1}{12}(1-Q)\right]^{k^i_{AA}+k^i_{BB}}\left[(1-\iota^*_M(p,Q,n))Q + \frac{1}{6}(1-Q)\right]^{k^i_{\phi\phi}}\right.$$
$$\left.\times\left[\frac{1}{6}(1-Q)\right]^{k^i_{\phi A}+k^i_{\phi B}}\left[\frac{1}{12}(1-Q)\right]^{k^i_{AB}+k^i_{A\phi}+k^i_{BA}+k^i_{B\phi}}\right\}$$

where $\iota^*_M(p, Q, n)$ is the equilibrium probability that an unbiased voter buys information and then votes according to the signal received in the subjective beliefs equilibrium model if the model parameters are (p, Q) and the committee size is n, using majority rule. The first term in the right-hand side corresponds to the event that subject i is biased, which happens with probability p. In that case, the subject intends not to acquire information and vote for A with probability 1/2, and not to acquire information and vote for B with probability 1/2, depending on the realization of the subject's bias. The subject does as intended with probability Q, and makes a mistake with probability $1 - Q$, in which case the subject does not acquire information and votes for A with probability 1/6, and similarly for B. Thus, the expression $[\frac{1}{2}Q + \frac{1}{6}(1-Q)]$ in the first term is equal to the probability that a biased subject does not acquire information and votes for A, and equal to the probability that a biased subject does not acquire information and votes for B. Other terms can be explained similarly.

Using $F(c^*) = \iota^*_M(p, Q, n)$, we get that $\iota^*_M(p, Q, n) = 10c^*$, where c^* is calculated as in Proposition 1. The log likelihood function is equal to the sum of $\log L_M(D^i|p, Q, n)$ across all the individuals in the sample. The estimation is then done by using Matlab to find the values of p and Q that maximize the log likelihood function.

4.2.2. *Likelihood function for unanimity rule*

The expression for the likelihood of D^i with unanimity is based on the equilibrium described in the last two columns of Table 2B. For any pair of parameters, (p, Q), and for group size n, the unanimity rule likelihood of D^i is given by:

$$L_U(D^i|p,Q,n) = p\left\{\left[\frac{1}{2}Q + \frac{1}{6}(1-Q)\right]^{k^i_{\phi A}+k^i_{\phi B}}\left[\frac{1}{6}(1-Q)\right]^{k^i_{\phi\phi}}\left[\frac{1}{12}(1-Q)\right]^{k^i_{AA}+k^i_{BB}+k^i_{AB}+k^i_{A\phi}+k^i_{BA}+k^i_{B\phi}}\right\}$$
$$+(1-p)\left\{\left[\frac{1}{2}\iota^*_U(p,Q,n)Q + \frac{1}{12}(1-Q)\right]^{k^i_{AA}+k^i_{BB}}\left[\frac{1}{2}(1-\iota^*_U(p,Q,n))Q + \frac{1}{6}(1-Q)\right]^{k^i_{\phi\phi}+k^i_{\phi B}}\right.$$
$$\left.\times\left[\frac{1}{6}(1-Q)\right]^{k^i_{\phi A}}\left[\frac{1}{12}(1-Q)\right]^{k^i_{AB}+k^i_{A\phi}+k^i_{BA}+k^i_{B\phi}}\right\}.$$

where $\iota^*_U(p, Q, n)$ is the equilibrium probability that an unbiased voter buys information and then votes according to the signal received in the subjective beliefs equilibrium model, if the model parameters are (p, Q) and the committee size is n, using unanimity rule. This expression is similar to the one for majority, except for how it deals with the probability that an unbiased subject does not get information and abstains, and the probability that an unbiased subject does not get information and votes for B. The reason is that under unanimity rule, for our parameter values, even for very small values of p, unbiased, uninformed voters can randomize in any way between voting for B and abstaining. We assume that such voters randomize with equal probability. Thus, the expression $[\frac{1}{2}(1-\iota^*_U(p, Q, n))Q + \frac{1}{6}(1-Q)]$ is equal to the probability that an unbiased subject does not acquire information and abstains, and is equal to the probability that an unbiased subject does not acquire information and votes for B.[9]

Using Eq. (4), we have

$$\iota^*_U(p,Q,n) = \left(\frac{10}{3}\right)(1-\nu_r)^{n-1} - \left(\frac{5}{3}\right)(1-\nu_W)^{n-1},$$

[9] We also explored an alternative estimation model that included an additional parameter for the probability an uninformed unbiased voter votes for B. That less parsimonious specification improves the fit somewhat, but leaves the estimates of p and Q unchanged.

Table 3
Estimation results for subject beliefs equilibrium model.

Treatment	Observations	$\hat{p}$	$\hat{Q}$	$-\ln L$
3M	1950	0.41	0.76	2889
7M	1554	0.45	0.76	2329
3U	1575	0.41	0.75	2497
7U	1575	0.10	0.78	2539
Pooled except 7U	5079	0.42	0.76	7716
All pooled	6654	0.30	0.76	10,282

Table 4
Comparison of action frequencies: estimated model vs. data.

Treatment		%I	AA	AB	Aϕ	BA	BB	Bϕ	ϕA	ϕB	$\phi\phi$
3M	Model	0.45	0.19	0.02	0.02	0.02	0.19	0.02	0.20	0.20	0.15
	Data	0.33	0.16	0.00	0.01	0.01	0.16	0.00	0.25	0.26	0.16
7M	Model	0.33	0.12	0.02	0.02	0.02	0.12	0.02	0.21	0.21	0.25
	Data	0.33	0.15	0.01	0.00	0.00	0.16	0.00	0.19	0.22	0.26
3U	Model	0.34	0.13	0.02	0.02	0.02	0.13	0.02	0.20	0.31	0.16
	Data	0.27	0.11	0.01	0.02	0.01	0.12	0.01	0.21	0.25	0.27
7U	Model	0.26	0.09	0.02	0.02	0.02	0.09	0.02	0.08	0.35	0.31
	Data	0.27	0.11	0.00	0.02	0.00	0.13	0.00	0.15	0.25	0.32

where

$$\nu_r = \frac{2}{3}(1-p)\iota_U^*(p, Q, n)Q + \frac{p}{2}Q + \frac{1}{3}(1-Q),$$
$$\nu_w = \frac{1}{3}(1-p)\iota_U^*(p, Q, n)Q + \frac{p}{2}Q + \frac{1}{3}(1-Q).$$

4.2.3. Estimation results

Table 3 reports the maximum likelihood estimates $(\hat{p}, \hat{Q})$. We report these estimates at two levels of aggregation:

(a) a separate estimate for each of the four treatments;
(b) an estimate for all treatments except 7U combined;
(c) an estimate for all treatments combined.

Table 3 also reports the number of observations, which is simply the number of subjects in each treatment times 25. Recall that for each subject we have a panel of 25 observations, each observation consisting of one of the nine possible pairs of actions listed above.[10]

The estimated values $(\hat{p}, \hat{Q})$ are very similar for the majority treatments and for the unanimity treatment with three member committees, with the proportion of biased subjects being approximately 40%, and error rates around 25%. In fact, these three treatments are statistically indistinguishable, based on a chi-square test of the difference between the pooled log likelihood and the sum of the three separately estimated log likelihoods.

The estimated error rate for the 7U treatment is about the same (22%) as the other treatments, but the estimated proportion of biased subjects is much lower (10% compared to 41%). This is the only of our treatments where we observe *too much* information acquisition. For example, when $p = 1/2$, the equilibrium probability of information acquisition in the 7U treatment is only 0.08, which is far below the observed frequency information acquisition of 0.27. To account for the observed frequency of information acquisition, the estimated model predicts that 90% of the subjects are unbiased. This however, leads to an estimated probability of voting for the status quo if uninformed that is half of what is observed in the data. A conjecture to explain the relatively high frequency of information acquisition is that, for 7U treatments, acquiring information above the best response cutoff is practically costless, so that in fact the fraction of unbiased subjects is similar to that in other treatments, but these individuals are doing more information acquisition than predicted by best response behavior.

Table 4 displays the fitted action probabilities corresponding to separate estimates for each treatment, and compares them to the frequencies observed in the data. Column 3 displays the percent of informed individuals, denoted *%I*. The predicted value for percent informed for each treatment is computed as

$$\frac{(1-\hat{Q})}{2} + (1-\hat{p})\hat{Q}\iota_R^*(\hat{p}, \hat{Q}, n)$$

[10] An exception is one of the majority sessions with seven member committees, where we have only 24 observations for each subject.

Table 5
Estimated model predictions vs. data.

Treatment (size, rule)	Estimated model				Data			
	3M	7M	3U	7U	3M	7M	3U	7U
Predicted probabilities of individual decisions								
Info acquisition	0.45	0.33	0.34	0.26	0.33	0.33	0.27	0.27
Vote A if uninformed	0.36	0.31	0.30	0.10	0.38	0.33	0.29	0.20
Vote B if uninformed	0.36	0.31	0.47	0.48	0.37	0.28	0.35	0.35
Abstain if uninformed	0.28	0.37	0.24	0.42	0.24	0.39	0.37	0.45
Vote A if signal s_A	0.82	0.76	0.75	0.72	0.96	0.96	0.80	0.82
Vote B if signal s_A	0.09	0.12	0.12	0.14	0.04	0.02	0.04	0.03
Abstain if signal s_A	0.09	0.12	0.12	0.14	0.00	0.02	0.16	0.15
Vote A if signal s_B	0.09	0.12	0.12	0.14	0.03	0.05	0.04	0.00
Vote B if signal s_B	0.82	0.76	0.75	0.72	0.97	0.93	0.89	0.89
Abstain if signal s_B	0.09	0.12	0.12	0.14	0.00	0.02	0.07	0.02
Predicted probability of group decision								
Correct decision	0.56	0.56	0.54	0.55	0.60	0.63	0.56	0.56

Table 6
Distribution of committee member types, based on modal behavior.

Behavioral type	3M	7M	3U	7U	All
Unbiased	0.44	0.59	0.52	0.73	0.56
Biased	0.53	0.40	0.41	0.21	0.40
Unclassified	0.03	0.01	0.07	0.06	0.04
Observations	77	63	63	63	266

for $R \in \{M, U\}$. Table 5 presents the predicted conditional probabilities of the estimated SBE model versus the observed frequencies in the data.

On the whole, the model fits the empirical distributions rather well in most cases. There are two notable exceptions. For both unanimity treatments, the model underestimates the frequency of abstention by uninformed voters and overestimates their frequency of voting for B. However, in no cases are these differences large in magnitude.

4.2.4. *Classification of individual subjects*

Using our estimates, we conduct a classification analysis based on individual behavior. We have 25 observations for each subject, except those in one majority session for whom we have only 24 observations. Each observation is one of the nine possible sequences of actions in $\{AA, AB, A\phi, BA, BB, B\phi, \phi A, \phi B, \phi\phi\}$. We compute, for each subject, the log-odds, λ_i, that the subject is a biased voter, calculated as the log of the ratio of the likelihood they are a biased and the likelihood they are unbiased, evaluated using the estimated parameters, $(\hat{p}, \hat{Q})$. We call λ_i subject i's *Lscore*. Thus, for example, for subject i in a 3M session, i's Lscore is computed as:

$$\lambda_i(D^i|\hat{p}_{3M}, \hat{Q}_{3M}) = \log\left\{\frac{\left[\frac{1}{2}\hat{Q}_{3M} + \frac{1}{6}(1-\hat{Q}_{3M})\right]^{k^i_{\phi A}+k^i_{\phi B}}\left[\frac{1}{6}(1-\hat{Q}_{3M})\right]^{k^i_{\phi\phi}}\left[\frac{1}{12}(1-\hat{Q}_{3M})\right]^{k^i_{AA}+k^i_{BB}+k^i_{AB}+k^i_{A\phi}+k^i_{BA}+k^i_{B\phi}}}{\left[\frac{1}{2}\iota^*_M(\hat{p}_{3M}, \hat{Q}_{3M}, 3)\hat{Q}_{3M} + \frac{1}{12}(1-\hat{Q}_{3M})\right]^{k^i_{AA}+k^i_{BB}}\left[(1-\iota^*_M(\hat{p}_{3M}, \hat{Q}_{3M}, 3))\hat{Q}_{3M} + \frac{1}{6}(1-\hat{Q}_{3M})\right]^{k^i_{\phi\phi}}}\right\}.$$

An Lscore greater than 0 corresponds to a subject who is more likely to be a biased type, and an Lscore below 0 indicates a subject more likely to be an unbiased type. Note that an Lscore above 3 or below −3 indicate that the odds (under the estimated model) are 20:1 that the subject is correctly classified as a biased or unbiased type, respectively. We use this 20:1 odds as our criterion for saying a subject is "classified" as a type. One can interpret 20:1 odds, for descriptive purposes, as indicating with 95% confidence that the subject is correctly classified by the estimated model.

Table 6 indicates, for each treatment, the percentage of subjects with Lscores below −3 (classified as "unbiased"), between −3 and 3 ("unclassified"), and above 3 (classified as "biased"). Across all sessions, 96% of subjects are classified, which suggests some support for our two-type mixture model.[11] Furthermore, 40% of subjects are classified as biased types, a percentage that corresponds almost exactly with the estimated value of $\hat{p}_{3M}$, $\hat{p}_{7M}$, and $\hat{p}_{3U}$.

Finally, Fig. 1 displays the action distribution for each subject, ordered by their Lscores, for each treatment. To avoid the figures becoming cluttered, we do not include all 9 possible sequences of actions in the figure, but condense these into 3 categories: ($IVote$, $\phi\phi$, $\phi Vote$) where $IVote = \{AA, AB, BA, BB\}$, $\phi\phi = \{\phi\phi\}$, and $\phi Vote = \{\phi A, \phi B\}$. The sequences $A\phi$ and $B\phi$ are very rare and are included in $IVote$.

[11] The classification rates are nearly as high if the classification criterion is considerably strengthened to 100:1 odds (92%), or even 1000:1 odds (90%).

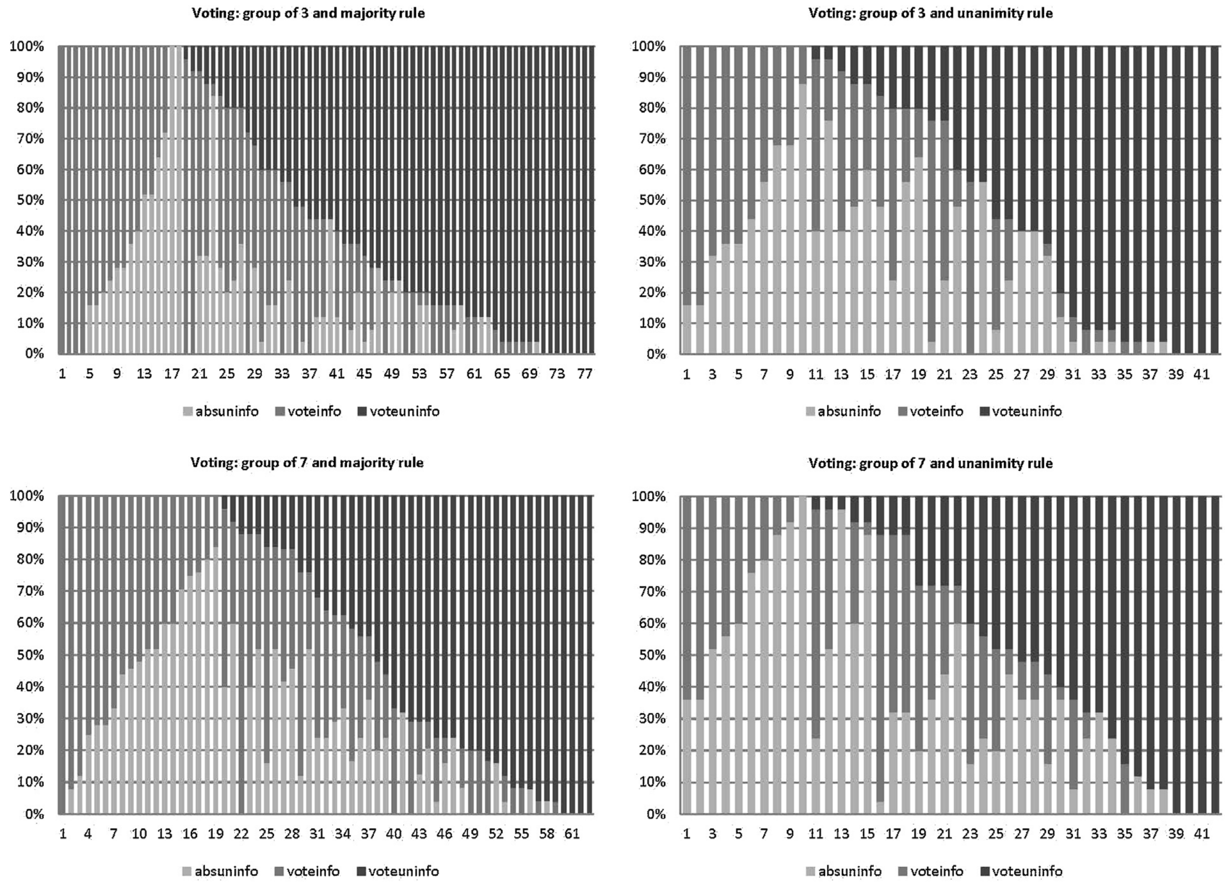

Fig. 1. Distribution of individual action profiles, ordered by Lscores.

As one can see from Fig. 1, for all four treatments, ordering of subjects by their Lscores is roughly the same as ordering them lexicographically by their relative frequencies of these three behavioral categories. In all treatments, individual Lscores are (almost) monotonically increasing in the probability of action ϕ*Vote*. Furthermore, in the majority rule committees, individual Lscores are (almost) monotonically decreasing in the probability of *IVote*. Thus, in these treatments, *IVote* is a strong indicator of an unbiased type. This indicator is weaker in the unanimity treatments. In fact, in the unanimity treatments, $\phi\phi$ is at least as strong a marker for an unbiased type as *IVote*, and is the strongest indicator for voters in 7U committees. Finally, voters with a roughly equal mix of three of the action categories are harder for the model to classify and therefore tend to have Lscores with lower absolute values (i.e., centrally located in these graphs).

4.3. Voting behavior of biased voters

We illustrate the behavior of biased voters in seven member committees in Fig. 2. (Behavior in three member committees is similar.) Fig. 2A and B is devoted to majority rule. As shown in Fig. 2A, on average, biased voters were equally likely to vote for either alternative when uninformed, although there is considerable individual variation, with some individuals behaving as “partisans.” As shown in Fig. 2B, unsurprisingly, on average, biased voters are roughly equally likely to “guess” correctly or incorrectly when voting without information, though a few individuals managed to guess incorrectly every time.

We turn to the behavior of biased voters under unanimity rule in Fig. 2C and D. As shown in Fig. 2C, biased voters are more likely to guess against the status quo, although there is considerable individual variation. Note that voting against the status quo when uninformed is not inconsistent with equilibrium behavior of an unbiased voter, so that it more difficult here to disentangle who is biased and who is not under unanimity. As shown in Fig. 2D, again unsurprisingly, on average, biased voters are roughly equally likely to “guess” correctly or incorrectly when voting without information, though a small number of individuals managed to guess incorrectly or correctly every time.

Note that biased voters do acquire information with some probability under both treatments, but are much less likely to do so, which is consistent with random mistakes.

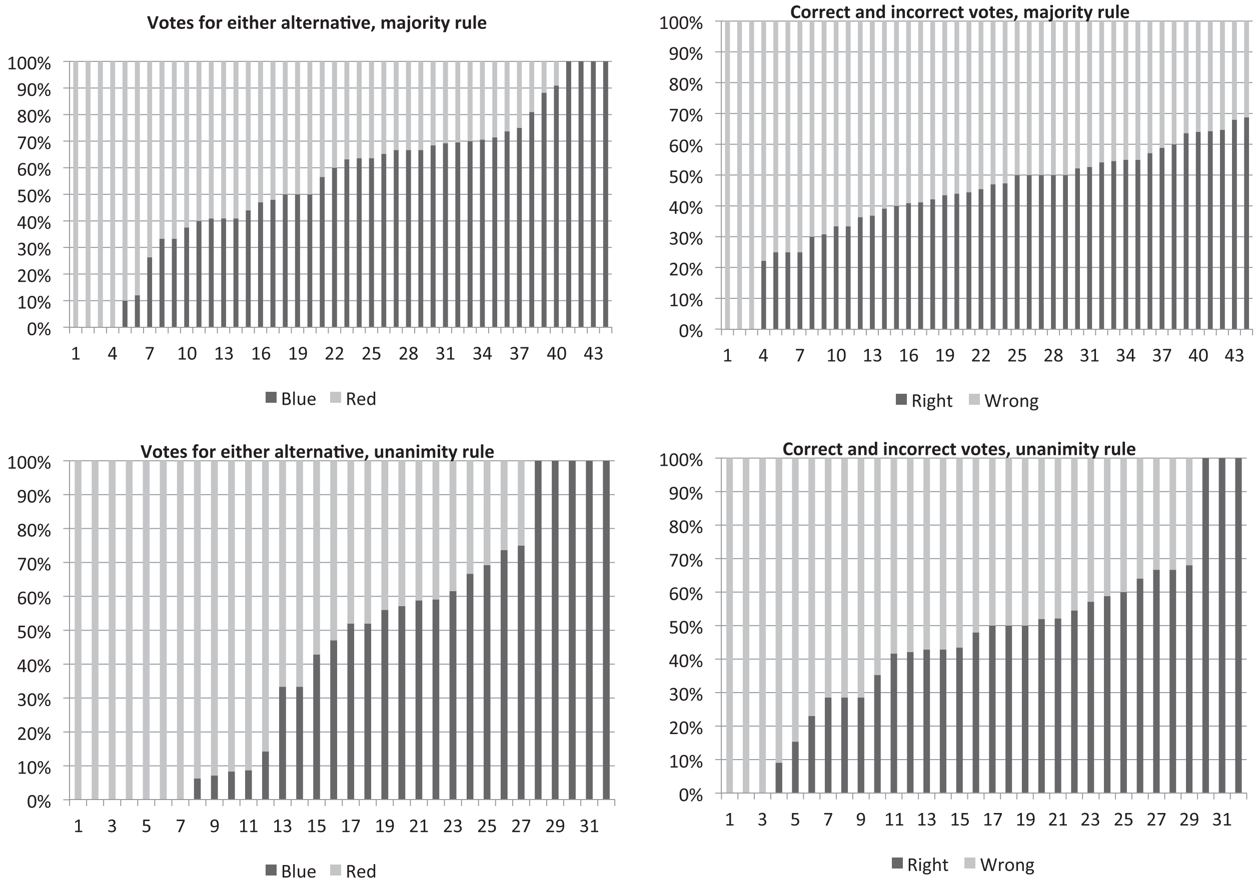

Fig. 2. Behavior of biased voters when uninformed in seven member committees.

5. Alternative behavioral theories

We developed the subjective belief equilibrium model only after carefully exploring alternative existing off-the-shelf behavioral theories that seemed promising at first blush, in terms of being able to explain our anomalous findings. Unfortunately, none of these alternative theories organize our data in a satisfactory way that is the same time consistent with extant data from earlier published experiments on the swing voter's curse. In this section, we briefly discuss here the implications of some of these alternative behavioral theories in the context of our collective decision problem, and the difficulties in organizing the data from this and previous experiments.

5.1. Cursed equilibrium

Cursed voters, as postulated by Eyster and Rabin (2005), predict correctly the distribution of action profiles of other voters, conditional on being uninformed or on having received a signal in favor of *A* or *B*, and hence would predict correctly the probability of being decisive, but would not take into account that other voters actions are affected by their private information about the state of the world. This has two different kinds of effects on behavior in our environment. First, because cursed voters think other voters' votes contain no information about the state of the world, when uninformed they would be indifferent between voting for either alternative or abstaining. A second, more subtle, effect is that cursed voters are more eager than rational voters to acquire information, because they act as if they implicitly believe that other voters do not gather information. That implicit belief implies that there is no free rider effect on the information gathering incentives for cursed voters. The first observation (uninformed voting) is consistent with our data, but inconsistent with past data. The second observation is inconsistent with our finding of underinvestment in information, coupled with the observation that of a large fraction of voters in our experiment (nearly forty percent) never purchase a signal yet always vote.

To flesh out the logic of overinvestment in information in a cursed equilibrium, consider a majority committee composed of three cursed voters who abstain or vote for either alternative with probability 1/3 each when uninformed. The probability that one of the other voters vote for A and the other for B is equal to

$$2\left(\frac{1}{3}+\left(\frac{1}{6}+q)F(c)\right)\right)\left(\frac{1}{3}+\left(\frac{1}{6}-q)F(c)\right)\right).$$

In this event, a cursed voter anticipates correctly that the value of information is qb, as the probability that the correct alternative gets chosen increases from 1/2 to $1/2+q$.

The probability that one of the other voters vote against the signal obtained by the voter and the other abstains is

$$2\left(\frac{1}{3}+\left(\frac{1}{6}-2q^2\right)F(c)\right)\frac{1}{3}(1-F(c)).$$

In this event, a cursed voter believes that the value of information is qb, as the probability that the correct alternative gets chosen increases with probability 1/2 from $1/2-q$ to $1/2+q$. Note that this calculation overestimates the true value of information since it ignores that the other voter who is not abstaining may be informed. Equating the perceived value of information to the cost as in Eq. (1), we obtain that in a fully-cursed equilibrium the frequency of information acquisition should be about 76%, far above both Bayesian equilibrium and observed frequency in the data.

5.2. Level-k reasoning

Suppose voters can be ordered in a cognitive hierarchy, as postulated by Stahl and Wilson (1995), Camerer et al. (2004), and others. First there is the thorny question of how to specify level-0 rationality, for which there are several candidate models in our experiment, none of which are particularly satisfactory. If one adopts the most common approach, where level-0 individuals randomize uniformly over available actions, such voters would acquire information and vote against the signal obtained with positive probability, and their information acquisition behavior would not be responsive to their idiosyncratic information costs. Both of these behaviors are at odds with observed data. Alternatively, one might assume that level 0 voters randomize uniformly over voting for either alternative or abstaining, and do not ever acquire information. This takes us one step closer to the data, as their behavior is identical to the behavior of our biased voters. In that case, in a three-member majority committee, for instance, level-1 voters would acquire information with probability near 74%, which is close to the prediction in our SBE model for unbiased voters when half the voters are biased, but on the other hand, these level-1 voters would randomize between voting for either alternative or abstaining when uninformed, which is at odds with the behavior of the unbiased subjects in the data. This specification of level 0 voters also has the flaw that it captures behavior of some voters in our experiment, but would predict high levels of uninformed voting in past swing voter curse experiments, where such behavior was not observed. Alternatively, one might assume that level 0 voters always get informed and always vote their signal. This would capture the behavior of a fraction of our unbiased subjects, but higher level types would not ever vote when they were uninformed. One could search for other specifications of level-0 behavior, but it seems that the most obvious candidates lead to implausible predictions about the distribution of aggregate behavior. The inability of this class of theories to explain data in swing voter curse experiments is also discussed in Battaglini et al. (2010).

5.3. Loss aversion

Treating the choice not to purchase information as a reference point, loss averse individuals, as postulated by Kahneman and Tversky (1983), would be less inclined than rational voters to acquire information, which is a loss in case the voter is not decisive or the information or the voter is decisive but wrong. This could account for the average small frequency of information acquisition in the data, but would leave unexplained uninformed voting. Something similar occurs with risk aversion for small stakes; it would deter voters from acquiring information, because the gain of being informed is probabilistic while the cost is certain, but, again, risk aversion cannot account for the very high frequency of uninformed voting.

5.4. Quantal response equilibrium

Stochastic models of equilibrium in games, such as the quantal response equilibrium of McKelvey and Palfrey (1995, 1998) could account for uninformed voting. However, given the very high rates of uninformed voting in our data, an implausibly high error rate would be required to fit the data.[12] But a very high error rate is inconsistent with the very systematic and rational behavior of voters who always vote their signal when buying information. This suggests that a heterogeneous QRE model (Rogers et al. 2009) might explain behavior very well, by simultaneously capturing the random voting behavior of the biased types in the SBE model and the highly rational behavior of the unbiased types. Unfortunately, the biased types in our model are not behaving randomly; they never buy information, while a random type would buy information half the

[12] In fact, the estimated value of lambda in a logit equilibrium is very close to 0.

time. That said, our estimation does include a stochastic choice parameter and in that sense captures to some degree the equilibrium effects of quantal choice behavior, and the estimated error parameter is significantly greater than zero.

5.5. Other behavioral theories of turnout

One may wonder whether uninformed voting may be motivated by a sense of *citizen duty*. Voting is compulsory in Mexico as in many other countries, but there is essentially no enforcement and turnout rates are similar to the US. Moreover, in our experiment abstention took the form of voting a blank ballot (in anticipation of this possible confound), which is allowed and common in Mexico as elsewhere. Recently, *ethical voting theories* proposed by Feddersen and Sandroni (2006) and others have provided a rationalization of duty-bound turnout. In our context, ethical voters would not vote when uninformed, which is detrimental to the group. *Expressive voting theory*, in the tradition of Brennan and Buchanan (1984), would not lead to uninformed voting unless uninformed voters are biased in favor of one of the alternatives, which is what our subjective beliefs notion postulates. That explanation also leaves open the question of why previous experiments related to ours did not observe expressive voting. *Conformist voting theories* which postulate a positive utility gain from the act of voting for the winning outcome also fails to organize the data. First, these theories are contradicted by earlier swing voter's curse experiments. Second, in our environment, where information acquisition is endogenous, conformist voters have a stronger incentive to purchase information than rational voters. Thus, to the extent that conformism might explain voting when uninformed, we would expect a higher rate of information acquisition than in the Bayesian equilibrium. But we observe significantly less information acquisition, and, considering our data at an individual level, it also fails to explain why most of these voters *never* gather information, even when their information costs are close to zero.

6. Final remarks

In this paper, we study theoretically and experimentally a group decision problem in which information is costly and therefore it may be rational to remain ignorant. The observed behavior in the experiment is inconsistent with several of the predictions from the standard common-prior Bayesian equilibrium model, and suggests an important role for noisy private beliefs. We develop and estimate a new model, subjective beliefs equilibrium that can account for the deviations from Bayesian equilibrium, and also provides a theoretical basis for understanding the observed heterogeneity in behavior. In the subjective beliefs equilibrium model, in addition to being rationally ignorant, some voters may be biased, in the sense of being subject to random, private shocks to prior beliefs. Therefore, the model proposed has a role both for *ignorance* and *bias* in collective choice environments with costly information, voting, and common values.

One must be cautious in interpreting our finding, as there are at least three critical questions that must be addressed. First, one of our key findings is the very large extent of uninformed voting, which appears to be a form of swing voter's curse. This observation is in stark contrast with the swing voter's curse experiments reported in Battaglini et al. (2008, 2010), where uninformed voting was rare and declined rapidly with experience. This raises the first question: How can we reconcile our findings with theirs, and what would the subjective beliefs equilibrium predict for the environment in those earlier studies?

A key difference between our environment and the Battaglini et al. studies is that private signals are imperfectly informative in our setup, but are perfectly informative in theirs. A second difference is that signals are costly in our framework, but are exogenously assigned in theirs. Because signals are weak in our design, sufficiently biased voters have no incentive to acquire information, even if it were costless. Furthermore, those same voters have subjective priors that swamp the weak signals of informed voters (if there are any), thus making them immune from the swing voter's curse. In a world of perfectly informative signals, as in the Battaglini et al. studies, neither of these effects are present: biased voters, even with very large biases, would still find it valuable to buy perfect information for sufficiently low information costs; and biased voters (even with large biases) would not believe that their prior swamps the information of the informed voters, and hence they would not think they are immune to the winner's curse. The bottom line is that *both* the results in Battaglini et al. and the behavior in our experiment are consistent with the subjective beliefs equilibrium model developed in this paper.

The second critical question to address is the following: Why develop a new model to explain these findings? Aren't the experimental results also consistent with other existing behavioral theories? Unfortunately, current behavioral theories cannot explain the data in a parsimonious way. The most obvious candidate is cursed equilibrium (Eyster and Rabin, 2005). Cursed or generalized cursed equilibria would indeed predict uninformed voting, however it cannot account for the fact that those subjects who vote when they are uninformed are also less likely to acquire information. In fact a considerable fraction of our subjects always or nearly always vote, yet never buy information. To the contrary, cursed voters would be more inclined than rational voters to obtain information, because they essentially act as if other voters are uninformed, so the free rider effects are smaller. Furthermore, cursed equilibrium would also predict uninformed voting in the Battaglini et al. studies, which is inconsistent with the huge difference in the prevalence of uninformed voting in our experiment and theirs.

There is also a natural alternative explanation for the under-acquisition of information in our experiment: loss aversion, as introduced by Kahneman and Tversky (1983). Loss averse voters would be less inclined to pay for information, since the cost of information is a certain loss in case the voter is not decisive, or even worse, of the imperfect signal the voter buys is incorrect and the voter is decisive. However, loss aversion fails to provide an explanation for the extensive uninformed voting we observe. In fact, loss aversion can increase the incentive for uninformed voters to abstain, in particular if the reference point is taken to be the decision reached by the other voters without one's own vote.

The third question concerns the broader applicability of our findings. Does subjective beliefs equilibrium have interesting new implications about behavior in a broader range of environments than the specific common-value voting problem we study? While a definitive answer to this last question is well beyond the scope of our paper, we are cautiously optimistic that the answer will ultimately be "yes." For many variations of the Condorcet jury environment there are likely to be significant effects, for example allowing for private values in addition to the common value component, continuous signals, sequential voting, etc. In the context of common value auction settings, biased bidders' behavior would be influenced not only by their private information but also by their noisy prior beliefs. This, in turn, would dampen the informational content of bids which could affect the degree to which bidders adjust to the winner's curse in equilibrium. In ascending bid auctions, it would change how bidders make inferences from the sequence of observed bids. Further examples that come to mind include, for example, sequential decision environments in which informational cascades are possible, lemon markets, asset trading, and signaling games. In all these instances, understanding the effects of the presence of biased players on the behavior of unbiased players requires careful equilibrium analysis. Looking forward, both theoretical and experimental work seem to be needed to achieve the goal of parsimonious, predictive game theory.

Appendix A.

Proof of Proposition 1. We start from the observation that under neutral strategies expected gain equations reduce to:

$$G(s_A|\epsilon,\sigma_{-i}) = \frac{1}{2}(q+\epsilon)D(0|\sigma_{-i}) + \frac{1}{2}\left(\frac{1}{2}+\epsilon\right)\left(\frac{1}{2}+q\right)D(-1|\sigma_{-i}) - \frac{1}{2}\left(\frac{1}{2}-\epsilon\right)\left(\frac{1}{2}-q\right)D(1|\sigma_{-i}),$$
$$G(s_B|\epsilon,\sigma_{-i}) = \frac{1}{2}(q-\epsilon)D(0|\sigma_{-i}) + \frac{1}{2}\left(\frac{1}{2}-\epsilon\right)\left(\frac{1}{2}+q\right)D(-1|\sigma_{-i}) - \frac{1}{2}\left(\frac{1}{2}+\epsilon\right)\left(\frac{1}{2}-q\right)D(1|\sigma_{-i}),$$
$$G(A|\epsilon,\sigma_{-i}) = \epsilon D(0|\sigma_{-i}) + \frac{1}{2}\left(\frac{1}{2}+\epsilon\right)D(-1|\sigma_{-i}) - \frac{1}{2}(\frac{1}{2}-\epsilon)D(1|\sigma_{-i}),$$
$$G(B|\epsilon,\sigma_{-i}) = -\epsilon D(0|\sigma_{-i}) + \frac{1}{2}\left(\frac{1}{2}-\epsilon\right)D(-1|\sigma_{-i}) - \frac{1}{2}\left(\frac{1}{2}+\epsilon\right)D(1|\sigma_{-i}).$$

The following lemma puts some bounds on what a voter can learn from being decisive, given that other voters play neutral strategies. In particular, the ratio of the probability that the correct alternative is ahead by one vote to the probability that the correct alternative is behind by one vote is bounded below by one, and is bounded above by the informativeness of a single signal, that is $(1/2+q)/(1/2-q)$. This result is useful because it implies that if priors are not too biased, then voters will prefer to abstain if uninformed and will prefer to vote for the alternative favored by signal received if informed. The idea of the proof is to match every vote profile of other voters in which the correct alternative is ahead by one vote with a vote profile in which the incorrect alternative is ahead by one vote, by reversing a single vote cast in favor of the correct alternative. The proof itself is an application of a theorem in graph theory that has been used in the analysis of networks but, to the best of our knowledge, never before in collective choice settings.

Lemma 3. *If other voters are playing neutral strategies, then*

$$1 \leq \frac{D(1|\sigma_{-i})}{D(-1|\sigma_{-i})} \leq \frac{(1/2)+q}{(1/2)-q},$$

where the lower bound is tight if and only if all other voters play uninformative strategies, and the upper bound is tight if and only if all other voters play informative strategies and vote when uninformed with probability zero and make no mistakes ($Q=1$).

Proof. Suppose voters other than i play neutral, informative strategies. From neutrality and symmetry of the distribution of ϵ, for each voter $i' \neq i$ the probability that the voter votes for A while uninformed is equal to the probability that the voter votes for B while uninformed, and the probability that the voter votes for A after receiving signal s_A is equal to the probability that the voter votes for B after receiving signal s_B, where these probabilities are calculated ex ante, taking into account the strategy of voter i' and the distribution of ϵ and c.

For each voter $i' \neq i$, let $\pi(\sigma_{i'})$ be the (ex ante) probability with which voter i' acquires information, $\rho(\sigma_{i'})$ the (ex ante) probability that the voter votes for alternative d after receiving signal s_d, and $\tau(\sigma_{i'})$ the (ex ante) probability that the voter votes for alternative d after not acquiring information, for $d=A, B$. Let $\nu_r(\sigma_{i'})$ and $\nu_w(\sigma_{i'})$ be the probabilities that voter i' votes for the correct and the incorrect alternative, respectively. We have

$$\frac{\nu_r(\sigma_i')}{\nu_w(\sigma_i')} = \frac{\pi(\sigma_{i'})((1/2)+q)\rho(\sigma_{i'}) + (1-\pi(\sigma_{i'}))\tau(\sigma_{i'})}{\pi(\sigma_{i'})((1/2)-q)\rho(\sigma_{i'}) + (1-\pi(\sigma_{i'}))\tau(\sigma_{i'})}.$$

Thus, for all $i' \neq i$,

$$\frac{\nu_r(\sigma_i')}{\nu_w(\sigma_i')} \geq 1,$$

with equality if and only if $\pi(\sigma_{i'})\rho(\sigma_{i'}) = 0$. Similarly, for all $i' \neq i$,

$$\frac{\nu_r(\sigma_i')}{\nu_w(\sigma_i')} \leq \frac{(1/2)+q}{(1/2)-q},$$

with equality if and only if $(1-\pi(\sigma_{i'}))\tau(\sigma_{i'}) = 0$.

Next, we claim that there exists a bijective mapping between the set of voting profiles such that the correct alternative wins by one vote and voting profiles such that the incorrect alternative wins by one vote, where only one voter needs to be switched from voting for the correct alternative to voting for the incorrect alternative to go from the profile where the correct alternative wins to the profile where the incorrect alternative wins. To see this, consider any subset of voters $C \subset \{1, \ldots, n\} \setminus \{i\}$ such that $|C|$ is odd. Define R_C to be the set of voting profiles such that voters abstain if and only if they are not in C, and the correct alternative wins by one vote. Similarly, define W_C to be the set of voting profiles such that voters abstain if and only if they are not in C, and the incorrect alternative wins by one vote.

Consider a graph where the vertices are the elements of $R_C \cup W_C$, and the edges are

$$E_C = \{(r, w) \in R_C \times W_C : r \text{ and } w \text{ differ by the vote of a single individual}\}.$$

Note that every element of R_C has $(|C|+1)/2$ edges incident to it, the same being true for every element of W_C. By Konig's Marriage Theorem (see Theorem 2.5 in Balakrishnan, 1995), there exists a perfect matching M_C, assigning to each voting profile in R_C a voting profile in W_C. Since this is true for every C, there exists a bijective mapping

$$f : \cup_C R_C \to \cup_C W_C$$

given by $f(r) = w$ such that $(r, w) \in M_C$ for any $r \in R_C$, and such that f assigns to each voting profile $r \in \cup_C R_C$ a unique profile $f(r) \in \cup_C W_C$ where a single voter switches to the "mistaken" side. Thus, the probability ratio between the two profiles is equal to the ratio between the probability of that voter being right and that voter being wrong, which, as we have already established, must be in the interval $[1, ((1/2)+q)/((1/2)-q)]$. Since, furthermore, $D(1|\sigma_{-i})$ is the sum of the probabilities of the voting profiles such that the correct alternative wins by one vote, and $D(-1|\sigma_{-i})$ is the sum of the probabilities of the voting profiles such that the incorrect alternative wins by one vote, the ratio of $D(1|\sigma_{-i})$ to $D(-1|\sigma_{-i})$ is also in that interval.□

We now put to work Lemmas 1 and 3. We claim that if other voters play neutral strategies, a best responding voter will play a neutral, informative strategy. To see this, assume all voters other than i play neutral strategies. Using Lemma 2, we get that if the voter is unbiased ($\epsilon = 0$),

$$G(s_A|0, \sigma_{-i}) > 0 \quad \text{and} \quad G(s_B|0, \sigma_{-i}) > 0.$$

Suppose voters other than i play informative strategies. Then, from Lemma 3, $D(1|\sigma_{-i}) > D(-1|\sigma_{-i})$. Thus, for unbiased voters,

$$G(A|0, \sigma_{-i}) < 0 \quad \text{and} \quad G(B|0, \sigma_{-i}) < 0.$$

That is, if other voters play informative strategies, an unbiased voter abstains when uninformed, and votes according to the signal received when informed. Moreover, from Lemma 1, the value of information for unbiased voters is

$$c(0, \sigma_{-i}) = qD(0|\sigma_{-i}) + \frac{1}{2}\left(\frac{1}{2}+q\right)D(-1|\sigma_{-i}) - \frac{1}{2}\left(\frac{1}{2}-q\right)D(1|\sigma_{-i}). \tag{6}$$

From Lemma 3, $c(0, \sigma_{-i}) > 0$. That is, the cutoff for information acquisition is strictly positive for unbiased voters if other voters play informative strategies.

Suppose instead that voters other than i play uninformative strategies. Since $G(s_A|0, \sigma_{-i}) > 0$ and $G(s_B|0, \sigma_{-i}) > 0$, an unbiased the voter will be indifferent if uninformed, and will vote for the alternative favored by their signal if informed. The value of information for this voter, using Lemma 1, is then

$$c(0, \sigma_{-i}) = q\left(D(0|\sigma_{-i}) + \frac{1}{2}D(-1|\sigma_{-i}) + \frac{1}{2}D(1|\sigma_{-i})\right)$$

as in the statement of Proposition 1. Note that $c(0, \sigma_{-i}) > 0$. That is, the cutoff for information acquisition is strictly positive for unbiased voters if other voters play uninformative strategies. Since by assumption $p < 1$, a symmetric, neutral equilibrium is necessarily informative. □

Proof of Proposition 2. It is straightforward to check that there are no equilibria in which voters acquire information with positive probability, vote for the alternative favored by the signal received, and abstain if uninformed (this echoes the result of Feddersen and Pesendorfer, 1996). The reason is that, if other voters adopt this strategy, $H(s_A|\epsilon, \sigma_{-i}) < 0$, then the best response would be to abstain rather than vote for A after signal s_A. Similarly, there are no equilibria in which voters acquire information with positive probability, vote for B after signal s_B, and abstain otherwise.

Next, it is straightforward to check that, in a symmetric strategy profile, $H(s_B|0, \sigma_j) \leq 0$ implies that best-responding voters who receive a signal s_A abstain with positive probability, which in turn implies $H(s_A|0, \sigma_j) \leq 0$. Thus, there is no informative equilibrium strategy such that $H(s_B|0, \sigma_j) \leq 0$. Similarly, $H(s_A|0, \sigma_j) < 0$ implies that best-responding voters who receive a signal s_A do not vote for A, which in turn implies $H(s_A|0, \sigma_j) > 0$, a contradiction. Thus, there is no informative equilibrium strategy such that $H(s_A|0, \sigma_j) < 0$.

Finally, $H(s_B|0, \sigma_j) > 0$ and $H(s_A|0, \sigma_j) > 0$ imply that best-responding voters vote for the alternative favored by the signal received, which in turn implies $H(s_A|0, \sigma_j) < 0$, unless uninformed voters vote for B with positive probability, corresponding to equilibria described in the second part of the theorem. The only remaining possibility is $H(s_B|0, \sigma_j) > 0$ and $H(s_A|0, \sigma_j) = 0$, corresponding to equilibria described in the first and second parts of the proposition.□

Appendix B: Experiment instructions.

What follows is a translation of the Spanish-language instructions for the experiment. Minor differences between treatments are shown in italics.

Instructions

Thank you for accepting to participate in this experiment about decision-making. During the experiment we shall require your complete attention and careful following of instructions. Furthermore, you will not be allowed to open other computer applications, talking with other participants, or doing other things that may distract your attention, such as using your cell phone, reading books, etc.

At the end of the experiment, you will be paid for your participation in cash. Different participants may earn different amounts. What you earn will depend, in part, on your decisions, in part on decisions of other participants, and in part on chance.

The experiment will be administered via computer terminals, and all the interaction between the participants shall happen through these computers. It is important that you do not talk or try to communicate in any manner with other participants during the experiment.

During the instruction period, you shall receive a complete description of the experiment and will be shown how to use the computers. If you have any questions during the instruction period, raise your hand and the question will be responded to in a loud voice, so that everybody can hear it. If you have any questions after the experiment has started, raise your hand and an experimenter will approach you and try to help you.

This experiment will continue for 25 periods. At the end of the experiment you shall be paid what you earned, plus a participation fee of $20 pesos. Everybody will receive his/her payment privately and will not be obliged to tell the others how much s/he earned. Your earnings during the experiment will be denominated in points. At the end of the experiment you will be paid $8 pesos for every 1000 points you earned.

Now we start a short instruction period, to be followed by a practice session. You will not be paid for the practice session. After the practice session, there will be a short comprehension test, to which you have to respond correctly before continuing to a session that will be paid for.

Rules of committee formation

We start the first period by dividing into seven (7) Committees of three (3) members each. <*Note: in two of the treatments, the instructions instead stated here "three (3) committees of seven (7) members each," with the appropriate substitution following through the rest of the text*>. Each one of you will be assigned by the computer to exactly one of these seven (7) Committees. You will not know the identity of the other members of the Committee to which you belong.

The committee decision task

Your Committee will have to decide between one of two options, which we shall call the Red Jar and the Blue Jar. The Committees will be making their decisions using the following voting procedure:

<*Instructions for majority treatments*>

The final decision of the Committee shall be the option which obtains the largest number of votes. In case of a draw in the number of votes, or voting in blank by all Committee members, the final decision of the Committee shall be determined randomly, with probability $\frac{1}{2}$ for each box color type.

<*Instructions for unanimity treatments*>

The Red Jar shall be elected by the Committee only when everybody who decides not to vote in blank, votes for the Red Jar, otherwise the Blue Jar shall be chosen by the Committee.

In other words, the Blue Jar shall be chosen by the Committee when at least one of the voters votes for the Blue Jar or everybody votes in blank.

Jar assignment by the computer

At the beginning of each period the computer will randomly assign one of two options as the correct Jar for your Committee. In each period there is a 50/50 chance that the Jar assigned is Red or Blue.

The computer will choose randomly the correct jar for each Committee and separately for each period. Therefore, the chance that your Committee is assigned a Red Jar or a Blue Jar shall not be affected by what happened in previous periods or by what is assigned to other Committees. The choice shall always be completely random in each period, with a probability of 50% for the Red Jar and 50% for the Blue Jar.

Buying information

You will not be informed of what is the correct Jar that is assigned to your Committee until after the Committee has chosen one of the options.

However, before the Committee decides on the option, each Committee member will have an opportunity to buy a piece of information about the color of the correct Jar assigned. During the practice period we shall explain exactly how this works.

If you decide to buy information about the color of the correct Jar assigned to your Committee, the cost of your purchase will be subtracted from your earnings. The cost of buying information which you would pay, which we shall call sampling cost, shall be equal to a number randomly chosen between 1 and 100 points.

You will be informed of the sampling cost before you decide whether to buy information, but you will not be informed about the costs of other members of your Committee.

These costs will be assigned randomly and independently for each of the Committee members and for each period. Any number between 1 and 100 points would have the same chance of being chosen.

Voting options

After all Committee members have decided, independently of each other, whether to buy information or not, every one of them will have to choose between:

- Voting for the Red Jar,
- Voting for the Blue Jar, or
- Voting in Blank.

<Next line of Instructions was read only for the unanimity treatments>

A Vote in Blank shall count in favor of the Blue Jar only if everybody decides to vote in blank.

After every member of your Committee has voted, the computer will count the votes in order to determine the final Committee decision.

Committee decision rule

<Instructions for majority treatments>

The Committee decision is determined using the majority rule.

<Instructions for unanimity treatments>

The Committee decision is determined using the unanimity rule, with <unanimity necessary> to decide in favor of the Red Jar; otherwise, the Blue Jar shall be chosen.

Payments for committee decisions

If your Committee decision is equal to the color of the Jar that was assigned, every Committee member will earn 1000 points.

If your Committee decision is not equal to the color of the Jar that was assigned, every Committee members will earn 0 points.

From the earnings of the Committee members who will have acquired information sampling costs will furthermore be subtracted.

Committee independence

Other Committees in the room will deal with similar problems, but the correct Jar assigned to each committee shall be different from that of other Committees. Remember that the Committees are completely independent and act independently.

After completing the first period, we proceed to the second period. You will be regrouped randomly into seven (7) new Committees and the process will repeat itself. This will continue for a total of 25 periods.

Description of the screen and the software

We now start the session and go to a practice period in order to familiarize ourselves with the experimental equipment.

During the practice period, please do not touch the keys, until you are asked to do so, and when you are instructed to enter certain information, please do exactly what you are asked to do.

Once more, you will not be paid for the practice period.

We now shall see the first experimental screen on the computer. You should see a similar screen in front of the room.

Please keep in mind that the screen shown in front is not necessarily identical to the screen that appears on your computer right now. The slides we show in front are only for illustration purposes. At the top left of the screen you will see your identification.

This ID shall be the same during the entire experiment. Please note it on the registration sheet that we have given to you.

Since this is the beginning of a period, you have been assigned by the computer to one of the seven committees of 3 members. This assignment will change every period.

At the top right of the screen you see the two Jars, each one containing exactly 12 balls. The Red Jar contains 8 Red balls and 4 Blue balls. The Blue Jar contains 8 Red balls and 4 Blue balls.

The Computer shall randomly assign one of the two jars to your Committee. In each period, the chance is 50/50 that the assigned Jar is Red or Blue. The assignment will be done 7 times, once for each Committee. Therefore, the seven committees in this period may have different jars.

You will not know whether the correct Jar for your Committee in this period is Red or Blue until after all the members of your Committee will have voted either Red, Blue or in Blank and the Committee decisions is determined. Before voting, one will have a chance to pay the cost and buy the information that may help you to determine the correct color of the Jar assigned to your group. Please wait and we will explain how to do it in a moment.

In front of the room you see a screen which shows how to determine the earnings. If the Committee decision coincides with the color of the Jar assigned to your Committee, you (and every one of the members of your Committee) will earn 1000 points for the period, and you will earn 0 points if the Committee decision does not coincide.

After the computer assigns a jar to each Committee, you shall see the following screen. Now you only see one Jar on the screen, but the colors of the balls are hidden, so at this point you cannot say which Jar has been assigned. This is the correct Jar assigned to your Committee. If it is the Red Jar it has 8 Red balls and 4 Blue balls; if it is the Blue Jar, it has 8 Blue balls and 4 Red balls.

Please keep in mind that the balls have been reordered randomly in each of the screens by the main computer, so that it is impossible to guess the location of the balls of each color and you cannot know which Jar has been assigned to your Committee.

At this point you will have an opportunity to pay a cost between 1 and 100 points to see the color of exactly one of the balls in the Jar assigned to your Committee. Your cost has been chosen randomly. Any cost between 1 and 100 points has equal probability of being assigned. The costs are assigned randomly and independently to each Committee member. These costs will also be randomly and independently chosen for each period. Your cost for this period is showing on the screen.

If you do not want to pay the cost, simply click the button that says “Do Not Observe”. In this case you will not obtain any information about the correct Jar assigned to your Committee. Otherwise, if you would like to pay the information cost, simply move the cursor to any of the balls in the Jar and click once. Please wait and do not click for the time being.

If you pay the sampling cost and click on one of the balls, we shall call it your “sample ball” for this period. This ball is your private information. Other Committee members will also have an opportunity to acquire a simple ball in the same manner, though locations of the balls in the Jar are ordered differently for each member, and different members normally may have different costs. Therefore, different members of the same Committee may be clicking on balls of different color even for the same Jar. Nevertheless, if the Jar is Red, the Red balls twice as likely to be chosen as the Blue balls, and if the Jar is Blue, the Blue balls are twice as likely to be chosen as the Red balls. The colors of other balls will stay hidden until the end of the period. You will not know how many other Committee members will have decided to buy a sample ball and how many decided not to buy. Now continue and make your choice, making a click on a ball, or clicking the “Do Not Observe” button.

We now go to the voting stage.

At this point you will have three options: vote in favor of the Red Jar, vote in favor of the Blue Jar, or cast a Blank vote. There are three buttons on the screen, which say “Red”, “Blue”, and “in Blank”. You can cast your vote by clicking on the corresponding button. Since this is a practice period, we shall not let you choose. Instead, we will ask that you vote according to your identification. If your identification number is between 0 and 7, please vote red. If your identification number is between 8 and 14, please vote blue. If your identification number is between 15 and 21, please vote in blank. Of course, during the periods played for money you will be making your own decisions.

<*Instructions for majority treatments*>

Remember that only Blue and Red votes will count for the Committee decisions, which shall be made by the majority. Ties will be resolved randomly.

<*Instructions for unanimity treatments*>

Remember that unanimity is needed to choose the Red Jar. The Committee decision will be the Red Jar only if everyone who decides not to vote in blank votes for the Red Jar, otherwise the Blue Jar will be chosen. In other words, if at least one vote is for the Blue Jar, or if everybody votes in blank, the Blue Jar will be chosen.

We are now ready for a short comprehension test. Everybody has to respond to all the questions correctly before we proceed to periods to be paid for. Also, during the test you must respond to all the questions on page 1 in order to move to

Fig. 3. Sample screen for the jar interface.

page 2. If you answer a question incorrectly, you will be asked to correct your answer. Please raise your hand if you have any questions during the test, so that we can come to your desk and respond to your question privately.

Once everybody has voted and finished the test, you will be informed of the final Committee decision, as well as of the correct Jar assigned to your Committee. Likewise, there will appear a small screen which will inform you of your earnings for the period, which shall be equal to zero, because we are in the practice period, which is not being paid for. Please close this little screen so that we may continue.

In the large screen at the end you will be shown how many votes were received by each Jar and how many voters decided to cast a blank vote. Also, please take into account that, at the end of the period the colors of all the balls in your Jar shall be revealed. This screen will mark the end of a period.

You are also shown your total earnings. Please click the “Accept” button to end the practice period.

COLUMN ONE shows the period number; COLUMN TWO shows your sampling cost; COLUMN THREE shows the color of the sample ball or says “Not Observed” if you decided not to buy the sample ball; COLUMN FOUR lists your vote; COLUMN FIVE provides the summary of the votes in the following order (RED JAR – BLUE JAR – IN BLANK); COLUMN SIX shows the Committee decision; COLUMN SEVEN shows the correct JAR assigned to the Committee; COLUMN EIGHT shows your earnings (do not appear now, because we are in practice period).

The table with columns in the bottom of the screen shows the history that includes all the key information for each period.

To sum up, please, remember the following important things.

<Instructions for the majority treatments>

The Committee decision is taken by the majority rule, with ties resolved randomly.

<Instructions for unanimity treatments>

The Committee decision is taken by the rule requiring unanimity to decide for the Red Jar. The Committee decision shall be Red Jar only if everybody who decided not to vote in Blank voted for the Red Jar; otherwise, it will be the Blue Jar. In other words, if at least one person voted for the Blue Jar, or if everybody votes in blank, the Blue Jar shall be chosen.

The Committee decisions are summed up in the text panel on the top left of the screen, and are also summed up in the history screen at the bottom of column five. While the experiment continues, the history screen will gradually show the information about all the previous periods in which you will have participated.

Are there any questions before we start the session that will be paid for?

We now start with the 25 paying periods of the experiment. If you have any problems or questions from now on, please raise your hand, and we will come by to help you in private.

References

Austen-Smith, D., Banks, J., 1996. Information aggregation, rationality, and the Condorcet jury theorem. Am. Polit. Sci. Rev. 90, 34–45.

Balakrishnan, V.K., 1995. Network Optimization. Chapman and Hall, London.

Battaglini, M., Morton, R., Palfrey, T.R., 2008. Information aggregation and strategic abstention in large laboratory elections. Am. Econ. Rev. 98 (2), 194–200.

Battaglini, M., Morton, R., Palfrey, T.R., 2010. The swing voter’s curse in the laboratory. Rev. Econ. Stud. 77, 61–89.

Bhattacharya, S., Duffy, J., Kim, S., 2014. Compulsory versus voluntary voting mechanisms: an experimental study. Games Econ. Behav. 84, 111–131.

Brennan, G., Buchanan, J., 1984. Voter choice. Am. Behav. Sci. 28, 185–201.

Camerer, C., Ho, T.-H., Chong, K., 2004. A cognitive hierarchy model of behavior in games. Q. J. Econ. 119, 861–898.

Camerer, C., Weigelt, K., 1988. Experimental tests of a sequential equilibrium reputation model. Econometrica 56, 1–36.

Condorcet, J.A.N.C.M., (1785) [1986]. Discours préliminaire de l’Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix, reprinted in Condorcet, Sur les Élections et autres textes. Libraire Artheme Fayard, Paris.

Che, Y.-K., Kartik, N., 2009. Opinions as incentives. J. Polit. Econ. 117, 815–860.
Daley, B., Sadowski, P., 2016. Magical thinking: a representation result. Theor. Econ. (forthcoming).
Esponda, I., Vespa, E., 2014. Hypothetical thinking and information extraction in the laboratory. Am. Econ. J. Microecon. 6 (4), 180–202.
Eyster, E., Rabin, M., 2005. Cursed equilibrium. Econometrica 73, 1623–1672.
Feddersen, T., Pesendorfer, W., 1996. The swing voter's curse. Am. Econ. Rev. 86, 408–424.
Feddersen, T., Pesendorfer, W., 1998. Convicting the innocent: the inferiority of unanimous jury verdicts under strategic voting. Am. Polit. Sci. Rev. 92, 23–35.
Feddersen, T., Sandroni, A., 2006. A theory of participation in elections. Am. Econ. Rev. 96, 1271–1282.
Friedman, J., Mezzetti, C., 2005. Random belief equilibrium in normal form games. Games Econ. Behav. 51, 296–323.
Gerardi, D., Yariv, L., 2008. Information acquisition in committees. Games Econ. Behav. 62, 436–459.
Gershkov, A., Szentes, B., 2009. Optimal voting schemes with costly information acquisition. J. Econ. Theory 144 (1), 36–68.
Goeree, J.K., Yariv, L., 2010. An experimental study of collective deliberation. Econometrica 79, 893–921.
Grosser, J., Seebauer, M., 2016. The curse of uninformed voting: an experimental study. Games Econ. Behav. 97, 205–226.
Guarnaschelli, S., McKelvey, R.D., Palfrey, T.R., 2000. An experimental study of jury decision rules. Am. Polit. Sci. Rev. 94, 407–423.
Kahneman, D., Tversky, A., 1983. Choices, values and frames. Am. Psychol. 39, 341–350.
Koriyama, Y., Szentes, B., 2009. A resurrection of the Condorcet jury theorem. Theor. Econ. 4, 227–252.
Martinelli, C., 2006. Would rational voters acquire costly information? J. Econ. Theory 129, 225–251.
Martinelli, C., 2007. Rational ignorance and voting behavior. Int. J. Game Theory 35, 315–335.
McKelvey, R.D., Palfrey, T.R., 1995. Quantal response equilibria for normal form games. Games Econ. Behav. 10, 6–38.
McKelvey, R.D., Palfrey, T.R., 1998. Quantal response equilibria for extensive form games. Exp. Econ. 1, 9–41.
Mukhopadhaya, K., 2005. Jury size and the free rider problem. J. Law Econ. Organ. 19, 24–44.
Nyarko, Y., Schotter, A., 2002. An experimental study of belief learning using elicited beliefs. Econometrica 70, 971–1005.
Oliveros, S., 2013. Abstention, ideology, and information acquisition. J. Econ. Theory 148 (3), 871–902.
Palfrey, T.R., Wang, S.W., 2009. On eliciting beliefs in strategic games. J. Econ. Behav. Organ. 71, 98–109.
Persico, N., 2004. Committee design with endogenous information. Rev. Econ. Stud. 71, 165–191.
Rogers, B.R., Palfrey, T.R., Camerer, C., 2009. Heterogeneous quantal response equilibrium. J. Econ. Theory. 144 (4), 1440–1467.
Stahl, D., Wilson, P., 1995. On players' models of other players: Theory and experimental evidence. Games Econ. Behav. 10, 154–218.
Sethi, R., Yildiz, M., 2012. Public disagreement. Am. Econ. J. Microecon. 4, 57–95.
Sethi, R., Yildiz, M., 2016. Communication with unknown perspectives. Econometrica (forthcoming).
Triossi, M., 2013. Costly information acquisition. Is it better to toss a coin? Games Econ. Behav. 82, 169–191.
Tyson, S.A., 2016. Information acquisition, strategic voting, and improving the quality of democratic choice. J. Polit. 78 (4), 1016–1031.