The Role of Information in Collective Decisions

Nicolás Figueroa, José-Alberto Guerra, Francisco Silva

THE ROLE OF INFORMATION IN COLLECTIVE DECISIONS

Nicolás Figueroa
Pontificia Universidad Católica de Chile, Instituto de Economía, Chile

José-Alberto Guerra
Universidad de los Andes, Department of Economics, Colombia

Francisco Silva
Deakin University, Department of Economics, Australia

Abstract
We study how members of a group vote for public information. We argue, both theoretically and experimentally, that voters are more likely to vote for information to be acquired relative to their own individual willingness to pay for information when ex-ante disagreement is higher and ex-post disagreement is lower. Ex-ante and ex-post disagreement refer to the disagreement among group members over the best policy for the group to follow before and after information is acquired, respectively. We discuss how the results inform the debate over the role of the State in fostering progress, and the value of the wisdom of the crowd. (JEL: D71, D83, C92)

1. Introduction

In a democracy, many of the decisions that shape society are made through voting: citizens vote for the president, parliament members vote on whether to approve the

The editor in charge of this paper was Nicola Pavoni.

Acknowledgments: We thank Santiago Neira and Maria Alejandra Paez for their outstanding research assistance. We thank audiences at BSE Summer Forum, ANZWEE, BEBES, Universidad de los Andes, and Pontificia Universidad Católica (PUC) de Chile, as well as the anonymous referees for their valuable suggestions. This project received financial support from ANID PIA/PUENTE AFB230002, Instituto Sistemas Complejos de Ingeniería. The experimental work of this paper was approved by Universidad de los Andes IRB (Act 1253 of 2020).

nicolasf@uc.cl (Figueroa); ja.guerra@uniandes.edu.co (Guerra); f.silva@deakin.edu.au (Silva)

Journal of the European Economic Association 2026 0(0):1–35
https://doi.org/10.1093/jeea/jvag001

budget, jurors vote on the guilt of a defendant, etc. If preferences are single peaked over the space of available policies, it is well understood which policies get selected (e.g., under simple majority rule, the winning policy is the most preferred by the median voter). However, it is often the case that voting is also used to determine the amount of information available to the group before a policy decision is made.

As an example, say that Parliament decides over whether to legalize the consumption of marijuana. Prior to the vote, Parliament has the option to obtain more information by authorizing a pilot project that legalizes marijuana only in part of the country.

From a conceptual point of view, voting on whether to legalize marijuana consumption is straightforward: Each parliament member only has to determine her preferences and then vote accordingly (assuming parliament members have independent private information). However, the vote on whether to authorize the pilot is not as straightforward; information is not valuable in and of itself, but rather a means to an end. Therefore, voters must not only consider the impact that information might have in shaping their own opinions but also its impact on other voters. A voter who is more pre-disposed to vote against the legalization of marijuana might prefer to vote against the pilot out of fear that support for the opposite view grows in light of the pilot results.

In this paper, we discuss the extent to which voting incentivizes information acquisition. Do group members tend to vote more or less in favor of information acquisition compared to what they would have done individually as if they were dictators? And what are the implications of these differences?

Ultimately, the extent to which voting incentivizes information acquisition is determined by whether each individual group member (agent) is more or less willing that the group acquires information. We argue, both theoretically and experimentally, that a key determinant in how willing each agent is for the group to obtain information is the role that information has in reducing or increasing the level of likely disagreement among the agents. Specifically, we construct a simple theoretical model of collective information acquisition and define the concept of *willingness to vote* (wtv) for information as the maximum (individual) cost a given agent would be willing to pay for the group to obtain information. We then prove that the difference between each agent's willingness to vote and willingness to pay (wtp) for information (which represents the maximum cost the agent would be willing to pay for information if she was deciding individually) is increasing with *ex-ante disagreement* (a measure of the likely disagreement between the various agents if information is not obtained) and decreasing with *ex-post disagreement* (a measure of the likely disagreement between the various agents if information is obtained). In the second part of the paper, we show the results of an experiment aimed at testing this theoretical prediction. We find that, if information is individually valuable (if $wtp > 0$), subjects behave in the manner predicted by the theory: $(wtv - wtp)$ is increasing with ex-ante disagreement and decreasing with ex-post disagreement.

We start by discussing an example in Section 2, which describes how each agent's lack of control over the group's final decision makes the problem strategically different from an individual problem. If there is substantial ex-ante disagreement, it is unlikely that the group will make the same decision as the individual in case information is not obtained, which makes agents more interested in the group obtaining information. By contrast, from the point of each agent, ex-post disagreement increases the likelihood of an undesirable choice after information has been obtained, which makes agents less willing to vote for information. Thus, ex-ante and ex-post disagreement determine whether agents are more or less interested in the group acquiring information relative to their own individual willingness to pay for information.

In Section 3, we present a decomposition result, which shows that the difference between each agent's wtv and wtp is linearly increasing on ex-ante disagreement and linearly decreasing on ex-post disagreement. The results hold for any information acquisition technology, for any decision rule (majority rule, unanimity, etc.) and if agents have private information (provided the agents' private information is independent of each other and of the public information that can be acquired by the group). We then test this result in Section 4 by conducting an experiment with subjects recruited from a set of college students from a top Colombian University.

In the experiment, each subject was placed in a group with two computerized players (bots) who behaved as to maximize their own expected payoffs. The experiment had several rounds and each round had two stages; an information acquisition stage and a final decision stage. In the information acquisition stage, the subject chose whether information should be obtained or not by the group. In the second stage, the group of players (the subject and the two bots) determined which policy to follow—"status quo" (Q) or "change" (C)—through voting. The winning alternative was the one with the majority of the votes. Payoffs depended on the chosen policy and on a binary public signal that would be revealed to the group if and only if information was obtained in the information stage.

In the first round, both bots' payoffs were known to be equal to the subject's payoffs. This made the subject's problem of whether to obtain information strategically equivalent to an individual problem (because the subject would be able to infer that the bots' votes would be the same as her's), which allowed us to recover each subject's wtp. In the following rounds, we generate ex-ante and/or ex-post disagreement by varying the bots' payoff functions while keeping the subject's constant. In the information stage of each round, each subject was asked whether they would like to obtain information by paying their own wtp. In that way, we recover whether each subject's wtv was larger or smaller than their own wtp in each round. The bots' payoff functions across rounds were designed to generate enough variation of ex-ante and ex-post disagreement to be able to estimate the impact of each type of disagreement on the probability that each subject's wtv was larger than or equal to their wtp.

We find that the results depend on whether information is individually valuable, that is, whether the signal's realization determines the subjects' preferred policy. If information is individually valuable, we find sufficient evidence to support the hypothesis that agents are more willing to vote for information acquisition when there is more ex-ante disagreement and less ex-post disagreement. If information is not individually valuable, the results are inconclusive, which suggests that information being individually valuable triggers the sort of strategic behavior the theory discusses.

We note that our results do not depend on the way in which information is obtained by the group: they hold if groups acquire information by paying a fee that is equal among the group members (e.g., Ginzburg and Guerra 2019), if groups acquire information by experimenting (e.g., Anesi and Bowen 2021; Gieczewski and Kosterina 2024; Khromenkova 2017; Strulovici 2010) and if groups acquire information by delaying a final decision (e.g., Anesi and Safronov 2023; Chan et al. 2018; Louis 2015). In this way, the paper makes a contribution to the literature on collective information acquisition by arguing that the way in which information is obtained by the group is not important in and of itself in determining the group members' willingness to vote for information; what is relevant is the ex-ante and ex-post disagreement. Take, as an example, the result in Strulovici (2010) and Khromenkova's (2017) collective experimentation papers that groups underacquire information—group members vote to stop experimentation sooner than any agent would have in the corresponding individual problem.[1] Our results demonstrate that this result of under-acquisition of information is not a feature of collective information acquisition per se, or even of collective experimentation, but rather a consequence of assuming that every group member is ex-ante the same (before experimentation begins), so that there is no ex-ante disagreement. The only disagreement is ex-post: after each period, some agents will have different preferences over the various options available. If there was ex-ante disagreement and the experimentation led to ex-post agreement one should expect the opposite result. Indeed, in Anesi and Safronov (2023) and Anesi and Bowen (2021), there is ex-ante disagreement that is reduced over time as information is acquired. In Chan et al. (2018) and Louis (2015), there is both ex-ante and ex-post disagreement.[2]

Our study also contributes to the experimental literature on collective information acquisition (see Martinelli and Palfrey 2020 for a recent survey). Ginzburg and Guerra (2019) find that groups are substantially more likely to vote against acquiring information whenever group members have conflicting public preferences. Freer, Martinelli, and Wang (2020) and Hudja (2019) experimentally test Strulovici's (2010) collective experimentation model. While the former shows that a simple majority voting rule reduces inefficient under-experimentation, the latter finds that groups under-experiment against equilibrium predictions. Reshidi et al. (2025) experimentally test static and sequential (as in Chan et al.'s (2018) model) information acquisition models where committee members have similar preferences.[3]

1. In Gieczewski and Kosterina (2024), a similar phenomenon occurs but is counteracted by the entry of more optimistic agents in the group who increase the group's willingness to experiment.

2. More tangentially related to our work is the literature on collective sequential search, where a group of agents vote on whether to select the available alternative or search for a new one (see Albrecht, Anderson, and Vroman 2010; Compte and Jehiel 2010). The main difference is the absence of recall. In these papers, the group is unable to select previously searched alternatives; the group either sticks with the current alternative or asks for a new alternative with random characteristics. Therefore, each voter's value of stopping the search is the same, regardless of whether the decision is made within a group or individually; when deciding collectively, it is only the option to search that is worsened by the fact that each agent does not have complete control over future decisions. That is why Albrecht, Anderson, and Vroman (2010) find that each agent has lower standards of acceptance than when deciding individually.

3. Our paper is also more loosely related to the work that experimentally tests collective search models (see Hizen, Kawata, and Sasaki 2013; Mak, Rapoport, and Seale 2014; Mak et al. 2019).

Our results have a number of implications. First, they suggest that democratic institutions might actually be halting innovation and progress. Progress requires change and change, by its very risky nature, is likely to be associated with ex-ante agreement against it. Indeed, sufficient support in favor of change usually requires research, testing and experimentation before enough trust is built on that policy. But if State institutions (like Parliament) are the ones determining whether information is obtained or not (through authorizing pilot programs for example), then our results suggest that the reluctance in obtaining information might lead to an inefficiently low rate of progress. Therefore, civil society and independent researchers might be more important in generating the information that precedes progress than democratic institutions with heterogeneous members.

Second, as we discuss in more detail in Section 5.1, our results also inform the debate over the value of the wisdom of the crowd. Is it better to rely on one expert's opinion or on the opinion of several non-experts to find out the right policy (assuming there is one)? The Condorcet jury theorem (as originally stated) analyzes this problem through statistics: if there are enough non-experts, their combined information will be better than the expert's. Our results allow us to go beyond the statistical comparison between the two options. In contexts where it is likely that a right answer will be found (like in a criminal trial), information reduces disagreement between group members. This implies that, even if the group of non-experts is less informed than the expert, it may be more likely to recommend the right policy, because each group member has more incentives to vote for information to be acquired. In that sense, our results can be interpreted as providing an argument in favor of juri trials as opposed to bench trials, where a single judge determines the guilt of the defendant.

Third, as we discuss in Section 5.2, over and under-acquisition of information may lead to Pareto inefficiencies: it might be that information is costless and valuable, but the group votes against it; it might also be that the group acquires information even when it is costly and has no value. If there is complete information, we prove that these inefficiencies can be ruled out whenever the agents' preferences have a non-reversal property that is typical of models of experimentation, delay and committee voting. However, they re-emerge if agents have private information. We also discuss how these inefficiencies can be mitigated by moving away from voting mechanisms and considering Vickrey-Clark-Groves mechanisms' instead, where group members report their private preferences rather than voting, and the mechanism determines the group's decision and the transfers paid/received by each group member.

2. Example

Jane must decide between two policies labeled Change (C) and Status Quo (Q). While the payoff from choosing Q is always 0, the payoff from choosing C depends on the realization of an unknown *state of the world* $\omega \in \{\omega^L, \omega^H\}$. Specifically, let us say that the payoff from choosing C is -1 if $\omega = \omega^L$ and 6 if $\omega = \omega^H$, so that Jane prefers Q if $\omega = \omega^L$ and C if $\omega = \omega^H$. How much should Jane pay for the opportunity to learn ω if both states are equally likely? Jane's *willingness to pay* (wtp) for information is given by

$$wtp = \frac{6}{2}\bar{\tau}(\omega^H) - \frac{1}{2}\bar{\tau}(\omega^L) - \frac{5}{2}\bar{\tau}^*,$$

where $\bar{\tau}(\omega)$ represents the probability that C is chosen by Jane once the state is revealed to be equal to $\omega = \omega^H, \omega^L$, and $\bar{\tau}^*$ represents the probability that C is chosen by Jane if information is not acquired. It trivially follows that $\bar{\tau}(\omega^H) = \bar{\tau}^* = 1$, and $\bar{\tau}(\omega^L) = 0$, so that $wtp = 1/2$. Jane's value of information is derived from increasing the probability that Jane's preferred choice is chosen in each state.

Let us now say that Jane is in a group with a few other people. The group faces the same decision: first, it decides on whether the state should be revealed, and then on whether C or Q should be chosen. Crucially, if the state is revealed, every group member learns it—group members cannot learn individually. Both decisions are made sequentially by majority rule—first, the group votes on whether the state should be revealed and then on the policy. If the group decides to acquire information, each group member pays some cost $c \geq 0$.

When does Jane vote in favor of information acquisition? If the difference between her expected payoff if information is acquired and her expected payoff if information is not acquired is greater than or equal to c. Therefore, for Jane to calculate her *willingness to vote* (wtv) for information, she must forecast the group's decision over C and Q in case information is and is not acquired. In particular,

$$wtv = \frac{6}{2}\tau(\omega^H) - \frac{1}{2}\tau(\omega^L) - \frac{5}{2}\tau^*,$$

where $\tau(\omega)$ and τ^* represent, respectively, the probability that C is chosen by the group in state ω if information is acquired, and the probability that C is chosen by the group if information is not acquired. Therefore,

$$wtv - wtp = \frac{6}{2}(\tau(\omega^H) - \bar{\tau}(\omega^H)) - \frac{1}{2}(\tau(\omega^L) - \bar{\tau}(\omega^L)) - \frac{5}{2}(\tau^* - \bar{\tau}^*).$$

As the reader will note, the difference between $\tau(\omega)$ and $\bar{\tau}(\omega)$ represents a measure of *ex-post* disagreement in state ω between Jane's favorite policy and the group's collective decision. The smaller it is, that is, the closer $\tau(\omega)$ is to $\bar{\tau}(\omega)$, the smaller is the difference between Jane's wtv and wtp. Likewise, the difference between τ^* and $\bar{\tau}^*$ represents a similar measure of disagreement but in the event that information is not acquired by the group—we call this *ex-ante* disagreement. Contrary to ex-post disagreement, the smaller is ex-ante disagreement, the larger is the difference between wtv and wtp.

Increases in ex-post disagreement in some state ω make Jane less willing to vote for the group to obtain information, because it makes it less likely that her favorite policy in state ω is chosen. By contrast, increases in ex-ante disagreement make her more willing to vote for information, because it is less likely that her favorite policy is chosen by the group if information is not acquired.

As an example, let us say that Jane is in a group with two other people—Anne and Bob. Table 1 displays Jane's wtv for a range of possible preferences.

TABLE 1. The payoff of policy Q is always 0 for Anne, Bob, and Jane in either state. The payoff of policy C is represented by a vector (ω_i^L, ω_i^H), where ω_i^L and ω_i^H represent the payoff of policy C in states ω^L and ω^H, respectively, for each agent $i = Anne, Bob, Jane$.

	Scenario 1	Scenario 2	Scenario 3
Anne	(−1,6)	(1,7)	(−4,2)
Bob	(−1,6)	(2,−4)	(2,−4)
Jane	(−1,6)	(−1,6)	(−1,6)
Jane's wtv	0.5	0	3

In scenario 1, Anne and Bob have the same preferences as Jane, so there is no disagreement. Therefore, Jane's wtv is equal to her wtp. In scenario 2, there is only disagreement if information is acquired and in state ω^L; the majority of the group members prefer C, while Jane prefers Q. As a result, Jane's wtv is lower than her wtp. By contrast, in scenario 3, there is only ex-ante disagreement—if information is acquired, there is always someone else who favors Jane's preferred policy in each state; however, if information is not acquired, Anne and Bob prefer Q. Therefore, Jane's wtv should exceed her wtp.

3. A General Decomposition Result

A group of agents collectively choose between policies C and Q. While policy Q returns a normalized payoff of 0 for every agent $i = 1, \ldots, I$, policy C returns a random and unknown payoff of $\omega_i \in \mathbb{R}$. We call vector $\omega = (\omega_1, \ldots, \omega_I)$ the *state of the world*.

Before choosing over policies C and Q, the group has the opportunity to obtain additional information about the state of the world. Information acquisition has an individual cost $c_i \in \mathbb{R}$ and produces a public signal $s \in S$. The joint distribution of vector (ω, s) is denoted by π and is known by every agent.

As the following two examples illustrate, this setting is general enough to allow for information to be obtained through experimentation and delay.[4]

EXAMPLE 1 (*Experimentation*). Prior to deciding over C and Q, Parliament decides over whether to run a pilot program aimed at testing policy C. If Parliament chooses to run the pilot, signal s is produced, which informs the agents about the state of the world. Due to the limited scope of the pilot (e.g., marijuana being legalized in part of the country), the direct payoffs that result from the pilot for each agent i are given by $\alpha\omega_i$ for some $\alpha \in [0, 1]$. In our model of collective information acquisition, this would be equivalent to assuming that $c_i = -\alpha E(\omega_i)$, so that the cost of obtaining information is (a fraction of) the expected cost of foregoing the status quo for each agent i. ■

4. As discussed in the Introduction, there is a large literature that has studied collective experimentation as a way to obtain information (Strulovici 2010, and the ensuing literature) and also a literature that views the cost of information as the cost of delaying the implementation of some project (Chan et al. 2018, and related papers).

EXAMPLE 2 (*Delay*). The Parliament of a democratic country decides whether to join the European Union (C) or not (Q). Parliament chooses between joining right away and waiting one more period to then choose whether to join. If Parliament chooses C right away, this is irreversible. If Parliament waits, the status quo remains in place for one extra period before a final decision is made. Waiting, however, produces information about the merits of joining, that is, signal s is realized. Therefore, one can interpret the cost of acquiring information as the cost of foregoing the possibility of implementing C right away, that is, $c_i = E(\omega_i)$ for each agent i.[5] ■

As in the Example, we define each agent i's individual willingness to pay for information as

$$wtp_i = \int_{s \in S} \pi(s) E(\omega_i|s)(\bar{\tau}_i(s) - \bar{\tau}_i^*) ds,$$

where

$$\bar{\tau}_i^* \equiv \mathbb{1}\{E(\omega_i) \geq 0\}$$

and

$$\bar{\tau}_i(s) \equiv \mathbb{1}\{E(\omega_i|s) \geq 0\}$$

represent the probability that agent i chooses policy C before and after observing signal s, respectively, while $\pi(s)$ represents the marginal probability of $s \in S$ being realized, and $E(\omega_i|s)$ represents the expected value of agent i's payoff in state ω, given signal s. It then follows that, when agent i makes both decisions individually, she acquires information if and only if $wtp_i \geq c_i$.

The group's decisions are made using (possibly signal-dependent) voting rules. Specifically, each group member starts by voting on whether information should be acquired or not. The vote is simultaneous and information is acquired (i.e., signal s is publicly revealed) if and only if at least $\bar{I}_{inf}$ of the group members vote in its favor. Then, regardless of the group's information acquisition decision, there is a second simultaneous vote on whether policy C should be implemented with an approval threshold that may depend on the realized signal s. Specifically, if information is not acquired, at least $\bar{I}^*_{pol}$ must vote for policy C for it to be implemented; if information is acquired and signal s is revealed, the analogous threshold is denoted by $\bar{I}^s_{pol}$.

As is well known, voting games generally have multiple Nash equilibria. Natural refinements involve the (iterated) elimination of (weakly) dominated strategies, which rule out undesirable equilibria where everybody votes for an option they do not like but deviations are irrelevant.[6] We follow the literature (e.g. Austen-Smith and Banks 1996;

5. From a strategic point of view, the two models are not exactly the same, because, with a model of delay, it is not possible not to acquire information and still choose Q. Nevertheless, the equivalence can be restored if one assumes that, if signal s is not known, it is common knowledge that C is preferred by every agent, that is, that $E(\omega_i) \geq 0$ for all $i = 1, \ldots, N$. In words, this means that, absent any new information, joining the European Union is preferred by every member of Parliament.

6. See Kalai and Samet (1984) or Brandenburger, Friedenberg, and Keisler (2008) for formal definitions.

Feddersen and Pesendorfer 1997) and consider a similar refinement. Specifically, when voting over C and Q, every agent is assumed to always vote for her preferred policy, given the available information; when voting for information, each agent takes as given that group members always vote for their preferred policy and votes as if they were pivotal. In this way, it follows that each agent i votes for information to be acquired by the group if and only if $wtv_i \geq c_i$, where

$$wtv_i \equiv \int_{s\in S} \pi(s)E(\omega_i|s)(\tau(s)-\tau^*)ds,$$

where

$$\tau^* \equiv \mathbb{1}\{\#\{i=1,\dots,I : E(\omega_i)\geq 0\} \geq \bar{I}^*_{pol}\}$$

and

$$\tau(s) \equiv \mathbb{1}\{\#\{i=1,\dots,I : E(\omega_i|s)\geq 0\} \geq \bar{I}^s_{pol}\}$$

represent the probability that the group decides in favor of policy C before and after signal s is revealed, respectively.

As in the Example, we interpret the absolute difference between what each individual agent would prefer to do and what the group decides to do in each circumstance as *disagreement*. Specifically, we define

$$\gamma_i^* \equiv |\tau^* - \bar{\tau}_i^*|$$

and

$$\gamma_i(s) \equiv |\tau(s) - \bar{\tau}_i(s)|$$

as measures of agent i's *ex-ante* and *ex-post* disagreement, respectively.

PROPOSITION 1. *For each agent* $i = 1, \dots, I$*,*

$$wtv_i - wtp_i = \delta_i^*\gamma_i^* + \int_{s\in S}\delta_i(s)\gamma_i(s)ds,$$

where

$$\delta_i^* \equiv |E(\omega_i)|$$

and

$$\delta_i(s) \equiv -\pi(s)|E(\omega_i|s)|,$$

Proof. It follows because, for each agent i,

$$E(\omega_i|s)\geq 0 \Rightarrow \tau(s) \leq \bar{\tau}_i(s)$$

for all $s \in S$ and

$$E(\omega_i)\geq 0 \Rightarrow \tau^* \leq \bar{\tau}_i^*.$$

□

Proposition 1 states that the difference between each agent's willingness to vote and willingness to pay for information is linearly increasing with the agent's perceived ex-ante disagreement and linearly decreasing with the agent's perceived ex-post disagreement for each signal that follows information acquisition.

3.1. Private Information

Proposition 1 refers to a setting with symmetric information: every agent has the same belief about everyone's payoff regardless of whether information is acquired or not. It is natural to wonder the extent to which the previous decomposition also holds when agents have heterogeneous (private) beliefs.

We extend the model by allowing each agent $i = 1, \ldots, I$ to observe a private signal $\theta_i \in \Theta_i$ prior to the information acquisition stage. In principle, each agent's belief about state ω depends on the realized θ_i so that, even when information is acquired and s is commonly observed, whether or not the group ends up voting for policy C is still uncertain for each individual agent i.

Analogously to the previous section, we consider Bayes-Nash equilibria with the same caveat that, if agent i knows that he is not pivotal in the information acquisition stage, he votes as if he is. Below, we confirm that the decomposition of Proposition 1 holds under the following two assumptions

Assumption A. Signal s is independent of $\theta = (\theta_1, \ldots, \theta_I)$.

Assumption B. For all $i = 1, \ldots, I$, ω_i is independent of θ_j for $j \neq i$.

Example 3. Each agent i's payoff ω_i is the sum of an idiosyncratic component θ_i that is independent across agents (Assumption B) and a common shock ε that influences every agent equally, that is, $\omega_i = \theta_i + \varepsilon$. Random variable s is interpreted as a signal of that shock, so that it is correlated with ε but not with θ (Assumption A). ■

As in the previous section, we can define ex-ante and ex-post agreement as perceived by agent i as the likelihood that the group will pick agent i's favorite policy given the available information, which includes agent i's private type θ_i, public signal s in the case of ex-post agreement, and that agent i's vote over whether information should be acquired is pivotal (if, in equilibrium, there is a positive probability that agent i is pivotal). Formally, ex-ante disagreement is defined as

$$\gamma_i^*(\theta_i) = |\tau_i^*(\theta_i) - \bar{\tau}_i^*(\theta_i)|,$$

where

$$\tau_i^*(\theta_i) \equiv Pr\{\#\{j = 1, \ldots I : E(\omega_j|\theta_j) \geq 0\} \geq \bar{I}_{pol}^*|\theta_i, Piv\}$$

represents the probability that the group chooses policy C if information is not acquired, conditional on agent i's type θ_i and on being pivotal, and

$$\bar{\tau}_i^*(\theta_i) \equiv \mathbb{1}\{E(\omega_i|\theta_i) \geq 0\}$$

represents the probability that agent i would have chosen policy C individually if information was not acquired.

Ex-post disagreement in signal $s \in S$ is defined as

$$\gamma_i(\theta_i)(s) = |\tau_i(\theta_i)(s) - \bar{\tau}_i(\theta_i)|,$$

where

$$\tau_i(\theta_i)(s) \equiv Pr\{\#\{j = 1, ..., I : E(\omega_j|\theta_j, s) \geq 0\} \geq \bar{I}_{pol}^s|\theta_i, Piv\}$$

represents the probability that the group chooses policy C if information is acquired and signal s is realized, conditional on agent i's type θ_i and on being pivotal, while

$$\bar{\tau}_i(\theta_i)(s) \equiv \mathbb{1}\{E(\omega_i|\theta_i, s) \geq 0\}$$

represents the probability that agent i would have chosen policy C individually if information was acquired and signal s was realized.

PROPOSITION 2. *For each agent $i = 1, ..., I$,*

$$wtv_i(\theta_i) - wtp_i(\theta_i) = \delta_i^*(\theta_i)\gamma_i^*(\theta_i) + \int_{s\in S} \delta_i(\theta_i)(s)\gamma_i(s)(\theta_i)ds,$$

where

$$\delta_i^*(\theta_i) \equiv |E(\omega_i|\theta_i)|$$

and

$$\delta_i(\theta_i)(s) \equiv -\pi(s)|E(\omega_i|\theta_i, s)|.$$

Proof. The proof follows because, under assumptions A and B, each agent i's willingness to pay and willingness to vote can be written as in the previous section:

$$wtp_i(\theta_i) = \int_{s\in S} \pi(s)E(\omega_i|\theta_i, s)(\bar{\tau}_i(\theta_i)(s) - \bar{\tau}_i^*(\theta_i))ds$$

and

$$wtv_i(\theta_i) = \int_{s\in S} \pi(s)E(\omega_i|\theta_i, s)(\tau_i(\theta_i)(s) - \tau_i^*(\theta_i))ds.$$

□

Assumptions A and B play a crucial role in the decomposition of Proposition 2. Combined, they imply that each agent does not care about any information held by any other agent when voting for the final policy. Not only does this allow us to rule out abstentions without loss of generality, but it also rules out the possibility that agents learn about their own payoffs from being pivotal. Were that not the case, ex-post (ex-ante) disagreement might *not* generate over-acquisition (under-acquisition) of information, as Example 4 demonstrates.

Example 4. There are two agents and two states $\omega^L = (-1, -1)$ and $\omega^H = (1, 4)$. Assume that state ω^L has an ex-ante probability of 3/4 and that acquiring information reveals the true state. While agent 1 is uninformed, agent 2 observes a signal $\theta_2 \in \{l, h\}$ such that $Pr\{\theta_2 = h|\omega = \omega^H\} = Pr\{\theta_2 = l|\omega = \omega^L\} = 1$, that is, agent 2 privately knows which state is realized. Assume that the rules of the group are that (1) information is acquired if and only if at least one agent votes for it and (2) policy C is selected if and only if both agents vote for C.

Let us focus on agent 1 and how her willingness to pay compares with her willingness to vote for information. Her willingness to pay is equal to 1/4—the value of information comes from correctly selecting policy C in state ω^H. As for her willingness to vote, let us first consider the second stage interaction between the agents. If information is acquired, both agents get to learn the state, so they both vote for C in state ω^H and Q in state ω^L. If information is not acquired, the informed agent also always votes for the correct policy. The uninformed agent knows this and would prefer to abstain. If abstention is not possible, an equivalent strategy is available: vote for C. In this way, the uninformed player essentially delegates the final decision to the informed player, ensuring that she maximizes her payoffs. But then, from the perspective of the uninformed player, whether the group obtains information or not is irrelevant; the group will always end up choosing C in state ω^H and Q in state ω^L, so that her willingness to vote is 0.

The failure of assumption A leads to agent 1 under-acquiring information ($wtp_1 > wtv_1$) despite there being ex-post agreement and ex-ante disagreement (if information is not acquired, agent 1 would select Q individually, while the group only selects Q with 75% probability). ■

4. Experiment

We test Propositions 1 and 2 in a series of online laboratory experiments, coded in oTree (Chen, Schonger, and Wickens 2016), with 159 students recruited through the Universidad de Los Andes Online Recruitment System for Economic Experiments (ORSEE) database (Greiner 2015). We measure how sensitive the difference between each agent's willingness to vote and willingness to pay is to (1) ex-ante disagreement and (2) ex-post disagreement. As Ginzburg and Guerra (2019) find that some subjects are unable to consider the behavior of others when making decisions about information acquisition (i.e., are strategically naive), especially in situations where information does not affect their preferred choice between C and Q, we also investigate whether situations where information is individually valuable trigger more strategic behavior.

4.1. Design

In each online session, each subject was assigned to a group with two bots. In each of four total rounds, the subject participated in a collective decision by the group by first voting on information revelation and then on whether alternative Q or C should be chosen.[7] In each round, each subject was initially endowed with 10 Experimental Tokens (ET, approximately 2.5USD). As in the model, alternative Q returned a constant payment of 10ET, while alternative C returned a random payment of $\omega_i = f^s(\theta_i)$, which depended on the privately known type θ_i of group member i and on the realization of a public signal s. We have assumed that the public signal was either High (H) or Low (L) with equal probability. Finally, while subjects had a fixed type $\theta_i = A$ for all rounds, bots were assigned a type from $\{A, B\}$. In the first round, each bot was assigned type A with probability 1; in the subsequent rounds, each bot was *privately* assigned each type with equal probability.

In each round, the initial decision of whether to learn the signal was made exclusively by the subject, as if she was pivotal. After that, the group decided by majority voting between alternatives Q and C. Subjects were informed that bots voted in order to maximize their expected payoffs, so that they could predict each bot's vote for each type.

There are two main departures from the theoretical model which were incorporated in order to simplify the subjects' problem. On the one hand, we use bots as additional group members, which serves two purposes. First, it alleviates concerns that the experimental results could be driven by social preferences. Specifically, while, in the theoretical model, we assume that group members are selfish, in practice, it could be that subjects have other-regarding preferences (Cooper and Kagel 2013). This would mean that, if groups were composed of human subjects, subjects would tend to condition their voting behavior to the payoffs other group members might receive. Replacing these by bots allows us to shut down this potential confounding channel in the experimental design.

Second, replacing other group members by bots also allowed us to reduce the strategic uncertainty that is common to every game. Indeed, the theoretical model assumes not only that every agent is rational, but also a common knowledge of rationality. While common knowledge of rationality seems compatible with the types of settings we have discussed in the Introduction, where agents face high-stakes' decisions, in an experimental setting with comparatively lower stakes, it would be possible for agents to doubt the rationality of the remaining group members. To account for that, we provide enough information for subjects to know exactly how each bot will behave for each type, with and without information. In this way, we restore common knowledge of rationality. As shown by Hanaki et al. (2016), by making the behavior of the bots perfectly predictable, we reduce strategic uncertainty to the point that, in our setup, it is unlikely that it explains the observed effect of ex-ante and ex-post disagreement on voting.

7. In the actual wording of the experiment, alternatives Q and C correspond, respectively, to Option X and Option Y. We have included an English translation of the instructions to subjects in Online Appendix B.4. An online session lasted for 50 minutes: filling out informed consent and reading instructions (10 minutes), making decisions (30 minutes), and filling out an exit questionnaire and payment receipt (10 minutes).

TABLE 2. Payoff function when alternative C is chosen given the signal, $(f^L(\theta_i), f^H(\theta_i))$.

Incentive level to choose C	Type, θ_i	Rounds 1	2	3	4
		***T*1. Information is individually valuable**			
Higher incentives	A	(9,16)			
	B	(9, 16)	(11, 17)	(6, 9)	(7, 12)
Lower incentives	A	(7,12)			
	B	(7, 12)	(6, 9)	(11, 17)	(9, 16)
		***T*2. Information is not individually valuable**			
Higher incentives	A	(11,17)			
	B	(11, 17)	(6,9)	(7, 12)	(9, 16)
Lower incentives	A	(6,9)			
	B	(6, 9)	(11, 17)	(9, 16)	(7, 12)

Note: Each cell shows $(f^L(\theta_i), f^H(\theta_i))$ for every treatment of individually valuable information (***T*1** or ***T*2**), for a given incentive level to choosing C (*Higher* or *Lower*), by type ($\theta_i \in \{A, B\}$), and across rounds.

In addition to incorporating bots, we have also made things easier for the subjects by explicitly telling them that they were the dictator in the information acquisition stage. Were that not the case, each subject would have to correctly calculate the voting equilibrium, in order to then form beliefs conditional on being pivotal. With a relatively low-stakes' experiment, it would not be realistic to expect subjects to be able to do this. Our design still allows us to test whether subjects are influenced by ex-ante and ex-post disagreement; telling them that they indeed are pivotal no matter what only makes the calculation of ex-ante and ex-post disagreement cognitively simpler.[8]

In Table 2, we show the payoff functions and the specific treatments implemented.[9] During the four rounds, the payoff function associated to type A remained constant while the payoff function linked to type B changed. We set these changes to have in some rounds ex-post disagreement under signal L or H, and in others ex-ante disagreement. Therefore, the variation across rounds in the payoff functions constitutes our main within-subjects experimental *treatment of ex-post or ex-ante disagreement*.[10] This treatment allows us to test whether the difference between willingness to vote and willingness to pay for information is increasing with ex-ante disagreement and decreasing with ex-post disagreement.

8. In theory, this makes each individual subject's decision of whether the group acquires information strategically equivalent to voting for information under the assumption that the bots' types get re-drawn in the second stage.

9. There was also a practice round with payoff levels that were $+1/-1$ deviations from the first round's. The practice round did not count towards the final experimental earnings.

10. Specifically, subjects face an ex-post disagreement treatment under signal s whenever $\mathbf{1}\{f^s(A) \geq 10\} \neq \mathbf{1}\{f^s(B) \geq 10\}$ and they face an ex-ante disagreement treatment whenever $\mathbf{1}\{1/2\left(f^L(A) + f^H(A)\right) \geq 10\} \neq \mathbf{1}\{1/2\left(f^L(B) + f^H(B)\right) \geq 10\}$.

To investigate deviations from the theory due to strategic naivety, we implemented a between-subjects *treatment of how individually valuable information* is to subjects. In treatment $T1$, information is individually valuable because the subject's payoff when the group chooses alternative C is not always larger or smaller than with alternative Q (see the upper panel of Table 2). In treatment $T2$, information is not individually valuable—the subject's payoff when alternative C is chosen is always larger or smaller than alternative Q (lower panel of Table 2). We distinguish between these two cases to study whether subjects behave differently when information is relevant to them individually. Within these two treatments, subjects were given different incentive levels to vote for alternative C—lower or higher incentives—to check whether our results are sensitive to incentive magnitudes.

In the first round, when the bots are known to be of type A just like the subject, we asked each subject whether they wished to pay a cost of c for the signal to be revealed. We retrieved each subject's responses for 11 values of c ranging from -0.5 to 3 ET and, to mitigate concerns about hypothetical bias, randomly selected one of these possible values for the first round. After implementing the information acquisition decision, the subject and the bots voted for either C or Q and the alternative with the most votes was chosen.

Notice that, by design, the problem the subject faced in this first round is equivalent to an individual problem. Therefore, we were able to recover each subject's (upper bound of her) willingness to pay for information $\bar{v}_i$, defined as the largest value of c (of our 11-value list) for which the subject chose the realization of the signal to be revealed.[11]

In the subsequent rounds, in the information acquisition stage, we asked each subject if they wanted to vote in favor of information being revealed by paying cost $\bar{v}_i$, that is, we recovered whether their willingness to vote for information was larger or smaller than their willingness to pay for information. We keep track of each subject i's answer per round r (denoted by $y_{i,r}$, which is equal to 1 if subject i chose to acquire information at round $r > 1$) as well as whether there is ex-post and/or ex-ante disagreement.

We have elected to elicit each subject's individual willingness to pay for information in a collective setting rather than using an individual task to ensure that subjects made decisions across the whole experiment within the same collective framework. In this way, any observed difference between each subject's willingness to pay and willingness to vote can be ascribed to distinct valuations of information, rather than changes in the elicitation method.

We conducted a total of seven sessions. In each session, we randomly allocated subjects to $T1$ where information is individually valuable (80 observations) or $T2$ where it is not individually valuable (79 observations). Subjects were also randomly given either higher or lower incentives to choosing C.

11. Table B.1, in the Online Appendix B.1, shows treatment specific values used for the cost of information in the choice-list task, which were calibrated based on risk aversion structural estimates found by Becerra and Guerra (2023).

Earnings were calculated in terms of ET and exchanged into Colombian pesos at a rate of 1 ET to Colombian Pesos (COP) 1,000. The total payment to each subject equaled her earnings from one randomly chosen round, plus a 2.5USD show-up fee. The average payment was approximately 5.4USD, equivalent to 7.5% of the subjects' average weekly expenses.[12] Payments were transferred to subjects' electronic money or bank accounts.

4.2. Experimental Results

In the experiment, we empirically observe whether each subject is more willing to vote for information to be acquired in each round than her stated willingness to pay for information (obtained in the first round). From Proposition 2, it follows that

$$y(\theta_i) = \mathbf{1}\left\{\delta^*(\theta_i)\gamma^*(\theta_i) + \sum_{s\in S}\delta(\theta_i)(s)\gamma(\theta_i)(s) \geq 0\right\},$$

where $y(\theta_i)$ represents whether an arbitrary agent of type i has a higher willingness to vote for information than her willingness to pay for information. The design of the experiment also allows us to observe the ex-ante disagreement $\gamma^*(\theta_i)$ and the ex-post disagreement $\gamma(\theta_i)(s)$ under signal $s \in \{L, H\}$ for each round. Our hypothesis is that $\delta^*(\theta_i) > 0$ for all θ_i and that $\delta(\theta_i)(s) < 0$ for all θ_i and s. We therefore want to regress the binary votes for information on ex-ante and ex-post disagreement to estimate $\delta^*(\theta_i)$ and $\delta(\theta_i)(s)$.

Given the structure of our experiment, we estimate a linear probability model with dependent variable $y_{i,r}$, where the set of main treatment variables includes indicator variables on whether in round r subject i faces ex-ante disagreement (Ex-ante$_{i,r}$), ex-post disagreement under signal $s = L$ (Ex-post$L_{i,r}$), and ex-post disagreement under signal $s = H$ (Ex-post$H_{i,r}$).[13] Additionally, we include subject i's upper bound of her willingness to pay for information ($\bar{v}_i$), which captures the subject's preferences and accounts for the incentive level to choosing C treatments. In all estimation specifications, we account for learning by controlling for the round the subject faced during the session (r). In the most comprehensive specification, we either control for observable characteristics ($\mathbf{z}_i$)—namely gender, socio-economic stratum (from 1 to 6), whether majoring in economics or business, semester, self-reported willingness to take risk, subject's guess from a beauty contest question as a proxy for rationality, and a constant—or subject fixed effects when we exploit the within-subjects treatment design. We cluster robust standard errors at the subject level and separately estimate equation (1) for $T1$ when information is individually valuable, and $T2$, when information is not individually valuable.

$$y_{i,r} = \delta\text{Ex-ante}_{i,r} + \delta_L\text{Ex-post}L_{i,r} + \delta_H\text{Ex-post}H_{i,r} + \boldsymbol{\beta}'\mathbf{z}_i + r + \bar{v}_i + \epsilon_{i,r} \qquad (1)$$

12. See descriptive statistics in Table A.1 in Appendix A.3.

13. In the Online Appendix B.2, we report the estimated average marginal effects from a logit model.

TABLE 3. Linear estimation of individual vote to acquire information.

	$T1$, Information is individually valuable		$T2$, Information is not individually valuable	
	(1)	(2)	(3)	(4)
Ex-ante Disagreement	0.422***	0.422***	−0.060	−0.058
	(0.105)	(0.118)	(0.072)	(0.081)
Ex-post L Disagreement	−0.160***	−0.160**	0.005	−0.043
	(0.059)	(0.066)	(0.068)	(0.080)
Ex-post H Disagreement	−0.202***	−0.202***	−0.113	−0.070
	(0.055)	(0.062)	(0.069)	(0.075)
Constant	1.368***	1.332***	1.121***	1.146***
	(0.178)	(0.071)	(0.161)	(0.053)
Round	Yes	Yes	Yes	Yes
$\bar{v}_i$	Yes	No	Yes	No
Controls†	Yes	No	Yes	No
Ind Fixed Effects	No	Yes	No	Yes
Observations/Clusters	320/80		316/79	
R-squared	0.317	0.521	0.274	0.506

Notes: *** $p<0.01$, ** $p<0.05$, and * $p<0.1$. Robust standard errors clustered at the subject level in parentheses. Coefficients come from an ordinary least squares regression. Columns (1) and (2) use the sample from $T1$ Individually valuable information treatment, while columns (3) and (4) restrict the sample to $T2$ Not individually valuable information treatment. Dependent variable is 1 if a subject voted for information. We include indicator variables on whether subject faces Ex-ante disagreement, Ex-post disagreement under signal $s = L$, and Ex-post disagreement under signal $s = H$. Round is a linear round trend. $\bar{v}_i$ is subject i's upper bound of her willingness to pay for information. †Controls include whether subject is female, socio-economic stratum (from 1 to 6), whether majoring in economics or business, semester, self-reported willingness to take risk, subject's guess from a beauty contest question as proxy for rationality, and a constant.

We report the estimation results from equation (1) in Table 3. In the first two columns, we use the sample from treatment $T1$, where information is individually valuable; in columns (3) and (4), we restrict the analysis to the sample of treatment $T2$, where it is not. For each individually valuable information treatment, the first column controls for round, the subject's upper bound of her willingness to pay for information, and the full set of observable characteristics. The second column incorporates individual fixed effects to control for unobserved characteristics. In this last specification, we fully exploit our within-subjects experimental design.

The estimated coefficients in Table 3 suggest that, whenever information is individually valuable (columns (1) and (2)), subjects are more willing to vote for information when there is more ex-ante disagreement (i.e., $\hat{\delta} > 0$) and less ex-post disagreement for either signal (i.e., $\hat{\delta}_L, \hat{\delta}_H < 0$). These effects hold even after controlling for individual fixed effects, which suggests that, after accounting for unobserved individual characteristic, the average subject is $42pp$ more willing to vote to acquire information when she faces ex-ante disagreement than when there is no disagreement. Additionally, when there is ex-post disagreement, subjects are between $16pp$ and $20pp$ less likely to vote for information acquisition. These results are in line with the theoretical predictions.

However, if the information is not individually valuable (columns (3) and (4)), the results do not support our theoretical channel, as no source of disagreement is relevant in explaining the observed voting behavior. Our experimental results suggest that the treatment where information is individually valuable is less prone to strategic naivety and triggers the sort of strategic behavior the theory discusses.

Next, we go beyond estimating the impact of ex-ante and ex-post disagreement on each subject's willingness to vote and study the extent to which our theory as a whole predicts the subjects' decisions. To that end, we calculate a measure of performance for the subjects: a ratio between the subjects' expected utility given their actual decisions and the expected utility they would have received had they made optimal decisions, as predicted by the theory.

To compute each subject's expected utility, we assume that subjects are expected utility maximizers with a constant relative risk aversion (CRRA) utility function. Each subject's utility function is based on the subject's estimated risk aversion coefficient, which we recover from her responses in the willingness-to-pay for information task.[14]

To ensure that the risk aversion parameter is well-identified, we restrict the analysis to consistent decision-makers (CDM). CDMs are identified using responses from the willingness-to-pay for information task in the first round. A subject qualifies as a CDM if she switches at most once from acquiring information to not acquiring information in the choice-list task. Additionally, when information is individually valuable (as in $T1$), a CDM's inferred willingness to pay for information must be positive (i.e., $\bar{v}_i \geq 0$); in contrast, when information is not individually valuable (as in $T2$), a CDM's elicited willingness to pay for information should be close to 0 (i.e., $\bar{v}_i \in [-0.1, 0.1]$).[15]

In Figure 1, we present the empirical distribution of these performance ratios by treatment group. We find that in a substantial proportion of rounds, subjects in each group made decisions fully aligned with theoretical predictions. Moreover, subjects in treatment T1 appear to perform better than subjects in treatment T2, even though we cannot reject the null hypothesis of equal means across both information treatments (p-value 0.156).[16]

Finally, our experiment also allows to compare the welfare generated by the voting mechanism with a few benchmark alternatives: The same voting mechanism under

14. Specifically, expected utilities are computed under CRRA preferences: $u_i(f) = \left(f^{1-\hat{\rho}_i}\right) / (1 - \hat{\rho}_i)$, where $\hat{\rho}_i$ is the estimated subject-specific risk aversion parameter. For subjects in treatment $T1$, where information is individually valuable, we infer risk aversion by matching their observed switching point in the individual willingness-to-pay task to theoretical indifference conditions under CRRA utilities. We then use a synthetic control approach to predict risk aversion for the full sample. We estimate a linear model $\rho_i = \boldsymbol{\varphi}'\mathbf{z}_i + \varepsilon_i$, where $\mathbf{z}_i$ includes the individual-level controls described above, and we impute fitted values for all subjects.

15. See Table B.2 for the distribution of $\bar{v}_i$ in the willingness to vote for information task by treatment.

16. An even more stringent test to the theory would have been to compare subjects' *recovered* willingness to vote for information in each round with a *reported* willingness to vote; the theoretical prediction is that these would be the same. However, this exercise would have substantially increased the number of decisions that each subject had to make—in each round, each subject would be faced with a list of costs and would have to think about which ones they would be willing to pay for information to be acquired—which would have reduced the reliability of our results on the impact of ex-ante and ex-post disagreement in voting.

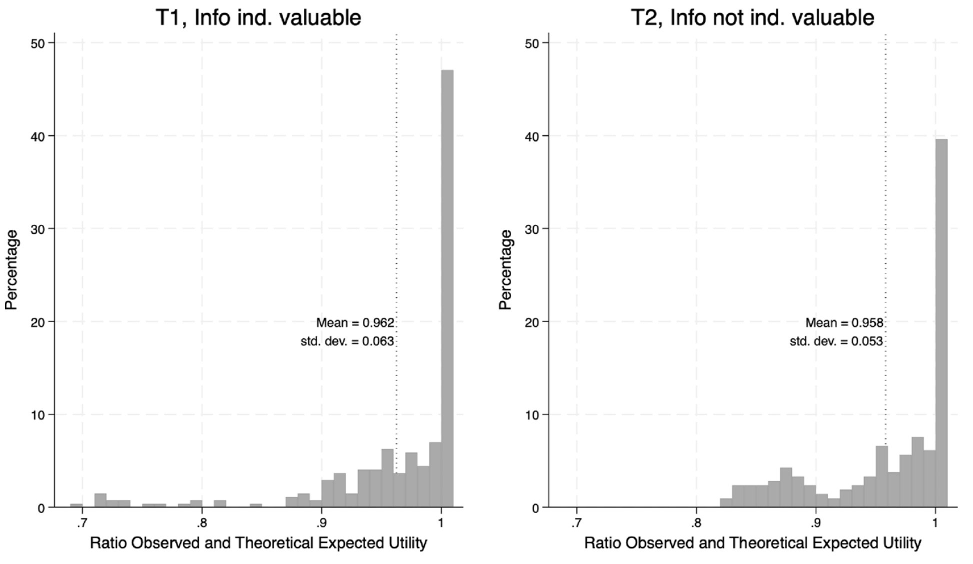

FIGURE 1. Ratio of observed and theoretical expected utility. Figure shows the histogram of the ratio between the subjects' expected utility given their actual decisions and the expected utility were subjects to make the optimal decision, as predicted by the theory, by rounds. To compute the expected utility, a constant relative risk aversion (CRRA) utility function is assumed for subjects based on the estimated risk aversion coefficient that we recover from the willingness-to-pay task. We report the mean and standard deviation for both $T1$ and $T2$ information treatments. The sample is restricted to consistent decision-makers—subjects who switched at most once from acquiring information to not acquiring information in the choice-list task, and for whom $\bar{v}_i \geq 0$ if in treatment $T1$ and $\bar{v}_i \in [-0.1, 0.1]$ if in treatment $T2$.

the assumption that subjects make optimal decisions (as predicted by the theory), a random dictator mechanism, where a randomly chosen member of the group gets to make both collective decisions, and the first best mechanism, which simply maximizes the sum of payoffs of all three members. The results are presented in Figure 2. As expected, the observed welfare is the smallest of the four, which is likely due to the subjects sub-optimal decisions. If any of the other mechanisms was experimentally tested, welfare measures would also be lower. Nevertheless, under the assumption that subjects do make optimal decisions, the voting mechanism does compare favorably to the random dictator mechanism.

4.3. Cognitive Limitations

Despite taking a number of steps to make the experiment cognitively simple, it is still the case that the main potential confounding channel to our results is the subjects' cognitive limitations (see Camerer, Ho, and Chong 2004). In order to gauge its impact, we estimated equation (1) by using a sample of consistent decision makers. We provide the results in Table 4. In columns (1) and (3), we restrict our sample to CDM, while in columns (2) and (4), we restrict it to inconsistent decision-makers. We note that our

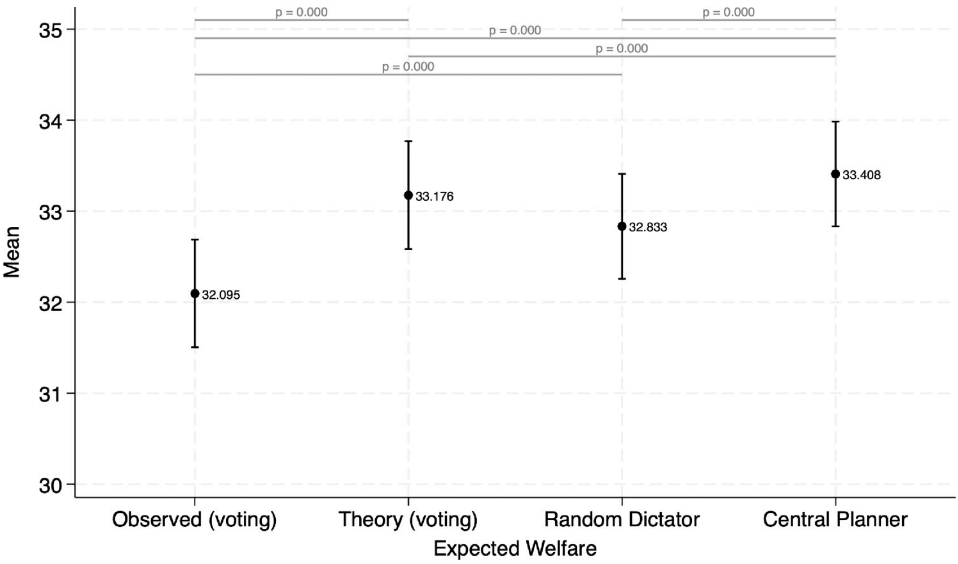

FIGURE 2. Expected Welfare under different collective decision making mechanisms. Figure shows the expected welfare generated by the voting mechanism given the subjects' actual decisions (*Observed* (*voting*)), the voting mechanism were subjects to make the optimal decision as predicted by the theory (*Theory* (*voting*)), a random dictator mechanism, where a randomly chosen member of the group gets to make both decisions (*Random Dictator*), and the first best mechanism, which simply maximizes the sum of payoffs of all three members (*Central Planner*). To compute the expected utilities, a constant relative risk aversion (CRRA) utility function is assumed and each risk aversion coefficient is recovered from the willingness-to-pay task. Bots are assumed to be risk-neutral. Vertical lines depict 95% robust confidence intervals clustered at the individual level. Horizontal lines test the null hypothesis of equality of means across welfare measures and report its corresponding *p-values*. The sample is restricted to consistent decision-makers—subjects who switched at most once from acquiring information to not acquiring information in the choice-list task, and for whom $\bar{v}_i \geq 0$ if in treatment $T1$ and $\bar{v}_i \in [-0.1, 0.1]$ if in treatment $T2$.

model is relatively better at explaining the behavior of subjects with greater strategic competence (columns (1) and (3)) than of non-CDM (columns (2) and (4)). For CDMs, the estimated coefficients $\left(\hat{\delta}, \hat{\delta}_L, \hat{\delta}_H\right)$ have the predicted signs in both $T1$ and $T2$, although not always significantly different than 0 when information is not individually valuable.

Additionally, to follow up on the suggestion (derived from the estimation results) that subjects in treatment $T2$, for whom information was not individually valuable, paid less attention to strategic incentives, we performed a (logit) Quantal Response Equilibrium (QRE) analysis (McKelvey and Palfrey 1995, 1998), where the probability that a subject votes for information is increasing in the difference in expected utility between voting for and against information acquisition.[17] The kernel density of

17. See more details of the maximum likelihood estimator in the Online Appendix B.3.

TABLE 4. Linear estimation of individual vote to acquire information, by consistent decision-making (CDM).

	Information is			
	$T1$ individually valuable		$T2$, not individually valuable	
	CDM‡		CDM‡	
	Yes (1)	No (2)	Yes (3)	No (4)
Ex-ante Disagreement	0.449***	0.156	0.063	−0.098
	(0.114)	(0.253)	(0.164)	(0.079)
Ex-post L Disagreement	−0.176***	−0.094	−0.108	0.043
	(0.063)	(0.176)	(0.146)	(0.080)
Ex-post H Disagreement	−0.250***	0.094	−0.282*	−0.062
	(0.058)	(0.139)	(0.160)	(0.075)
Constant	1.346***	1.219***	1.147***	1.148***
	(0.055)	(0.105)	(0.037)	(0.033)
Round	Yes	Yes	Yes	Yes
$\bar{v}_i$	Yes	Yes	Yes	Yes
Controls†	Yes	Yes	Yes	Yes
Ind Fixed Effects	No	No	No	No
Observations/Clusters	272/68	48/12	76/19	240/60
R-squared	0.310	0.188	0.246	0.227

Notes: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. Robust standard errors clustered at the subject level in parentheses. Coefficients come from an ordinary least squares regression. Columns (1) and (2) use the sample from $T1$ Individually valuable information treatment, while columns (3) and (4) restrict the sample to $T2$ Not individually valuable information treatment. Only consistent decision makers (CDM) from the willingness to pay for information task in round 1 are included in columns (1) and (3), while non-consistent decision makers are included in columns (2) and (4). ‡CDM in treatment $T1$ and $T2$ are subjects who switched at most once from acquiring information to not acquiring information in the choice-list task, and for whom $\bar{v}_i \geq 0$ if in treatment $T1$ and $\bar{v}_i \in [-0.1, 0.1]$ if in treatment $T2$. Dependent variable is 1 if a subject voted for information. We include indicator variables on whether subject faces Ex-ante disagreement, Ex-post disagreement under signal $s = L$, and Ex-post disagreement under signal $s = H$. Round is a linear round trend. $\bar{v}_i$ is subject i's upper bound of her willingness to pay for information. †Controls include whether subject is female, socio-economic stratum (from 1 to 6), whether majoring in economics or business, semester, self-reported willingness to take risk, subject's guess from a beauty contest question as proxy for rationality, and a constant.

the estimated precision parameter λ (see Figure 3) reveals a wide range of values across subjects, with some subjects displaying low responsiveness to expected utility differences and a median above 1.00. This suggests that some of the deviations from the theoretical predictions may be due to limited strategic sensitivity. More importantly, the mean of the estimated precision parameter λ is found to be 11.5% higher for subjects for whom information is individually valuable than for those for whom it is not (p-value 0.036), and we can reject the null hypothesis that both λ distributions come from the same data generating process. This suggests that subjects facing the information treatment $T1$ present a larger estimated precision parameter consistent with more sensitivity to expected utility differences and strategic incentives than subjects who faced treatment $T2$.

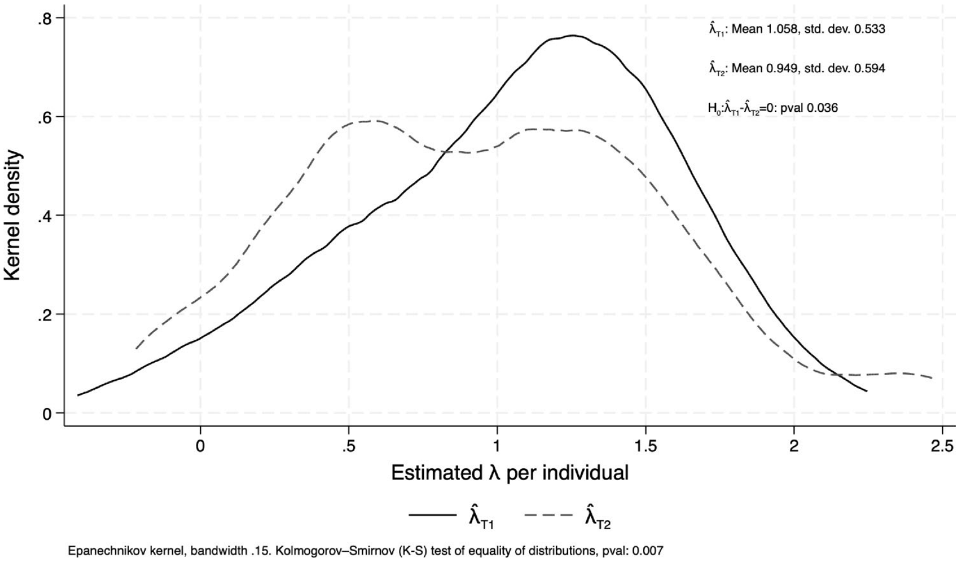

FIGURE 3. Estimated λ from Quantal Response Equilibrium (QRE). The figure displays the kernel densities for the estimated precision λ_i from the QRE maximum likelihood estimator described in Online Appendix B.3 allowed to vary by information treatment $T1$ and $T2$. We report the mean, standard deviation of the precision parameter for each information treatment. We test for distributional differences in the kernel densities using the Kolmogorov–Smirnov test and compare means using robust regressions clustered at the subject level.

4.4. External Validity

We expect the model to fit the observed behavior best when subjects are more similar to members of real-life decision-making bodies (DMB). From the experimental end-line questionnaire, we gather the number of DMB (in high school, university student councils, or other decision boards) in which the subject participated. In Table 5, apart from distinguishing between $T1$ and $T2$ treatments, we split the sample between subjects with experience in DMB (columns (1) and (3)) and those with no experience whatsoever (columns (2) and (4)), and estimate equation (1) for both samples.

The estimated coefficients confirm that our theoretical channel is good at explaining the behavior of subjects facing the treatment where information is individually valuable, independently of the experience in DMB (columns (1) and (2)). However, when subjects face the treatment where information is not individually valuable, our results suggest that subjects with more experience in actual collective decision-making are more likely to act in line with the model (i.e., coefficients in column (3) present the expected signs) than subjects with no experience (column (4)). These results speak to the model's external validity and relevance when subjects resemble members of real-life committees.

Overall, the model is better at explaining the behavior of subjects who have experience in DMB, who are consistent-decision makers, and who face situations where information is individually valuable.

TABLE 5. Linear estimation of individual vote to acquire information, by experience in decision-making bodies (DMB).

	Information is			
	$T1$ individually valuable		$T2$, not individually valuable	
	Experience in DMB‡		Experience in DMB‡	
	Yes (1)	No (2)	Yes (3)	No (4)
Ex-ante Disagreement	0.502***	0.357**	0.059	−0.171*
	(0.166)	(0.136)	(0.107)	(0.096)
Ex-post L Disagreement	−0.124	−0.190**	−0.111	0.092
	(0.084)	(0.085)	(0.095)	(0.103)
Ex-post H Disagreement	−0.218**	−0.190**	−0.083	−0.124
	(0.086)	(0.073)	(0.090)	(0.110)
Constant	1.360***	1.310***	1.121***	1.171***
	(0.078)	(0.073)	(0.038)	(0.039)
Round	Yes	Yes	Yes	Yes
$\bar{v}_i$	Yes	Yes	Yes	Yes
Controls†	Yes	Yes	Yes	Yes
Ind Fixed Effects	No	No	No	No
Observations/Clusters	152/38	168/42	156/39	160/40
R-squared	0.317	0.335	0.198	0.296

Notes: *** $p<0.01$, ** $p<0.05$, and * $p<0.1$. Robust standard errors clustered at the subject level in parentheses. Coefficients come from an ordinary least squares regression. Columns (1)–(2) use the sample from $T1$ Individually valuable information treatment, while columns (3)–(4) restrict the sample to $T2$ Not individually valuable information treatment. ‡Subjects with experience in DMB (such as high school or university student councils) are included in columns (1) and (3), while subjects with no experience whatsoever in DMB are included in columns (2) and (4). Dependent variable is 1 if a subject voted for information. We include indicator variables on whether subject faces Ex-ante disagreement, Ex-post disagreement under signal $s = L$, and Ex-post disagreement under signal $s = H$. Round is a linear round trend. $\bar{v}_i$ is subject i's upper bound of her willingness to pay for information. †Controls include whether subject is female, socio-economic stratum (from 1 to 6), whether majoring in economics or business, semester, self-reported willingness to take risk, subject's guess from a beauty contest question as proxy for rationality, and a constant.

5. Discussion

5.1. Condorcet Jury Theorem and the Wisdom of the Crowd

Does a crowd of non-experts make better judgments than an individual expert? That is, say that a group of people chooses between two alternatives through voting. It is common knowledge that one of the alternatives is better for everyone but each group member has a prior opinion about which alternative is better. Is the group more likely to guess the correct alternative than the expert?

The Condorcet jury theorem says that it is, provided the group is large enough, that is, the wisdom of the crowd wins if the crowd is sufficiently large. The basis for the result is statistical: if agents have statistically independent private signals about which alternative is correct and vote sincerely, the probability that the best alternative receives more votes converges to 1 as the number of voters increases. This is true even if each individual signal is only marginally informative; with enough agents, the combined signals become more informative than the expert's signal. Feddersen and Pesendorfer (1997) prove the result still holds even in equilibrium, where voters may not vote sincerely.

However, the Condorcet jury theorem assumes that each voter's information is exogenous. If each voter can, at a cost, acquire private information prior to the vote, a free-riding problem emerges; because the probability of each voter being pivotal reduces as the number of voters increases, each individual voter has little incentive to obtain information in large groups (Li 2001; Persico 2004). Therefore, in general, it is no longer always the case that larger groups produce more accurate decisions (Persico 2004). This suggests that a single expert should be preferred over a group of non-experts even if their collective signal is as informative as the expert's.

By contrast, our results are based on a model where learning is collective: agents do not learn privately; instead, a public signal is produced and everyone pays for its cost. Therefore, there is no incentive to free-ride. Interestingly, in these settings, the exact opposite result emerges. The group of non-experts actually has *more* incentive to acquire information; not less. Indeed, as we illustrate below, it might be that, statistically, the expert is more informed than the group of non-experts, and yet the group of non-experts may still be more likely to guess the correct alternative. We make this point more precisely below.

Let us say that a defendant has been accused of a crime and is being judged by a group of I jurors. The defendant can be guilty ($\omega = \omega^{gu}$) or innocent ($\omega = \omega^{in}$). Jurors can choose to punish the defendant (policy C) or acquit him (policy Q). In the spirit of the Condorcet jury theorem, we assume that $\omega_i^{gu} > 0$ and $\omega_i^{in} < 0$ for all $i = 1, \ldots, I$, that is, everyone agrees that, if the defendant is guilty, he should be punished, and if he is innocent, he should be acquitted. However, jurors value wrong punishments and wrong acquittals differently, so that $E(\omega_i|\theta_i)$ may be positive or negative depending on θ_i and/or on index i. If information is obtained, a public signal s is realized (e.g., whether more time is needed to go through the evidence) and jurors pay some (opportunity) cost (e.g., time of deliberation).

If decisions are made by a single judge—the expert—the decision rule is straightforward: The judge acquires signal s if her willingness to pay (wtp) for information exceeds its cost. Let us compare how accurate the judge's decision is with the decisions made by a group of jurors who, individually, are less informed, but, as a group, are equally informed as the judge. Specifically, let us say that, if information is acquired by the group of jurors, the same signal s is realized. For simplicity, assume that s is sufficiently informative so that, for all $s' \in S$, either $E(\omega_i|\theta_i, s') > 0$ for all $\theta_i \in \Theta_i$ and $i = 1, \ldots, I$ or $E(\omega_i|\theta_i, s') < 0$ for all $\theta_i \in \Theta_i$ and $i = 1, \ldots, I$. Using the terms defined in the previous sections, there is ex-ante disagreement and ex-post agreement for each agent i.[18]

18. Given that s is such that there is ex-post agreement, if the deliberation produced individual signals privately observed by each juror, they would be incentivized to share them with everyone else. Signal s

Proposition 2 tells us that, for any judge with preference parameter θ_i such that $wtp_i(\theta_i) \geq c_i$, there is a juror with the same preference parameter for whom $wtv_i(\theta_i) \geq c_i$. That is, if the preferences of judges and jurors are drawn from the same distribution and have the same costs of acquiring information, jurors are more likely to vote for information to be acquired than judges are of acquiring information. This insight, combined with the standard assumption that jury trials use the unanimity rule (where information is acquired unless every juror objects), leads to the result that, if a group of jurors is capable of generating the same information than an individual judge, the former is more likely to reach the correct outcome, because they have more incentives to obtain information.[19]

5.2. *On Efficiency*

We now turn our attention to studying how voting for information can produce inefficient outcomes. We first examine the role of preference reversal and then the general possibility of inducing efficient outcomes if a planner is allowed to freely choose a mechanism and transfers are possible.

5.2.1. Inefficient Outcomes under Voting. Dynamic voting models, where uncertainty unravels over time, are known to cause Pareto inefficient outcomes (Roessler, Shelegia, and Strulovici 2018). Our results show how ex-ante and ex-post disagreement may generate these inefficiencies, as Example 5, adapted from Roessler, Shelegia, and Strulovici (2018), illustrates.

EXAMPLE 5. Three committee members decide whether to expand (C) or not (Q) the subway of the city of Rome. Before voting on expansion, the committee votes on whether to order an archaeological study aimed at identifying the likelihood (**H**igh and **L**ow) that no historical landmark is destroyed if the expansion is completed. Both votes are decided by majority rule. Let us say that $\omega^H = (-2, 3, 1)$ and $\omega^L = (1, 2, -2)$, so that, if the study is completed, the committee votes for expansion regardless of the study results. However, if each state is equally likely, committee members 1 and 3 prefer not to expand in the event that the study is not completed.

This means that the unique equilibrium outcome is for the committee to vote against information acquisition even when it is costless, and then voting against expansion. Indeed, this would still be the unique equilibrium outcome if each group member was paid up to 1/2 (i.e., if $c_i \geq -1/2$) in case information was acquired. In that case, every group member could be made better off if the group was to commit to obtaining information and selecting not to expand no matter the results of the study. ■

is the product of this information sharing in line with the concept of the wisdom of the crowd—even if individual signals are not very informative, the collective signal is as informative as the one obtained by the expert.

19. Not only are jurors more willing to obtain information as judges with the same preference type, but also it is enough for one juror to want to obtain information.

Example 5 has the unnatural feature that group member 1 actually prefers to expand the subway if there is a high probability of finding a historical landmark but prefers not to if the probability is low. In many settings, this inconsistency of preferences is problematic: one would think that, if group member 1 is more predisposed to expand the subway than group member 3 if the probability of finding a historical landmark is high, the same should be true if the probability is low. Indeed, in Appendix A.1, we prove a general result: If there is complete information and payoffs have the non-reversal property described below, the equilibrium is Pareto efficient.

We define the following "no-preference-reversal" property: If agent i is more favorable to C relative to Q than another agent i' in some state, then agent i should be more favorable to C than agent i' in every state.[20] If it holds, one can order group members by their relative preference for policy C. In this way, if the same voting rule is used to determine the winning policy regardless of what signal (if any) is produced (e.g., simple majority rule), there is only one agent, labeled i^*, who might be pivotal in the policy decision (e.g., the agent with the median preference for C in the case of the simple majority rule). As this is common knowledge, agent i^* votes as if he is a dictator and cannot be made worse off, as long as the cost function c_i does not break the order across agents induced by the non-reversal property. Specifically, in Appendix A.1, we show that the group decides as if agent i^* is a dictator if either (1) the cost of obtaining information is (weakly) larger for agents with a stronger preference for Q, and, absent information, agent i^* prefers Q; or (2) the cost of obtaining information is (weakly) larger for agents with a stronger preference for C and, absent information, agent i^* prefers C.

As the reader will note, condition (1) includes experimentation models with a known ex-ante preference for the status quo (similar to the example over the legalization of marijuana), while condition (2) includes delay models with a known ex-ante preference for change (similar to the example of a country that considers joining the European Union). Moreover, both include the case of constant cost functions, which are standard when committees make decisions to purchase information using common funds. Our result shows that, in these settings, very much like what would happen with any other policy, majority rule leads to the group acquiring exactly as much information as the "pivotal voter" would have acquired, had he been a dictator, so that the equilibrium outcome is Pareto efficient.

By contrast, if agents have private information, the identity of the pivotal voter is not known a priori, even if preferences satisfy the no-preference reversal property and a simple majority rule is used for all decisions. The uncertainty over her identity may make the pivotal voter herself vote differently from when she decides individually (by over or underacquiring information depending on her levels of ex-ante and ex-post disagreement), which may lead to inefficient outcomes. Example 6 illustrates this.

20. Formally, if there is ω' such that $\omega'_i > \omega'_{i'}$, then $\omega_i > \omega_{i'}$ for all $\omega \in \Omega$.

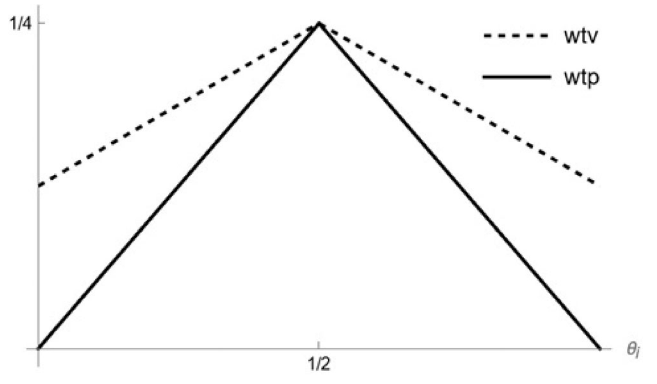

FIGURE 4. The full and dashed line represent each agent i's willingness to pay and willingness to vote for information, respectively, as a function of θ_i.

EXAMPLE 6. Consider the following variation of Example 5. Continue to assume that there are three agents but, now, each agent observes an independent private signal $\theta_i \in [0, 1]$ drawn from a uniform distribution. Each agent i's payoff is $\omega_i = \theta_i + s$, where public signal $s \in \{-1, 0\}$ and takes each value with equal probability. All agents' have the same cost $c_i = c \in (0, 1/4)$ of acquiring information.

It directly follows that each agent i's willingness to pay is given by

$$wtp_i(\theta_i) = \left\{ \begin{array}{ll} \frac{\theta_i}{2} - (\theta_i - 1/2) & \text{if } \theta_i \geq 1/2, \\ \frac{\theta_i}{2} & \text{if } \theta_i < 1/2. \end{array} \right\}.$$

Each agent i's willingness to vote is given by

$$wtv_i(\theta_i) = \frac{\theta_i}{2} - (\theta_i - 1/2)Pr\{median(\theta_1, \theta_2, \theta_3) \geq 1/2|\theta_i, Piv\}.$$

In Appendix A.2, we show that, for every $c \in (0, 1/4)$, there is an equilibrium, where

$$Pr\{median(\theta_1, \theta_2, \theta_3) \geq 1/2|\theta_i, Piv\} = \begin{cases} 1/4 & \text{if } \theta_i < 1/2 \\ 3/4 & \text{if } \theta_i \geq 1/2. \end{cases}$$

Therefore, each agent i's willingness to vote can be written as[21]

$$wtv_i(\theta_i) = \begin{cases} \frac{\theta_i}{2} - (\theta_i - 1/2)(3/4) & \text{if } \theta_i \geq 1/2, \\ \frac{\theta_i}{2} - (\theta_i - 1/2)(1/4) & \text{if } \theta_i < 1/2. \end{cases}$$

Figure 4 represents both wtp_i and wtv_i as a function of θ_i.

As is implied by Proposition 2, each agent's willingness to vote for information exceeds her willingness to pay, because, in this example, there is ex-post agreement and

21. It is worth noting that, in general, an agent's willingness to vote is a function of the costs of all agents, because these influence the voting equilibrium, which, in turn, influences the beliefs of each agent i over the identity of the pivotal voter i^*, conditional on agent i himself being pivotal in the information acquisition stage.

only ex-ante disagreement. In particular, for any vector of preferences $\theta = (\theta_1, \theta_2, \theta_3)$, the agent with the median value of θ_i over-acquires information relative to what would maximize her payoff if she knew that she was the median voter (unless the median θ_i is exactly equal to $1/2$). Indeed, it is easy to find realizations of vector θ for whom all three agents could be made better off.[22] ■

5.2.2. Optimal Mechanisms. As the previous analysis confirms, unlike what happens with static voting models, voting over information may produce inefficient outcomes. While these can be easily resolved by replacing voting by a dictatorial mechanism, where a single agent is given complete control over the group's decisions (which would ensure Pareto efficiency), this alternative is still unsatisfying, as a good measure of welfare should take into account every group member's well-being.

One possible alternative way of measuring the quality of the mechanism is by verifying whether it produces *majority efficient* outcomes, defined as outcomes for which no other outcomes exist that are preferred by a majority of the population. Roessler, Shelegia, and Strulovici (2018) study whether such mechanisms exist in a dynamic model where, just like in our model, information unravels over time as a function of the group's decisions. Their results are mostly negative; if the equilibrium outcome of the dynamic voting game is majority inefficient, any alternative mechanism will also produce a majority inefficient outcome. Applied to our setting, this suggests that dynamic majority voting mechanisms are hard to improve over if what one cares about is majority efficiency.

Another popular alternative is to consider the sum of payoffs as a measure of welfare. Indeed, in settings with transferable utility, efficiency is achieved if and only if it maximizes the sum of the agents' payoffs. The classic literature on mechanism design has shown how to design efficient mechanisms even when agents have private information. A principal elicits reports from the agents and incentivizes them to report truthfully by designing appropriate transfer schemes between them. The classic Vickrey-Clark-Groves family of mechanisms provides incentives for truthful reporting to be a dominant strategy by making agents pay for the externality they create by misreporting (Vickrey 1961; Clarke 1971; Groves 1973). While ensuring voluntary participation, this family of mechanisms does not guarantee budget-balance, and can therefore require subsidies or induce unnecessary levels of spending. By contrast, the expected externality mechanism (Arrow 1979; d'Aspremont and Gérard-Varet 1979), which makes agents pay for the expected externality they generate, guarantees budget-balance at the expense of voluntary participation constraints. In the context of voting procedures, where we assume that agents are forced to participate, budget-balance

22. For example, take $\theta = (0.3, 0.31, 0.32)$. All agents agree on the policy ex-ante and ex-post. Their willingness to pay for information ranges from 0.15 to 0.16 but their willingness to vote for information ranges from 0.2 to 0.205. Therefore, for any individual cost from 0.16 to 0.2 the group selects to acquire information following a unanimous vote, but all group members would have actually been better off had the group not acquired information.

becomes a primary concern and, therefore, we focus on this class. Below, we describe it in the context of our model.

The allocation space is $X = \{S, C\} \cup \{h : S \to \{Q, C\}\}$, where $\Theta = \Theta_1 \times \ldots \times \Theta_I$, and represents the decision of the group—Q and C represent choosing Status Quo and Change, respectively, without acquiring further information, while h represents the choice rule over Q and C for each signal $s \in S$ after acquiring information. A mechanism is then given by the allocation rule $x : \Theta \to X$ and the transfer vector $t : \Theta \to \mathbb{R}^I$, where each t_i represents the payment made by agent i. Transfers enter linearly in the agents' payoff functions.

To find the efficient allocation rule x^*, we start by considering the principal's choice, conditional on acquiring information, which we label by $h^*(\theta) : S \to \{Q, C\}$:

$$h^*(\theta)(s) = \begin{cases} Q & \text{if } \sum_{i=1}^{I} E(\omega_i|s, \theta_i) \geq 0, \\ C & \text{otherwise.} \end{cases}$$

The efficient allocation rule x^* compares $h^*(\theta)$ with the options of immediately implementing Q or C:

$$x^*(\theta) = \begin{cases} Q & \text{if } \sum_{i=1}^{I} E(\omega_i|\theta_i) \leq 0 \text{ and} \\ & E_s\left[\max\left\{\sum_{i=1}^{I} E(\omega_i|s, \theta_i), 0\right\}\right] - \sum_{i=1}^{I} c_i \leq 0, \\ C & \text{if } \sum_{i=1}^{I} E(\omega_i|\theta_i) > 0 \text{ and} \\ & E_s\left[\max\left\{\sum_{i=1}^{I} E(\omega_i|s, \theta_i), 0\right\}\right] - \sum_{i=1}^{I} c_i < \sum_{i=1}^{I} E(\omega_i|\theta_i), \\ h^*(\theta) & \text{otherwise.} \end{cases}$$

To build the expected externality transfer vector t^*, start by defining each agent i's expected externality from reporting θ_i as

$$z_i(\theta_i) \equiv -E_{\theta_{-i}}\left[\sum_{j \neq i} v_j^*(\theta_i, \theta_{-i})\right],$$

where

$$v_j^*(\theta) = \begin{cases} 0 & \text{if } x^*(\theta) = Q, \\ E(\omega_j|\theta_j) & \text{if } x^*(\theta) = C, \\ E[\mathbb{1}\{h^*(\theta)(s) = C\}\omega_j|\theta_j] - c_j & \text{if } x^*(\theta) = h^*(\theta). \end{cases}$$

represents each agent j's payoffs from the efficient allocation absent any transfers. Then define

$$t_i^*(\theta) = z_i(\theta_i) - \frac{1}{I-1}\sum_{j \neq i} z_j(\theta_j) \; \forall i.$$

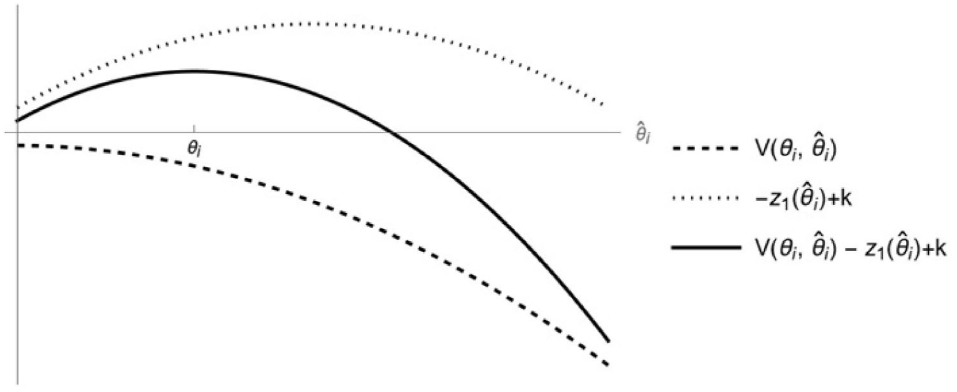

FIGURE 5. Incentive analysis in the expected externality mechanism: when agent 1 is asked to pay for the expected externality his report creates, his optimal strategy is to report truthfully ($\hat{\theta}_1 = \theta_1$).

Agent i has no incentive to misreport because his expected payoff is given by

$$E_{\theta_{-i}}\left[\sum_{j=1}^{I} v_j^*(\theta_i, \theta_{-i})\right] + \frac{1}{I-1}\sum_{j\neq i} z_j(\theta_j),$$

which is maximized by truthfully reporting θ_i; the design of the transfer makes agent i internalize the expected externality caused from misreporting. It can be verified that the sum of transfers always equals 0 for any vector θ, which ensures efficiency.

To get a sense of how the expected externality solves the principal's problem and eliminates over and under acquisition of information, consider again Example 6. Start by noting that

$$h^*(\theta)(s) = \begin{cases} Q & \text{if } s = -1, \\ C & \text{if } s = 0, \end{cases}$$

and

$$x^*(\theta) = \begin{cases} Q & \text{if } \theta_1 + \theta_2 + \theta_3 \leq 6c \\ h^* & \text{if } \theta_1 + \theta_2 + \theta_3 \in [6c, 3 - 6c] \\ C & \text{otherwise.} \end{cases}$$

Consider the incentives of agent 1 in the absence of transfers. The ex-post payoff for agent 1 when reporting $\hat{\theta}_1$ is given by

$$\hat{v}_1(\theta, \hat{\theta}_1) \equiv \begin{cases} 0 & \text{if } \hat{\theta}_1 + \theta_2 + \theta_3 \leq 6c, \\ \frac{\theta_1}{2} - c & \text{if } \hat{\theta}_1 + \theta_2 + \theta_3 \in [6c, 3 - 6c], \\ \theta_1 - 1/2 & \text{otherwise.} \end{cases}$$

Therefore, agent 1 would have reported $\hat{\theta}_1$ in order to maximize $V_1(\hat{\theta}_1, \theta_1) \equiv E_{\theta_2,\theta_3}(\hat{v}(\theta, \hat{\theta}_1))$, which is displayed in Figure 5 for the case of $c = 0.16$ and $\theta_1 = 0.3$ (as in footnote 21). As the reader will note, agent 1 has an incentive to under-report his type. The reason is straightforward: of the three options available—Q, C and getting information, Q is the most preferred whenever $\theta_1 = 0.3$. Agent 1 maximizes the chances that Q is selected by reporting $\hat{\theta}_1 = 0$.

The expected externality transfer eliminates incentives to misreport, thereby eliminating all inefficiencies, including those generated by over or under acquisition of information. In Example 6, agent 1's expected externality is

$$z_1(\hat{\theta}_1) = -E_{\theta_2,\theta_3}(v_2(\hat{\theta}_1, \theta_2, \theta_3) + v_3(\hat{\theta}_1, \theta_2, \theta_3)),$$

which is depicted in Figure 5.[23] As Figure 5 shows, the incentives to over-report in order to reduce the expected externality payment exactly counteract the incentives to under-report in order to maximize the payoff from the allocation, making reporting truthfully optimal.

Appendix A

A.1. No-Preference Reversal under Complete Information

Assume that the no-preference reversal defined in the text holds and assume that all voting is decided by majority rule.[24] For convenience, re-label each agent's index i such that agents with larger indexes have a stronger preference for C in any state, that is, for any $\omega \in \Omega$, $\omega_i > \omega_{i'}$ if and only if $i > i'$. Then, denote by $i^* \equiv \frac{I+1}{2}$ as the median agent (under the assumption that the number of agents is odd). Notice that it is common knowledge that, in the second period voting, agent i^* is pivotal no matter what information is produced. Finally, define the following two conditions: (1) c_i is weakly decreasing in i and $E(\omega_{i^*}) \leq 0$, and (2) c_i is weakly increasing in i and $E(\omega_{i^*}) \geq 0$.

Proposition A.1. *Assume that either* (1) *or* (2) *holds. Then, the group acquires information if and only if* $wtp_{i^*} \geq c_{i^*}$.

Proof. For each $i = 1, \ldots, I$, define $\widehat{wtv}_i \equiv wtv_i - c_i$. Notice that the group acquires information if and only if

$$median\left(\widehat{wtv}_1, \widehat{wtv}_2, \ldots, \widehat{wtv}_I\right) \geq 0.$$

We start by proving that, if either (1) and (2) hold, $\widehat{wtv}_i$ is weakly monotone in i. To that purpose, we define $i^s = 0, \ldots, I$ to be such that agent i (weakly) prefers C over Q if signal $s \in S$ is realized if and only if $i > i^s$. Likewise, define $\bar{i} = 0, 1\ldots, I$ to be such that agent i (weakly) prefers C over Q ex-ante if and only if $i > \bar{i}$.

23. Parameter $k \in \mathbb{R}$ was included for illustration purposes as it does not influence the set of maximizers.

24. The result also applies for any other rule as long as it applies for both stages.

Let us start by assuming that (1) holds. In that case, we can write

$$wtv_i = E[y^s E(\omega_i|s)],$$

where

$$y^s = \begin{cases} 1 \text{ if } i^* > i^s \\ 0 \text{ otherwise} \end{cases}$$

so that wtv_i is weakly increasing in i. Given that c is weakly decreasing in i, it follows that $\widehat{wtv}_i$ is weakly increasing in i. When (2) holds, we can write

$$wtv_i = E[z^s E(\omega_i|s)],$$

where

$$z^s = \begin{cases} -1 \text{ if } i^* \leq i^s \\ 0 \text{ otherwise} \end{cases}$$

so that wtv_i is weakly decreasing in i. Given that c is weakly increasing in i, it follows that $\widehat{wtv}$ is weakly decreasing in i.

The fact that $\widehat{wtv}_i$ is weakly monotone in i implies that

$$median\left(\widehat{wtv}_1, \widehat{wtv}_2, ..., \widehat{wtv}_I\right) = \widehat{wtv}_{i^*}.$$

Finally, the fact that agent i^* knows she is pivotal implies that $wtv_{i^*} = wtp_{i^*}$, so that

$$\widehat{wtv}_{i^*} = wtp_{i^*} - c_{i^*}.$$

□

A.2. Example 6 (continued)

From the calculation of wtv_i in the text, it follows that, for any cost $c \in (1/8, 1/4)$, there is some threshold $\tilde{\theta} \in (0, 1/2)$ such that agent $i = 1, 2, 3$ votes in favor of information acquisition if and only if $\theta_i \in [\tilde{\theta}, 1 - \tilde{\theta}]$ (see Figure 4). Without loss of generality, we discuss how to calculate the probability of the median value of θ_i being greater than $1/2$ conditional on agent 1 being pivotal in the information stage and on his type.

For agent 1 to be pivotal in the information acquisition stage, it must be that one of the other agents' type (say agent 2) lies in interval $[\tilde{\theta}, 1 - \tilde{\theta}]$, while the other agent's type (say agent 3) lies outside of it. That means that, if $\theta_1 \geq 1/2$, the median θ_i is larger than $1/2$ if either agent 2's type lies in interval $[1/2, 1 - \tilde{\theta}]$, or if agent 3's type lies in interval $[1 - \tilde{\theta}, 1]$. The probability that one of these two events happens, conditional on agent 2 voting for information acquisition and on agent 3 voting against it, is $1 - 1/4 = 3/4$.

By contrast, if $\theta_1 < 1/2$, for the median θ_i to be above $1/2$ it would be necessary that both θ_2 and θ_3 were to on the right half of their intervals, that is, θ_2 would have to lie on interval $[1/2, 1 - \tilde{\theta}]$ and θ_3 would have to lie on interval $[1 - \tilde{\theta}, 1]$. The probability of this happening, conditional on agent 2 voting for information acquisition and on agent 3 voting against it, is $1/4$.

If $c \leq 1/8$, the equilibrium behavior would be for every agent to always vote in favor of acquiring information. Because there would be a 0 probability of being pivotal, agents calculate the probability of C being the group's decision if information is not acquired using their priors. Therefore, each agent i will still place a probability of 1/4 of the group choosing C if $\theta_i < 1/2$ (because the other two agents would have to be above 1/2) and of 3/4 if $\theta_i \geq 1/2$ (because only one agent would have to be above 1/2).

A.3. Descriptive Statistics

TABLE A.1. Descriptive statistics.

	Mean	Sd	Valuable information treatment		p-value:
			T1, individually	T2, not individually	$H_0 : T1 = T2$
Female	0.50	0.50	0.59	0.42	0.03
Semester	6.35	2.90	6.60	6.10	0.28
Econ/Business	0.13	0.34	0.16	0.10	0.26
SES stratum	3.57	1.21	3.69	3.46	0.23
Weekly expenses USD	72.16	97.81	67.42	77.08	0.74
Beauty Contest	37.96	21.61	35.35	40.59	0.13
WT risk	6.33	1.82	6.41	6.25	0.58
Observations	159		80	79	

Notes: Female is dummy indicating if subject is female participant, Semester is subject's academic term currently in, Econ/Business is 1 if subject is studying Economics or Business Administration, Socioeconomic status (SES) stratum is 1 for poorest and 6 for richest households, Beauty contest is subject's guess from a beauty contest question as proxy for rationality, WT risk is self reported willingness to take risk 1 "not at all willing" 10 "Always willing". Treatment refers to how individually valuable information is to subjects; in $T1$ information is individually valuable, in $T2$ information is not individually valuable.

References

Albrecht, James, Axel Anderson, and Susan Vroman (2010). "Search by Committee." *Journal of Economic Theory*, 145, 1386–1407.

Anesi, Vincent and T. Renee Bowen (2021). "Policy Experimentation in Committees: A Case Against Veto Rights under Redistributive Constraints." *American Economic Journal: Microeconomics*, 13, 124–62.

Anesi, Vincent and Mikhail Safronov (2023). "Deciding When to Decide: Collective Deliberation and Obstruction." *International Economic Review*, 64, 757–781.

Arrow, Kenneth J. (1979). "The Property Rights Doctrine and Demand Revelation under Incomplete Information." In *Economics and Human Welfare*, pp. 23–39. Elsevier.

Austen-Smith, David and Jeffrey S. Banks (1996). "Information Aggregation, Rationality, and the Condorcet Jury Theorem." *American Political Science Review*, 90, 34–45.

Becerra, Oscar and José-Alberto Guerra (2023). "Personal Safety First: Do Workers Value Safer Jobs?" *Journal of Economic Behavior & Organization*, 212, 996–1016.

Brandenburger, Adam, Amanda Friedenberg, and H. Jerome Keisler (2008). "Admissibility in Games 1." *Econometrica*, 76, 307–352.

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong (2004). "A Cognitive Hierarchy Model of Games." *Quarterly Journal of Economics*, 119, 861–898.

Chan, Jimmy, Alessandro Lizzeri, Wing Suen, and Leeat Yariv (2018). “Deliberating Collective Decisions.” *Review of Economic Studies*, 85, 929–963.

Chen, Daniel L., Martin Schonger, and Chris Wickens (2016). “oTree—An Open-Source Platform for Laboratory, Online, and Field Experiments.” *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Clarke, Edward H. (1971). “Multipart Pricing of Public Goods.” *Public Choice*, 11, 17–33.

Compte, Olivier and Philippe Jehiel (2010). “Bargaining and Majority Rules: A Collective Search Perspective.” *Journal of Political Economy*, 118, 189–221.

Cooper, David J. and John H. Kagel (2013). “Other-Regarding Preferences.” In *The Handbook of Experimental Economics*, vol. 2, edited by J. Kagel and A. Roth, chap. 4, pp. 217–289. Princeton University Press, Princeton.

d’Aspremont, Claude and Louis-André Gérard-Varet (1979). “Incentives and Incomplete Information.” *Journal of Public Economics*, 11, 25–45.

Feddersen, Timothy and Wolfgang Pesendorfer (1997). “Voting Behavior and Information Aggregation in Elections with Private Information.” *Econometrica*, 65, 1029–1058.

Freer, Mikhail, César Martinelli, and Siyu Wang (2020). “Collective Experimentation: A Laboratory Study.” *Journal of Economic Behavior & Organization*, 175, 365–379.

Gieczewski, Germán and Svetlana Kosterina (2024). “Experimentation in Endogenous Organizations.” *Review of Economic Studies*, 91, 1711–1745.

Ginzburg, Boris and José-Alberto Guerra (2019). “When Collective Ignorance Is Bliss: Theory and Experiment on Voting for Learning.” *Journal of Public Economics*, 169, 52–64.

Greiner, Ben (2015). “Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE.” *Journal of the Economic Science Association*, 1, 114–125.

Groves, Theodore (1973). “Incentives in Teams.” *Econometrica*, 41, 617–631.

Hanaki, Nobuyuki, Nicolas Jacquemet, Stéphane Luchini, and Adam Zylbersztejn (2016). “Cognitive Ability and the Effect of Strategic Uncertainty.” *Theory and Decision*, 81, 101–121.

Hizen, Yoichi, Keisuke Kawata, and Masaru Sasaki (2013). “An Experimental Test of a Committee Search Model.” *European Economic Review*, 61, 59–76.

Hudja, Stanton (2019). “Voting for Experimentation: A Continuous Time Analysis.” Working paper, SSRN 3473426.

Kalai, Ehud and Dov Samet (1984). “Persistent Equilibria in Strategic Games.” *International Journal of Game Theory*, 13, 129–144.

Khromenkova, Daria (2017). “Collective Experimentation with Breakdowns and Breakthroughs.”Working paper, University of Mannheim.

Li, Hao (2001). “A Theory of Conservatism.” *Journal of Political Economy*, 109, 617–636.

Louis, Philippos (2015). “Learning Aversion and Voting Rules in Collective Decision Making.” *Economics Letters*, 133, 24–26.

Mak, Vincent, Amnon Rapoport, and Darryl A. Seale (2014). “Sequential Search by Groups with Rank-Dependent Payoffs: An Experimental Study.” *Organizational Behavior and Human Decision Processes*, 124, 256–267.

Mak, Vincent, Darryl A. Seale, Amnon Rapoport, and Eyran J. Gisches (2019). “Voting Rules in Sequential Search by Committees: Theory and Experiments.” *Management Science*, 65, 4349–4364.

Martinelli, César and Thomas R. Palfrey (2020). “Communication and Information in Games of Collective Decision: A Survey of Experimental Results.” In *Handbook of Experimental Game Theory*, pp. 348–375. Edward Elgar Publishing.

McKelvey, Richard D. and Thomas R. Palfrey (1995). “Quantal Response Equilibria for Normal form Games.” *Games and Economic Behavior*, 10, 6–38.

McKelvey, Richard D. and Thomas R. Palfrey (1998). “Quantal Response Equilibria for Extensive form Games.” *Experimental Economics*, 1, 9–41.

Persico, Nicola (2004). “Committee Design with Endogenous Information.” *Review of Economic Studies*, 71, 165–191.

Reshidi, Pëllumb, Alessandro Lizzeri, Leeat Yariv, Jimmy H. Chan, and Wing Suen (2025). “Sequential Sampling by Individuals and Groups: An Experimental Study.” *American Economic Review: Insights*, 7, 480–496.

Roessler, Christian, Sandro Shelegia, and Bruno Strulovici (2018). "Collective Commitment." *Journal of Political Economy*, 126, 347–380.
Strulovici, Bruno (2010). "Learning While Voting: Determinants of Collective Experimentation." *Econometrica*, 78, 933–971.
Vickrey, William (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, 16, 8–37.

Supplementary Data

Supplementary data are available at *JEEA* online.

Journal of the European Economic Association 2026 0(0):1–35
https://doi.org/10.1093/jeea/jvag001