Public versus Secret Voting in Committees*

Andrea Mattozzi† Marcos Y. Nakaguma‡
EUI and MOVE EESP-FGV

March 2019

Abstract

In this paper we study the effect of transparency on voting in committees when members are heterogeneous in competence and bias, they are career-concerned and they can abstain. We show that public voting attenuates the biases of competent members and secret voting attenuates the biases of incompetent members. Public voting leads to better decisions when the magnitude of the bias is large, while secret voting performs better otherwise. We discuss evidence from the lab and from the field that is consistent with our theory.

Keywords: Committees, Voting, Career-Concern, Transparency.
JEL Classification Codes: D72, C92, D71.

*We thank Veronica Rattini, Hanno Kase, Chiara Santantonio and Riccardo Ghidoni for excellent research assistance. We thank the editor, three excellent anonymous referees as well as Marina Agranov, Paulo Furquim de Azevedo, Maria Bigoni, Stefania Bortolotti, Marco Casari, Alessandra Casella, Vincenzo Denicolò, Juan Dubra, Francesco Feri, Piero Gottardi, Daniela Iorio, Philipp Kircher, Gilat Levy, Massimo Morelli, Salvatore Nunnari, Emanuel Ornelas, Luigi Proietti, Debraj Ray, Aniol Llorente-Saguer, Arthur Schram, Justin Valasek and seminar audiences at Bologna, CORS-RWIO, EESP-FGV, Erasmus University, EUI, Insper, King's College, Leicester, Ridge Forum, Royal Holloway, SBE, St.Andrews, University of Oslo, USP, Warwick,Venice 4th Political Economy conference for useful conversations, comments, and encouragement.

†Department of Economics, EUI, Florence, Italy. Email: andrea.mattozzi@eui.eu
‡Sao Paulo School of Economics, FGV, Sao Paulo, Brazil. Email: marcos.nakaguma@fgv.br

1 Introduction

Committee decision-making is a central feature of many political and economic organizations, including government agencies, legislative bodies, central banks, law courts and private companies.[1] The issues confronted by committees are typically multi-faceted and complex, and may involve a variety of conflicts and personal interests. Furthermore, committee members are usually motivated by the desire to advance their own careers and, therefore, care about being perceived as competent decision-makers. This is the case, for example, of regulatory agencies such as the U.S. Food and Drug Administration (FDA) and legislative committees. Finally, committee members may have different competences and it is not unusual to observe situations where some members abstain when unable to form a firm conviction about a particular issue.

This paper studies a committee decision-making problem which combines all elements described above. Specifically, committee members in our model are heterogeneous in their level of competence, they are biased towards different alternatives, they care about their reputation for competence, and they may vote or abstain. In this context, we investigate how the degree of transparency of the committee, i.e. whether individual votes are observed or not, affects equilibrium voting behavior and the quality of the decisions. From a positive point of view, our goal is to understand how the incentives for a committee member to abstain, vote for her bias or against it depend on career-concerns and the degree of transparency of the committee. From a normative point of view, our main objective is to characterize the circumstances under which voting should be public or secret. To the best of our knowledge, this is the first paper investigating how competence, individual biases and career-concerns interact in shaping individuals' voting behaviors in a committee, and how this interaction is affected by transparency.

In our model a committee takes a decision over a binary agenda by simple majority and committee members can vote for either alternative or abstain. The payoff of a member depends on three components: i) a common value, i.e. whether the committee adopts the correct decision; ii) whether the decision matches the member's bias; and iii) the ex-post perceived competence of the member. Competence and bias are private information. Our analysis highlights that the interaction between career-concerns and

[1] A widespread view in the literature is that voting in committees provides an efficient way to aggregate disperse information and contributes to mitigate the interference of individual biases in the decision. See Gerling et al [15] and Li and Suen [23] for reviews of this literature.

transparency leads to qualitatively different implications depending on the member's level of competence and the magnitude of her bias relative to the common value component. We show that, when committee members are relatively biased, transparency acts to "correct" the vote of competent members who would have otherwise simply voted in accordance with their personal interests. On the other hand, when committee members are relatively unbiased, transparency induces incompetent members to vote, even though they would have otherwise preferred to abstain.

Intuitively, competent members know which alternative is the correct one, so that transparency creates an incentive for them to vote correctly. Conversely, incompetent members are uncertain about which alternative is correct, so that transparency simply creates an incentive for them to vote, either for their biases or for the ex-ante more likely alternative. In the absence of career-concerns and when the common value is sufficiently large, it is optimal for incompetent members to abstain, since by doing so they delegate the decision to the competent members. This is the well-known swing voter's curse, first studied by Feddersen and Pesendorfer [10]. In the presence of career-concerns, however, such behavior affects perceived competence negatively, since abstentions can be interpreted as a sign of incompetence in equilibrium. Hence, while transparency attenuates the pre-existing biases of competent members, it may actually exacerbate the pre-existing biases of incompetent members. While these incentives exist everywhere in the parameters' space, our analysis shows that they may lead to actual changes in observable voting behavior in different situations depending on the magnitude of the bias.

We show that public voting should be preferred when the magnitude of the bias is large relatively to the common value, in which case transparency helps to mitigate the influence of private interests on the decisions. Conversely, secret voting should be preferred when the magnitude of the bias is relatively small, in which case the non-observability of the individual votes helps reducing the incentives for incompetent members to "gamble" and vote just to avoid revealing their lack of competence. Notice that, under secrecy, the effect of an individual's correct vote on her own reputation is diluted across all members. Since this dilution effect is proportional to the size of the committee, it follows that the choice between secret and public voting becomes more relevant as the size of the committee increases.

We extend our basic model to allow for a behind closed-doors deliberation stage prior to voting, where committee members may choose to share their private information. In this environment, we show that information is not always aggregated and we identify situations where competent members may have an incentive to strategically withhold information and then vote correctly in order to separate themselves from incompetent members. Furthermore, we show that the observability of individual votes might lead to a trade-off between quality of information aggregation at the deliberation stage and quality of the decision taken at the voting stage. Under certain conditions secrecy may actually make it more likely that information about the state of the world is revealed at the deliberation stage, while transparency creates an incentive for the informed individuals to vote correctly at the voting stage.

We also consider a repeated version of the benchmark model, where committee members vote on a sequence of different independent issues and individual competence may either be iid across periods or persistent over time. We show that repeated interaction does not necessarily lead to better outcomes. In fact, while repetition improves outcomes when competence is iid, it has the opposite effect when competence is persistent. Intuitively, repetition helps disciplining the behavior of incompetent members in the current period only if they may turn out to be competent in the future. Otherwise, it actually increases the incentives for them to vote in order to hide incompetence. We also show that the choice between secret and public voting is still relevant in the context of an infinitely repeated game. Indeed transparency affects both the rewards associated with a correct vote today and the expected punishment associated with an incorrect vote or abstention in the future.

Our analysis has implications for the design of committee decision-making rules. The basic model suggests that voting should be public in committees where members are highly influenced by ideological or self-interested motives such as congressional committees. Conversely, voting should be kept secret when the dissent among members due to individual biases is relatively small, as it is perhaps the case of committees of experts and top bureaucrats responsible for technical decisions.

The premises of our theory and several of its main implications apply to a number of real-world settings. We first discuss two striking cases from national legislatures in Brazil and Italy that illustrate well how a change in transparency might completely change voting behavior and outcomes in the context of a committee of politicians.[2]

[2] Hansen et al [19], Meade and Stasavage [26] and Swank et al [41] exploit the decision of the Federal Open Market Committee (FOMC) to make the transcripts of its meeting publicly available to show that transparency changed the nature of deliberation in the Committee. In particular, they find that the dissent among members decreased significantly after the move to transparency.

Moreover, using detailed information about the meetings of the FDA's advisory committees, we document that abstentions do occur among experts. We further provide evidence showing that abstentions are determined, to a considerable extent, by the desire to delegate the decision to members with more expertise in a particular area.

Given the difficulties involved in evaluating the impact of secret versus public voting using observational data only, we test the main theoretical predictions of our model by means of a controlled laboratory experiment. The experimental setting allows us to control for the level of information and biases of committee members as well as to impose a structure on the rewards associated with career-concerns. These characteristics are rarely observed in field data, but are nonetheless critical for testing the mechanisms underlying models based on asymmetric information. Furthermore, as will become clear later in the paper, there are regions of the parameters where our model features multiple equilibria with different properties. From this perspective, a controlled experiment can help to inform whether individuals coordinate on certain equilibria and not on others.

Consistently with our theory, the experimental results show that transparency improves information aggregation when the bias is high and that secrecy performs better when the bias is low. Furthermore, when the bias is low, approximately half of the incompetent subjects abstain under secret voting and this proportion falls dramatically when we move to public voting. There are almost no abstentions in the case of high bias. Finally, when there are multiple equilibria, while our theory is unclear about which equilibrium players should coordinate on, the experimental results suggest that subjects gradually learn to coordinate on the efficient equilibrium.

The rest of the paper is organized as follows. In the next section we document a number of stylized facts pertaining to committee decision-making. In Section 3, we describe the theoretical model. We solve for the equilibrium and present comparative static results in Section 4. Section 5 discusses a number of possible extensions. We describe the experimental design in Section 6 and present the empirical results in Section 7. Finally, we discuss the related literature and conclude in Section 8.

2 Stylized Facts

In this section we present novel stylized facts about decision-making in committees. We exploit data from the national legislatures in Brazil and Italy and from the FDA's advisory committees to illustrate the relevance of our theory in a number of real-world settings.

Evidence from Legislatures: Secret versus Public Voting. In November 2013, the Brazilian Congress approved a constitutional amendment that changed the procedure to be employed in cases concerning the expulsion of congressmen from secret to public voting. Brazil's House of Representatives is composed of 513 deputies and a member can only be expelled from it if a petition requesting his expulsion receives the support of a majority of representatives. Figure 1 depicts the outcomes of the seventeen expulsion votes that occurred in the House during the period 2005-2018, where the vertical bar indicates the date of approval of the constitutional amendment. Observe that the shift from secret to public voting is associated with a major change in voting behavior. Under secret voting, the average number of votes in favor of an expulsion was 231 and only 4 out of 14 representatives (28%) were expelled from the House. After the change to public voting, on the other hand, the average number of votes in favor of an expulsion increased to 425 and 3 out of 3 representatives (100%) were expelled from the House.[3] It is also interesting to note that, in the four months following the change in the voting rule, five other deputies resigned from the House in order to avoid an almost certain expulsion.

Another interesting case that illustrates well the importance of the degree of transparency for collective decision-making comes from Italian politics. In June 2017, the Italian Parliament voted on a proposal to change the electoral law of one of its twenty districts, the so-called Biancofiore amendment (from the name of the MP who proposed it). The vote was supposed to be secret, but due to a technical error, all individual votes were shown, for a few seconds as they were being cast, in the front panel of the parliament. The mistake was realized by the president of the House within six seconds of the beginning of the vote, at which point she shouted "*it's a secret vote, it's a secret vote*" and, after a few more seconds, the vote was suspended.[4] This case provides an interesting opportunity to examine whether the behavior of legislators is impacted by a completely unexpected change in transparency. An inspection of the video record of the session shows that at least 62 members switched their votes in a span of less than

[3] We exclude abstentions from this analysis, since they are used, in practice, as a way to help the member being subjected to the expulsion vote.

[4] Apparently, the technician in charge of the system forgot to switch the panel's mode from public to secret vote.

eight seconds, which represents around 15% of the total number votes which had been cast in that short period.[5] While we do not claim that this evidence is causal, these rapid and numerous changes suggest that, under certain circumstances, legislators do behave sharply differently under public and secret voting.

Evidence from the FDA: Abstentions. Perhaps one of the most important regulatory agencies of the United States, the Food and Drug Administration (FDA) relies heavily on a number of advisory committees to make decisions about new drug applications and other major public health issues. We collected detailed information available in the transcripts of all meetings held during 2009-2017 of eighteen of the main FDA's advisory committees to establish novel stylized facts about abstention among experts. Our data set contains information about individual votes, which are always public, the justification provided by each member for her vote and a number of individual characteristics, such as educational background and professional affiliation.[6]

The committees are composed by independent experts and relevant stakeholders, including patient and consumer representatives. In our data set, 63% of the members have a MD degree, 28% have a PhD degree and 60% are university professors. Furthermore, career-concerns seem to play an important role on the voting behavior of committee members. For example, by exploiting a change from sequential to simultaneous vote implemented in 2007, Newham and Midjord [34] show that on average 46% of the members take into consideration the sequence of previous votes when casting their own vote, which is consistent with models of reputational herding.[7]

An analysis of the data shows that abstentions are not uncommon in the FDA's advisory committees. They correspond to 2.4% of the individual votes and one or more abstentions are observed in about 25% of committee votes (see Table 1, Panel A).[8] Furthermore, the majority of justifications provided for an abstention are related

[5] The video record of the session is available at https://bit.ly/2HjCefy or upon request. The changes in votes happened in both directions, from "Favorable" to "Against" and vice-versa, most likely to conform to constituents' interests and party lines.

[6] A committee meeting begins with a presentation from the FDA and the sponsor company's staffs followed by a discussion. The committee then votes simultaneously on one or more questions (one at a time) through an electronic system and members have the option to vote yes or no or abstain. Our data set consists of 246 committee meetings, 654 votes on different issues and $9,675$ individual votes. There are on average 2.56 votes per meeting and the mean committee size is 14.79.

[7] Regarding potential biases of FDA's committee members, the existing empirical literature has found a weak relationship between votes and financial ties of members. See Lurie et al. [24], Camara and Kyle [4] and Cooper and Golec [3].

[8] Interestingly, Newham and Midjord [34] note that abstentions increased from 1.2% to 2.7% as a

to lack of adequate data and evidence, and there is a significant number of individuals who acknowledge their own lack of expertise to justify their decision to abstain.[9] A detailed analysis of the transcripts further suggests that specialists without everyday clinical experience, such as PhDs and statisticians without medical background, are more likely to abstain.[10] Indeed, there exists a systematic relationship between academic background and abstention rates: MDs abstain significantly less than the average, while members with a PhD and statisticians are significantly more likely to abstain (see Panel B of Table 1). In particular, the abstention rates among PhDs and statisticians without medical background are about 1.5% larger than the other members, which represents an abstention rate almost 60% above the average.[11]

Overall, our analysis shows not only that abstentions in the FDA do occur, but also that they are motivated, to a considerable extent, by the desire to delegate the decision to members with more expertise, an evidence consistent with the swing voter's curse. This desire to delegate, together with how it is affected by transparency, is at the core of our theoretical model, which we describe next.

3 The Model

We consider a committee of n members, with $n \geq 3$ odd, that must decide between two alternatives, A and B. There are two states of the world, $\omega \in \{A, B\}$, with $\Pr(\omega = A) = q \in (0,1)$.[12] While the true state is a priori unknown, committee members receive a signal about it $s_i \in \{A, \emptyset, B\}$. A member may be either competent, $\mathtt{c}$, in which case he receives a perfectly informative signal $s_i \in \{A, B\}$, or incompetent,

---

result of the shift from sequential to simultaneous voting.

[9] Examples of such statements are *"I abstained only because I don't have the expertise − I yield to my other colleagues"* or *"I abstained because I don't feel I'm competent to make that judgment."*

[10] One member justified his decision as follows: *"I voted to abstain because I'm not in front line to deal with patients. It's really hard to appreciate. I don't feel comfortable either way."* A cardiologist participating as a temporary member in the Peripheral and Central Nervous System Drugs Advisory Committee stated: *"I abstain because [...] this is just sufficiently far away from my area of expertise so that I frankly don't think I add anything."* Other examples are *"I abstained because I am a statistician, and it is outside my expertise to define such things"* or *"I abstained only because I really just have never seen this product. I worked as a nurse for many years, and even then I never saw it. And I just don't know enough to say yes or no."*

[11] These results are robust to a regression analysis controlling for meeting fixed-effects, with standard errrors clustered at the meeting level. Column (4) in Panel B of Table 1 reports the estimated coefficients for each group's dummy obtained from separate regressions. All results are statistically significant, with the estimated differences actually increasing in magnitude.

[12] Our model extends the setting analyzed by Nakaguma [33] to an asymmetric environment.

$\texttt{nc}$, in which case he receives an uninformative signal $s_i = \emptyset$. We assume that each member knows her own competence type $\tau_i \in \{\texttt{c},\texttt{nc}\}$ and the distribution of other members' competences, which is given by $\Pr(\tau_i = \texttt{c}) = \sigma \in (0,1)$. After observing their private signals, all members decide simultaneously whether to vote for A or B or to abstain, $v_i \in \{A, \emptyset, B\}$, where abusing notation we denote abstention by $v_i = \emptyset$. The final decision $x \in \{A, B\}$ is determined by simple majority rule and ties are broken randomly.

Committee members care about making correct decisions and receive a common value $\alpha > 0$ whenever the final decision is equal to the state of the world, $x = \omega$. Additionally, every member is biased towards either A or B and knows her own bias type, $\beta_i \in \{A, B\}$, as well as the distribution of other members' biases, $\Pr(\beta_i = A) = p \in (0,1)$, which we assume to be common knowledge. A member with bias β_i receives an extra payoff $\gamma > 0$ when alternative $x = \beta_i$ is chosen by the committee, regardless of the state of the world.

Committee members are also concerned about building a reputation for competence and making correct decisions. In particular, we assume the existence of an external evaluator, whose only task is to update his beliefs about the likelihood that each member is competent and voted correctly, i.e. voted for the state of the world, conditional on the state and on any other relevant information that might be available to him. The state of the world is always revealed ex-post and, under *public voting*, the evaluator is also able to observe the individual votes of all members, while under *secret voting* he is able to observe only the aggregate number of votes for each alternative.[13] The posterior probability that a committee member i is competent and voted correctly is, therefore, given by:

$$r_i^{\omega,\lambda} \equiv \Pr(\tau_i = \texttt{c}, v_i = \omega | \omega, \mathcal{I}^\lambda), \tag{1}$$

where ω is the state of the world, $\lambda \in \{\texttt{p}, \texttt{s}\}$ denotes whether voting is public or secret, and $\mathcal{I}^\lambda$ represents all relevant information available under λ.

Given the state of the world ω and the committee's decision x, the utility of a member i biased towards β_i under voting rule λ is given by:

$$u_i^{\beta_i,\lambda}(x,\omega) = \phi r_i^{\omega,\lambda} + \mathbb{I}_{\{x=\omega\}}\alpha + \mathbb{I}_{\{x=\beta_i\}}\gamma, \tag{2}$$

[13] We assume that committee members are unable to reveal their votes truthfully ex-post, otherwise in our model all individuals who voted correctly would have an incentive to do so and voting would become de facto public.

where $\phi > 0$ is the weight assigned to career-concerns and $\mathbb{I}_{\{\cdot\}}$ is an indicator function that is equal to one if the condition inside brackets is satisfied and zero otherwise.

Remark. Our model makes a number of simplifying assumptions which deserve to be discussed in more detail. In Online Appendix A, we undertake the following extensions and robustness checks: (i) we examine the case where competent and incompetent members receive signals of positive but imperfect precisions, (ii) we show that the model can be extended to allow for the existence of unbiased members and for the possibility of correlation between competence and bias, (iii) we discuss the role of the assumption that the state of the world is always observed ex-post, (iv) we show that our main results are robust to assuming that only the final decision of the committee is observed under secret voting, (v) we show that our main qualitative results remain unchanged when we use the more standard notion of career-concerns that is based only on the posterior probability that the member is competent, $r_i^{\omega,\lambda} \equiv \Pr(\tau_i = \mathsf{c}|\omega, \mathcal{I}^\lambda)$ and, finally, (vi) we characterize the institutional preferences of committee members of different types between secret and public voting. Furthermore, we consider a version of the model with deliberation in Subsection 5.1 and we allow for repeated interaction in Subsection 5.2.

4 Equilibrium Analysis

We solve the model for symmetric pure-strategy equilibria, where committee members of the same type (with the same bias and competence level) choose identical strategies. We assume that members do not use weakly-dominated strategies. In equilibrium, each committee member chooses a voting strategy that maximizes his expected utility given the equilibrium strategies of other members and the beliefs of the external evaluator. At the same time, the external evaluator's beliefs must be consistent with the members' strategies and computed by Bayes' rule.

4.1 Basic Properties

We begin our analysis by providing a general characterization of the basic properties of the equilibria. Let μ_i denote the conjecture held by a committee member i about the behavior of other members and the beliefs of the external evaluator. Suppose first that member i observes the state of the world prior to voting, i.e. he receives a perfectly informative signal. Given the conjecture μ_i and the state of the world ω, player i's action v_i induces a probability distribution over final outcomes, which is represented by the mapping $\rho^{\omega}_{\mu_i} : \{A, \emptyset, B\} \to [0,1]$, where $\rho^{\omega}_{\mu_i}(v_i)$ denotes the probability (as perceived by the member) that the committee's decision is A when his choice is v_i, given μ_i and ω. Note that we must have $\rho^{\omega}_{\mu_i}(B) \leq \rho^{\omega}_{\mu_i}(\emptyset) \leq \rho^{\omega}_{\mu_i}(A)$ since a vote for A can never lead to a lower probability that the committee's decision is A relative to the case where the member abstains or votes for B.[14]

Let μ_e be the external evaluator's beliefs about the behavior of committee members. Under public voting, all individual votes are observed ex-post, so that career-concern rewards depend only on each member's own vote in accordance with the following expression

$$r^{\omega,\mathtt{p}}_{i,\mu_e} = \Pr_{\mu_e}(\tau_i = \mathtt{c} | v_i = \omega)\mathbb{I}_{\{v_i=\omega\}}, \tag{3}$$

where $\Pr_{\mu_e}(\tau_i = \mathtt{c}|v_i = \omega)$ is computed based on the external evaluators' beliefs about the behavior of members. Under secret voting, on the other hand, only the aggregate voting outcome is observed ex-post, so that career-concern rewards depend on the total number of correct votes, $V^c \equiv \sum_i \mathbb{I}_{\{v_i=\omega\}}$, in accordance with the following expression

$$r^{\omega,\mathtt{s}}_{i,\mu_e} = \Pr_{\mu_e}(\tau_i = \mathtt{c}|v_i = \omega)\tfrac{V^c}{n}, \tag{4}$$

where V^c/n represents the probability that a particular member voted correctly. Note that, in this case, the career-concern rewards are the same across all members and equal to the average expected competence in the committee conditional on V^c.

In equilibrium, each committee member correctly anticipates the voting behavior of other members as well as the beliefs of the external evaluator. Before casting a vote, a member forms an expectation about the career-concern reward that she will receive as a function of her vote. Suppose, first, that the state of the world is observed by the member. Under public voting, the expected career-concern reward is given by

$$\widetilde{r}^{\omega,\mathtt{p}}_i(v_i) = \Pr(\tau_i = \mathtt{c}|v_i = \omega)\mathbb{I}_{\{v_i=\omega\}}, \tag{5}$$

where we omit the index for the evaluator's beliefs to simplify the notation. Under

[14]Observe that the probability $\rho^{\omega}_{\mu_i}(v_i)$ already takes into account the uncertainty related to the realization of types of all other members of the committee. Note also that the inequalities are weak, since for certain beliefs μ_i the member might expect to be pivotal with zero probability.

secret voting, on the other hand, the expected career-concern reward is given by

$$\widetilde{r}_i^{\omega,\mathsf{s}}(v_i) = \Pr(\tau_i = \mathsf{c}|v_i = \omega)\tfrac{1}{n}(\mathbb{I}_{\{v_i=\omega\}} + \mathbb{E}(\textstyle\sum_{j\neq i}\mathbb{I}_{\{v_j=\omega\}})), \tag{6}$$

where $\mathbb{E}(\sum_{j\neq i}\mathbb{I}_{\{v_j=\omega\}})$ represents the number of correct votes expected to be cast by all other committee members. Hence, under secret voting, the impact of a member correct vote on her own career-concern reward is diluted in proportion to the size of the committee.

Finally, when the state of the world is not observed, each member computes her expected career-concern reward as follows

$$\widetilde{r}_i^{\lambda}(v_i) = q\widetilde{r}^{\omega=A,\lambda}(v_i) + (1-q)\,\widetilde{r}^{\omega=B,\lambda}(v_i). \tag{7}$$

Assuming that the state of the world is A, the expected utility of a competent member biased towards β_i can be expressed as a function of his vote v_i as follows

$$U^{\beta_i=A,\lambda}(v_i, s_i = A) = \phi\widetilde{r}_i^{\omega=A,\lambda}(v_i) + \rho^{\omega=A}(v_i)(\alpha+\gamma) \tag{8}$$

and

$$U^{\beta_i=B,\lambda}(v_i, s_i = A) = \phi\widetilde{r}_i^{\omega=A,\lambda}(v_i) + \rho^{\omega=A}(v_i)\alpha + (1-\rho^{\omega=A}(v_i))\gamma, \tag{9}$$

depending on whether the member is biased towards A or B, respectively. Similar expressions can be derived for the case where $\omega = B$.

The next lemma provides a general characterization of the behavior of competent members.[15]

Lemma 1. *The behavior of competent members is characterized by the following properties:*

i. Both abstaining and voting against the bias are weakly dominated strategies for a competent member whose bias is equal to the signal, $s_i = \beta_i$;

ii. Abstaining is a weakly dominated strategy for a competent member whose bias is different from her signal, $s_i \neq \beta_i$.

Intuitively, competent members observe the state of the world and, as a consequence, are not subject to the "swing voter's curse" (Feddersen and Pesendorfer [10]).

[15] All proofs can be found in Online Appendix D.

Therefore, there is no reason for them to abstain. Lemma 1 implies that a competent member who receives a signal equal to her bias always (weakly) prefers to vote for the state. Instead, a competent member who receives a signal different from her bias may either (weakly) prefer to vote for the state or for her bias. Note also that Lemma 1 guarantees that, in any equilibrium, every competent members biased towards the state of the world votes correctly. Thus, by Bayes rule, the likelihood that a member is competent given that she voted correctly is strictly positive in any equilibrium, $\Pr(t = \mathsf{c}|v = \omega) > 0$. The next lemma follows as an implication of this result.

Lemma 2. *In equilibrium, a member's expected career-concern reward is always strictly larger when she votes correctly than when she abstains or votes incorrectly.*

Based on the above two results, we are now able to characterize the equilibrium behavior of incompetent members in the next lemma.

Lemma 3. *There exists no equilibrium where a competent member who receives a signal different than her bias votes against the signal and an incompetent member abstains. Furthermore, if in equilibrium a competent member with bias β_i votes for her bias when the signal is $s_i \neq \beta_i$, then all incompetent members with bias β_i must vote for their bias.*

Intuitively, incompetent members are always more inclined to follow their biases relative to competent members. Note that when a competent individual decides to vote against the signal, she is certain to be casting an incorrect vote, while an incompetent member always attributes positive probability to the event that her vote is correct. Specifically, Lemma 3 guarantees that if, for instance, in equilibrium, competent members biased towards A vote for their bias when the state is B, then all incompetent members biased towards A must vote for their bias. As for the behavior of incompetent members with bias B in this case, the lemma just says that they will never abstain − they might vote either for their bias or for the ex-ante more likely alternative (against their bias).[16]

Finally, we can show that it is possible to classify the equilibria of the model into three classes.

[16] Note that voting against the bias might be optimal for an incompetent member if the prior q is very asymmetric.

Proposition 1. *The set of symmetric pure-strategy equilibria of the model can be categorized into one of the following classes:*

i. A fully-competent equilibrium, where all competent members vote in accordance with the signal and all incompetent members abstain;

ii. A partially-competent equilibrium, where all competent members vote in accordance with the signal and not all types of incompetent members abstain;

iii. A biased equilibrium, where not all types of competent members vote in accordance with the signal and all incompetent members vote.

A fully-competent equilibrium completely pins down the behavior of all committee members, while both a partially-competent and a biased equilibrium allow for a variety of different behaviors within each class. A partially-competent equilibrium pins down only the behavior of competent members and is consistent with either both types of incompetent members voting – either for the ex-ante more likely alternative or for their bias – or with one type voting (e.g., A biased) and the other abstaining (e.g., B biased). A biased equilibrium is consistent with either both types of competent members always voting for their biases or with one type always voting for the state and the other always voting for the bias. As for the incompetent members, they never abstain in a biased equilibrium and, as prescribed by Lemma 3, they vote for their bias if the competent members of the same type are doing so as well.

Proposition 1 helps organizing the set of all possible equilibria by grouping them in terms of key qualitative features of voters' behavior.[17] Importantly, our characterization holds under both public and secret voting, although the region of parameters where each class of equilibrium exists does depend on the level of transparency of the voting rule.

4.2 Main Comparative Statics Results

This subsection provides a general characterization of the regions of parameters where it is possible to sustain each class of equilibrium under secret and public voting. The following proposition summarizes the main properties of a fully-competent equilibrium.

[17] Furthermore, as we shall discuss in Subsection 4.3, under symmetry $q = p = 1/2$ there exists a unique equilibrium in each class.

Proposition 2. *There exists a unique threshold* $\overline{\gamma}^{\lambda}_{full}(\alpha, \phi, \sigma, n) < \alpha$ *such that a fully-competent equilibrium can be sustained if and only if* $\gamma \leq \overline{\gamma}^{\lambda}_{full}(\alpha, \phi, \sigma, n)$. *Furthermore* $\overline{\gamma}^{\mathtt{s}}_{full}(\alpha, \phi, \sigma, n) > \overline{\gamma}^{\mathtt{p}}_{full}(\alpha, \phi, \sigma, n)$.

A fully-competent equilibrium can be sustained only if the magnitude of the bias is small relatively to the common value and it is more likely to exist under secret voting. The binding constraint for this class of equilibrium is that on the behavior of incompetent members and the interaction between transparency and career-concerns creates an incentive for incompetent members to vote, since abstaining perfectly reveals their lack of competence in equilibrium.

The next proposition characterizes the bounds on parameters such that a partially-competent equilibrium exists.

Proposition 3. *There exist unique thresholds* $\underline{\gamma}^{\lambda}_{part}(\alpha, \phi, \sigma, n)$ *and* $\overline{\gamma}^{\lambda}_{part}(\alpha, \phi, \sigma, n)$, *with* $\underline{\gamma}^{\lambda}_{part}(\alpha, \phi, \sigma, n) < \alpha < \overline{\gamma}^{\lambda}_{part}(\alpha, \phi, \sigma, n)$, *such that a partially-competent equilibrium can be sustained if and only if:*

$$\underline{\gamma}^{\lambda}_{part}(\alpha, \phi, \sigma, n) \leq \gamma \leq \overline{\gamma}^{\lambda}_{part}(\alpha, \phi, \sigma, n).$$

Furthermore $\underline{\gamma}^{\mathtt{p}}_{part}(\alpha, \phi, \sigma, n) < \underline{\gamma}^{\mathtt{s}}_{part}(\alpha, \phi, \sigma, n) < \overline{\gamma}^{\mathtt{s}}_{part}(\alpha, \phi, \sigma, n) < \overline{\gamma}^{\mathtt{p}}_{part}(\alpha, \phi, \sigma, n)$.

A partially-competent equilibrium can be sustained even if the magnitude of the bias is large relatively to the common value and it is more likely to exist under public voting. In particular, the region of parameters where a partially-competent equilibrium can be supported under secret voting is strictly contained in the region where a partially-competent equilibrium can be supported under public voting. Observe that transparency acts to counter-balance the effect of the biases for competent members by creating an incentive for them to vote correctly in order to signal their competence. At the same time, it provides incentive for incompetent members to vote rather than to abstain.

We emphasize that a partially-competent equilibrium is consistent with a number of different behaviors by incompetent members, so that the condition above simply guarantees that a partially-competent equilibrium of "some sort" exists. To be clear, a move from secret to public voting (or vice-versa) might cause the equilibrium to change from one type of partially-competent equilibrium to another (e.g. it might cause one group of incompetent members to change from abstaining to voting). The results presented in Proposition 3 apply broadly to the class of partially-competent equilibrium, as we do distinguish between different subclasses. We will be able to derive more specific predictions about the behavior of committee members in Subsection 4.3, where we analyze the symmetric version of the model.

Finally, the next proposition characterizes the bound on parameters such that a biased equilibrium exists.

Proposition 4. *There exists a unique threshold* $\underline{\gamma}^{\lambda}_{bias}(\alpha, \phi, \sigma, n) > \alpha$ *such that a biased equilibrium can be sustained if and only if* $\underline{\gamma}^{\lambda}_{bias}(\alpha, \phi, \sigma, n) \leq \gamma$. *Furthermore,* $\underline{\gamma}^{\mathtt{s}}_{bias}(\alpha, \phi, \sigma, n) < \underline{\gamma}^{\mathtt{p}}_{bias}(\alpha, \phi, \sigma, n)$.

A biased equilibrium can be sustained only if the bias is large enough and is more likely to exist under secret voting. In particular, the region of parameters where a biased equilibrium can be supported under public voting is strictly contained in the region of parameters where a biased equilibrium can be supported under secret voting. Intuitively, secrecy reduces the career-concern rewards associated with a correct vote, which makes competent members more willing to disregard their information about the state of the world and vote in accordance with their biases. We also observe that a biased equilibrium is consistent a number of different behaviors by competent and incompetent members, so that the same caveats discussed above apply to this case.

Finally, it is important to note that there will generally be an overlap between the regions of parameters where a fully-competent and a partially-competent equilibria can be supported as well as between the regions of parameters where a partially-competent and a biased equilibrium can be supported. Overall, our analysis highlights the fact that transparency affects the behavior of competent and incompetent members in different ways. On the one hand, transparency *attenuates* the preexisting biases of competent members by inducing them to vote correctly, even if the state of the world contradicts their biases. On the other hand, transparency *exacerbates* the preexisting biases of incompetent members by inducing them to vote in order to avoid exposing their lack of competence. While these incentives exist everywhere in the parameters' space, our analysis shows that they may lead to actual changes in observable voting behavior in different situations. Specifically, when the magnitude of the biases is relatively large, transparency may induce competent members to vote correctly rather than incorrectly (an attenuation effect) – while incompetent members vote anyway. Alternatively, when the magnitude of the biases is relatively small, transparency may induce incompetent members to vote rather than abstain (an exacerbation effect) – while competent members vote correctly anyway.

4.3 The Symmetric Case

In this subsection, we assume that the distributions of both the prior and the biases are symmetric, i.e. $q = p = 1/2$. The symmetric prior assumption guarantees that when an incompetent member decides to vote she always votes for her bias. Moreover, the assumption on the distribution of biases further simplifies the analysis by making symmetric the incentives of members with the same competence but different biases. Together, these assumptions also imply that there exists a unique equilibrium in each class and that behavior is completely pinned-down. Specifically, the unique partially-competent equilibrium is such that all competent members vote correctly and all incompetent members vote for their biases, while the unique biased equilibrium is such that all members votes for their biases. Under symmetry we can explicitly solve for the thresholds defined in Propositions 2, 3 and 4.

Proposition 5. *Suppose that $q = p = 1/2$, then:*

i. A fully-competent equilibrium can be supported if and only if:

$$\gamma \leq \overline{\gamma}^{\lambda}_{full}\left(\alpha, \phi, \sigma, n\right) \equiv \frac{\left(n-1\right)\sigma}{2+\left(n-3\right)\sigma}\alpha - \frac{\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathtt{s}\}}\right)\phi}{\left(1+\frac{n-3}{2}\sigma\right)\left(1-\sigma\right)^{n-2}}$$

ii. A partially-competent equilibrium can be supported if and only if:

$$\gamma \leq \overline{\gamma}^{\lambda}_{part}\left(\alpha, \phi, \sigma, n\right) \equiv \alpha + \frac{2^{n}\sigma\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathtt{s}\}}\right)\phi}{\binom{n-1}{(n-1)/2}\left(1+\sigma\right)^{\frac{n+1}{2}}\left(1-\sigma\right)^{\frac{n-1}{2}}}$$

iii. A biased equilibrium can be supported if and only if:

$$\gamma \geq \underline{\gamma}^{\lambda}_{bias}\left(\alpha, \phi, \sigma, n\right) \equiv \alpha + \frac{2^{n-1}\sigma\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathtt{s}\}}\right)\phi}{\binom{n-1}{(n-1)/2}}$$

Furthermore, $\overline{\gamma}^{\lambda}_{full} < \underline{\gamma}^{\lambda}_{bias} < \overline{\gamma}^{\lambda}_{part}$, $\overline{\gamma}^{\mathtt{p}}_{full} < \overline{\gamma}^{\mathtt{s}}_{full}$, $\overline{\gamma}^{\mathtt{p}}_{part} > \overline{\gamma}^{\mathtt{s}}_{part}$, and $\underline{\gamma}^{\mathtt{p}}_{bias} > \underline{\gamma}^{\mathtt{s}}_{bias}$.

The term $((n-1)/n)\cdot\mathbb{I}_{\{\lambda=\mathtt{s}\}}$ appearing in the expressions above captures the impact of the dilution of career-concern rewards under secret voting.[18] Note that a change from public to secret voting is qualitatively equivalent to a reduction in the weight on reputation ϕ. Figure 2 shows the values of the parameters α and γ for which each class of equilibrium can be sustained, for given transparency λ and for fixed values of ϕ, σ and n.

Observe that since $\overline{\gamma}^{\lambda}_{full} < \overline{\gamma}^{\lambda}_{part}$, the region of parameters where a fully-competent equilibrium exists is contained inside the region where a partially-competent equilibrium exists. Recall that the main reason for an incompetent member to abstain is to avoid adding "noise" to the decision process. However, a coordination issue arises in the region where the two equilibria overlap in that abstaining can only be optimal for an incompetent member if she expects other incompetent members to abstain as well. Similarly, since $\underline{\gamma}^{\lambda}_{bias} < \overline{\gamma}^{\lambda}_{part}$, there exists a region of parameters where both a partially-competent and a biased equilibrium can be sustained simultaneously. The multiplicity of equilibria arises in this case due to the existence of a coordination issue among competent members who are biased against the state of the world. In the region where the two equilibria overlap, voting in accordance with one's bias can only be optimal if the member expects other competent members of the same type to do the same. The reason is that an individual is less likely to be pivotal when she is the only competent member voting against the state, in which case she would prefer to vote correctly in order to guarantee a larger career-concern reward for herself.

Figure 3 summarizes the main comparative static results of the model. Observe that in region I, where $\overline{\gamma}^{\mathsf{s}}_{part} < \gamma < \overline{\gamma}^{\mathsf{p}}_{part}$, a partially-competent equilibrium can be sustained under public but not under secret voting. Instead, in region II, where $\overline{\gamma}^{\mathsf{p}}_{full} < \gamma < \overline{\gamma}^{\mathsf{s}}_{full}$, a fully-competent equilibrium can be sustained under secret but not under public voting. When the magnitude of the bias is relatively large, as in region I, incompetent members always vote in accordance with their biases, but public voting may induce competent members to vote correctly rather than incorrectly. On the other hand, when the magnitude of the bias is relatively small, as in region II, competent members always vote correctly, but secret voting may induce incompetent members to abstain rather than to vote.

For each class of equilibrium, it can be shown that the probability of a correct

[18] Note that the threshold $\underline{\gamma}^{\lambda}_{part}$ which appears in Proposition 3 is strictly negative under symmetry, meaning that incompetent members always prefer to vote for their biases rather than to abstain. Intuitively, given that all other incompetent members are voting in equilibrium, it becomes optimal for each of them to vote as well.

decision is given by

$$\Pi_{full}(\sigma,n) = 1-\tfrac{1}{2}(1-\sigma)^n \tag{10}$$

$$\Pi_{part}(\sigma,n) = \textstyle\sum_{i=(n+1)/2}^{n} \binom{n}{i} \left(\sigma+\tfrac{1}{2}(1-\sigma)\right)^i \left(\tfrac{1}{2}(1-\sigma)\right)^{n-i} \tag{11}$$

$$\Pi_{bias}(\sigma,n) = \tfrac{1}{2}, \tag{12}$$

with $\Pi_{full}(\sigma,n) > \Pi_{part}(\sigma,n) > \Pi_{bias}(\sigma,n)$ for any $0<\sigma<1$ and $n \geq 3$.[19] We are, therefore, able to rank public and secret voting in welfare terms, based on the expected quality of the decisions.

Proposition 6. *Suppose that $q = p = 1/2$. In equilibrium, we have that:*

i. If $\overline{\gamma}^{\mathsf{s}}_{part}(\alpha,\phi,\sigma,n) < \gamma < \overline{\gamma}^{\mathsf{p}}_{part}(\alpha,\phi,\sigma,n)$ then the probability of a correct decision under public voting is at least as large as under secret voting.

ii. If $\overline{\gamma}^{\mathsf{p}}_{full}(\alpha,\phi,\sigma,n) < \gamma < \overline{\gamma}^{\mathsf{s}}_{full}(\alpha,\phi,\sigma,n)$ then the probability of a correct decision under secret voting is at least as large as under public voting.

Note that because of the existence of multiple equilibria, as discussed above, we are only able to rank public and secret voting weakly in terms of welfare. We complement our characterization of the equilibria by providing additional comparative statics results based on the expressions derived in Proposition 5.

Proposition 7. *The following comparative static results hold:*

i. Career-concerns (ϕ). For any $\alpha \geq 0$, $\phi > 0$, $0<\sigma<1$, $n \geq 3$, we have that:

$$\frac{\partial \overline{\gamma}^{\lambda}_{full}(\alpha,\phi,\sigma,n)}{\partial \phi} < 0, \quad \frac{\partial \overline{\gamma}^{\lambda}_{part}(\alpha,\phi,\sigma,n)}{\partial \phi} > 0 \quad \textit{and} \quad \frac{\partial \underline{\gamma}^{\lambda}_{bias}(\alpha,\phi,\sigma,n)}{\partial \phi} > 0$$

ii. Competent members (σ). There exists $\overline{n} \in \mathbb{R}$ such that for any $\alpha \geq 0$, $\phi > 0$, $0<\sigma<1$ and $n \geq \overline{n}$, we have that:

$$\frac{\partial \overline{\gamma}^{\lambda}_{full}(\alpha,\phi,\sigma,n)}{\partial \sigma} < 0 \quad \textit{and} \quad \frac{\partial \overline{\gamma}^{\lambda}_{part}(\alpha,\phi,\sigma,n)}{\partial \sigma} > 0$$

[19] Observe that the probability of a correct decision is smaller than 1 even under a fully-competent equilibrium, since with probability $(1-\sigma)^n$ all committee members are incompetent.

Furthermore, for any $n \geq 3$, we have:

$$\frac{\partial \underline{\gamma}^{\lambda}_{bias}(\alpha, \phi, \sigma, n)}{\partial \sigma} > 0$$

The comparative static results with respect to ϕ are intuitive in light of our previous discussion. Both a fully-competent and a biased equilibrium become more difficult to sustain as the importance of career-concerns increases, while a partially-competent equilibrium becomes easier to sustain. In fact, for an arbitrarily large ϕ only a partially-competent equilibrium exists. On the other hand, the comparative static results with respect to σ are somewhat more subtle. We can guarantee that, for n sufficiently large, a fully-competent equilibrium becomes less likely to exist as the proportion of competent members increases. Indeed, as σ goes up, the likelihood that an incompetent member is pivotal when she casts an incorrect vote decreases, which gives her a stronger incentive to vote. Furthermore, a partially-competent equilibrium becomes more likely to exist as σ increases, provided that n is large enough. Note that in this case an increase in σ reduces the likelihood that a competent member is pivotal when she casts an incorrect vote, which gives her a stronger incentive to vote correctly. Finally, a biased equilibrium is always less likely to exist as σ increases. The general intuition here is that an increase in σ raises the opportunity cost of voting against the state, given that the career-concern rewards associated with a correct vote are increasing in the fraction of competent members.

In the next proposition we analyze what happens to the equilibrium thresholds as the size of the committee gets arbitrarily large.

Proposition 8. *For any $\alpha \geq 0$, $\phi > 0$, $0 < \sigma < 1$, we have that:*

i. Under public voting:

$$\lim_{n\to\infty} \overline{\gamma}^{\mathsf{p}}_{full}(\cdot) = -\infty, \quad \lim_{n\to\infty} \overline{\gamma}^{\mathsf{p}}_{part}(\cdot) = +\infty \text{ and } \lim_{n\to\infty} \underline{\gamma}^{\mathsf{p}}_{bias}(\cdot) = +\infty$$

ii. Under secret voting:

$$\lim_{n\to\infty} \overline{\gamma}^{\mathsf{s}}_{full}(\cdot) = -\infty, \quad \lim_{n\to\infty} \overline{\gamma}^{\mathsf{s}}_{part}(\cdot) = +\infty \text{ and } \lim_{n\to\infty} \underline{\gamma}^{\mathsf{s}}_{bias}(\cdot) = \alpha$$

As n gets arbitrarily large a fully-competent equilibrium can never be supported. Indeed, the probability that an incompetent member is pivotal in a fully-competent equilibrium converges to zero as $n \to \infty$, so that incompetent members have a large incentive to vote. Thus, contrarily to Feddersen and Pesendorfer (1996), information is never fully aggregated in large elections.[20] Furthermore, note that a partially-competent equilibrium exists everywhere in the parameters' space under both public and secret voting for large n. Finally, a biased equilibrium can only exist under secrecy. Thus, overall, our analysis shows that in large elections with career-concerns, transparency is expected to lead to (weakly) better decisions

5 Extensions

In this section we examine the implications of allowing for information sharing prior to the voting stage and we also consider a dynamic extension of our model. We cover other generalizations and extensions of the basic model in the Online Appendix A.

5.1 Model with Deliberation

In our benchmark model we assume that signals are privately observed and that competent members are unable to share information with others. In a common value environment, Coughlan (2000) showed that voters have a strong incentive to share information truthfully. However, given the presence of bias and career-concerns, the direction of incentives is less clear in our setting.

We analyze an extension of the symmetric version of the model with two stages: (i) a deliberation stage and (ii) a voting stage. At the beginning of the game, each member receives a signal $s_i \in \{A, \emptyset, B\}$ about the state, as in the benchmark model. We focus our analysis on the case where, at the deliberation stage, members are able to exchange hard information: each member sends a message $m_i \in \{s_i, \emptyset\}$ to the committee, which may reveal or not his private signal truthfully.[21] At the voting stage, after having observed all messages $m = (m_1, ..., m_n)$, the members of the committee vote and the decision is taken by majority. As before, under public voting, all votes are observed by the external evaluator, while under secret voting only the vote tally is observed. We

[20] This result holds for any positive ϕ, even if arbitrarily small.

[21] This framework captures the idea that experts' opinions must be backed up with facts and evidence, which cannot be fabricated by incompetent individuals.

assume that in both cases the messages exchanged at the deliberation stage are never observed by the evaluator.

The sequential nature of the model gives rise to a variety of possible equilibria. We focus on the characterization of two general classes of equilibria. First, the set of *responsive equilibria*, where on the equilibrium path all committee members vote correctly when the state is revealed. Within this class, we distinguish between *full revelation equilibrium*, where all competent members reveal their signals at the deliberation stage and *partial revelation equilibrium*, where only the competent members with bias equal to the state reveal their signals.[22] Second, within the set of *non-responsive equilibria* we focus on one particular subclass, which we refer to as *irrelevant deliberation equilibrium*, where committee members always vote for their biases irrespective of whether the state is revealed or not at the deliberation stage.[23] We are now ready to summarize the main insights that emerge from adding a stage of deliberation to our benchmark model.[24]

A preliminary observation is that in any responsive equilibrium the revelation of the state of the world by a single member is sufficient to induce an unanimous correct decision. In equilibrium, the external evaluator takes this fact into account and understands that a unanimous correct decision implies that with high probability there is at least one competent member in the committee.

Full Revelation Equilibrium. In a full revelation equilibrium, all competent members reveal their signals at the deliberation stage. Then, at the voting stage, if the state was revealed, all members vote correctly. Otherwise, everyone votes for their biases. Note that, in this case, career-concern rewards are positive only if the committee's decision is unanimously correct. Therefore, the structure of incentives under public and secret voting is exactly the same.

Notice that behavior at the voting stage is straightforward. If information about the state is revealed during the deliberation process, then at the voting stage no member ever wants to vote incorrectly, since the probability of being pivotal is zero. On the

[22] An equilibrium where competent members never reveal their signals does not exist within the class of responsive equilibria. Indeed, note that in this case a competent member with bias equal to the state would always have an incentive to share information.

[23] While there are other types of non-responsive equilibria, all of them involve peculiar behavior, such as a competent member voting against the state when it is revealed and for the state when it is not revealed. We do not consider these equilibria in our analysis given their limited applied interest.

[24] We provide a detailed characterization of the equilibria in the Online Appendix B.1.

other hand, if information is not revealed on the equilibrium path, then it becomes common knowledge among committee members that everyone is incompetent and it is, therefore, optimal for all members to simply vote for their biases.

At the deliberation stage, the relevant incentive constraint is that on competent members biased against the state. These members are willing to share information only if the magnitude of the bias is not too large. Formally, we show that there exists a threshold $\overline{\gamma}_{fr}^{\lambda}(\alpha, \phi, \sigma, n) > \alpha$, such that a full revelation equilibrium can be sustained if and only if $\gamma \leq \overline{\gamma}_{fr}^{\lambda}(\alpha, \phi, \sigma, n)$. Furthermore, the existence condition is the same under public and secret voting, $\overline{\gamma}_{fr}^{\mathtt{s}}(\alpha, \phi, \sigma, n) = \overline{\gamma}_{fr}^{\mathtt{p}}(\alpha, \phi, \sigma, n)$.

Partial Revelation Equilibrium. In a partial revelation equilibrium, only the competent members with bias equal to the state reveal their signals, yet everyone votes correctly when information is shared. Note that, conditional on the state not being revealed during deliberation, the equilibrium is consistent with a variety of voting behaviors. We focus on the more interesting case: a partial revelation equilibrium with "partially-competent voting" when no information is revealed.[25]

A partial revelation equilibrium captures situations in which competent members might have an incentive to strategically withhold information and then vote correctly in order to look "smart" in the eyes of the external evaluator. Note that the binding constraints for the existence of such equilibrium are those on the behavior of competent members biased against the state. In particular, the magnitude of the bias should be large enough for them not to reveal information, but not too large so that they are willing to vote correctly when information is not shared. Formally, we show that there are thresholds $\underline{\gamma}_{pr}^{\lambda}(\alpha, \phi, \sigma, n) > \alpha$ and $\overline{\gamma}_{pr}^{\lambda}(\alpha, \phi, \sigma, n) > \alpha$, such that a partial revelation equilibrium with partially-competent voting can be sustained if and only if $\underline{\gamma}_{pr}^{\lambda}(\alpha, \phi, \sigma, n) \leq \gamma \leq \overline{\gamma}_{pr}^{\lambda}(\alpha, \phi, \sigma, n)$. Furthermore, the equilibrium is always more likely to exist under transparency, since $\underline{\gamma}_{pr}^{\mathtt{p}} < \underline{\gamma}_{pr}^{\mathtt{s}}$ and $\overline{\gamma}_{pr}^{\mathtt{p}} > \overline{\gamma}_{pr}^{\mathtt{s}}$. Intuitively, the career-concern rewards associated with the strategy of withholding information and then voting correctly are larger under transparency. Thus, our analysis suggests that there might be situations in which secret voting may actually lead to better outcomes by stimulating more information sharing at the deliberation stage.

[25] A partial revelation equilibrium with "fully-competent voting" does not exist under public voting. Indeed, a competent member with bias equal to the state would also have an incentive to withhold information and vote correctly in an attempt to separate himself from incompetent members. The case of a partial revelation equilibrium with "biased voting" has qualitative features that make it very similar to an irrelevant deliberation equilibrium, which we analyze below in detail.

Irrelevant Deliberation Equilibrium. Within the set of non-responsive equilibria, an irrelevant deliberation equilibrium is such that all committee members vote for their biases regardless of whether the state was revealed or not at the deliberation stage. The structure of an irrelevant deliberation equilibrium is very similar to that of a biased equilibrium. In fact, it is possible to show that the equilibrium can only exist if the bias term is sufficiently large and that it is always more likely to exist under secret voting. Formally, there exists a threshold $\underline{\gamma}_{irr}^{\lambda}(\alpha, \phi, \sigma, n) > \alpha$ such that an irrelevant deliberation equilibrium can be sustained if and only if $\gamma \geq \underline{\gamma}_{irr}^{\lambda}(\alpha, \phi, \sigma, n)$, where $\underline{\gamma}_{irr}^{\mathsf{p}} > \underline{\gamma}_{irr}^{\mathsf{s}}$.

Comparative Results. While it is difficult to derive general comparative static results in the model with deliberation, our analysis highlights the fact that the level of transparency still matters, although the nature of the trade-off is slightly different in this case.[26] First, in line with Coughlan (2000), we conclude that when $\alpha > \gamma$ a full deliberation equilibrium always exists under both public and secret voting. In this case the state is revealed at the deliberation stage $-$ provided that there is at least one competent member in the committee $-$ and the correct decision is taken by unanimity. Interestingly, the same force that induced incompetent members to abstain in a model without deliberation now generates an incentive for competent members to share information in equilibrium.

Conversely, when the bias term is relatively large, $\alpha < \gamma$, transparency might involve a trade-off between the quality of information aggregation at the deliberation stage and the quality of the decisions at the voting stage. Under certain conditions, secrecy may actually make it more likely that information about the state is revealed at the deliberation stage, while transparency creates an incentive for informed members to vote correctly. In particular, note that while public voting always makes the irrelevant deliberation equilibrium less likely, there are some regions of the parameters where a change from public to secret voting might lead to a shift from a partial deliberation to a full deliberation equilibrium.[27] Altogether, these results reinforce our previous conclusions, highlighting another dimension in which the level of transparency might

[26] Some qualitative insights about the effect of transparency on information sharing and voting behavior can be obtained by evaluating the existence conditions for each class of equilibrium studied above for specific parameters' values. We discuss a simple example for the case where $n = 3$, $\sigma = 1/2$ and $\phi = 1$ in the Online Appendix B.1.

[27] See the Online Appendix B.1 for a quantitative example and further discussion.

be relevant for the quality of the decisions of the committee.[28]

5.2 Dynamic Model

In this subsection, we study an infinitely repeated version of our benchmark model, where committee members vote on a sequence of different independent issues. Sequential interaction is an important feature of many real world settings and it is important to understand how the main results of our model are influenced by it. For tractability, we assume that the state of the world is iid across periods ($q = 1/2$) and that each member's biases are iid across periods ($p = 1/2$). We consider two polar cases for competence. Individual competence may either be iid across periods or persistent (fixed) over time. Both the magnitude of the common value α and of the bias term γ are assumed to be constant to facilitate comparison with the static model. The discount factor is given by $\delta \in (0, 1)$.

We focus our analysis on the characterization of the conditions under which a *dynamic fully-competent equilibrium* can be sustained. We provide formal definitions below. As shall become clear, this class of equilibrium serves as a good benchmark for comparative analysis, since it is associated with the best outcomes in expectation. We are particularly interested in the following questions: Does repeated interaction necessarily lead to better outcomes? Is the difference between secret and public voting still relevant in the context of an infinitely repeated game?

5.2.1 The IID Case

We begin our analysis by considering the case where competence is iid across periods, with $\Pr(\tau = \mathsf{c}) = \sigma$.[29] Intuitively, this case can be viewed as capturing situations where a committee member may be an expert in some issues but not in others. Since both competence and bias are iid over time, the structure of the game in any period is very similar to that of the static model, except that the external evaluator is now able to condition the career-concern rewards on the entire history of observed votes. In this case a *dynamic fully-competent equilibrium* is defined as an equilibrium where committee members vote correctly in all periods in which they are competent and abstain in all

[28] Our characterization also highlights the idea present in Swank and Visser [39] that committee members have an incentive to show an united front whenever information is revealed during the deliberation process, since unanimity signals more strongly the competence of the committeee.

[29] A detailed derivation of the results presented in this section can be found in the Online Appendix B.2.

periods in which they are incompetent. Thus, on the equilibrium path, no member ever votes incorrectly and a correct vote in a given period perfectly reveals competence in that particular period. Outside of the equilibrium path, after an incorrect vote is observed, we assume that committee members switch to voting in accordance with a partially-competent equilibrium.[30]

Under public voting, the career-concern rewards in any period are such that, on the equilibrium path, if a member has never cast an incorrect vote before, his career-concern gain is equal to one if her vote is correct and zero otherwise. Furthermore, we assume that a member who votes incorrectly in one period is forever punished with zero reputation afterwards.

Under secret voting, on the other hand, the career-concern rewards in any period are such that, on the equilibrium path, the reputation of the committee is equal to the total number of correct votes in that period divided by the size of the committee. We assume that when an incorrect vote is observed, and in every period thereafter, the external evaluator simply attributes a reward that is equal to the prior probability that a member is competent σ multiplied by a factor $(n-1)/n$. The scaling captures the idea that the external evaluator "randomly" punishes one individual with zero reputation. Thus, as before, punishment is diluted among committee members.[31] The main results of our analysis are summarized in the following proposition:

Proposition 9. *There exists a threshold $\overline{\gamma}^{\lambda}_{dfull-i}(\alpha,\phi,\sigma,n,\delta)$ such that a dynamic fully-competent equilibrium can be sustained if and only if:*

$$\gamma \leq \min\{\alpha, \overline{\gamma}^{\lambda}_{dfull-i}(\alpha,\phi,\sigma,n,\delta)\}$$

Furthermore, we have that:

i. $\overline{\gamma}^{\lambda}_{dfull-i}(\alpha,\phi,\sigma,n,\delta)$ *is strictly increasing in* δ.

ii. There exists $\overline{\delta} \in (0,1)$ *such that* $\overline{\gamma}^{\mathsf{p}}_{dfull-i}(\alpha,\phi,\sigma,n,\delta) \geq \overline{\gamma}^{\mathsf{s}}_{dfull-i}(\alpha,\ \phi,\sigma,n,\delta)$ *if and only if* $\delta \geq \overline{\delta}$.

[30]Note that committee members will have an incentive to behave in accordance with a partially-competent equilibrium outside of the equilibrium path as long as $\gamma \leq \alpha$. The main insights of our analysis are robust to assuming different behaviors off the equilibrium path.

[31]Given the one-shot deviation principle, it is enough to consider situations where, out of equilibrium, a single member votes incorrectly. Our qualitative results are robust to specifying harsher punishments outside of the equilibrium, as long as the external evaluator does not punish the entire committee with zero career-concern rewards as a consequence of an incorrect vote.

Our analysis shows that it is easier to sustain behavior consistent with a fully-competent equilibrium when the game is played repeatedly and competence is iid across periods. Formally, we can show that $\overline{\gamma}^{\lambda}_{full}(\alpha, \phi, \sigma, n) < \overline{\gamma}^{\lambda}_{dfull-i}(\alpha, \phi, \sigma, n, \delta)$ for any $\delta \in (0, 1)$, so that the condition for the existence of a dynamic fully-competent equilibrium is easier to satisfy than the condition for a fully-competent equilibrium in the static model. Indeed, the fact that an incompetent member may turn out to be competent in the future helps to discipline her behavior in the current period.[32]

Regarding the comparison between public and secret voting (point *ii.* above), a dynamic fully-competent equilibrium is more likely to be sustained under public voting if and only if the discount factor is large enough. Note that a novel trade-off between public and secret voting emerges in the dynamic version of the model. Besides the benefit associated with voting in the current period, there now exists the expectation of punishments in future periods. Transparency increases both the incentive for an incompetent member to vote today and the expected punishment associated with an incorrect vote, so that the overall effect depends on the discount factor. When the discount factor is large enough, committee members place a higher weight on the future stream of payoffs, so that a harsher expected punishment makes it easier to sustain a dynamic fully-competent equilibrium. In this case, public voting is expected to lead to better voting outcomes in terms of information aggregation. On the other hand, when the discount factor is small, committee members place a higher weight on the current payoff, so that lower career-concern rewards associated with voting in the current period make it easier to sustain a dynamic fully-competent equilibrium. In this case, secret voting is expected to lead to better outcomes.

5.2.2 Persistent Competence

We now consider the case where the competence of committee members is realized at the beginning of the game, with $\Pr(\tau = \mathtt{c}) = \sigma$, and remains fixed over time. As before, we maintain the assumption that both the state of the world and the individual biases are iid across periods. A dynamic fully-competent equilibrium in this case is defined as an equilibrium where competent members vote correctly and incompetent members abstain in all periods $-$ unless it becomes common knowledge that all members are

[32] Observe that the term α, which appears inside the minimum operator in the first inequality of Proposition 9, is required only to guarantee that a behavior consistent with a partially-competent equilibrium can be sustained out of equilibrium.

incompetent, in which case everyone votes for their biases.[33]

Let $v_{it}^{c} \equiv \mathbb{I}_{\{v_{it}=\omega_t\}}$ denote an indicator variable that equals one if the vote of member i in period t is correct and define $V_t^c \equiv \sum_{i=1}^{n} v_{it}^{c}$ as the total number of correct votes in period t. Under public voting, the external evaluator observes the entire history of individual votes and we assume that the career-concern reward of member i in period t is given by $r_{it}^{\mathsf{p}} = \min\left\{v_{i1}^{c}, v_{i2}^{c}, ..., v_{it}^{c}\right\}$. Note that this formulation assumes that both an incorrect vote and an abstention in a given period are forever punished with zero reputation. Under secret voting, on the other hand, the external evaluator observes only the history of aggregate votes, in which case we assume that the career-concern reward of the committee is given by $r_t^{\mathsf{s}} = 1/n \cdot \min\left\{V_1^c, ..., V_t^c\right\}$, where we suppose that the external evaluator attributes reputation based on the minimum number of correct votes in order to punish deviations from the equilibrium strategies. The main results of our analysis are summarized in the following proposition:

Proposition 10. *There exists a threshold $\overline{\gamma}_{dfull-p}^{\lambda}(\alpha, \phi, \sigma, n, \delta) < \alpha$ such that a dynamic fully-competent equilibrium can be sustained if and only if:*

$$\gamma \leq \overline{\gamma}_{dfull-p}^{\lambda}(\alpha, \phi, \sigma, n, \delta)$$

Furthermore, we have that:

i. $\overline{\gamma}_{dfull-p}^{\lambda}(\alpha, \phi, \sigma, n, \delta)$ *is strictly decreasing in* δ.

ii. A dynamic fully-competent equilibrium is always more likely to exist under secret voting, i.e. $\overline{\gamma}_{dfull-p}^{\mathsf{p}}(\alpha, \phi, \sigma, n, \delta) < \overline{\gamma}_{dfull-p}^{\mathsf{s}}(\alpha, \phi, \sigma, n, \delta)$.

Contrary to the iid case, when competence is persistent it is harder to sustain behavior consistent with a fully-competent equilibrium when the game is played repeatedly. Intuitively, when types are fixed, incompetent members have a larger incentive to vote, since an abstention in a given period is forever punished with zero career-concern rewards. Note that there now exists a dynamic gain associated with not revealing their types in the current period. Finally, regarding the comparison between secret and

[33] For simplicity, we assume that out of equilibrium, when an incorrect vote is observed, committee members continue behaving in accordance with a fully-competent equilibrium. As will become clear, while this assumption simplifies the characterization of the equilibrium considerably, the main mechanism behind the results does not depend on it.

public voting, our analysis highlights the fact that both the static and the dynamic incentives for an incompetent member to vote are larger under transparency. Therefore, when competence is persistent, it is always easier to sustain a dynamic fully-competent equilibrium under secret voting.

6 Experimental Design

In this section we test the main theoretical predictions of our basic model by means of a controlled laboratory experiment. Since the choice of adopting secret or public voting may be endogenous to the composition of the committee as well as to the types of decisions being made, it is particularly difficult to evaluate the impact of transparency on voting outcomes using non-experimental data.[34] A lab experiment allows us to both collect data on individuals' behaviors and compare the quality of the decisions under public and secret voting, while controlling for the degree of information and biases of committee members. Furthermore, since our model features multiple equilibria with different information aggregation properties, the experimental results can inform on whether subjects select a particular equilibrium.

For the experimental implementation, we amend the basic model imposing two simplifying assumptions on the structure of the career-concern rewards. First, we assume that the career-concern rewards associated with a correct vote are exogenous under both public and secret voting. Specifically, before voting, each committee member knows, and is guaranteed to receive, a certain payoff $R^{\lambda} > 0$ when she votes correctly. Note that this simplification maintains all basic features of the original model, except that now the updating process of the external evaluator is not being explicitly modelled.[35] Second, while it is natural to suppose that $R^{\mathsf{p}} > R^{\mathsf{s}}$, we further assume that $R^{\mathsf{s}} = 0$, i.e. the career-concern reward associated with a correct vote is zero under secret voting. We make this assumption in order to sharpen the contrast between the two treatments. The rest of the model remains unchanged. In particular, the same three classes of equilibria exist, there are multiple equilibria in some regions of the

[34] We consider the incentives of different types of members to choose between secret and public voting in the Online Appendix A.

[35] In this way we do not need a human evaluator whose role would be to guess the competence of committee members. Both Fehrler and Hughes [11] and Meloso and Ottaviani [27] find that experimental subjects have a hard time updating beliefs correctly in the lab. In particular, Meloso and Ottaviani [27] show that human evaluations tend to be so noisy that they considerably dampen the incentives of other participants, especially in treatments where there are multiple equilibria.

parameters' space and all previous comparative static results hold.

We consider committees of three members with uniform prior ($q = 1/2$) and symmetric distribution of both biases ($p = 1/2$) and competence ($\sigma = 1/2$). Recall that in this case there exists a unique equilibrium in each class.[36] Under this parametrization, the conditions for the existence of each class of equilibria are given by

$$\gamma \leq \overline{\gamma}^{\lambda}_{full} \equiv \frac{1}{2}\alpha - 2R^{\lambda} \tag{13}$$

$$\gamma \leq \overline{\gamma}^{\lambda}_{part} \equiv \alpha + \frac{8}{3}R^{\lambda} \tag{14}$$

and

$$\gamma \geq \overline{\gamma}^{\lambda}_{bias} \equiv \alpha + 2R^{\lambda}, \tag{15}$$

with the usual notation.[37] We concentrate our analysis on regions of the parameters' space where a change in the degree of transparency is expected to lead to a change in observed behavior. The choice of parameters as well as the equilibrium predictions associated with each of the four treatments considered in the experiments are reported in Table 2. The common value is set to $\alpha = 10$ in all treatments, while the magnitude of the bias can be either low, $\gamma = 1$, or high, $\gamma = 14$. Furthermore, the career-concern reward is $R^{\mathrm{p}} = 9$ under public voting and $R^{\mathrm{s}} = 0$ under secret voting. Accordingly, there are four treatments labelled Low/Secret, Low/Public, High/Secret and High/Public, with the Low/Secret treatment being consistent with both a fully-competent and a partially-competent equilibrium.[38]

The experiments were conducted at the Bologna Laboratory for Experiments in Social Science (BLESS) with registered undergraduates from the University of Bologna. We run the experiments in 6 sessions, each consisting of 2 parts with a different treatment being tested in each part (within-subject design). Each treatment was repeated for 32 rounds, the first two being practice non-paid rounds. In every session, the value of the bias was held fixed and only the parameter corresponding to the career-concern reward (public or secret voting) changed from one part to the other. Table 3 reports

[36] The fully-competent equilibrium is such that all competent members vote correctly and all incompetent abstain, the unique partially-competent equilibrium is such that all competent members vote correctly and all incompetent members vote for their biases, and the the unique biased equilibrium is such that all members votes for their biases.

[37] See the Online Appendix E for the derivation of these conditions.

[38] Since there are multiple equilibria in the Low/Secret treatment, in principle one could observe no difference in voting behavior and percentage of correct decisions between Low/Secret and Low/Public.

the sequence of treatments and number of participants in each session. In total, 144 different subjects took part in the experiments.

Thes experiments were implemented via computer terminals and programmed in z-Tree. In every session, instructions were read aloud at the beginning of each part, after which a short comprehension quiz was administered in order to check basic understanding of the rules.[39] Subjects were randomly divided into groups of three members and were re-assigned, in every period, to different groups using a random matching procedure. The task of each group was to choose between two colors, blue or yellow. The "group's color" (i.e. the state of the world) was ex-ante unknown and could be either one of the two colors with equal probability.

Before voting, each subject received a message about the group's color that could be either perfectly informative or non-informative with equal probability.[40] Specifically, subjects were told that the messages would be randomly assigned so that, among all participants in a given session, half of them would receive a perfectly informative message saying either "*blue*" or "*yellow*" depending on the group's color, and the other half would receive an uninformative message saying "*blue or yellow with equal probability*", in which case no new information would be added to what was previously known.[41] At this point, we were explicit in emphasizing that this procedure did not guarantee that there would always be an informed member in every group and that, in fact, the number of informed members in a given committee could be anything between zero and three. Then, each subject was informed about his or her "role" (i.e. bias), which could be either blue or yellow with equal probability. The procedure used to assign individual colors was the same used to assign group colors. After observing their messages and roles, each subject had to choose whether to vote for blue or yellow or to abstain. The "group's decision" was taken by majority rule and ties were broken

[39] All participants were provided with a copy of the instructions that they could consult at any moment during the experiment. See the Online Appendix F for a version of the instructions translated into English.

[40] In our discussion of the experiment, we will refer to subjects who receive informative messages (competent) as "informed" and to subjects who receive non-informative messages (incompetent) as "uninformed".

[41] This distribution procedure was adopted in order to make the experiment as transparent as possible. Note, however, that it introduces a minor correlation in the distribution of messages in that if, for instance, a subject receives an informative message, then it is slightly less likely that another participant will receive an informative message as well. As a consequence, the conditions for the existence of each class of equilibria are slightly different than (13)-(15). However, for the number of participants and parameters' values used in each session, all of our equilibrium predictions remain unchanged.

randomly. At the end of each period, subjects were provided with information about their group's color, the decision taken and the number of members of the group that voted for blue, yellow or abstained.

The final payoff in a given period was such that if the group's decision was equal to the group's color, then each member of the group received 10 points. Moreover, if the group's decision was equal to the role of one of its members, then she received 1 extra point in the low bias treatments and 14 extra points in the high bias treatments. Finally, under public voting treatments, subjects were also given an additional payoff of 9 points if her vote was equal to the group's color, while no points were given to a correct vote under secret voting. The points obtained during the experiment were converted to Euros at a rate of 1€ per 80 points and participants were paid the sum of their earnings over the 60 paid rounds at the end of the experiment. The average earning was €13.9, including a show-up fee of €2, with each session lasting for approximately 60 minutes.

7 Experimental Results

7.1 Committee Decisions

Table 4 reports the fraction of correct decisions observed under each treatment, alongside with the model predictions. The quality of the decisions is slightly higher under Low/Secret (85.56%) than Low/Public (84.31%), whereas the fraction of correct decisions under High/Secret (59.58%) is significantly lower than under High/Public (81.53%), as expected.[42]

7.2 Individual Choices

Table 5 summarizes the aggregate choices of uninformed subjects. When the magnitude of the bias is low, uninformed subjects are much more likely to abstain under secret (44.17%) than public voting (18.98%), while being significantly more likely to vote in accordance with their biases under public (64.81%) than secret voting (46.20%). On the other hand, when the magnitude of the bias is high, the vast majority of uninformed subjects vote in accordance with their biases under both secret (87.96%) and

[42] The χ^2 statistic for the difference between Low/Secret and Low/Public is 0.43, with $p = 0.50$, and the χ^2 statistic for the difference between High/Secret and High/Public is 83.4, with $p = 0.00$.

public voting (84.26%).[43] These results are all in line with our theoretical comparative static predictions. It should be noted that while 18.98% of subjects abstain under Low/Public, this number decreases substantially when we account for sequencing effects (see the Online Appendix C). We also observe between 3% and 16% of uninformed subjects voting *against* their biases depending on the treatment. Interestingly, the incentive to vote against the bias seems to be larger under public voting, which may be interpreted as evidence that some subjects do so in an attempt to guess the state of the world.[44]

In Table 6 we summarize the behavior of informed subjects who received a signal different than their biases and, therefore, face a trade-off between voting correctly and voting for their biases. Observe that, as predicted by the theory, when the magnitude of the bias is high, these subjects are much more inclined to vote correctly under public (84.60%) than secret voting (21.86%), while when the magnitude of the bias is low, the vast majority of them vote correctly under both secret (95.96%) and public voting (97.71%).[45] The percentage of subjects who vote correctly under High/Secret (21.86%) and the percentage of subjects who vote in accordance with their biases under High/Public (11.94%) are larger than expected. We note, however, that these proportions tend to decrease when we account for learning and sequencing effects.[46] We also observe that a fraction of informed subjects abstain under High/Secret (14.70%). This result is puzzling given that, in theory, abstaining is weakly dominated for members of this type. A possible explanation for this result could be attributed to the fact that both the common value (10 points) and the bias (14 points) are relatively close in magnitude in this case, which may lead some informed subjects to simply prefer to abstain.

Finally, we complement our analysis of individual choices by classifying subjects in accordance with their overall behavior during a session. Table 7 reports the distribution of the types of strategies used by uninformed subjects in Low/Secret and Low/Public treatments. Consistently with our theoretical predictions, the majority of subjects (44.44%) vote for their biases more than any other alternative in both treatments

[43] All these differences are significant at the 99% confidence level.

[44] This finding is consistent with experimental results previously obtained by Elbittar et al [7], who argue that a large proportion of uninformed subjects vote based on "hunches" (subjective beliefs). Similar findings are also in Guarnaschelli et al [18] and Bouton et al [2].

[45] The χ^2 statistic for the difference in correct votes when the bias is high is 434.0, with $p = 0.00$, and the χ^2 statistic for the difference in correct votes when the bias is small is 2.6, with $p = 0.11$.

[46] See the Online Appendix C for a detailed discussion.

and a substantial proportion of subjects (20.83%) mostly abstain under Low/Secret and vote for their biases under Low/Public. We also observe a considerable fraction of subjects (18.06%) abstaining more than any other choice in both Low/Secret and Low/Public treatments. Next, Table 8 reports the most frequent strategies adopted by informed subjects in High/Secret and High/Public treatments when they receive a signal different than their biases. As expected, we find that the vast majority of subjects (65.38%) mostly vote for their biases under High/Secret and vote for their signals under High/Public.[47]

While our results are consistent with the main comparative static predictions about the behavior of uninformed voters, still the fraction of subjects who change from voting to abstaining as a result of a change from public to secret voting is significantly below one. Given that there are multiple equilibria under Low/Secret, it would be interesting to investigate why uninformed voters do not select the efficient equilibrium, although, from a theoretical point of view, it is unclear which equilibrium players should coordinate on.[48] A potential explanation for the relatively low levels of abstention is that, while some subjects may have recognized the potential benefits of abstaining, they were discouraged from doing so by the fact that other uninformed subjects were not abstaining as well.[49] We explore this possibility in the Online Appendix C and provide empirical evidence that a negative feedback in one period (i.e. realizing that at least one other committee member "distorted" the decision by voting for the wrong alternative) significantly impacts the subsequent decisions to abstain in future periods and limits the convergence of voting behavior towards the efficient equilibrium.

7.3 Voting Profiles

We conclude our analysis by examining the frequency with which the observed voting profiles are exactly in accordance with one of the three classes of theoretical equilibria. In order to do so, we restrict the sample to include only decisions that involve at

[47] All the above results are robust to performing multivariate regressions including a rich set of controls (including subject performance in the comprehension quiz), individual fixed effects and standard errors clustered at the individual level. Furthermore clustering by session and adjusting the standard errors to account for the small number of clusters using a procedure proposed by Ibragimov and Müller [21] does not change any of our main results. See the Online Appendix C for a detailed discussion.

[48] Previous studies by Elbittar et. al [7], and Grosser and Seebauer [17] found, in a setting with common values, that a substantial proportion of subjects vote even though they have no information about the state of the world.

[49] Indeed, the optimal behavior for an uninformed subject is to vote in accordance with her bias if she believes that other uniformed subjects are voting in accordance with their biases.

least one uninformed subject and one informed subject who received a signal different than her bias. This restriction allows us to associate each voting profile with a single class of equilibria. As shown in Table 9, the proportion of voting profiles that are consistent with a fully-competent equilibrium decreases, as expected, from 33.23% under Low/Secret to 15.73% under Low/Public. This reduction is accompanied by a proportional increase in the profiles compatible with a partially-competent equilibrium from 35.00% under Low/Secret to 51.96% under Low/Public. Furthermore, the fraction of voting profiles consistent with a biased equilibrium drops significantly from 48.71% under High/Secret to 8.56% under High/Public. Again, this reduction is accompanied by an increase in the profiles compatible with a partially-competent equilibrium from 17.47% under High/Secret to 63.47% under High/Public.[50]

Finally, we find evidence (not reported in Table 8) that the percentage of voting profiles consistent with a fully-competent equilibrium under Low/Secret, a treatment in which there are multiple equilibria, increases substantially within the treatment. This result suggests that subjects were gradually learning to coordinate on the fully-competent equilibrium. Indeed, the percentage of voting profiles that are exactly in line with a fully-competent equilibrium increases from 27.11% in periods 1-10 to 29.31% in periods 11-20 to finally reach 44.33% in periods 21-30.

8 Concluding Remarks

We presented a new model of voting in committees where members are heterogeneous in competence and bias, they are career-concerned and can abstain. We identified a novel trade-off: transparency of individual votes attenuates the pre-existing biases of competent members and exacerbates the biases of incompetent members. Public voting leads to better decisions when the magnitude of the bias is large, while secret voting performs better otherwise. We presented new stylized facts about decision making in committees to illustrate the relevance of our theory in a number of real-life settings. Finally, we provided experimental evidence that is consistent with the main predictions of the model.

A number of papers in the literature have shown that transparency in decision-making is not always advisable since it creates incentives for agents to distort their behavior in order to convey information about their types. This has been investigated

[50] All these differences are significant at the 99% confidence level.

for single decision-makers and only recently − and partly following a trend towards increased procedural transparency in central banking − the literature has started focusing on the effects of the transparency of voting procedures on decision-making in committees. None of the existing papers has investigated how competence, individual biases and career-concerns interact in shaping individuals' voting behaviors in a committee, and how this interaction is affected by transparency.

Gersbach and Hahn [13] and Levy [22] examine models where agents care about acquiring a reputation for competence and show that secret voting may reduce distortions arising from signalling. In particular, Levy [22] identifies a tendency for "conformity" under secrecy in that committee members are more likely to vote for the alternative that is favored by the prior. This is not the case in our model. The combination of a common value component and the possibility of abstention lead to a different form of conformity: secret voting creates an incentive for incompetent members to abstain and it, therefore, attenuates their pre-existing biases. In this respect, our model uncovers an interaction between Levy's conformity effect and the swing voter's curse of Feddersen and Pesendorfer [10].[51]

Gersbach and Hahn [12] and Stasavage [38] analyze a setting where committee members may be misaligned with the interests of society, but also care about being perceived as "unbiased" to the extent that this enhances their reelection prospects. They show that transparency induces biased agents to act in accordance with the public interest. Conversely, in single decision-makers models, Ely and Välimäki [8], Morris [29] and, more recently, Shapiro [37] argue that transparency and career-concerns create an incentive for an unbiased agent to ignore her private information and choose the alternative that makes her look impartial.[52] Our model can help reconcile these seemingly opposing results: Transparency leads to better decisions when the biases are large, and secrecy leads to better decisions when the biases are small. Furthermore, our model does not assume that individual biases are per se punished.

In addition to these papers, Gersbach and Hahn [14] show that transparency induces agents to exert more effort in order to improve their chances of reappointment, Dal Bo [6] and Felgenhauer and Gruner [9] argue that public voting makes the committee more vulnerable to the influence of special interest groups, and Swank and Visser [39] show that career-concerns create an incentive for committees to conceal internal

[51] For models that focus only on the moral hazard aspect of secret voting see Seidmann [36] and references therein.

[52] For single decision-makers with career-concerns see also Maskin and Tirole [25] and Prat [35].

disagreements and show a united front in public.[53] Finally, Midjord et al [28] point out that career-concerns induce experts to be too conservative in order not to put their reputation at risk, and Gradwohl [16] shows that transparency leads to a trade-off between the accuracy of the decisions and the welfare of agents in a model where committee members have privacy concerns.

As for the experimental literature on committee decision making, a closely related paper to ours is Fehrler and Hughes [11]. As in our paper, they focus on the effect of transparency on committee decision making where agents are career-concerned. Differently from us, they concentrate on committees of two individuals, members are unbiased, and the experimental focus is on deliberation.[54] Also related is Battaglini et al [1] who provided the first test of the swing voters' curse in a laboratory setting.[55]

[53]See also Swank and Visser [40] for a model that combines public's demand for transparency, and committee members' opposition to it.

[54]See also Morton and Ou [30] for an empirical investigation of whether secret voting leads to less pro-social voting behavior than public voting.

[55]See also Morton and Tyran [31] and [32] for related experiments and Herrera et al [20] for a theory on strategic abstention in proportional elections.

References

[1] Battaglini, M., R. Morton and T. Palfrey. 2010. "The Swing Voter's Curse in the Laboratory." *Review of Economic Studies*, 77(1): 61-89.

[2] Bouton, L., M. Castanheira and A. Llorente-Saguer. 2016. "Divided Majority and Information Aggregation: Theory and Experiment." *Journal of Public Economics*, 134: 114-128.

[3] Cooper, J. and J. Golec. 2017. "Conflicts of Interest on Expert Committees: The Case of FDA Drug Advisory Committees." University of Connecticut School of Business Research Paper No. 17-02; George Mason Law & Economics Research Paper No. 18-10.

[4] Camara, F. and M. Kyle "Experts and Financial Ties: Evidence from the FDA Advisory Committees." Mimeo.

[5] Coughlan, P. 2000. "In Defense of Unanimous Jury Verdicts: Mistrials, Communication, and Strategic Voting." *American Political Science Review*, 94(2): 375-393.

[6] Dal Bo, E. 2007. "Bribing Voters." *American Journal of Political Science*, 51(4): 789-803.

[7] Elbittar, A., A. Gomberg, C. Martinelli and T. Palfrey. 2014. "Ignorance and Bias in Collective Decisions." Working Paper, California Institute of Technology.

[8] Ely, J. C. and J. Välimäki. 2003. "Bad Reputation." *Quarterly Journal of Economics*, 118(3): 785-814.

[9] Felgenhauer, M. and H. P. Grüner. 2008. "Committees and Special Interests." *Journal of Public Economic Theory*, 10(2): 219-43.

[10] Feddersen, T. and W. Pesendorfer. 1996. "The Swing Voter's Curse." *American Economic Review*, 86(3): 408-24.

[11] Fehler, S. and N. Hughes. 2018. "How Transparency Kills Information Aggregation: Theory and Experiments." *American Economic Journal: Microeconomics*, 10(1): 181-209.

[12] Gersbach, H. and V. Hahn. 2004. "Voting Transparency, Conflicting Interests, and the Appointment of Central Bankers." *Economics and Politics*, 16(3): 321-45.

[13] Gersbach, H. and V. Hahn. 2008. "Should the Individual Voting Records of Central Bankers be Published?" *Social Choice and Welfare*, 30: 655-83.

[14] Gersbach, H. and V. Hahn. 2012. "Information Acquisition and Transparency in Committees." *International Journal of Game Theory*, 41(2): 427-453.

[15] Gerling, K., H. P. Grüner, A. Kiel and E. Schulte. 2005. "Information Acquisition and Decision Making in Committees: A Survey." *European Journal of Political Economy*, 21: 563-97.

[16] Gradwohl, R. 2018. "Voting in the Limelight." *Economic Theory*, 66(1): 65-103.

[17] Grosser, J. and M. Seebauer. 2015. "The Curse of Uninformed Voting: An Experimental Study." Working Paper.

[18] Guarnaschelli, S., R. McKelvey and T. Palfrey. 2000. "An Experimental Study of Jury Decision Rules." *The American Political Science Review*, 94(2): 407-423.

[19] Hansen, S., M. McMahon and A. Prat. 2018. "Transparency and Deliberation within the FOMC: A Computational Linguistic Approach." *Quarterly Journal of Economics*, 133(2): 801-870.

[20] Herrera, H., J. McMurray and A. Llorente-Saguer. 2016. "The Marginal Voter's Curse." Working Paper.

[21] Ibragimov, R. and U. Müller. 2010. "t-Statistic Based Correlation and Heterogeneity Robust Inference." *Journal of Business and Economic Statistics*, 28: 453-468.

[22] Levy, G. 2007. "Decision Making in Committees: Transparency, Reputation and Voting Rules." *American Economic Review*, 97(1): 150-68.

[23] Li, H. and W. Suen. 2009. "Decision-making in Committees." *Canadian Journal of Economics*, 42(2): 359-92.

[24] Lurie P., C. Almeida, N. Stine, A. Stine, and S. Wolfe. 2006. "Financial Conflict of Interest Disclosure and Voting Patterns at Food and Drug Administration Drug Advisory Committee Meetings." *JAMA*, 295(16):1921?1928.

[25] Maskin, E. and J. Tirole. 2004. “The Politician and the Judge.” *American Economic Review*, 94(4): 1034-54.

[26] Meade, E. and D. Stasavage. 2008. “Publicity of Debate and the Incentive to Dissent: Evidence from the US Federal Reserve.” *Economic Journal*, 118: 695-717.

[27] Meloso, D. and M. Ottaviani. 2015. "The Mechanics of Reputational Cheap Talk: An Experiment with Crystal Balls." Mimeo.

[28] Midjord, R., T. Rodriguez-Barraquer and J. Valasek. 2016. “Over-Caution of Large Committees of Experts." Mimeo.

[29] Morris, S. 2001. “Political Correctness.” *Journal of Political Economy*, 109(2): 231-65.

[30] Morton, R. and K. Ou. 2015. “Public Voting and Prosocial Behavior.” Working Paper.

[31] Morton, R. and J. Tyran. 2011. “Let the Experts Decide? Asymmetric Information, Abstention, and Coordination in Standing Committees.” *Games and Economic Behavior*, 72(2): 485-509.

[32] Morton, R. and J. Tyran. 2015. “Corruption in Committees: An Experimental Study of Information Aggregation through Voting.” *Journal of Public Economic Theory*, 17(4): 553-579.

[33] Nakaguma, M. Y. 2012. Essays on Political Institutions and Constitutional Design. Ph.D. Thesis, Columbia University.

[34] Newham, M. and R. Midjord. 2018 “Herd Behavior in FDA Committees: A Structural Approach.” DIW Berlin Discussion Paper No. 1744.

[35] Prat, A. 2005. “The Wrong Kind of Transparency.” *American Economic Review*, 95(3): 862-77.

[36] Seidmann, D. 2011. “A Theory of Voting Patterns and Performance in Private and Public Committees.” *Social Chioce and Welfare*, 36(1): 49-74.

[37] Shapiro, J. M. 2016. "Special Interests and the Media: Theory and an Application to Climate Change." *Journal of Public Economics*, 144: 91-108.

[38] Stasavage, D. 2004. "Public versus Private Deliberation in a Responsive Democracy." Working Paper.

[39] Swank, O. H. and B. Visser. 2007. "On Committees of Experts." *Quarterly Journal of Economics*, 122(1): 337-72.

[40] Swank, O. H. and B. Visser. 2013. "Is Transparency To No Avail?." *Scandinavian Journal of Economics*, 115(4): 967-94.

[41] Swank, J., O. H. Swank and B. Visser. 2008. "How Committees of Experts Interact with the Outside World: Some Theory and Evidence from the FOMC." *Journal of the European Economic Association*, 6(2): 478-86.

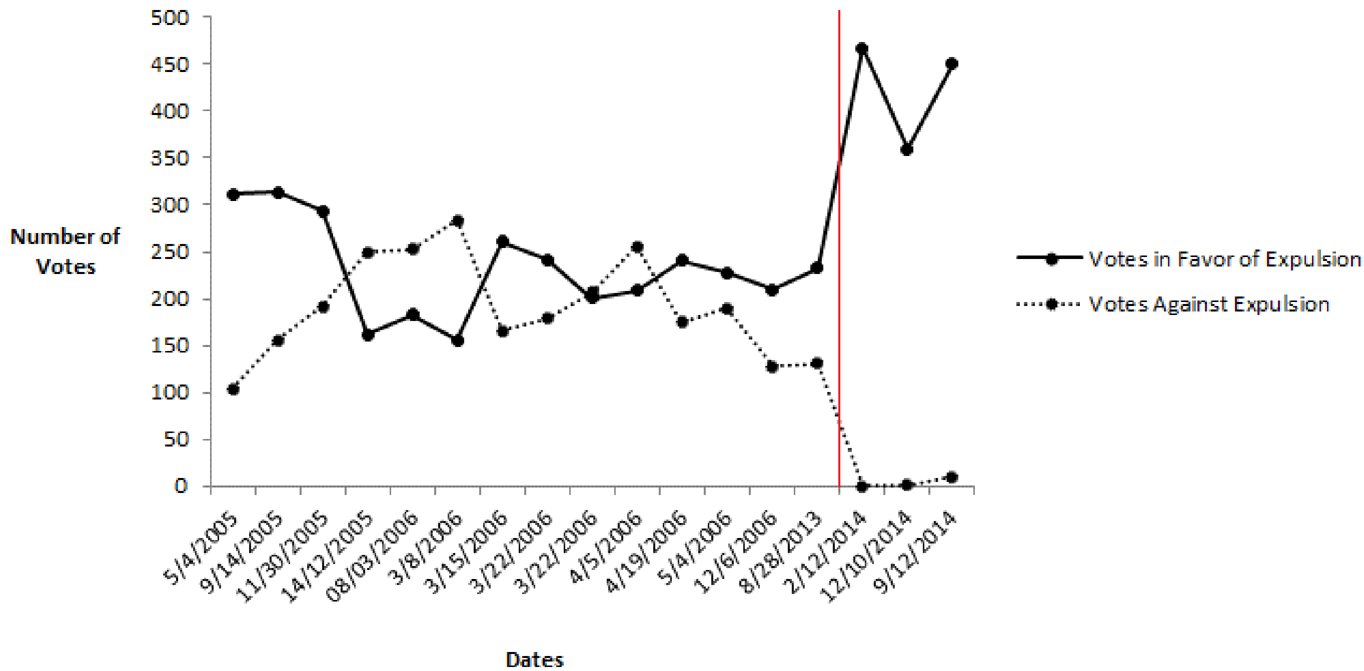

Notes. This figure plots the number of votes in favor and against an expulsion in seventeen expulsion votes that occurred in the Brazilian House of Representatives during the period 2005-2018. The vertical bar represents the date of approval of a constitutional amendment that changed the vote from secret to public.

Figure 1. Secret versus Public Vote: Brazil's House of Representatives

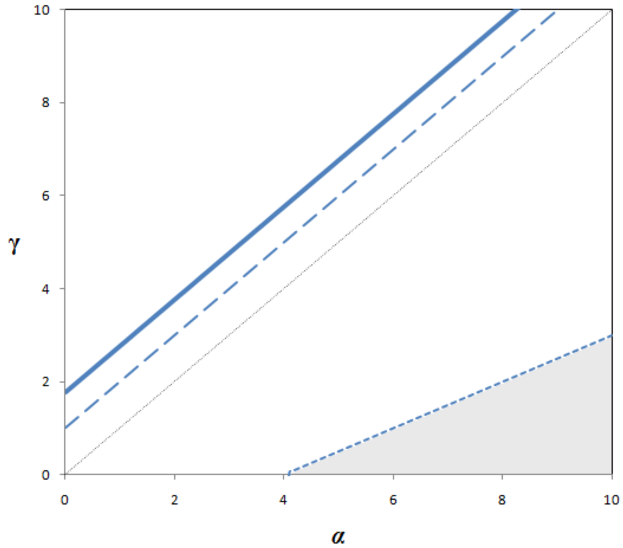

Panel a. Fully Competent Equilibrium

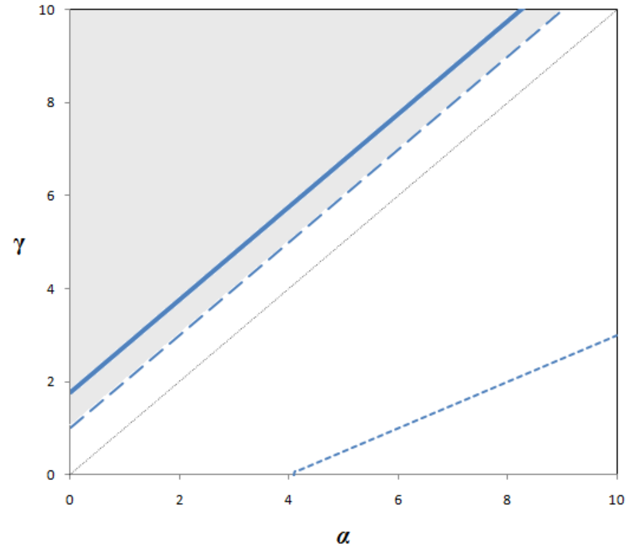

Panel c. Biased Equilibrium

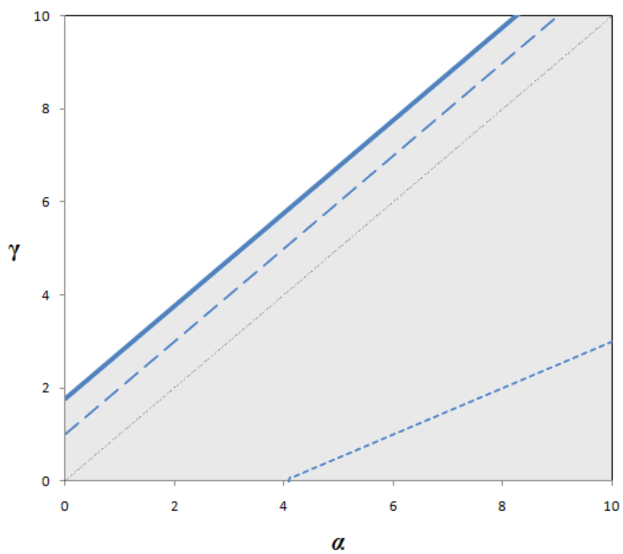

Panel b. Partially Competent Equilibrium

Notes. This figure illustrates the region of the parameters where each class of equilibrium can be sustained under a given voting rule λ, as derived in Proposition 5. The structure of the equilibria looks similar under secret and public voting, although the exact regions where each class of equilibrium can be sustained differ. *Panel a* represents, shaded in grey, the region of the parameters where a fully-competent equilibrium can be sustained. P*anel b* represents, shaded in grey, the region of the parameters where a partially-competent equilibrium can be sustained. Finally, *panel c* represents, shaded in grey, the region where a biased equilibrium exists. Observe that the shaded areas may overlap in some regions, representing the existence of multiple equilibria. The 45-degree line is depicted as a small dotted line. The parameter values assumed for the construction of this graph were: $p=0.5$, $q=0.5$, $n=3$, $\phi=1$ and $\sigma=0.5$.

Figure 2. Equilibria: The Symmetric Case

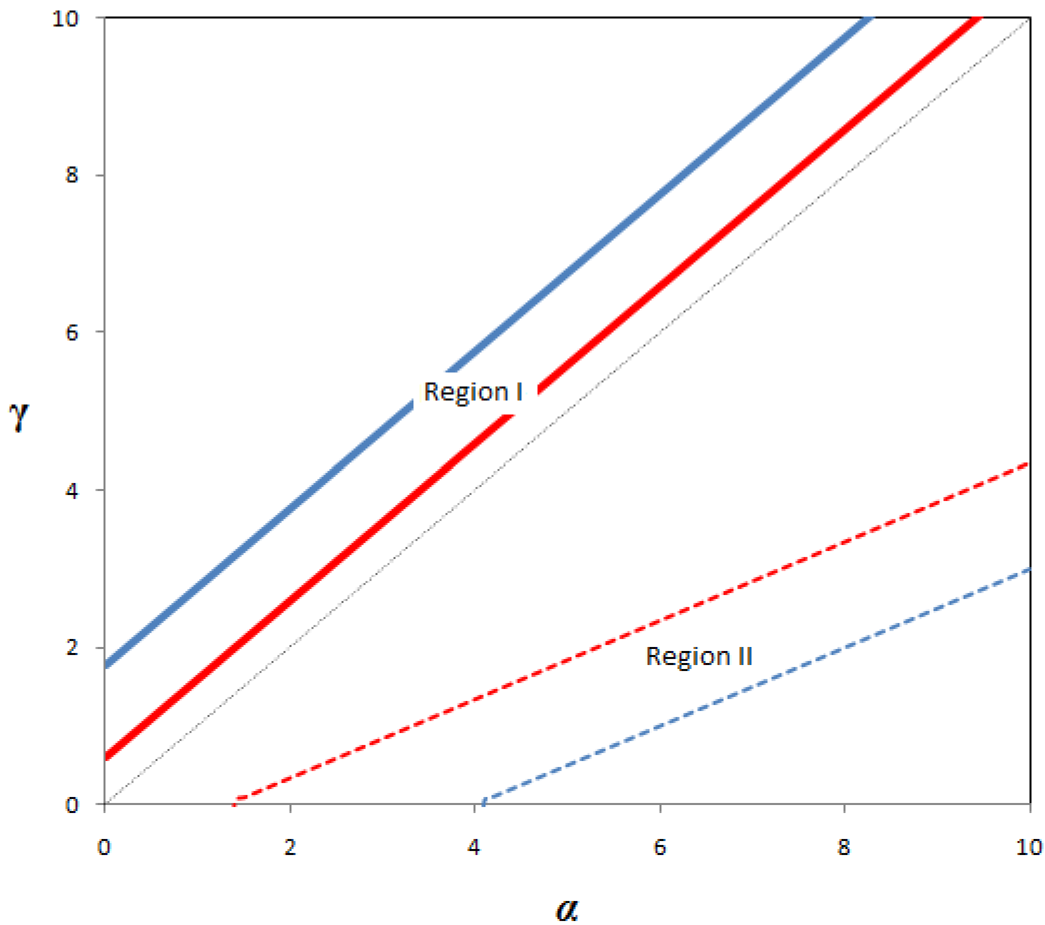

Notes. This figure provides a comparison of the parameter regions where a fully-competent and a partially-competent equilibrium can be supported under each voting rule. The relevant thresholds for the public and secret voting rules are depicted in blue and red, respectively. Region I represents the region of parameters where a partially-competent equilibrium can be sustained under public but not under secret voting, while region II represents the region of parameters where a fully-competent equilibrium can be sustained under secret but not under public voting. The 45-degree line is depicted as a small dotted line. The parameter values assumed for the construction of this graph were: p=0.5, q=0.5, n=3, ϕ=1 and σ=0.5.

Figure 3. Comparative Static Results

Panel A. Abstentions

	Obs	Mean	SD
Individual Votes: Abstention	9,675	0.024	0.153
Committee Votes: One or More Abstentions	654	0.2522	0.434

Panel B. Abstentions by Academic Backgroud

	Means			FE Regressions
Academic Background	Yes	No	Diff.	Diff.
	(1)	(2)	(3)	(4)
MD	0.021	0.029	-0.0079**	-0.0129***
	N = 6,066	N = 3,609	(0.0032)	(0.0033)
PhD	0.030	0.021	0.0086**	0.0143***
	N = 2671	N = 7004	(0.0034)	(0.0046)
Statistician	0.033	0.022	0.0111**	0.0127**
	N = 1212	N = 8463	(0.0047)	(0.0060)
PhD without MD	0.037	0.021	0.0165***	0.0215***
	N = 1782	N = 7893	(0.0040)	(0.0055)
Statistician without MD	0.038	0,022	0,0155***	0.018***
	N = 972	N = 8,703	(0.0051)	(0.0063)

Notes. Panel A reports summary statistics for abstentions at the individual and committee levels for votes held in the FDA's Advisory Committees. Panel B reports abstention rates for members with different academic backgrounds, namely: members with MD, PhDs, statisticians, members with PhD but without MD and statisticians without MD. Columns (1) and (2) report abstention rates for members who, respectively, belong and not belong to each of these groups, while Column (3) reports the difference in abstention rates, with standard errors in parentheses. Finally, Column (4) reports estimates for the differences in abstention rates between groups controlling for 246 meeting fixed effects, with standard errors clustered at the meeting level. ***, ** and * denote significance at 1%, 5% and 10% respectively.

Table 1. FDA's Advisory Committees: Abstention Rates

Treatment	alpha	gamma	career-concerns	Predicted Equilibria
Low/Secret	10	1	0	Fully-Competent and Partially-Competent
Low/Public	10	1	9	Partially-Competent
High/Secret	10	14	0	Biased
High/Public	10	14	9	Partially-Competent

Notes. This table presents the choice of parameters' values and the equilibrium predictions associated with each treatment. A fully-competent equilibrium is such that all informed subjects vote in accordance with the signal and all uninformed subjects abstain, a partially-competent equilibrium is such that all informed subjects vote in accordance with the signal and all uninformed subjects vote for their biases, and a biased equilibrium is such that all subjects vote for their biases.

Table 2. Treatments

Session	Sequence	Subjects
1	Low/Secret – Low/Public	24
2	Low/Secret – Low/Public	30
3	Low/Public – Low/Secret	18
4	High/Secret – High/Public	24
5	High/Secret – High/Public	24
6	High/Public – High/Secret	24

Notes. This table summarizes the sequence of treatments and number of participants in each experimental session.

Table 3. Sequence of Treatments

Treatment	Obs	Correct Decisions (%)	Predicted (%)
Low/Secret	720	85.56	93.00 / 84.00
Low/Public	720	84.31	84.00
High/Secret	720	59.58	50.00
High/Public	720	81.53	84.00

Notes. This table presents the proportion of correct decisions observed under each treatment, alongside with the predictions of the theoretical model.

Table 4. Decisions

		Uninformed Voters		
Treatment	Obs	Abstention (%)	Bias (%)	Against-Bias (%)
Low/Secret	1080	44.17	46.20	9.63
Low/Public	1080	18.98	64.81	16.20
High/Secret	1080	9.35	87.96	2.69
High/Public	1080	5.83	84.26	9.91

Notes. This table reports the aggregate choices of uninformed subjects by treatment. The model predicts that all uninformed subjects vote for their biases in treatments Low/Public, High/Secret and High/Public. The Low/Secret treatment is consistent with both a fully-competent equilibrium, where all uninformed subjects abstain, and a partially-competent equilibrium, where all uninformed subjects vote for their biases.

Table 5. Individual Choices: Uninformed Subjects

		Informed Voters with Signal ≠ Bias		
Treatment	Obs	Signal (%)	Bias (%)	Abstention (%)
Low/Secret	520	95.96	1.54	2.50
Low/Public	524	97.71	2.29	0.00
High/Secret	517	21.86	63.44	14.70
High/Public	578	84.60	11.94	3.46

Notes. This table reports the aggregate choices of informed subjects who received a signal different than their biases by treatment. The model predicts that all of these subjects vote in accordance with the signal in treatments Low/Secret, Low/Public and High/Public and that all of them vote for their biases in treatment High/Secret.

Table 6. Individual Choices: Informed Subjects

			Low/Secret	
		Abstainers (%)	Bias-Followers (%)	Against-Bias (%)
	Abstainers (%)	18.06	0.00	0.00
Low/Public	Bias-Followers (%)	20.83	44.44	1.39
	Against-Bias (%)	6.94	4.17	4.17

Notes. This table reports the distribution of subjects who can be classified as "abstainers", "bias-followers" and "against-bias" when uninformed in treatments Low/Secret and Low/Public. A subject is considered to be an abstainer if she abstains more than any other alternative, a bias-follower if she votes for the bias more than any other alternative and an against-bias if she votes against the bias more than any other alternative. The model predicts that uninformed subjects vote for their biases under Low/Public and is consistent with both them abstaining or voting for their biases under Low/Secret.

Table 7. Strategies: Uninformed Subjects

		High/Secret		
		Signal-Followers (%)	Bias-Followers (%)	Abstainers (%)
	Signal-Followers (%)	19.44	65.28	12.50
High/Public	Bias-Followers (%)	0.00	0.00	1.39
	Abstainers (%)	0.00	0.00	1.39

Notes. This table reports the distribution of subjects who can be classified as "signal-followers", "bias-followers" and "against bias-followers" when informed and the signal is different than their biases in treatments High/Secret and High/Public. A subject is considered to be a signal-follower if she votes for the signal more than any other alternative, a bias-follower if she votes for the bias more than any other alternative and an abstainer if she abstains more than any other alternative. The model predicts that informed subjects vote for the signal under High/Public and for their biases under High/Secret.

Table 8. Strategies: Informed Subjects

Treatment	Obs	Fully Competent (%)	Partially Competent (%)	Biased (%)	Other (%)
Low/Secret	340	33.23	35.00	0.00	31.77
Low/Public	356	15.73	51.96	0.00	32.31
High/Secret	349	0.00	17.47	48.71	33.82
High/Public	397	2.77	63.47	8.56	25.20

Notes. This table reports the proportion of voting profiles that are consistent with a fully-competent equilibrium, a partially-competent equilibrium, a biased equilibrium or none by treatment. The sample is restricted to include only decisions that involved at least one uninformed subject and one informed subject who received a signal different than her bias. The model predicts either a fully-competent or a partially-competent equilibrium under Low/Secret, a partially-competent equilibrium under both Low/Public and High/Public and a biased equilibrium under High/Secret.

Table 9. Voting Profiles

Online Appendix: Public versus Secret Voting in Committees

(NOT FOR PUBLICATION)

This Online Appendix is organized as follows: Section A discusses a number of extensions to our benchmark model, Section B provides a detailed analysis of the versions of the model with deliberation and repeated interaction, Section C presents additional experimental results omitted from the main text, Section D collects the proofs of the propositions of the paper, Section E presents the derivation of the version of the model tested in the lab and, finally, Section F presents the English version of the experiment instructions.

Appendix A. Discussion and Extensions

This section discusses a number of assumptions which we have made throughout our main analysis as well as some possible extensions to the basic model.

A.1 Precision of Signals

Following Feddersen and Pesendorfer [4] and Battaglini et al [1], our analysis assumed that competent members receive perfectly informative signals about the state of the world, while incompetent members received no information at all. Although our main results rely on the hypothesis that the precision of signals received by competent and incompetent model be sufficiently different, the extreme assumption of perfectly informative and non-informative signals is not essential. Formally, our main comparative static results concerning the impact of transparency on the behavior of committee members would still hold in an environment where competent members received signals with precision $\Pr(s=\omega|\omega)=1-\varepsilon$, while incompetent members received signals with precision $\Pr(s=\omega|\omega)=\frac{1}{2}+\delta$, for $\omega\in\{A,B\}$, $\varepsilon>0$ and $\delta>0$, provided that ε and δ are relatively small. In particular, the set of possible equilibria would still consist of the same three classes of equilibria described in Proposition 1, although the precise conditions for the existence of each class would have to be adjusted in order to take into account the fact that competent members may now vote for the incorrect alternative even if they follow their signals.

A.2 Unbiased Members

Although our basic model assumes that all committee members are biased towards one of the alternatives, the main qualitative results of our analysis are robust to allowing for the existence of unbiased members. In fact, note that unbiased competent members would always have an incentive to follow their signals, since they care only about the common value and the career-concern rewards associated with a correct vote, while unbiased incompetent members would always be more willing to abstain relatively to biased incompetent members. Next, with these observations in mind, it would be interesting to consider what would happen if we allowed for the existence of correlation between competence and bias. Suppose, for instance, that competent members were ideologically more neutral and consider, in particular, the extreme case where all competent members were unbiased, while incompetent members could be either biased or unbiased. Observe that in this case competent members would always have an incentive to vote for the correct alternative, so that the degree of transparency would have no impact on their behavior. For incompetent members, on the other hand, public voting would always make them more willing to vote, so that we should expect secret voting to lead to better decisions in general. Conversely, if competent members were either biased or unbiased and all incompetent members were unbiased, then none of our main comparative static results would change. Observe that unbiased incompetent members would still have an incentive to vote due to career-concerns, although they would not have a "preferred" alternative in this case. Therefore, the basic trade-off between public and secret voting would remain unchanged.

A.3 Ex-Post Observability of the State of the World

An important assumption in our model is that the external evaluator always observes the state of the world ex-post. This feature guarantees that, under transparency, voting for the correct alternative is always associated with strictly positive career-concern rewards, whereas an incorrect vote is not rewarded in equilibrium. Note that if the evaluator did not observe the state of the world, then the role played by career-concerns in providing incentives for members to vote correctly would be weakened. In particular, as emphasized by Canes-Wrone et al [2], the desire to acquire reputation could create an incentive for committee members to ignore whatever information they might have about the state of the world and simply vote for the alternative which the evaluator believes is more likely to be correct.[1] Furthermore, as in Swank and Visser [11], there would be an incentive for the committee members to show "internal agreement", since competent members always receive the same signal. The incentive to pander to the evaluator's opinion makes transparency in committees less appealing in general, a result also emphasized by Stasavage [10].

A.4 Voting Rule and Degree of Transparency

Throughout the analysis, we assumed that the main difference between public and secret voting is that while all individual votes are observed under public voting, only the vote tally is revealed under secret voting. Note that, in this case, neither the final decision of the committee nor the size of the majority required for an alternative to be chosen has any impact on the external evaluator's posterior beliefs, given that the observation of the aggregate vote alone provides more information about the behavior of members than knowledge about the committee's decision and voting rule. As a consequence, a change in the size of the majority required for an alternative to be approved would have no significant impact on our main qualitative results. If, on the other hand, we had assumed, as Levy [7] and Swank and Visser [11], that only the final decision of the committee was observed

[1] See Morris [9] and Maskin and Tirole [8] for studies that also emphasize the importance of pandering in principal-agent models.

under secrecy, then the voting rule would have played a more important role in determining how much information is conveyed to the evaluator. Nonetheless, our basic comparative static results would remain unchanged, since the dilution of career-concern rewards, the key mechanism behind our results, would still exist under secret voting.

A.5 Career-Concern Rewards

We assumed in our basic model that the career-concern rewards of a committee member are given by the conditional probability that the member is competent *and* voted correctly (see equation (1)). However, our main qualitative results would remain the same even if we allowed for these rewards to be based only on the posterior probability that the member is competent, $r_i^{\omega,\lambda} \equiv \Pr(\tau_i = \mathsf{c}|\omega, \mathcal{I}^\lambda)$. In particular, both the fully-competent and partially-competent equilibria would still be characterized by Propositions 2 and 3, respectively, and all comparative static results concerning these two classes of equilibria would remain unchanged. The intuition is that in both of these cases the career-concern rewards associated with a correct vote are strictly larger than those associated with an abstention or an incorrect vote, since all competent members vote correctly in equilibrium.[2] It is in this sense that we can say that our basic conclusion that transparency attenuates the biases of competent members while exacerbating the biases of incompetent members is robust to how career-concerns are defined.

The main implication of relaxing the assumption that career-concern rewards are attributed only in connection with a correct vote is that it is now possible to sustain a larger set of equilibria than those described in Proposition 1. In particular, there may also exist equilibria featuring: (*i*) competent members with bias equal to the state of the world voting against the state and (*ii*) competent members abstaining. Note, however, that any equilibrium involving one of these behaviors require a very particular structure of incentives, namely: a member who abstains or votes incorrectly must be seen as more likely to be competent than a member who votes correctly. There is an aspect of self-fulfilling prophecy involved in such equilibria in that whatever the external evaluator expects competent members to do may end up occurring, provided that the weight attached to career-concerns is sufficiently large. We believe that this element is not likely to be crucial in most applications of our model and this is one reason why we view our initial assumption that career-concerns are related to the joint probability that the member is competent and voted correctly as a reasonable form of refinement.

In any case, it is possible to show that the equilibria discussed above can only exist in a very specific region of the parameter space, where the sum of the common value and the bias term is small. Note that, for a competent member with bias equal to the state of the world, voting correctly would increase the likelihood that the member gets a payoff of $\alpha + \gamma$. Therefore, for such member to have an incentive to either abstain or vote against the state of the world, both the common value and the bias term must be sufficiently small.

[2] This result follows directly from Bayes' rule, since if all competent members vote correctly in equilibrium then it must be the case that $\Pr(t_i = \mathsf{c}|v_i \neq \omega) = 0$.

A.6 Institutional Preferences

Which level of transparency would the members of the committee prefer if, prior to voting, they could choose between public and secret voting? The choice between public and secret voting affects the payoffs of members both in terms of how the career-concern rewards are distributed across members and in terms of the likelihood that the correct decision is taken. Observe that, overall, due to the dilution effect, competent members are more likely to prefer public voting, whereas incompetent members are more likely to prefer secret voting. There are, however, some exceptions to this general point. First, if the weight associated to career-concerns is small and the common value is high relatively to the bias, then competent members may actually prefer secret voting, since secrecy is more likely to lead to better decisions in this case. Furthermore, whenever a biased equilibrium is expected to prevail regardless of the voting rule, then competent members biased against the state actually prefer secret voting, since they always get zero career-concern rewards when voting is public in this case.

Overall, our discussion highlights the fact that the choice of voting rule may be endogenous to the composition of the committee as well as to the types of decisions being taken, a result that has important implications for the empirical evaluation of the impact of transparency on voting outcomes and individual behavior using non-experimental data.

Appendix B. Deliberation and Dynamics

This section presents a detailed discussion about the versions of the model with deliberation and with dynamics.

B.1 Model with Deliberation

We solve the model described in the main text for symmetric pure-strategy equilibria. Conditional on the member's competence and bias, a strategy is now characterized by a message, m_i, and a voting function, $v_i(m)$, which specifies a vote as a function of the message vector m.

A few preliminary remarks are in order before we proceed to the characterization of the equilibria. First, note that, at the deliberation stage, the incompetent members have no option but to keep silent, so that it only remains to determine the conditions under which the competent members are willing to share their private information. Second, if the state of the world is revealed at the deliberation stage, then all members become informed, rendering irrelevant the distinction between competent and incompetent types. Third, if all members are informed about the state of the world, then by Lemma 1 we can restrict attention to two types of outcomes at the voting stage: either (i) all vote for the correct alternative, or (ii) all vote for their biases.[3]

[3] Note that Lemma 1 implies that: (i) abstaining is weakly dominated for any informed member and (ii) voting against the state of the world is weakly dominated for any informed member with bias equal to the state.

Responsive Equilibria. We begin by characterizing the set of responsive equilibria, where, on the equilibrium path, all committee members vote correctly when the state of the world is revealed. Recall that in any responsive equilibrium revelation by a single member is sufficient to induce an unanimous correct decision. Therefore, in equilibrium the external evaluator knows that an unanimous correct decision means that with high likelihood there is at least one competent member in the committee.

Formally, given a system of beliefs about the behavior of committee members and conditional on the state of the world as well as on all information disclosed about the outcome of the vote under voting rule λ, the external evaluator computes the posterior probability that the state of the world was revealed (R) or not revealed (NR) at the deliberation stage, $\Pr(R|\omega, \mathcal{I}^{\lambda})$ and $\Pr(NR|\omega, \mathcal{I}^{\lambda})$. Then, conditional on each of these two events, the external evaluator computes the posterior probability that a member is competent and voted correctly. Thus, the carreer-concerns rewards are given by:

$$r_i^{\omega,\lambda} = \sum_{k \in \{R, NR\}} \Pr(k|\omega, \mathcal{I}^{\lambda}) \Pr(\tau_i = \mathrm{c}, v_i = \omega | \omega, \mathcal{I}^{\lambda}, k)$$

The conditions for the existence of a particular class of equilibria now involve constraints on the equilibrium behavior of members both at the deliberation stage and the voting stage. In what follows, we provide a detailed discussion about the construction of each type of equilibrium as well as their main properties.

In a *full revelation equilibrium*, all competent members reveal their signals at the deliberation stage. Then, at the voting stage, if the state of the world was revealed, all members vote correctly; otherwise, if the state was not revealed (which can only happen on the equilibrium path if all members are incompetent), everyone votes for their biases. In this case, career-concern rewards are positive only if the committee's decision is unanimously correct, since a non-unanimous decision perfectly reveals that there are no competent members in the committee.

In particular, under a full revelation (fr) equilibrium, the career-concern reward associated with an unanimous correct decision is given by:

$$r_{fr}^{un} = \frac{1 - (1-\sigma)^n}{(1 - (1-\sigma)^n) + (1-\sigma)^n \left(\frac{1}{2}\right)^n} \left(\sigma + \frac{1-\sigma}{n}\right),$$

where the first ratio captures the probability that the state was revealed at the deliberation stage conditional on an unanimous correct decision, while the second ratio represents the conditional probability that a member is competent, given that the state was revealed at the deliberation stage.[4] Note that, in a full revelation equilibrium, a member is competent with zero probability conditional on the state of the world not being revealed at the deliberation stage.

Since, in this case, information about the unanimity of a correct decision is all that matters for career-concerns, the structure of incentives under public and secret voting are exactly the

[4] Intuitively, this second ratio captures the conditional probability that an agent is competent given that there is at least one competent member in the committe, which can be written as $\frac{1+(n-1)\sigma}{n} = \sigma + \frac{1-\sigma}{n}$.

same. At the deliberation stage, competent members with bias equal to the state of the world are always willing to reveal their private information, so that the relevant incentive constraint is that on competent members biased against the state of the world. Furthermore, if information about the state is revealed during the deliberation process, then, at the voting stage, no committee member would ever want to vote against it: if all other members are voting correctly, then nobody is pivotal and a non-unanimous decision is always associated with zero career-concerns rewards. Conversely, if information about the state is not revealed on the equilibrium path, then it becomes common knowledge inside the committee that all members are incompetent, so that it is indeed an equilibrium for everyone to vote for their biases.

The following proposition provides a general characterization of the conditions under which a full revelation (fr) equilibrium exists.

Proposition B.1. *A full revelation equilibrium can be supported if and only if:*

$$\gamma \leq \overline{\gamma}_{fr}^{\lambda}\left(\alpha, \phi, \sigma, n\right),$$

where $\overline{\gamma}_{fr}^{\lambda}\left(\alpha, \phi, \sigma, n\right) > \alpha$*. Furthermore, the condition for the existence of a fullrevelation equilibrium is the same under public and secret voting,* $\overline{\gamma}_{fr}^{\mathsf{s}}\left(\alpha, \phi, \sigma, n\right) = \overline{\gamma}_{fr}^{\mathsf{p}}\left(\alpha, \phi, \sigma, n\right).$

The proof is straightforward and follows directly from the incentive compatibility constraint on competent members biased against the state. In equilibrium, a competent member biased against the state of the world is willing to share information only if the magnitude of his bias is not too large. Furthermore, it can be shown that the threshold $\overline{\gamma}_{fr}^{\lambda}$ is increasing in ϕ, meaning that a full revelation equilibrium is more likely to exist the larger the importance attached to career-concerns.

In a *partial revelation equilibrium*, not all competent members share their private information, yet everyone votes correctly when the state is revealed. A few general remarks about this class of equilibrium are in order before we proceed. First, note that the members who might want to withhold information in equilibrium are those biased against the state of the world. Second, since both the prior and the distribution of biases are symmetric, no additional information about the state of the world can be inferred when the state is not revealed at the deliberation stage, even considering the fact that signals are disclosed strategically. Third, conditional on the state of the world not being revealed, committee members may vote in accordance with either: (i) a fully-competent equilibrium, (ii) a partially-competent equilibrium, or (iii) a biased equilibrium.

We begin our analysis by observing that a partial revelation equilibrium with "fully-competent voting" cannot exist. The reason is that in such equilibrium a competent member with bias equal to the state of the world would also want to withhold information and then vote correctly in an attempt to distinguish herself from the incompetent members. Note that, by doing so, a competent member would still be able to guarantee that the correct decision is taken, while also obtaining higher career-concern rewards in expectation.

We, therefore, focus our analysis on the characterization of a partial revelation equilibrium with "partially-competent voting", which is the more interesting case for our study.[5] Observe that the relevant incentive constraint in this case is that on competent members biased against the state of the world. Specifically, at the deliberation stage, members of this type must prefer not to reveal information, while at the voting stage, conditional on the state not being revealed, they must prefer to vote correctly. The next proposition provides a general characterization of the conditions for a partial revelation equilibrium with partially-competent voting (pr).

Proposition B.2. *A partial revelation equilibrium with partially-competent voting can be supported if and only if:*

$$\underline{\gamma}_{pr}^{\lambda}\left(\alpha,\phi,\sigma,n\right)\leq\gamma\leq\overline{\gamma}_{pr}^{\lambda}\left(\alpha,\phi,\sigma,n\right),$$

where $\underline{\gamma}_{pr}^{\lambda}\left(\alpha,\phi,\sigma,n\right)>\alpha$ *and* $\underline{\gamma}_{pr}^{\lambda}\left(\alpha,\phi,\sigma,n\right)>\alpha$*. Furthermore, we have* $\underline{\gamma}_{pr}^{\mathsf{s}}\left(\alpha,\phi,\sigma,n\right)>\underline{\gamma}_{pr}^{\mathsf{p}}\left(\alpha,\phi,\sigma,n\right)$ and $\overline{\gamma}_{pr}^{\mathsf{s}}\left(\alpha,\phi,\sigma,n\right)<\overline{\gamma}_{pr}^{\mathsf{p}}\left(\alpha,\phi,\sigma,n\right)$.

Proof. In a partial revelation equilibrium with partially-competent voting, the career-concern reward associated with an unanimous correct decision is given by:

$$r_{pr-part}^{un}=\chi_{pr-part}\left(\sigma+\frac{1-\sigma}{n}\right)+\left(1-\chi_{pr-part}\right)\varrho_{pr-part}$$

where $\chi_{pr-part}$ is the probability that the state was revealed at the deliberation stage conditional on an unanimous correct decision and $\varrho_{pr-part}$ is the probability that a member is competent given that the state was not revealed and conditional on an unanimous correct decision. In equilibrium, we have:

$$\chi_{pr-part}=\frac{1-\left(1-\frac{1}{2}\sigma\right)^{n}}{\left(1-\left(1-\frac{1}{2}\sigma\right)^{n}\right)+\left(1-\frac{1}{2}\sigma\right)^{n}\left(\frac{1}{2-\sigma}\right)^{n}}$$

and

$$\varrho_{pr-part}=\frac{\frac{1}{2}\sigma}{\frac{1}{2}\sigma+\frac{1}{2}\left(1-\sigma\right)}=\sigma$$

Note that probability that a member is competent conditional on a correct vote and on a non-unanimous decision is $r_{pr-part}^{*}=\varrho_{pr-part}=\sigma$.

Suppose that the state of the world was not revealed at the deliberation stage. Without loss of generality, assume that the state is A. At the voting stage, a competent member biased towards B must prefer to vote correctly rather than to vote for his bias, that is:

$$\begin{aligned}&\phi\widetilde{r}_{pr-part}^{\omega=A,\lambda}\left(A\right)+\rho_{pr-part}^{\omega=A}\left(A\right)\alpha+\left(1-\rho_{pr-part}^{\omega=A}\left(A\right)\right)\gamma\\&\qquad\geq\phi\widetilde{r}_{pr-part}^{\omega=A,\lambda}\left(B\right)+\rho_{pr-part}^{\omega=A}\left(B\right)\alpha+\left(1-\rho_{pr-part}^{\omega=A}\left(B\right)\right)\gamma,\end{aligned}$$

where $\widetilde{r}_{pr-part}^{\omega=A,\lambda}\left(v_{i}\right)$ is the expected career-concern reward of member i when he chooses v_i at the

[5] The analysis of a partial revelation equilibrium with "biased voting" is available upon request. This type of equilibrium has some qualitative features that make it very similar to an irrelevant deliberation equilibrium, which we analyze later in detail.

voting stage, conditional on the state not being revealed before. Remember that $\rho^{\omega=A,\lambda}_{pr-part}(v_i)$ represents the probability that the committee's decision is A when the member chooses v_i. Re-arranging, we get:

$$\gamma \leq \alpha + \frac{\phi(\widetilde{r}^{\omega=A,\lambda}_{pr-part}(A) - \widetilde{r}^{\omega=A,\lambda}_{pr-part}(B))}{\rho^{\omega=A}_{pr-part}(A) - \rho^{\omega=A}_{pr-part}(B)}$$

Furthermore, note that $\widetilde{r}^{\omega=A,\mathtt{p}}_{pr-part}(A) - \widetilde{r}^{\omega=A,\mathtt{p}}_{pr-part}(B) > \widetilde{r}^{\omega=A,\mathtt{s}}_{pr-part}(A) - \widetilde{r}^{\omega=A,\mathtt{s}}_{pr-part}(B)$, so that the above condition is more likely to be satisfied under transparency.

Next, at the deliberation stage, a competent member biased towards B must prefer not to reveal information and then vote correctly:

$$\begin{aligned}\phi r^{un}_{pr-part} + \alpha &\leq \left(1 - \left(1 - \frac{1}{2}\sigma\right)^{n-1}\right)\left(\phi r^{un}_{pr-part} + \alpha\right) \\ &\quad + \left(1 - \frac{1}{2}\sigma\right)^{n-1}\left(\phi\widetilde{r}^{\omega=A,\lambda}_{pr-part}(A) + \rho^{\omega=A}_{pr-part}(A)\,\alpha + \left(1 - \rho^{\omega=A}_{pr-part}(A)\right)\gamma\right)\end{aligned}$$

which can be re-expressed as:

$$\phi r^{un}_{pr-part} + \alpha \leq \phi\widetilde{r}^{\omega=A,\lambda}_{pr-part}(A) + \rho^{\omega=A}_{pr-part}(A)\,\alpha + \left(1 - \rho^{\omega=A}_{pr-part}(A)\right)\gamma$$

$$\gamma \geq \alpha + \frac{\phi(r^{un}_{pr-part} - \widetilde{r}^{\omega=A,\lambda}_{pr-part}(A))}{1 - \rho^{\omega=A}_{pr-part}(A)}$$

Furthermore, note that $\widetilde{r}^{\omega=A,\mathtt{p}}_{pr-part}(A) > \widetilde{r}^{\omega=A,\mathtt{s}}_{pr-part}(A)$, so that the above condition is more likely to be satisfied under transparency. ■

A partial revelation equilibrium with partially-competent voting is always more likely to exist under transparency. In fact, it can be shown that for certain parameter values this class of equilibrium exists under public, but not secret voting. Intuitively, the expected career-concern rewards for a competent member associated with the strategy of withholding information and then voting correctly are larger under transparency. Thus, our analysis suggests that, under certain conditions, secret voting may actually lead to better outcomes by stimulating more information sharing at the deliberation stage. In fact, as we show below, a move towards secrecy may induce a change from a partial deliberation to a full deliberation equilibrium.

Finally, in a *no revelation equilibrium*, competent members never share information. Then, at the voting stage, members may vote in accordance with either (*i*) a fully-competent equilibrium, (*ii*) a partially-competent equilibrium or (*iii*) a biased equilibrium. Consistently with our focus on responsive equilibria, we concentrate our characterization of the equilibria on the case where committee members vote correctly whenever the state of the world is revealed, which in this case can only happen outside of the equilibrium path.[6] It is relatively straightforward to see that a no

[6] The alternative case, where all members vote for their biases when the state is revealed would only make intuitive sense if, on the equilibrium path, members voted in accordance with a biased equilibrium. However, this case is similar

revelation equilibrium cannot exist. To see why, note that, on the equilibrium path, the external evaluator always computes the posterior probability that a member is competent assuming that no information was shared prior to the voting stage. But if that is the case, then a competent member with bias equal to the state of the world would always have an incentive to reveal his signal, since by doing so he would guarantee that the correct decision is made without causing a dilution of his career-concern rewards.

Non-Responsive Equilibria Within the set of non-responsive equilibria, we focus on one particular class of equilibrium, the *irrelevant deliberation equilibrium*, where all committee members vote for their biases independently of whether the state of the world was revealed or not at the deliberation stage. For simplicity, we assume that, on the equilibrium path, all competent members reveal their signals, though nothing essential in the characterization of the equilibrium would change if we assumed otherwise. As in the case of a biased equilibrium, the binding constraint here is that of the members biased against the state of the world and requires that they are willing to vote in accordance with their biases. The next proposition provides a general characterization of the existence condition for an irrelevant deliberation (irr) equilibrium.

Proposition B.3. *An irrelevant deliberation equilibrium can be supported if and only if:*

$$\gamma > \underline{\gamma}_{irr}^{\lambda}\left(\alpha, \phi, \sigma, n\right),$$

where $\underline{\gamma}_{irr}^{\lambda}\left(\alpha, \phi, \sigma, n\right) > \alpha$. *Furthermore, we have* $\underline{\gamma}_{irr}^{\mathrm{p}}\left(\alpha, \phi, \sigma, n\right) > \underline{\gamma}_{irr}^{\mathrm{s}}\left(\alpha, \phi, \sigma, n\right)$.

The structure of an irrelevant deliberation equilibrium is the essentially same as that of a biased equilibrium, the only difference being that all members are informed in this case. Furthermore, both equilibria are more likely to exist under secret voting. Therefore, it is still the case that, when the bias term is sufficiently large, transparency is expected to lead to better voting outcomes by inducing members to vote correctly.

Public versus Secret Voting: Example Consider the case of $n = 3$, $\phi = 1$ and $\sigma = 0,5$. Under public voting, the conditions for the existence of a full revelation (fr), partial revelation (pr) and irrelevant deliberation (irr) equilibria are, respectively, the following:

$$\gamma \leq \overline{\gamma}_{fr}^{\mathrm{p}} \equiv \alpha + 0.87$$

$$\underline{\gamma}_{pr}^{\mathrm{p}} \equiv \alpha + 0.68 \leq \gamma \leq \overline{\gamma}_{pr}^{\mathrm{p}} \equiv \alpha + 1.26$$

and

$$\gamma \geq \underline{\gamma}_{irr}^{\mathrm{p}} \equiv \alpha + 1$$

to an irrelevant deliberation equilibrium, which we discuss later in detail.

Under secret voting, the conditions are the following:

$$\gamma \leq \overline{\gamma}^{\mathrm{s}}_{fr} \equiv \alpha + 0.87$$

and

$$\gamma \geq \underline{\gamma}^{\mathrm{s}}_{irr} \equiv \alpha + 0,33,$$

where a partial revelation equilibrium does not exist under secrecy. Figure B.1 depicts the regions of the parameters where each class of equilibrium exists under public and secret voting for a given value of α. As before, an important feature of the equilibrium structure is the existence of multiple equilibria.

Observe that, consistently with our previous discussion, an irrelevant deliberation equilibrium is always more likely to exist under secret voting. In particular, there is a region of parameters, $\alpha + 0.33 < \gamma < \alpha + 1$, where this class of equilibrium exists under secret, but not under public voting. Conversely, a partial revelation equilibrium is always more likely to exist under public voting. Furthermore, there is a region of parameters, $\alpha + 0.685 < \gamma < \alpha + 0.873$, where a change from public to secret voting may lead to more information aggregation and better voting outcomes, provided that the equilibrium shifts from a partial deliberation to a full deliberation equilibrium. Note, however, that this may not always be the case since both a full deliberation and an irrelevant deliberation equilibrium exist in this region under secret voting, so that the impact of a change from transparency to secrecy is actually ambiguous in this case.

B.2 Dynamic Model

B.2.1 The IID Case

Under public voting, the career-concern rewards in any period t are such that, on the equilibrium path, if a member has never cast an incorrect vote before, his reputational gain is given by $r^{\omega_t,\mathrm{p}}_{i,t} = \mathbb{I}_{\{v_{it}=\omega_t\}}$, i.e. it equals one if his vote is correct. Otherwise, if a committee member has cast an incorrect vote in a previous period, we assume that his career-concern reward is $\widehat{r}^{\omega_t,\mathrm{p}}_{i,t} = 0$, i.e. a member who votes incorrectly in one period is forever punished with zero reputation afterwards, where we use the "hat" notation to denote objects outside of the equilibrium path. Under secret voting, on the other hand, the career-concern rewards in any period t are such that, on the equilibrium path, if an incorrect vote has never been observed, the reputational gain of the committee is given by $r^{\omega_t,\mathrm{s}}_t = \frac{V^c_t}{n}$, where V^c_t represents the number of correct votes in period t. Otherwise, if an incorrect vote has been observed, we assume that the career-concern reward of the committee is $\widehat{r}^{\omega_t,\mathrm{s}}_t = \frac{n-1}{n}\sigma$.

For a dynamic fully-competent equilibrium to be sustained, we must guarantee that in any period, after any history of the game on the equilibrium path, an incompetent member prefers to abstain rather than to vote, while a competent member biased against the state of the world prefers to vote correctly rather than incorrectly. To do so, we apply the one-shot deviation principle.

Consider, first, the behavior of a member who is incompetent in period t and take a history of the game where no incorrect vote has been cast before. Suppose that all other members are following the equilibrium strategies. Let $u^{\tau,\lambda}(v)$ denote the expected utility on the equilibrium path of an individual of type τ under voting rule λ who chooses $v \in \{A, \emptyset, B\}$ and, similarly, let $\widehat{u}^{\tau,\lambda}(v)$ denote the same object outside of the equilibrium path, after an incorrect vote has been observed. We use $v = \beta$ to denote a vote in accordance with one's bias and $v = s$ to represent a vote in accordance with one's signal.

Observe that the expected discounted value of following the equilibrium strategy for an incompetent member under voting rule λ is given by:

$$W^{\mathrm{nc},\lambda} = u^{\mathrm{nc},\lambda}(\emptyset) + \sum_{k=1}^{\infty} \delta^k [\sigma u^{\mathrm{c},\lambda}(s) + (1-\sigma) u^{\mathrm{nc},\lambda}(\emptyset)],$$

where the term inside brackets captures the fact that a member may be competent in future periods with probability σ. On the other hand, the expected discounted value of a one-shot deviation in the current period is given by:

$$\widehat{W}^{\mathrm{nc},\lambda} = u^{\mathrm{nc},\lambda}(\beta) + \sum_{k=1}^{\infty} \delta^k \left\{ \frac{1}{2}[\sigma u^{\mathrm{c},\lambda}(s) + (1-\sigma) u^{\mathrm{nc},\lambda}(\emptyset)] + \frac{1}{2}[\sigma \widehat{u}^{\mathrm{c},\lambda}(s) + (1-\sigma) \widehat{u}^{\mathrm{nc},\lambda}(\beta)] \right\},$$

where the expression inside curly brackets captures the idea that a member is "caught" deviating only with probability $\frac{1}{2}$, i.e. the probability that his vote turns out to be incorrect; otherwise, the game remains on the equilibrium path.

The condition for an incompetent member not to have an incentive to deviate is $W^{\mathrm{nc},\lambda} \geq \widehat{W}^{\mathrm{nc},\lambda}$, which can be expressed as:

$$u^{\mathrm{nc},\lambda}(\beta) - u^{\mathrm{nc},\lambda}(\emptyset) \leq \sum_{j=1}^{\infty} \frac{\delta^k}{2} [\sigma (u^{\mathrm{c},\lambda}(s) - \widehat{u}^{\mathrm{c},\lambda}(s)) + (1-\sigma) (u^{\mathrm{nc},\lambda}(\emptyset) - \widehat{u}^{\mathrm{nc},\lambda}(\beta))]$$

Observe that the term on the left-hand side represents the benefit of voting relative to abstaining in the current period, while the right-hand side represents the relative gains obtained from following the equilibrium strategy in the long-run. Note that, in the static version of the model, the condition for a fully-competent equilibrium is given by $u^{\mathrm{nc},\lambda}(\beta) - u^{\mathrm{nc},\lambda}(\emptyset) \leq 0$, which can be obtained from the above expression by letting $\delta \to 0$.

Given our assumptions on the behavior of players outside of the equilibrium path, we have:

$$u^{\mathrm{c},\lambda}(s) - \widehat{u}^{\mathrm{c},\lambda}(s) = \left(1 - \frac{n-1}{n} \mathbb{I}_{\{\lambda = \mathrm{s}\}}\right) \phi + \zeta(n,\sigma) \alpha$$

$$u^{\mathrm{nc},\lambda}(\emptyset) - \widehat{u}^{\mathrm{nc},\lambda}(\beta) = \zeta(n,\sigma) \alpha,$$

where $\zeta(n,\sigma) \equiv 1 - \sum_{j=0}^{(n-1)/2} \binom{n-1}{j} \left(\sigma + \frac{1}{2}(1-\sigma)\right)^{n-1-j} \left(\frac{1}{2}(1-\sigma)\right)^j$ and $0 < \zeta(n,\sigma) < 1$ for any n and σ. Intuitively, the term $\zeta(n,\sigma)$ captures the probability difference with which the correct decision is taken on and off the equilibrium path, while the term $\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathtt{s}\}}\right)\phi$ captures the difference in career-concern rewards for a competent member on and off the equilibrium path.

Thus, the condition for an incompetent member not to have an incentive to deviate is given by:

$$\begin{aligned}\gamma \quad \leq \quad \overline{\gamma}^{\lambda}_{dfull-i}\left(\alpha,\phi,\sigma,n,\delta\right) &\equiv \left(\frac{\left(n-1\right)\sigma}{2+\left(n-3\right)\sigma}+\frac{\delta}{1-\delta}\zeta\left(n,\sigma\right)\right)\alpha \\ &\quad -\left(\frac{1}{\left(1+\frac{n-3}{2}\sigma\right)\left(1-\sigma\right)^{n-2}}-\frac{\sigma\delta}{2\left(1-\delta\right)}\right)\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathtt{s}\}}\right)\phi,\end{aligned} \tag{B.1}$$

where the only differences with respect to the condition for the existence of a fully-competent equilibrium in the static version of the model are the terms $\frac{\delta}{1-\delta}\zeta\left(n,\sigma\right)$ and $\frac{\sigma\delta}{1-\delta}$. We, therefore, conclude that it is always easier to sustain behavior consistent with a fully-competent equilibrium when competence is iid across periods and the game is played repeatedly. Furthermore, the threshold $\overline{\gamma}^{\lambda}_{dfull-i}$ is strictly increasing in δ. Finally, note that the condition for the existence of a dynamic fully-competent equilibrium must be expressed as $\gamma \leq \min\{\alpha, \overline{\gamma}^{\lambda}_{dfull-i}\left(\alpha,\phi,\sigma,n,\delta\right)\}$, since it is necessary that $\gamma \leq \alpha$ in order for it to be possible to sustain behavior consistent with a partially-competent equilibrium outside of the equilibrium path.[7]

Regarding the comparison between public and secret voting, it follows from the above expression that a dynamic fully-competent equilibrium is more likely to be sustained under public voting if, and only if:

$$\frac{1}{\left(1+\frac{n-3}{2}\sigma\right)\left(1-\sigma\right)^{n-2}}-\frac{\sigma\delta}{1-\delta}\leq 0 \;\; \Rightarrow \delta \geq \overline{\delta}\left(\sigma,n\right),$$

where $\overline{\delta}\left(\sigma,n\right) \equiv \frac{1}{1+\sigma\left(1+\frac{n-3}{2}\sigma\right)\left(1-\sigma\right)^{n-2}} \in (0,1)$. That is, transparency is expected to lead to better voting outcomes if, only if, the discount factor is large enough.

B.2.2 Persistent Competence

Let $v^c_{it} \equiv \mathbb{I}_{\{v_{it}=\omega_t\}}$ be a binary variable that equals one if the vote of member i in period t is correct and define $V^c_t \equiv \sum_{i=1}^n v^c_{it}$ as the total number of correct votes in the committee in period t. Note that, on the equilibrium path, as long as there exists at least one competent member in the committee, we must have $v^c_{it} = v^c_{it'}$ for any i and any two periods $t \neq t'$and $V^c_t = V^c_{t'}$ any $t \neq t'$, since competent members always vote correctly and incompetent members always abstain. Furthermore, when there are no competent members, we have $v^c_{i1} = 0$ for all i and $V^c_1 = 0$, since, on the equilibrium path, incompetent members must abstain at least in the first period.

The decision problem of a player i in a given period t depends on two state variables that completely characterize the conditions under which a member makes a decision in the game. First, define $x^c_{it} \equiv \min\left\{v^c_{i1},...,v^c_{it-1}\right\}$. This variable summarizes the history of individual votes of member i and indicates whether or not he voted correctly in all previous periods. Second, let $V^c_{-it} \equiv \sum_{j\neq i} v^c_{jt}$ be the number of correct votes in the committee excluding player i and define $X^c_{-it} \equiv$

[7] It follows immediately from our analysis of the static version of the model that a competent member biased against the state of the world will always have an incentive to vote correctly, provided that condition (B.1) is satisfied.

$\min\left\{V^c_{-i1}, ..., V^c_{-it}\right\}$. This variable summarizes the history of the votes of all committee members other than player i and, in equilibrium, provides information about the number of competent individuals among those $n-1$ members. We assume, for completeness, that $x^c_{i1} = \emptyset$ and $X^c_{-i1} = \emptyset$. Finally, note that the vector of state variables (x^c_{it}, X^c_{-it}) is always observed by all members in any period under both voting rules.

Before we proceed, let us define two additional important objects for our analysis. First, let $u^{\tau,\lambda}_{it}\left(v|x^c_i, X^c_{-it}\right)$ denote the expected utility under voting rule λ of an individual of competence type τ when he chooses $v \in \{A, \emptyset, B\}$ and the state is (x^c_{it}, X^c_{-it}). Similarly, let $\widetilde{r}^{\tau,\lambda}_{it}\left(v|x^c_i, X^c_{-it}\right)$ be the expected career-concern reward under voting rule λ of an individual of competence τ when he chooses v and the state is (x^c_{it}, X^c_{-it}). We want to derive the conditions under which a dynamic fully-competent equilibrium can be sustained and we start by considering the behavior of an incompetent member in period $t > 1$, assuming that all other members are following their equilibrium strategies. Observe that, in this case, $X^c_{-it'} = X^c_{-it}$ for any two periods $t \neq t'$, with $t, t' > 1$. Note that the state variable X^c_{-it} provides perfect information about the number of competent members in the committee.

Let us, first, look at a subgame where $x^c_{it} = 0$, i.e. the member either abstained or voted incorrectly in a previous period. Note that, in this case, $x^c_{it'} = 0$ for any $t' > t$, by definition. Furthermore, observe that, under both voting rules, the expected career-concern reward of voting and abstaining are the same, i.e. $\widetilde{r}^{\texttt{nc},\lambda}_{it}(\beta|0, j) = \widetilde{r}^{\texttt{nc},\lambda}_{it}(\emptyset|0, j)$ for any $j \in [0, n-1]$. Intuitively, voting can never lead to any reputational gain if a member has already voted incorrectly or abstained in the game. Indeed, note that $\widetilde{r}^{\texttt{nc,p}}_{it}(\emptyset|0, j) = \widetilde{r}^{\texttt{nc,p}}_{it}(\beta|0, j) = 0$ under public voting and $\widetilde{r}^{\texttt{nc,s}}_{it}(\emptyset|0, j) = \widetilde{r}^{\texttt{nc,s}}_{it}(\beta|0, j) = \frac{j}{n}$ under secret voting. Given these observations, we can characterize the optimal behavior of incompetent members when $x^c_{it} = 0$.

Lemma B.1. *Let $t > 1$ and suppose $x^c_{it} = 0$. Assume that all other $n-1$ members are following a strategy consistent with a dynamic fully-competent equilibrium. In this case, we have:*

i. An incompetent member prefers to abstain rather than to vote for his bias when $X^c_{-it} \geq 1$, provided that $\gamma \leq \alpha$;

ii. An incompetent member always prefers to vote for his bias rather than to abstain whenever $X^c_{-it} = 0$.

Intuitively, the way in which a member behaves does not affect his career-concern rewards when $x^c_{it} = 0$. The first part of the Lemma above simply says that, in this case, an incompetent member is better off by abstaining when there is at least one competent member in the committee, provided that the common value is larger than the bias term. The second part, in turn, states that it is optimal for an incompetent member to vote for his bias whenever all other members are incompetent. We proceed our analysis under the assumption that $\gamma \leq \alpha$, in which case it follows that the value function of incompetent members, $W^{\texttt{nc},\lambda}(x^c_{it}, X^c_{-it})$, when $x^c_{it} = 0$ is characterized by:

$$W^{\mathtt{nc},\lambda}(0,j)=\frac{1}{1-\delta}u_{it}^{\mathtt{nc},\lambda}(\emptyset|0,j) \qquad \text{for any } j\geq 1$$

and

$$W^{\mathtt{nc},\lambda}(0,0)=\frac{1}{1-\delta}u_{it}^{\mathtt{nc},\lambda}(\beta|0,0)$$

Next, let us look at a subgame in which $x_{it}^{c}=1$, i.e. the member voted correctly in all previous periods. Observe that a stage in the game where $x_{it}^{c}=1$ can only be reached by an incompetent member if he has deviated from the equilibrium strategies in all previous periods. The main question here is how should the member behave in this situation. Should he continue voting or abstain? The following Lemma provides a partial characterization.

Lemma B.2. *Let $t>1$ and suppose $x_{it}^{c}=1$. Assume that all other $n-1$ members are following a strategy consistent with a dynamic fully-competent equilibrium. In this case, an incompetent member always prefers to vote for his bias rather than to abstain whenever $X_{-it}^{c}\geq 2$ or $X_{-it}^{c}=0$.*

Note, first, that when $X_{-it}^{c}\geq 2$, i.e. there are two or more competent members in the committee, the incompetent member is never pivotal. In this case, he is strictly better off by voting, since by doing so he obtains a larger expected career-concern reward without affecting the outcome in the wrong direction. On the other hand, when $X_{-it}^{c}=0$, i.e. all committee members are incompetent, there is no point in abstaining. Note that, since $x_{it}^{c}=1$, all other individuals believe that there is one competent member in the committee and abstain. Therefore, by voting the member not only obtains a larger expected career-concern reward, but also guarantees that the outcome will be consistent with his own bias.

Finally, there remains to consider the case where $X_{-it}^{c}=1$, i.e. there is a single competent member in the committee. The next lemma provides an useful result.

Lemma B.3. *Let $t>1$ and suppose $x_{it}^{c}=1$. Assume that all other $n-1$ members are following a strategy consistent with a dynamic fully-competent equilibrium. In this case, a dynamic fully-competent equilibrium can never be supported if an incompetent member prefers to vote for his bias rather than to abstain when $X_{-it}^{c}=1$.*

Note that when $X_{-it}^{c}=1$, an incompetent member is pivotal with probability one whenever he casts an incorrect vote. Observe that this is exactly the situation where incompetent members have the strongest incentive to abstain in order to avoid the swing voters' curse. Intuitively, if an incompetent individual does not abstain in this case, then he will necessarily not have an incentive to abstain in the first period, where the number of competent members in the committee is still uncertain. Therefore, for a dynamic fully-competent equilibrium to exist under voting rule λ, we must have $u_{it}^{\mathtt{nc},\lambda}(\emptyset|1,1)\geq u_{it}^{\mathtt{nc},\lambda}(\beta|1,1)$. We will proceed our analysis under the assumption that this condition holds, but, as observed above, this will not the binding constraint on the behavior of incompetent members.

Given these results, we can show that when $x_{it}^{c}=1$ the value function of incompetent members is characterized by:

$$W^{\mathtt{nc},\lambda}(1,j)=\frac{1}{1-\frac{\delta}{2}}\left[u_{it}^{\mathtt{nc},\lambda}(\beta|1,j)+\frac{\delta}{2}W^{\mathtt{nc},\lambda}(0,j)\right]\qquad \text{for any } j\geq 2$$

$$W^{\mathtt{nc},\lambda}(1,1)=W^{\mathtt{nc},\lambda}(0,1)$$

and

$$W^{\mathtt{nc},\lambda}(1,0)=\frac{1}{1-\frac{\delta}{2}}\left[u_{it}^{\mathtt{nc},\lambda}(\beta|1,0)+\frac{\delta}{2}W^{\mathtt{nc},\lambda}(0,0)\right]$$

Note that $W^{\mathtt{nc},\lambda}(1,j)-W^{\mathtt{nc},\lambda}(0,j)\geq 0$, with strict inequality when $j\geq 2$ or $j=0$, i.e. the value of concealing his type is positive for an incompetent member. Furthermore, it is possible to show that $W^{\mathtt{nc},\mathtt{p}}(1,j)-W^{\mathtt{nc},\mathtt{p}}(0,j)\geq W^{\mathtt{nc},\mathtt{s}}(1,j)-W^{\mathtt{nc},\mathtt{s}}(0,j)$, with strict inequality when $j\geq 2$ or $j=0$, i.e. the value of concealing incompetence is larger under transparency.

Next, we must guarantee that an incompetent member prefers to abstain rather than to vote in the first period, taking into account the dynamic implications of his decision. Note that this condition can be expressed as:

$$\begin{aligned}&\textstyle\sum_{j=0}^{n-1}\Pr\left(X_{-i}^{c}=j\right)\left[u_{it}^{\mathtt{nc},\lambda}(\emptyset|1,j)+\delta W^{\mathtt{nc},\lambda}(0,j)\right]\ \geq\\ &\qquad\textstyle\sum_{j=0}^{n-1}\Pr\left(X_{-i}^{c}=j\right)\left[u_{it}^{\mathtt{nc},\lambda}(\beta|1,j)+\delta\left(\frac{1}{2}W^{\mathtt{nc},\lambda}(1,j)+\frac{1}{2}W^{\mathtt{nc},\lambda}(0,j)\right)\right]\end{aligned}$$

where $u_{it}^{\mathtt{nc},\lambda}(\emptyset|1,j)$ and $u_{it}^{\mathtt{nc},\lambda}(\beta|1,j)$ represent the expected utility of abstaining and voting in the first period, conditionally on there being j competent members in the committee. Re-arranging, we can express the above inequality as:

$$\begin{aligned}&\textstyle\sum_{j=0}^{n-1}\Pr\left(X_{-i}^{c}=j\right)\left(u_{it}^{\mathtt{nc},\lambda}(\emptyset|1,j)-u_{it}^{\mathtt{nc},\lambda}(\beta|1,j)\right)\ \geq\\ &\qquad\frac{\delta}{2}\textstyle\sum_{j=0}^{n-1}\Pr\left(X_{-i}^{c}=j\right)\left(W^{\mathtt{nc},\lambda}(1,j)-W^{\mathtt{nc},\lambda}(0,j)\right)\end{aligned}$$

Note that, when $\delta\to 0$, the term on the right-hand side goes to zero, so that this constraint converges to the condition for the existence of a fully-competent equilibrium in the static model. Moreover, since $W^{\mathtt{nc},\lambda}(1,j)-W^{\mathtt{nc},\lambda}(0,j)\geq 0$ for any j, the right-hand side must be strictly positive, meaning that, when competence is persistent, it is actually harder to sustain behavior consistent with a fully-competent equilibrium when the game is played repeatedly. In particular, note that the right-hand side is strictly increasing in the discount factor, so that a dynamic fully-competent equilibrium is less likely to sustain when δ increases.

Regarding the comparison between secret and public voting, we already knew from our previous analysis that incompetent members have a larger static incentive to vote under transparency. Furthermore, since $W^{\mathtt{nc},\mathtt{p}}(1,j)-W^{\mathtt{nc},\mathtt{p}}(0,j)\geq W^{\mathtt{nc},\mathtt{s}}(1,j)-W^{\mathtt{nc},\mathtt{s}}(0,j)$ for any j, we conclude that it is always easier to sustain a dynamic fully-competent equilibrium under secret voting when competence is persistent.

Finally, the above condition can be re-expressed as $\gamma \leq \overline{\gamma}^{\lambda}_{dfull-p}(\alpha, \phi, \sigma, n, \delta)$, with $\overline{\gamma}^{\lambda}_{dfull-p}(\alpha, \phi, \sigma, n, \delta) < \alpha$, where $\overline{\gamma}^{\lambda}_{dfull-p}(\alpha, \phi, \sigma, n, \delta)$ is strictly decreasing in δ and $\overline{\gamma}^{\mathsf{p}}_{dfull-p}(\alpha, \phi, \sigma, n, \delta) < \overline{\gamma}^{\mathsf{s}}_{dfull-p}(\alpha, \phi, \sigma, n, \delta)$.

Appendix C. Additional Experimental Results

This section presents a number of additional experimental results omitted from the main text.

C.1 Learning Effects

This subsection investigates whether learning within a treatment affects the behavior of subjects. In fact, as subjects become more familiar with the structure of the game, we would expect their choices to converge towards the theoretical predictions of the model. In order to test whether this is the case, we compare the aggregate behavior of subjects across periods 1-10, 11-20 and 21-30 and check whether any pattern emerges from the data. Table C.1 reports the aggregate choices of uninformed subjects. Note, first, that abstentions under Low/Secret are significantly higher in later periods, increasing from 39.17% in periods 1-10 to 48.33% in period 21-30.[8] Furthermore, we observe an increase in the percentage of uninformed subjects who vote for their biases under High/Secret from 85.83% in periods 1-10 to 90.83% in periods 21-30.[9] Both of these results are consistent with the learning hypothesis in that they show that the observed behavior tends to converge towards the predictions of the model.

Next, Table C.2 reports separately for periods 1-10, 11-20 and 21-30 the aggregate choices of informed subjects who received a signal different than their biases. Note, first, that the percentage of informed subjects who vote in accordance with their signals under High/Secret decreases from 25.88% in periods 1-10 to 16.57% in periods 21-30.[10] We also observe a significant reduction in the proportion of subjects who vote for their biases under High/Public from 17.20% in periods 1-10 to 5.00% in periods 21-30.[11] While these result are consistent with the learning hypothesis, the percentage of abstentions under High/Secret increases slightly from 12.94% in periods 1-10 to 18.86% in periods 21-30.[12] As conjectured in section 8.2 of the paper, this result could be due to the fact that both common and private values are relatively close to each other in our setting. Thus, it is possible that some informed subjects may have simply decided to abstain as a result of being "practically" indifferent between the two alternatives.

Finally, Figure C.1 plots the dynamics of average voting behavior across all periods for each treatment together with a local polynomial smooth. Note that, consistently with our previous discussion, abstention rates increase steadily across periods under Low/Secret, while remaining

[8] The χ^2 statistic for this difference is 6.14, with $p = 0.01$.
[9] The χ^2 statistic for this difference is 4.36, with $p = 0.04$.
[10] The χ^2 statistic for this difference is 4.48, with $p = 0.03$.
[11] The χ^2 statistic for this difference is 14.8, with $p = 0.00$.
[12] Note, however, that this difference is only marginally significant, with $\chi^2 = 2.25$ and $p = 0.13$.

relatively stable under Low/Public. Moreover, while the percentage of correct votes decreases under High/Secret, it increases under High/Public. Importantly, voters' behavior seems to stabilize towards the end of each treatment, as evidenced by the fact that the learning curves become flatter in the last ten periods.

C.2 Sequencing Effects

This subsection investigates whether the main comparative results presented in section 8.3 are robust to the sequence of treatments. Table C.3 summarizes the behavior of uninformed subjects by sequence and treatment. Observe that, consistently with previous results, the percentage of abstentions is significantly higher under Low/Secret than under Low/Public irrespective of the order of treatments; that is, when the magnitude of the bias is low, abstentions are always higher under secret voting. However, the order of treatments does seem to affect the behavior of uninformed subjects in one dimension, namely the proportion of abstentions is significantly higher when the session starts with Low/Secret.[13] Thus, it seems that once a subjects "learns" to behave in a certain way (e.g. abstaining or voting for his bias), he tends to repeat the same behavior in later treatments even though it is no longer optimal for him to do so. Nonetheless, it is interesting to observe that the reduction in abstentions associated with a change from Low/Secret to Low/Public is almost identical in both sequences and approximately equal to 25%. Thus, while the order of treatments affects the baseline abstention rate, it has no impact on the size of the treatment effect itself.

Next, Table C.4 reports the behavior of informed subjects broken down by sequence and treatment, focusing, as before, on the subjects who received a signal different than their biases. Observe that our main comparative result is robust to the order of treatments, namely: under both sequences, when the magnitude of the bias is high, the proportion of informed subjectss who vote in accordance with their signals is significantly higher under public voting. However, it should be noted that the proportion of subjects who vote correctly under High/Public is larger when the session starts with High/Public (89.62%) than when it starts with High/Secret (82.28%).[14] Furthermore, a change from High/Secret to High/Public leads to an increase of 68.79% (=82.28% − 13.49%) in the percentage of correct votes when the session starts with High/Secret in comparison with an increase of 51.55% (=89.62% − 38.07%) when the session starts with High/Public. Thus, it seems that a change in behavior from voting incorrectly to voting correctly is more likely to occur than the opposite.

[13] Note that the percentage of subjects who abstain under Low/Secret is 47.65% when the session starts with Low/Secret and 33.70% when the session starts with High/Secret − the χ^2 statistic associated with this difference is 15.9, with $p = 0.00$. Similarly, the percentage of subjects who abstain under Low/Public is 22.59% when the session starts with Low/Secret but only 8.15% when the session starts with Low/Public − the χ^2 statistic associated with this difference is 27.4, with $p = 0.00$.

[14] The χ^2 statistic for this difference is 5.17, with $p = 0.02$.

C.3 Regression Analysis

We now present a detailed regression analysis of the results of the experiment. The fact that the same subjects were exposed to two different treatments, allows us to perform a rigorous analysis controlling for individual fixed effects.[15] We start by examining the determinants of a correct vote by informed subjects. Table C.5 presents the results of linear probability models where the dependent variable is a dummy that equals one if the subject voted correctly in a given period and zero otherwise. The sample is restricted to subject-period observations where the subject received a signal different than his bias. Furthermore, we focus only on high bias treatments, i.e. High/Secret and High/Public, since these are the cases where we expect a change in the degree of transparency to have an impact on voting behavior. All standard errors were clustered at the individual level.[16]

We begin by presenting in column [1] the results of a simple OLS regression of correct vote on High/Secret. Consistently with previous findings, a change from public to secret voting leads to a significant 62.7 percentage points (p.p.) decrease in the likelihood that an informed subject votes correctly. Note that, as shown in column [2], this result is very robust to controlling for individual fixed effects, as can be observed by the fact that the estimated coefficient remains almost unchanged.[17] Next, in column [3], we estimate the impact of High/Secret on the likelihood of a correct vote separately in periods 1-10, 11-20 and 21-30.[18] We find that a change from public to secret voting reduces the probability of a correct vote by 56.5 p.p. in periods 1-10, 60.4 p.p. in periods 11-20 and 68.3 p.p. in periods 21-30, which corroborates the existence of a strong learning effect for informed subjects.[19]

Finally, we create a dummy variable that captures whether a subject performed poorly in the comprehension quizzes administered before the beginning of each treatment.[20] We interpret a low performance in these tests as evidence that either the subject did not fully understand a particular aspect of the treatment or, perhaps more likely, that he or she did not put enough effort to think through the questions. The results reported in column [4] shows that subjects who performed poorly in the comprehension quiz are less responsive to changes in the degree of transparency; in particular, they are 26.4 p.p. more likely to vote correctly under High/Secret, a treatment in which we would expect all informed subjects to vote in accordance with their biases.

We now proceed to examine the determinants of abstention by uninformed subjects. Table C.6

[15] The results remain unchanged when we control for random effects instead of fixed effects.

[16] Clustering by session and adjusting the standard errors to account for the small number of clusters using a procedure proposed by Ibragimov and Müller [6] does not change any of our main results.

[17] Note that the individual fixed effects already control for all session specific characteristics, including the order of the treatments and general characteristics of the pool of participants.

[18] Our results are robust to an alternative specification where we include an interaction between High/Secret and a continuous period variable that assumes values between 1 and 30.

[19] The null hypothesis that the coefficients for these three dummies are identical is rejected at 5% confidence level ($F = 3.21$).

[20] Before the beginning of each treatment, and immediately after instructions were read aloud, subjects were asked to answer a short comprehension quiz consisting of several multiple choice questions. While these questions were simple in general, most of them required calculation of hypothetical payoffs under various scenarios. An individual is defined to have performed poorly in the comprehension quiz if the number of questions he or she got wrong was above average. Our results are robust to alternative definitions of bad performance.

presents the results of linear probability models where the dependent variable is a dummy that equals one if the subject abstained in a given period and zero otherwise. The sample is restricted to subject-period observations where the subject did not receive any information about the state of the world. The analysis focuses only on low bias treatments, i.e. Low/Secret and Low/Public. All standard errors were clustered at the subject level.[21] We, first, present in column [1] the results of a simple OLS regression of abstention on Low/Secret. The estimates confirm our previous findings that uninformed subjects are more likely to abstain under secret voting. In particular, a change from public to secret voting leads to a 25.1 p.p increase in the probability that an uninformed subject abstains. Moreover, as shown in column [2], this result is very robust to the inclusion of individual fixed effects in the regression. Next, in column [3], we estimate the impact of the Low/Secret treatment on the likelihood of abstention separately in periods 1-10, 11-20 and 21-30. The results corroborate the previous evidence that there is substantial learning occurring within a treatment, even after controlling for individual fixed effects. Specifically, the impact of a change from public to secret voting on the probability that an uninformed subject abstains is 20.5 p.p. in periods 1-10, 24.7 p.p. in periods 11-20 and 27.6 p.p. in periods 21-30.[22]

Overall, the above results are consistent with our main comparative predictions about the behavior of uninformed subjects. Still, the fraction of subjects who change from voting to abstaining as a result of a change from public to secret voting is significantly below one. Given that there are multiple equilibria under Low/Secret, it would be interesting to better understand why uninformed subjects do not select the efficient equilibrium, which involves all of them abstaining in order to let the "experts" decide. Notice that previous studies by Elbittar et. al [3], and Grosser and Seebauer [5] also found, in a setting with common values, that a substantial proportion of subjects vote even though they have no information about the state of the world.

One possible explanation for this finding can be attributed to the fact that some subjects may simply have failed to recognize the advantages associated with abstaining. Indeed, some degree of sophistication is required to understand that, under some circumstances, "doing nothing" may be better than trying to influence the voting outcome (Feddersen and Pesendorfer [4]). In order to investigate this hypothesis, we run a fixed effect regression including the interaction between Low/Secret and the dummy for poor performance in the comprehension quiz. The results reported in column [4] show that subjects who perform badly in the quiz tend to be much less responsive to changes in the degree of transparency. In particular, our estimates imply that these subjects are approximately 16.4 p.p. less likely to abstain under Low/Secret.

An alternative explanation for the relatively low levels of abstention is that, while some subjects may have recognized the potential benefits of abstaining, they were discouraged from doing so by the fact that other uninformed subjects were not abstaining as well. Indeed, the optimal behavior for an uninformed subject is to vote in accordance with his bias if he believes that other uniformed

[21] As before, clustering by session and adjusting the standard errors to account for the small number of clusters does not change any of our main results.

[22] The null hypothesis that the coefficients for these three dummies are identical is rejected at 6% confidence level ($F = 2.92$)

subjects are also voting in accordance with their biases. In order to examine whether a negative feedback in one period impacts the subsequent decisions of subjects, we define a "bad abstention" as a situation where an uninformed subjects abstains, but the decision of his or her group is incorrect, meaning that at least one other committee member “distorted" the decision by voting for the wrong alternative. We count the number of bad abstentions experienced by each subject during the first ten periods of Low/Secret and add the interaction of this variable with the Low/Secret dummy in a fixed effects regression. In doing so, we restrict the estimation sample to include only observations from the last twenty periods of each treatment (periods 11-30). We also control for the number of times that each subject abstained when uninformed in the first ten rounds of Low/Secret, given that a subject who abstains in the beginning of the treatment is more likely to continue doing so. The results reported in column [5] show that ceteris paribus a bad abstention in the first ten periods reduces the probability of an abstention in subsequent rounds by 13.9 p.p., suggesting that coordination problems among uninformed subjects may, indeed, significantly limit the convergence of voting behavior towards the efficient equilibrium.

Appendix D. Proofs

D.1 Lemma 1

Suppose, without loss of generality, that the state of the world is $\omega = A$. (By symmetry, all arguments are valid for the opposite case, where $\omega = B$.) Consider, first, the behavior of a competent member whose signal, $s_i = A$, is *equal* to his bias, $\beta_i = A$. Fix the beliefs of the external evaluator and the strategies of other players. The expected payoffs associated with each of the members' pure strategies, $v_i \in \{A, \emptyset, B\}$, are the following:

$$U^{\beta_i=A,\lambda}(v_i = A, s_i = A) = \phi\widetilde{r}_i^{\omega=A,\lambda}(v_i = A) + \rho^{\omega=A}(v_i = A)(\alpha+\gamma)$$

$$U^{\beta_i=A,\lambda}(v_i = \emptyset, s_i = A) = \phi\widetilde{r}_i^{\omega=A,\lambda}(v_i = \emptyset) + \rho^{\omega=A}(v_i = \emptyset)(\alpha+\gamma)$$

$$U^{\beta_i=A,\lambda}(v_i = B, s_i = A) = \phi\widetilde{r}_i^{\omega=A,\lambda}(v_i = B) + \rho^{\omega=A}(v_i = B)(\alpha+\gamma),$$

where $\rho^{\omega=A}(v_i = B) \leq \rho^{\omega=A}(v_i = \emptyset) \leq \rho^{\omega=A}(v_i = A)$ and $\widetilde{r}_i^{\omega=A,\lambda}(v_i = \emptyset) = \widetilde{r}_i^{\omega=A,\lambda}(v_i = B) \leq \widetilde{r}_i^{\omega=A,\lambda}(v_i = A)$, i.e. voting for A leads to a larger probability that the committee's decision is A and is also associated with a (weakly) larger career-concern reward. Thus, it follows that:

$$\max\left\{U^{\beta_i=A,\lambda}(v_i = \emptyset, s_i = A), U^{\beta_i=A,\lambda}(v_i = B, s_i = A)\right\} \leq U^{\beta_i=A,\lambda}(v_i = A, s_i = A) \tag{D.1}$$

Note that this inequality is valid for any belief of the external evaluator and any profile of strategies for the other players. Therefore, voting against the signal and abstaining are both weakly dominated strategies for a competent member whose signal is equal to his bias.

Next, consider the behavior of a competent member whose signal, $s_i = A$, is *different* than his bias, $\beta_i = B$. Fix the beliefs of the external evaluator and the strategies of other players. The expected payoffs associated with each of the member's pure strategies, $v_i \in \{A, \emptyset, B\}$, are the following:

$$U^{\beta_i=B,\lambda}(v_i = A, s_i = A) = \phi\widetilde{r}_i^{\omega=A,\lambda}(v_i = A) + \rho^{\omega=A}(v_i = A)\,\alpha + \left(1 - \rho^{\omega=A}(v_i = A)\right)\gamma$$

$$U^{\beta_i=B,\lambda}(v_i = \emptyset, s_i = A) = \phi\widetilde{r}_i^{\omega=A,\lambda}(v_i = \emptyset) + \rho^{\omega=A}(v_i = \emptyset)\,\alpha + \left(1 - \rho^{\omega=A}(v_i = \emptyset)\right)\gamma$$

$$U^{\beta_i=B,\lambda}(v_i = B, s_i = A) = \phi\widetilde{r}_i^{\omega=A,\lambda}(v_i = B) + \rho^{\omega=A}(v_i = B)\,\alpha + \left(1 - \rho^{\omega=A}(v_i = B)\right)\gamma$$

Note that, in this case, if $\alpha \geq \gamma$, then:

$$U^{\beta_i=B,\lambda}(v_i = \emptyset, s_i = A) \leq U^{\beta_i=B,\lambda}(v_i = A, s_i = A), \tag{D.2}$$

whereas if $\alpha < \gamma$, then:

$$U^{\beta_i=B,\lambda}(v_i = \emptyset, s_i = A) \leq U^{\beta_i=B,\lambda}(v_i = B, s_i = A) \tag{D.3}$$

As before, observe that these conditions are valid for any belief of the external evaluator and any profile of strategies for the other players. Therefore, abstaining is a weakly dominated strategy for a competent member whose signal is different than his bias. ■

D.2 Lemma 2

In any equilibrium where committee members do not use weakly dominated strategies, it must be the case that every competent member whose signal is equal to his bias votes correctly, $v_i = \omega$ (Lemma 1, part *i*). Therefore, by the Bayes' rule, the probability that a member is competent given that he voted correctly is strictly positive:

$$\Pr(\tau_i = \mathtt{c}|v_i = \omega) > 0$$

Given the beliefs of the external evaluator, it follows from equations (5) and (6) in the paper that the expected career-concern rewards associated with a correct vote under public and secret voting are, respectively, given by:

$$\widetilde{r}_i^{\omega,\mathtt{p}}(v_i = \omega) = \Pr(\tau_i = \mathtt{c}|v_i = \omega) \tag{D.4}$$

and

$$\widetilde{r}_i^{\omega,\mathtt{s}}(v_i = \omega) = \Pr(\tau_i = \mathtt{c}|v_i = \omega).\tfrac{1}{n}(1+\mathbb{E}(\textstyle\sum_{j\neq i}\mathbb{I}_{\{v_j=\omega\}})), \tag{D.5}$$

while the expected career-concern gains associated with an abstention or an incorrect vote under public and secret voting are, respectively, given by:

$$\widetilde{r}_i^{\omega,\mathtt{p}}(v_i \neq \omega) = 0 \tag{D.6}$$

and

$$\widetilde{r}_i^{\omega,\mathsf{s}}(v_i \neq \omega) = \Pr(\tau_i = \mathsf{c}|v_i = \omega).\tfrac{1}{n}\mathbb{E}(\textstyle\sum_{j\neq i}\mathbb{I}_{\{v_j=\omega\}})) \tag{D.7}$$

Therefore, since $\Pr(\tau_i = \mathsf{c}|v_i = \omega) > 0$, we must have that:

$$\widetilde{r}_i^{\omega,\lambda}(v_i = \omega) > \widetilde{r}^{\omega,\lambda}(v_i \neq \omega) \blacksquare$$

D.3 Lemma 3

Suppose for concreteness, and without loss of generality, that the state of the world is $\omega = A$. Assume that, in equilibrium, all competent members biased towards B vote against the signal when the state is A. Fix the beliefs of the external evaluator and the strategies of all players. In this case, we must have:

$$U^{\beta_i=B,\lambda}(v_i = B, s_i = A) \geq U^{\beta_i=B,\lambda}(v_i = A, s_i = A), \tag{D.8}$$

so that, by equation (9) in the paper, it follows that:

$$\phi\widetilde{r}_i^{\omega=A,\lambda}(B) + \rho^{\omega=A}(B)\alpha + (1-\rho^{\omega=A}(B))\gamma \geq \phi\widetilde{r}_i^{\omega=A,\lambda}(A) + \rho^{\omega=A}(A)\alpha + (1-\rho^{\omega=A}(A))\gamma$$

Note that since $\widetilde{r}_i^{\omega=A,\lambda}(B) < \widetilde{r}_i^{\omega=A,\lambda}(A)$ by Lemma 2 and $\rho^{\omega=A}(B) \leq \rho^{\omega=A}(A)$, the above inequality can be satisfied if, and only if, $\gamma > \alpha$.

Next, observe that Lemma 1, part ii, together with the fact that $\gamma > \alpha$ imply:

$$U^{\beta_i=B,\lambda}(v_i = B, s_i = A) \geq U^{\beta_i=B,\lambda}(v_i = \emptyset, s_i = A), \tag{D.9}$$

and

$$U^{\beta_i=A,\lambda}(v_i = A, s_i = B) \geq U^{\beta_i=A,\lambda}(v_i = \emptyset, s_i = B) \tag{D.10}$$

Let us now consider the behavior of an incompetent member biased towards B. We want to show that it can never be optimal for a member of this type to abstain. Note that his expected utility is given by:

$$U^{\beta_i=B,\lambda}(v_i, s_i = \emptyset) = qU^{\beta_i=B,\lambda}(v_i, s_i = A) + (1-q)\,U^{\beta_i=B,\lambda}(v_i, s_i = B)$$

Observe that voting for B is preferred to abstaining, since $U^{\beta_i=B,\lambda}(v_i = B, s_i = A) \geq U^{\beta_i=B,\lambda}(v_i = \emptyset, s_i = A)$ by (D.9), and $U^{\beta_i=B,\lambda}(v_i = B, s_i = B) > U^{\beta_i=B,\lambda}(v_i = \emptyset, s_i = B)$ by Lemma 1. Thus, for any prior $q \in (0,1)$, it follows that:

$$U^{\beta_i=B,\lambda}(v_i = B, s_i = \emptyset) \geq U^{\beta_i=B,\lambda}(v_i = \emptyset, s_i = \emptyset) \tag{D.11}$$

Similarly, we can show:

$$U^{\beta_i=A,\lambda}\left(v_i=A,s_i=\emptyset\right)\geq U^{\beta_i=A,\lambda}\left(v_i=\emptyset,s_i=\emptyset\right), \tag{D.12}$$

Therefore, abstaining can never optimal for an incompetent member.

Moreover, for an incompetent member biased towards B, we can also show that:

$$U^{\beta_i=B,\lambda}\left(v_i=B,s_i=\emptyset\right)\geq U^{\beta_i=B,\lambda}\left(v_i=A,s_i=\emptyset\right), \tag{D.13}$$

since $U^{\beta_i=B,\lambda}(v_i=B,s_i=A)\geq U^{\beta_i=B,\lambda}(v_i=A,s_i=A)$ by assumption (see (D.8)) and the fact that $U^{\beta_i=B,\lambda}(v_i=B,s_i=B)\geq U^{\beta_i=B,\lambda}(v_i=A,s_i=B)$. Therefore, from (D.11) and (D.13) it follows that an incompetent member biased towards B will necessarily vote for his bias in this case.

Finally, since none of the above results depend on the prior probability q, the same arguments can be applied to the opposite case where the competent members biased towards A vote against the signal when the state of the world is B. ■

D.4 Proposition 1

We focus on symmetric pure-strategy equilibria where players do not use weakly dominated strategies. From Lemma 1, it follows that competent members never abstain in equilibrium. Therefore, we can divide their possible equilibrium strategies into two categories: either (a) they all vote in accordance with the signal; or (b) not all types vote in accordance with the signal. Next, from Lemma 3, it follows that incompetent members never abstain when a competent member votes against the state of the world, which corresponds to the situation described in case (b) above. Therefore, combining the results in Lemmas 1 and 3, the result follows. ■

D.5 Proposition 2-4: Preliminaries

For a given system of beliefs μ and voting rule λ, let us define the following objects, which will be used in the proofs of Propositions 2-4:

$$\Lambda^{\lambda}_{1,\mu}\equiv\frac{\phi(\widetilde{r}^{\omega=B,\lambda}_{\mu}(B)-\widetilde{r}^{\omega=B,\lambda}_{\mu}(A))}{\rho^{\omega=B}_{\mu}(A)-\rho^{\omega=B}_{\mu}(B)}\geq 0 \tag{D.14}$$

$$\Lambda^{\lambda}_{2,\mu}\equiv\frac{\phi(\widetilde{r}^{\omega=A,\lambda}_{\mu}(A)-\widetilde{r}^{\omega=A,\lambda}_{\mu}(B))}{\rho^{\omega=A}_{\mu}(A)-\rho^{\omega=A}_{\mu}(B)}\geq 0 \tag{D.15}$$

$$\Lambda^{\lambda}_{3,\mu}\equiv\frac{q\phi(\widetilde{r}^{\omega=A,\lambda}_{\mu}(A)-\widetilde{r}^{\omega=A,\lambda}_{\mu}(\emptyset))}{q(\rho^{\omega=A}_{\mu}(A)-\rho^{\omega=A}_{\mu}(\emptyset))+(1-q)\left(\rho^{\omega=B}_{\mu}(A)-\rho^{\omega=B}_{\mu}(\emptyset)\right)}\geq 0 \tag{D.16}$$

$$\Lambda^{\lambda}_{4,\mu}\equiv\frac{(1-q)\,\phi(\widetilde{r}^{\omega=B,\lambda}_{\mu}(B)-\widetilde{r}^{\omega=B,\lambda}_{\mu}(\emptyset))}{q(\rho^{\omega=A}_{\mu}(\emptyset)-\rho^{\omega=A}_{\mu}(B))+(1-q)\left(\rho^{\omega=B}_{\mu}(\emptyset)-\rho^{\omega=B}_{\mu}(B)\right)}\geq 0 \tag{D.17}$$

and

$$\Gamma_{1,\mu} \equiv \frac{(1-q)\left(\rho_\mu^{\omega=B}(A)-\rho_\mu^{\omega=B}(\emptyset)\right)-q(\rho_\mu^{\omega=A}(A)-\rho_\mu^{\omega=A}(\emptyset))}{q(\rho_\mu^{\omega=A}(A)-\rho_\mu^{\omega=A}(\emptyset))+(1-q)\left(\rho_\mu^{\omega=B}(A)-\rho_\mu^{\omega=B}(\emptyset)\right)} \tag{D.18}$$

$$\Gamma_{2,\mu} \equiv \frac{q(\rho_\mu^{\omega=A}(\emptyset)-\rho_\mu^{\omega=A}(B))-(1-q)\left(\rho_\mu^{\omega=B}(\emptyset)-\rho_\mu^{\omega=B}(B)\right)}{q(\rho_\mu^{\omega=A}(\emptyset)-\rho_\mu^{\omega=A}(B))+(1-q)\left(\rho_\mu^{\omega=B}(\emptyset)-\rho_\mu^{\omega=B}(B)\right)} \tag{D.19}$$

Note that the terms $\Gamma_{1,\mu}$ and $\Gamma_{2,\mu}$ do not depend on the voting rule λ. Observe also that the terms $\Lambda_{1,\mu}^\lambda$, $\Lambda_{2,\mu}^\lambda$, $\Lambda_{3,\mu}^\lambda$ and $\Lambda_{4,\mu}^\lambda$ are all between 0 and 1. Finally, although we cannot determine the sign of $\Gamma_{1,\mu}$ and $\Gamma_{2,\mu}$, it must be the case that $-1 \leq \Gamma_{1,\mu} \leq 1$ and $-1 \leq \Gamma_{2,\mu} \leq 1$.

D.6 Proposition 2

The conditions for the existence of a fully-competent equilibrium are the following. First, every competent member who receives a signal different than his bias must prefer to vote in accordance with the signal:

$$U_{full}^{\beta_i=A,\lambda}(v_i=B, s_i=B) \geq U_{full}^{\beta_i=A,\lambda}(v_i=A, s_i=B) \tag{D.20}$$

and

$$U_{full}^{\beta_i=B,\lambda}(v_i=A, s_i=A) \geq U_{full}^{\beta_i=B,\lambda}(v_i=B, s_i=A) \tag{D.21}$$

Second, all incompetent members must prefer to abstain rather than to vote for either one of the alternatives:

$$U_{full}^{\beta_i=A,\lambda}\left(v_i=\emptyset, s_i=\emptyset\right) \geq \max\{U_{full}^{\beta_i=A,\lambda}\left(v_i=A, s_i=\emptyset\right), U_{full}^{\beta_i=A,\lambda}\left(v_i=B, s_i=\emptyset\right)\} \tag{D.22}$$

and

$$U_{full}^{\beta_i=B,\lambda}\left(v_i=\emptyset, s_i=\emptyset\right) \geq \max\{U_{full}^{\beta_i=B,\lambda}\left(v_i=A, s_i=\emptyset\right), U_{full}^{\beta_i=B,\lambda}\left(v_i=B, s_i=\emptyset\right)\}, \tag{D.23}$$

where we assume that the beliefs of all players, including that of the external evaluator, are consistent with the equilibrium strategies.

After some algebra, it is possible to re-express the conditions on the behavior of competent members (D.20) and (D.21), respectively, as:

$$\gamma \leq \alpha + \Lambda_{1,full}^\lambda \tag{D.24}$$

$$\gamma \leq \alpha + \Lambda_{2,full}^\lambda \tag{D.25}$$

Furthermore, the condition on the behavior of incompetent members (D.22) can be shown to imply that the following two conditions must hold:

$$\gamma \leq \alpha\Gamma_{1,full} - \Lambda_{3,full}^\lambda \tag{D.26}$$

$$\gamma \geq -\alpha\Gamma_{2,full} + \Lambda_{4,full}^\lambda \tag{D.27}$$

Similarly, condition (D.23) implies that:

$$\gamma \leq \alpha\Gamma_{2,full} - \Lambda^{\lambda}_{4,full} \tag{D.28}$$

$$\gamma \geq -\alpha\Gamma_{1,full} + \Lambda^{\lambda}_{3,full}, \tag{D.29}$$

where the terms $\Lambda^{\lambda}_{1,full}$, $\Lambda^{\lambda}_{2,full}$, $\Lambda^{\lambda}_{3,full}$, $\Lambda^{\lambda}_{4,full}$, $\Gamma_{1,full}$ and $\Gamma_{2,full}$ are defined in accordance with expressions (D.14) – (D.19).

In equilibrium, conditions (D.24) – (D.29) must hold simultaneously, but the characterization of the equilibrium is greatly simplified by the following two remarks. First, we can show that:

Remark 1. *If* (D.26) *is satisfied, then* (D.24) *and* (D.29) *must also hold.*

Proof. Note that if condition (D.26) is satisfied, then we must have $\alpha\Gamma_{1,full} - \Lambda^{\lambda}_{3,full} > 0$, since $\gamma > 0$, which in turn implies that $\Gamma_{1,full} > 0$. Thus, since $0 < \Gamma_{1,full} \leq 1$ and $\Lambda^{\lambda}_{1,full}, \Lambda^{\lambda}_{3,full} \geq 0$, condition (D.24) must hold, given that $\alpha\Gamma_{1,full} - \Lambda^{\lambda}_{3,full} < \alpha + \Lambda^{\lambda}_{1,full}$. Furthermore, from $\alpha\Gamma_{1,full} - \Lambda^{\lambda}_{3,full} > 0$ it follows immediately that $-\alpha\Gamma_{1,full} + \Lambda^{\lambda}_{3,full} < 0$, which means that condition (D.29) is also necessarily satisfied. ■

Next, following a similar argument, we can show that:

Remark 2. *If* (D.28) *is satisfied, then* (D.25) *and* (D.27) *must also hold.*

Therefore, the binding constraints for the existence of a fully-competent equilibrium are given by (D.26) and (D.28). Intuitively, what we have shown here is that if, for instance, an incompetent member biased towards A prefers to abstain rather than to vote for A, then: (a) all incompetent members biased towards B also prefer to abstain rather than to vote for A, and (b) all competent members biased towards A always prefer to vote for the signal rather than to follow their bias.

Observe that we can express conditions (D.26) and (D.28) more compactly as:

$$\gamma \leq \overline{\gamma}^{\lambda}_{full}(\alpha, \phi, \sigma, n), \tag{D.30}$$

where $\overline{\gamma}^{\lambda}_{full}(\alpha, \phi, \sigma, n) \equiv \min\{\alpha\Gamma_{1,full} - \Lambda^{\lambda}_{3,full}, \alpha\Gamma_{2,full} - \Lambda^{\lambda}_{4,full}\}$. Note that $\overline{\gamma}^{\lambda}_{full}(\alpha, \phi, \sigma, n) < \alpha$, since $\Lambda^{\lambda}_{3,full}, \Lambda^{\lambda}_{4,full} \geq 0$ and $-1 < \Gamma_{1,full}, \Gamma_{2,full} < 1$.

Moreover, we must have:

$$\overline{\gamma}^{\mathtt{p}}_{full}(\alpha, \phi, \sigma, n) < \overline{\gamma}^{\mathtt{s}}_{full}(\alpha, \phi, \sigma, n), \tag{D.31}$$

since $\Lambda^{\mathtt{p}}_{3,full} > \Lambda^{\mathtt{s}}_{3,full}$ and $\Lambda^{\mathtt{p}}_{4,full} > \Lambda^{\mathtt{s}}_{4,full}$ in equilibrium, which follow, respectively, from:

$$\widetilde{r}^{\,\omega=A,\mathtt{p}}_{full}(A) - \widetilde{r}^{\,\omega=A,\mathtt{p}}_{full}(\emptyset) > \widetilde{r}^{\,\omega=A,\mathtt{s}}_{full}(A) - \widetilde{r}^{\,\omega=A,\mathtt{s}}_{full}(\emptyset)$$

and

$$\widetilde{r}^{\,\omega=B,\mathtt{p}}_{full}(B) - \widetilde{r}^{\,\omega=B,\mathtt{p}}_{full}(\emptyset) > \widetilde{r}^{\,\omega=B,\mathtt{s}}_{full}(B) - \widetilde{r}^{\,\omega=B,\mathtt{s}}_{full}(\emptyset)$$

Intuitively, the relative career-concern reward associated with a correct vote is larger under public voting, so that the incompetent members have less incentive to abstain under transparency. ■

D.7 Proposition 3

A partially-competent equilibrium requires that all competent members vote correctly and that incompetent members of at least one type vote – either for their biases or for the ex-ante more likely alternative. Thus, a partially-competent equilibria is consistent with a variety of different behaviors on the part of incompetent members. The next remark helps us to put some order in the set of all possible equilibria.

Remark 1. *If, in equilibrium, an incompetent member biased towards B votes for A, then all incompetent members biased towards A must vote for A. A similar argument holds for the opposite case: if, in equilibrium, an incompetent member biased towards A votes for B, then all incompetent members biased towards B must vote for B.*

Proof. Suppose without loss of generality that, in equilibrium, an incompetent member biased towards B votes for A. In this case, we must have that:

$$U^{\beta_i=B,\lambda}\left(v_i=A,s_i=\emptyset\right)\geq U^{\beta_i=B,\lambda}\left(v_i=\emptyset,s_i=\emptyset\right)$$

Note that the above inequality can be re-written as:

$$\begin{aligned}q(U^{\beta_i=B,\lambda}\left(v_i=A,s_i=A\right)-U^{\beta_i=B,\lambda}\left(v_i=\emptyset,s_i=A\right))\geq\\ -\left(1-q\right)\left(U^{\beta_i=B,\lambda}\left(v_i=A,s_i=B\right)-U^{\beta_i=B,\lambda}\left(v_i=\emptyset,s_i=B\right)\right)\end{aligned}$$

Next, using expressions (8) and (9) in the paper, and after some algebra, it is possible to show that:

$$U^{\beta_i=A,\lambda}\left(v_i=A,s_i=A\right)-U^{\beta_i=A,\lambda}\left(v_i=\emptyset,s_i=A\right)>U^{\beta_i=B,\lambda}\left(v_i=A,s_i=A\right)-U^{\beta_i=B,\lambda}\left(v_i=\emptyset,s_i=A\right)$$

$$U^{\beta_i=A,\lambda}\left(v_i=A,s_i=B\right)-U^{\beta_i=A,\lambda}\left(v_i=\emptyset,s_i=B\right)>U^{\beta_i=B,\lambda}\left(v_i=A,s_i=B\right)-U^{\beta_i=B,\lambda}\left(v_i=\emptyset,s_i=B\right)$$

Thus, it follows that:

$$\begin{aligned}q(U^{\beta_i=A,\lambda}\left(v_i=A,s_i=A\right)-U^{\beta_i=A,\lambda}\left(v_i=\emptyset,s_i=A\right))>\\ -\left(1-q\right)\left(U^{\beta_i=A,\lambda}\left(v_i=A,s_i=B\right)-U^{\beta_i=A,\lambda}\left(v_i=\emptyset,s_i=B\right)\right),\end{aligned}$$

which implies:

$$U^{\beta_i=A,\lambda}\left(v_i=A,s_i=\emptyset\right)>U^{\beta_i=A,\lambda}\left(v_i=\emptyset,s_i=\emptyset\right)\tag{D.32}$$

To complete the argument, note that if an incompetent member biased towards B votes for A, then we must also have that:

$$U^{\beta_i=B,\lambda}\left(v_i=A,s_i=\emptyset\right)\geq U^{\beta_i=B,\lambda}\left(v_i=B,s_i=\emptyset\right)$$

and it is immediate to show that the above condition implies:

$$^{\beta_i=A,\lambda}\left(v_i=A,s_i=\emptyset\right)>U^{\beta_i=A,\lambda}\left(v_i=\emptyset,s_i=\emptyset\right) \tag{D.33}$$

Therefore, from (D.32) and (D.33), it follows that an incompetent member biased towards A would also vote for A. ∎

Observe that the above result allows us to rule out the existence of any equilibrium where, for instance, an incompetent member biased towards B votes for A and an incompetent member biased towards A abstain. Thus, it follows that in a partially-competent equilibrium we must always have incompetent members of at least one type voting for their bias. We, therefore, focus our analysis on the following three subclasses of partially-competent equilibria: (i) incompetent members biased towards A vote for A and incompetent members biased towards B either vote or abstain, (ii) incompetent members biased towards B vote for B and incompetent members biased towards A either vote or abstain, and (iii) both types of incompetent members vote for their biases. This classification allows us to focus on types of equilibria for which the binding constraint on the behavior of incompetent members is different.

Fix a subclass of partially-competent equilibrium and denote it by $part_j$. For this equilibrium to exist, the following conditions must hold. First, every competent member who receives a signal different than his bias must prefer to vote in accordance with the state of the world, that is:

$$U_{part_j}^{\beta=A,\lambda}(v_i=B,s_i=B)\geq U_{part_j}^{\beta=A,\lambda}(v_i=A,s_i=B) \tag{D.34}$$

and

$$U_{part_j}^{\beta=B,\lambda}(v_i=A,s_i=A)\geq U_{part_j}^{\beta=B,\lambda}(v_i=B,s_i=A) \tag{D.35}$$

Second, incompetent members of at least one type must prefer to vote for their biases:

$$U_{part_j}^{\beta_i=A,\lambda}\left(v_i=A,s_i=\emptyset\right)\geq U_{part_j}^{\beta_i=A,\lambda}\left(v_i=\emptyset,s_i=\emptyset\right) \tag{D.36}$$

or/and

$$U_{part_j}^{\beta_i=B,\lambda}\left(v_i=B,s_i=\emptyset\right)\geq U_{part_j}^{\beta_i=B,\lambda}\left(v_i=\emptyset,s_i=\emptyset\right) \tag{D.37}$$

Note here that, depending on the equilibrium, we may have one of the following cases: (i) only condition (D.36) is satisfied, (ii) only condition (D.37) is satisfied, or (iii) both conditions are satisfied.

After some algebra, it is possible to re-express the conditions on the behavior of competent members, respectively, as:

$$\gamma \leq \alpha + \Lambda^{\lambda}_{1,part_j}$$

and

$$\gamma \leq \alpha + \Lambda^{\lambda}_{2,part_j},$$

with $\Lambda^{\lambda}_{1,part_j}$, $\Lambda^{\lambda}_{2,part_j} \geq 0$; while the conditions on the behavior of incompetent members can be expressed as:

$$\gamma \geq \alpha\Gamma_{1,part_j} - \Lambda^{\lambda}_{3,part_j}$$

and/or

$$\gamma \geq \alpha\Gamma_{2,part_j} - \Lambda^{\lambda}_{4,part_j},$$

with $\Lambda^{\lambda}_{3,part_j}$, $\Lambda^{\lambda}_{4,part_j} \geq 0$, $-1 \leq \Gamma_{1,part_j} \leq 1$ and $-1 \leq \Gamma_{2,part_j} \leq 1$, where the terms $\Lambda^{\lambda}_{1,part_j}$, $\Lambda^{\lambda}_{2,part_j}$, $\Lambda^{\lambda}_{3,part_j}$, $\Lambda^{\lambda}_{4,part_j}$, $\Gamma_{1,part_j}$ and $\Gamma_{2,part_j}$ are defined in accordance with (D.14)$-$(D.19).

Given the inequalities above, it follows that the condition for the existence of a partially-competent equilibrium $part_j$ can be expressed as:

$$\underline{\gamma}^{\lambda}_{part_j}(\alpha, \phi, \sigma, n) \leq \gamma \leq \overline{\gamma}^{\lambda}_{part_j}(\alpha, \phi, \sigma, n), \tag{D.38}$$

for some thresholds $\underline{\gamma}^{\lambda}_{part_j}(\alpha, \phi, \sigma, n) < \alpha$ and $\overline{\gamma}^{\lambda}_{part_j}(\alpha, \phi, \sigma, n) > \alpha$, where $\overline{\gamma}^{\lambda}_{part_j}(\alpha, \phi, \sigma, n) \equiv \min\{\alpha + \Lambda^{\lambda}_{1,part_j}, \alpha + \Lambda^{\lambda}_{2,part_j}\}$, while $\underline{\gamma}^{\lambda}_{part_j}(\alpha, \phi, \sigma, n)$ depends on which constraint on the behavior of incompetent members is binding in equilibrium. Moreover, we must also have that:

$$\underline{\gamma}^{\mathtt{p}}_{part_j}(\alpha, \phi, \sigma, n) < \underline{\gamma}^{\mathtt{s}}_{part_j}(\alpha, \phi, \sigma, n), \tag{D.39}$$

since $\Lambda^{\mathtt{p}}_{3,part_j} > \Lambda^{\mathtt{s}}_{3,part_j}$ and $\Lambda^{\mathtt{p}}_{4,part_j} > \Lambda^{\mathtt{s}}_{4,part_j}$; and

$$\overline{\gamma}^{\mathtt{s}}_{part_j}(\alpha, \phi, \sigma, n) < \overline{\gamma}^{\mathtt{p}}_{part_j}(\alpha, \phi, \sigma, n), \tag{D.40}$$

since $\Lambda^{\mathtt{p}}_{1,part_j} > \Lambda^{\mathtt{s}}_{1,part_j}$ and $\Lambda^{\mathtt{p}}_{2,part_j} > \Lambda^{\mathtt{s}}_{2,part_j}$. The argument here is similar to the one used in the proof of Proposition 2.

Next, define:

$$\underline{\gamma}^{\lambda}_{part}(\alpha, \phi, \sigma, n) \equiv \min_{part_j} \underline{\gamma}^{\lambda}_{part_j}(\alpha, \phi, \sigma, n) \tag{D.41}$$

and

$$\overline{\gamma}^{\lambda}_{part}(\alpha, \phi, \sigma, n) \equiv \max_{part_j} \overline{\gamma}^{\lambda}_{part_j}(\alpha, \phi, \sigma, n) \tag{D.42}$$

The region of parameters where a partially-competent equilibrium of any sort can be sustained is given by:

$$\underline{\gamma}^{\lambda}_{part}(\alpha, \phi, \sigma, n) \leq \gamma \leq \overline{\gamma}^{\lambda}_{part}(\alpha, \phi, \sigma, n), \tag{D.43}$$

where $\underline{\gamma}^{\lambda}_{part}(\alpha, \phi, \sigma, n) < \alpha$ and $\overline{\gamma}^{\lambda}_{part}(\alpha, \phi, \sigma, n) > \alpha$. Intuitively, for any equilibrium $part_j$ it must be the case that (D.43) is satisfied, since (D.38) implies (D.43). Conversely, we can guarantee that, in any region of parameters where (D.43) is satisfied, it is possible to sustain an equilibrium which belongs to a subclass of partially-competent equilibrium.

Finally, it follows from definition (D.41) and (D.42) together with the relationships expressed in inequalities (D.39) and (D.40) that:

$$\underline{\gamma}^{\mathsf{p}}_{part}\left(\alpha,\phi,\sigma,n\right) < \underline{\gamma}^{\mathsf{s}}_{part}\left(\alpha,\phi,\sigma,n\right) \tag{D.44}$$

and

$$\overline{\gamma}^{\mathsf{s}}_{part}\left(\alpha,\phi,\sigma,n\right) < \overline{\gamma}^{\mathsf{p}}_{part}\left(\alpha,\phi,\sigma,n\right) \ \blacksquare \tag{D.45}$$

D.8 Proposition 4

A biased equilibrium requires that: (a) competent members of at least one type vote against the signal when the state is different than their bias and (b) all incompetent members vote. Note, however, that Lemma 3 guarantees that (a) implies (b), so that we do not have to worry about the conditions on the behavior of incompetent members in order to characterize the equilibrium. We, therefore, focus our analysis on the following three subclasses of biased equilibrium in which, when the signal is different than the bias: (i) competent members biased towards A vote for A and competent members biased towards B may or may not vote for their bias, (ii) competent members biased towards B vote for B and competent members biased towards A may or may not vote for their bias, and (iii) both types of competent members vote for their biases. This classification allows us to focus on types of equilibria for which the binding constraint on the behavior of competent members is different.

Fix a subclass of biased equilibrium and denote it by $bias_j$. For such an equilibrium to exist, one of the following conditions must be satisfied:

$$U^{\beta=A,\lambda}_{bias_j}(v_i = B, s_i = B) \leq U^{\beta=A,\lambda}_{bias_j}(v_i = A, s_i = B) \tag{D.46}$$

and/or

$$U^{\beta=B,\lambda}_{bias_j}(v_i = A, s_i = A) \leq U^{\beta=B,\lambda}_{bias_j}(v_i = B, s_i = A) \tag{D.47}$$

Note here that, depending on the equilibrium, we may have one of the following cases: (i) only condition (D.46) is satisfied, (ii) only condition (D.47) is satisfied, or (iii) both conditions are satisfied. After some algebra, it is possible to re-expressed the above conditions, respectively, as:

$$\gamma \geq \alpha + \Lambda^{\lambda}_{1,bias_j} \tag{D.48}$$

and/or

$$\gamma \geq \alpha + \Lambda^{\lambda}_{2,bias_j}, \tag{D.49}$$

with $\Lambda^{\lambda}_{1,bias_j}$,$\Lambda^{\lambda}_{2,bias_j} \geq 0$, where the terms $\Lambda^{\lambda}_{1,bias_j}$ and $\Lambda^{\lambda}_{2,bias_j}$ are defined in accordance with expressions (D.14) and (D.15).

From the above inequalities, it follows that the condition for the existence of a biased equilibrium $bias_j$ can be expressed as:

$$\gamma \geq \underline{\gamma}^{\lambda}_{bias_j}\left(\alpha, \phi, \sigma, n\right), \tag{D.50}$$

for some threshold $\underline{\gamma}^{\lambda}_{bias_j}\left(\alpha, \phi, \sigma, n\right) > \alpha$, where the exact expression for $\underline{\gamma}^{\lambda}_{bias_j}\left(\alpha, \phi, \sigma, n\right)$ depends on which constraint on the behavior of competent members is binding in equilibrium. Furthermore, it must also be the case that:

$$\underline{\gamma}^{\mathtt{s}}_{bias_j}\left(\alpha, \phi, \sigma, n\right) < \underline{\gamma}^{\mathtt{p}}_{bias_j}\left(\alpha, \phi, \sigma, n\right), \tag{D.51}$$

since $\Lambda^{\mathtt{p}}_{1,bias} > \Lambda^{\mathtt{s}}_{1,bias}$ and $\Lambda^{\mathtt{p}}_{2,biast} > \Lambda^{\mathtt{s}}_{2,bias}$.

Next, define:

$$\underline{\gamma}^{\lambda}_{bias}\left(\alpha, \phi, \sigma, n\right) \equiv \min_{bias_j} \underline{\gamma}^{\lambda}_{bias_j}\left(\alpha, \phi, \sigma, n\right) \tag{D.52}$$

The region of parameters where a biased equilibrium of any sort can be sustained is characterized by:

$$\gamma \geq \underline{\gamma}^{\lambda}_{bias}\left(\alpha, \phi, \sigma, n\right), \tag{D.53}$$

where $\underline{\gamma}^{\lambda}_{bias}\left(\alpha, \phi, \sigma, n\right) > \alpha$. Finally, it follows from definition (D.52) together with the relationship expressed in inequality (D.51) that:

$$\underline{\gamma}^{\mathtt{s}}_{bias}\left(\alpha, \phi, \sigma, n\right) < \underline{\gamma}^{\mathtt{p}}_{bias}\left(\alpha, \phi, \sigma, n\right) \quad \blacksquare \tag{D.54}$$

D.9 Proposition 5

We start by deriving the conditions for the existence of a fully-competent equilibrium under symmetry. Assuming that all competent members vote correctly and all incompetent members abstain, we have:

$$\rho^{\omega=A}_{full}(A) - \rho^{\omega=A}_{full}(\emptyset) = \rho^{\omega=B}_{full}(\emptyset) - \rho^{\omega=B}_{full}(B) = \frac{1}{2}\left(1-\sigma\right)^{n-1}$$

$$\rho^{\omega=A}_{full}(\emptyset) - \rho^{\omega=A}_{full}(B) = \rho^{\omega=B}_{full}(A) - \rho^{\omega=B}_{full}(\emptyset) = \frac{1}{2}\left(1-\sigma\right)^{n-1} + \frac{1}{2}\left(n-1\right)\left(1-\sigma\right)^{n-2}\sigma$$

Moreover, note that in this case:

$$\widetilde{r}^{\,\omega=A,\mathtt{p}}_{full}(A) = \widetilde{r}^{\,\omega=B,\mathtt{p}}_{full}(B) = 1$$

$$\widetilde{r}^{\,\omega=A,\mathtt{p}}_{full}(\emptyset) = \widetilde{r}^{\,\omega=B,\mathtt{p}}_{full}(\emptyset) = 0$$

$$\widetilde{r}^{\,\omega=A,\mathtt{s}}_{full}(A) = \widetilde{r}^{\,\omega=B,\mathtt{s}}_{full}(B) = \frac{1}{n}\left(1 + \mathbb{E}\left(\ \textstyle\sum_{j\neq i}\ \mathbb{I}_{\{v_j=\omega\}}\right)\right)$$

$$\widetilde{r}^{\,\omega=A,\mathtt{s}}_{full}(\emptyset) = \widetilde{r}^{\,\omega=B,\mathtt{p}}_{full}(\emptyset) = \frac{1}{n}\mathbb{E}\left(\ \textstyle\sum_{j\neq i}\ \mathbb{I}_{\{v_j=\omega\}}\right)$$

Therefore, from (D.30), it follows that:

$$\overline{\gamma}^{\lambda}_{full}\left(\alpha,\phi,\sigma,n\right) \equiv \frac{\left(n-1\right)\sigma}{2+\left(n-3\right)\sigma}\alpha - \frac{1}{\left(1+\frac{n-3}{2}\sigma\right)\left(1-\sigma\right)^{n-2}}\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathtt{s}\}}\right)\phi \tag{D.55}$$

Next, we proceed to derive the conditions for the existence of a partially-competent equilibrium under symmetry. Assuming that all competent members vote correctly and all incompetent members vote for their biases, we have:

$$\rho^{\omega}_{part}(A) - \rho^{\omega}_{part}(B) = \binom{n-1}{\left(n-1\right)/2}\left(\sigma+\frac{1}{2}\left(1-\sigma\right)\right)^{\frac{n-1}{2}}\left(\frac{1}{2}\left(1-\sigma\right)\right)^{\frac{n-1}{2}}$$

$$\rho^{\omega}_{part}(A) - \rho^{\omega}_{part}(\emptyset) = \rho^{\omega}_{part}(\emptyset) - \rho^{\omega}_{part}(B) = \frac{1}{2}\binom{n-1}{\left(n-1\right)/2}\left(\sigma+\frac{1}{2}\left(1-\sigma\right)\right)^{\frac{n-1}{2}}\left(\frac{1}{2}\left(1-\sigma\right)\right)^{\frac{n-1}{2}},$$

for $\omega \in \{A,B\}$, where the term $\sigma+\frac{1}{2}\left(1-\sigma\right)$ represents the proportion of committee members that are expected to vote for the correct alternative in equilibrium. Note, also, that:

$$\widetilde{r}^{\,\omega=B,\mathtt{p}}_{part}(B) = \widetilde{r}^{\,\omega=A,\mathtt{p}}_{part}(A) = \frac{\sigma}{\sigma+\frac{1}{2}\left(1-\sigma\right)}$$

$$\widetilde{r}^{\,\omega=A,\mathtt{p}}_{part}(B) = \widetilde{r}^{\,\omega=B,\mathtt{p}}_{part}(A) = \widetilde{r}^{\,\omega=A,\mathtt{p}}_{part}(\emptyset) = \widetilde{r}^{\,\omega=B,\mathtt{p}}_{part}(\emptyset) = 0$$

$$\widetilde{r}^{\,\omega=B,\mathtt{s}}_{part}(B) = \widetilde{r}^{\,\omega=A,\mathtt{s}}_{part}(A) = \frac{1}{n}\frac{\sigma}{\sigma+\frac{1}{2}\left(1-\sigma\right)}\left(1+\mathbb{E}\left(\sum_{j\neq i}\mathbb{I}_{\{v_j=\omega\}}\right)\right)$$

$$\widetilde{r}^{\,\omega=A,\mathtt{s}}_{part}(B) = \widetilde{r}^{\,\omega=B,\mathtt{s}}_{part}(A) = \widetilde{r}^{\,\omega=A,\mathtt{s}}_{part}(\emptyset) = \widetilde{r}^{\,\omega=B,\mathtt{s}}_{part}(\emptyset) = \frac{1}{n}\frac{\sigma}{\sigma+\frac{1}{2}\left(1-\sigma\right)}\mathbb{E}\left(\sum_{j\neq i}\mathbb{I}_{\{v_j=\omega\}}\right)$$

Therefore, from equations (D.38), it follows that:

$$\underline{\gamma}^{\lambda}_{part}\left(\alpha,\phi,\sigma,n\right) < 0 \tag{D.56}$$

and

$$\overline{\gamma}^{\lambda}_{part}\left(\alpha,\phi,\sigma,n\right) = \alpha + \frac{2^n\sigma}{\binom{n-1}{(n-1)/2}\left(1+\sigma\right)^{\frac{n+1}{2}}\left(1-\sigma\right)^{\frac{n-1}{2}}}\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathtt{s}\}}\right)\phi, \tag{D.57}$$

where the first expression follows from the fact that $\alpha\Gamma^{\lambda}_{1,part} - \Gamma^{\lambda}_{2,part} = \alpha\Gamma^{\lambda}_{3,part} - \Gamma^{\lambda}_{4,part} < 0$, since $\Gamma^{\lambda}_{1,part} = \Gamma^{\lambda}_{3,part} = 0$ and $\Gamma^{\lambda}_{2,part}, \Gamma^{\lambda}_{4,part} > 0$.

Finally, let us derive the conditions for the existence of a biased equilibrium. Assuming that all members vote in accordance with their biases, we have:

$$\rho^{\omega}_{bias}(A) - \rho^{\omega}_{bias}(B) = \binom{n-1}{(n-1)/2}\left(\frac{1}{2}\right)^{\frac{n-1}{2}}\left(\frac{1}{2}\right)^{\frac{n-1}{2}},$$

for $\omega \in \{A,B\}$. Observe that, in this case, the proportion of members expected to vote for each of the alternatives is exactly $\frac{1}{2}$. Note, also, that:

$$\widetilde{r}_{bias}^{\omega=A,\mathtt{p}}(A)=\widetilde{r}_{bias}^{\omega=B,\mathtt{p}}(B)=\sigma$$

$$\widetilde{r}_{bias}^{\omega=A,\mathtt{p}}(B)=\widetilde{r}_{bias}^{\omega=B,\mathtt{p}}(A)=0$$

$$\widetilde{r}_{bias}^{\omega=A,\mathtt{s}}(A)=\widetilde{r}_{bias}^{\omega=B,\mathtt{s}}(B)=\frac{\sigma}{n}\left(1+\mathbb{E}(\ \textstyle\sum_{j\neq i}\ \mathbb{I}_{\{v_j=\omega\}})\right)$$

$$\widetilde{r}_{bias}^{\omega=A,\mathtt{s}}(B)=\widetilde{r}_{bias}^{\omega=B,\mathtt{s}}(A)=\frac{\sigma}{n}\left(\mathbb{E}(\ \textstyle\sum_{j\neq i}\ \mathbb{I}_{\{v_j=\omega\}})\right)$$

Therefore, from equation (D.53), it follows that:

$$\underline{\gamma}_{bias}^{\lambda}\left(\alpha,\phi,\sigma,n\right)=\alpha+\frac{2^{n-1}\sigma}{\binom{n-1}{(n-1)/2}}\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathtt{s}\}}\right)\phi \tag{D.58}$$

Finally, note that:

$$0\leq\frac{(n-1)\,\sigma}{2+(n-3)\,\sigma}\leq 1,$$

since $n\geq 3$ and $\sigma\in(0,1)$; and

$$\frac{2^{n}\sigma}{\binom{n-1}{(n-1)/2}\left(1+\sigma\right)^{\frac{n+1}{2}}\left(1-\sigma\right)^{\frac{n-1}{2}}}>\frac{2^{n-1}\sigma}{\binom{n-1}{(n-1)/2}},$$

since $2>(1+\sigma)^{\frac{n+1}{2}}(1-\sigma)^{\frac{n-1}{2}}$.[23] Therefore, comparing equations (D.55), (D.57) and (D.58), we have:

$$\overline{\gamma}_{full}^{\lambda}\left(\alpha,\phi,\sigma,n\right)<\underline{\gamma}_{bias}^{\lambda}\left(\alpha,\phi,\sigma,n\right)<\overline{\gamma}_{part}^{\lambda}\left(\alpha,\phi,\sigma,n\right)$$

Furthermore, form the inspection of these expressions, it is immediate to see that:

$$\overline{\gamma}_{full}^{\mathtt{p}}\left(\alpha,\phi,\sigma,n\right)<\overline{\gamma}_{full}^{\mathtt{s}}\left(\alpha,\phi,\sigma,n\right)$$

$$\overline{\gamma}_{part}^{\mathtt{p}}\left(\alpha,\phi,\sigma,n\right)>\overline{\gamma}_{part}^{\mathtt{s}}\left(\alpha,\phi,\sigma,n\right)$$

and

$$\underline{\gamma}_{bias}^{\mathtt{p}}\left(\alpha,\phi,\sigma,n\right)>\underline{\gamma}_{bias}^{\mathtt{s}}\left(\alpha,\phi,\sigma,n\right)\ \blacksquare$$

D.10 Proposition 6

Note that if $\overline{\gamma}_{part}^{\mathtt{s}}\left(\alpha,\phi,\sigma,n\right)<\gamma<\overline{\gamma}_{part}^{\mathtt{p}}\left(\alpha,\phi,\sigma,n\right)$, then a partially-competent equilibrium can be sustained under public but not under secret voting. Furthermore, for this range of parameters, a biased equilibrium always exists under secret voting, but may or may not exist under public voting. Therefore, the probability of a correct decision under public voting is at least as large as under

[23]Note that $2>(1+\sigma)^{\frac{n+1}{2}}(1-\sigma)^{\frac{n-1}{2}}\leftrightarrow 2>(1+\sigma)(1+\sigma)^{\frac{n-1}{2}}(1-\sigma)^{\frac{n-1}{2}}\leftrightarrow 2^{\frac{2}{n-1}}>(1+\sigma)^{\frac{2}{n-1}}(1-\sigma^2)$. Observe that the last inequality always holds for any $n\geq 3$ and $\sigma\in(0,1)$, since $2^{\frac{2}{n-1}}>(1+\sigma)^{\frac{2}{n-1}}$ and $1-\sigma^2<1$.

secret voting, i.e.:

$$\Pi^{\mathsf{p}} = \min\left\{\sum_{i=(n+1)/2}^{n}\binom{n}{i}\left(\sigma+\tfrac{1}{2}\left(1-\sigma\right)\right)^{i}\left(\tfrac{1}{2}\left(1-\sigma\right)\right)^{n-i},\tfrac{1}{2}\right\} \geq \Pi^{\mathsf{s}} = \tfrac{1}{2}$$

Next, observe that if $\overline{\gamma}^{\mathsf{p}}_{full}\left(\alpha,\phi,\sigma,n\right) < \gamma < \overline{\gamma}^{\mathsf{s}}_{full}\left(\alpha,\phi,\sigma,n\right)$, then a fully-competent equilibrium can be sustained under secret but not under public voting. Note that for this range of parameters, a partially-competent equilibrium always exists under both secret and public voting. Thus, the probability of a correct decision under secret voting is at least as large as under public voting, i.e.:

$$\Pi^{\mathsf{s}} \min\left\{1-\tfrac{1}{2}\left(1-\sigma\right)^{n},\sum_{i=(n+1)/2}^{n}\binom{n}{i}\left(\sigma+\tfrac{1}{2}\left(1-\sigma\right)\right)^{i}\left(\tfrac{1}{2}\left(1-\sigma\right)\right)^{n-i}\right\}$$

$$\geq \Pi^{\mathsf{p}} = \sum_{i=(n+1)/2}^{n}\binom{n}{i}\left(\sigma+\tfrac{1}{2}\left(1-\sigma\right)\right)^{i}\left(\tfrac{1}{2}\left(1-\sigma\right)\right)^{n-i} \quad \blacksquare$$

D.11 Proposition 7

The derivatives of the equilibrium thresholds with ϕ are given by:

$$\frac{\partial\overline{\gamma}^{\lambda}_{full}\left(\alpha,\phi,\sigma,n\right)}{\partial\phi} = -\frac{\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathsf{s}\}}\right)}{\left(1+\frac{n-3}{2}\sigma\right)\left(1-\sigma\right)^{n-2}} < 0$$

$$\frac{\partial\overline{\gamma}^{\lambda}_{part}\left(\alpha,\phi,\sigma,n\right)}{\partial\phi} = \frac{2^{n}\sigma\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathsf{s}\}}\right)}{\binom{n-1}{(n-1)/2}\left(1+\sigma\right)^{\frac{n+1}{2}}\left(1-\sigma\right)^{\frac{n-1}{2}}} > 0$$

and

$$\frac{\partial\overline{\gamma}^{\lambda}_{bias}\left(\alpha,\phi,\sigma,n\right)}{\partial\phi} = \frac{2^{n-1}\sigma\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathsf{s}\}}\right)}{\binom{n-1}{(n-1)/2}} > 0$$

Next, the derivatives of the equilibrium thresholds with respect to σ are given by:

$$\frac{\partial\overline{\gamma}^{\lambda}_{full}\left(\alpha,\phi,\sigma,n\right)}{\partial\sigma} = \frac{2\left(n-1\right)\alpha}{\left(2+\left(n-3\right)\sigma\right)^{2}} - \frac{2\left(n-1\right)\left(1-\sigma\right)\left(1+\left(n-3\right)\sigma\right)\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathsf{s}\}}\right)\phi]}{\left(2+\left(n-3\right)\sigma\right)^{2}\left(1-\sigma\right)^{n}}$$

$$\frac{\partial\overline{\gamma}^{\lambda}_{part}\left(\alpha,\phi,\sigma,n\right)}{\partial\sigma} = \frac{2^{n}\left(1-\sigma\left(1-\left(n-1\right)\sigma\right)\right)\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathsf{s}\}}\right)\phi}{\binom{n-1}{(n-1)/2}\left(1+\sigma\right)^{\frac{n+3}{2}}\left(1-\sigma\right)^{\frac{n+1}{2}}},$$

where it is possible to show that for n large enough, we have $\frac{\partial\overline{\gamma}^{\lambda}_{full}(\alpha,\phi,\sigma,n)}{\partial\sigma} < 0$ and $\frac{\partial\overline{\gamma}^{\lambda}_{part}(\alpha,\phi,\sigma,n)}{\partial\sigma} > 0$. Furthermore, we have:

$$\frac{\partial\overline{\gamma}^{\lambda}_{bias}\left(\alpha,\phi,\sigma,n\right)}{\partial\sigma} = \frac{2^{n-1}\left(1-\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathsf{s}\}}\right)\phi}{\binom{n-1}{(n-1)/2}} > 0 \quad \blacksquare$$

D.12 Proposition 8

By the L'Hospital rule, it follows immediately that $\lim_{n\to\infty}\overline{\gamma}_{full}^{\lambda}(\cdot)=-\infty$ for $\lambda=$p, s. Furthermore, we can show that:

$$\lim_{n\to\infty}(1+\sigma)^{\frac{n+1}{2}}(1-\sigma)^{\frac{n-1}{2}}=0$$

and, using Stirling's approximation, $\binom{n-1}{(n-1)/2}\underset{n\to\infty}{\sim}\frac{2^{(n-1)}}{\sqrt{\pi\frac{n-1}{2}}}$, we get:

$$\lim_{n\to\infty}\frac{2^n\sigma\phi}{\binom{n-1}{(n-1)/2}}=+\infty \quad \text{and} \quad \lim_{n\to\infty}\frac{2^n\sigma\frac{\phi}{n}}{\binom{n-1}{(n-1)/2}}=0$$

Thus, it follows that $\lim_{n\to\infty}\overline{\gamma}_{part}^{\lambda}(\cdot)=-\infty$ for $\lambda=$p, s. Finally, we have that:

$$\lim_{n\to\infty}\frac{2^{n-1}\sigma\phi}{\binom{n-1}{(n-1)/2}}=+\infty \quad \text{and} \quad \lim_{n\to\infty}\frac{2^{n-1}\sigma\frac{\phi}{n}}{\binom{n-1}{(n-1)/2}}=0,$$

so that $\lim_{n\to\infty}\overline{\gamma}_{bias}^{\mathtt{p}}(\cdot)=+\infty$ and $\lim_{n\to\infty}\overline{\gamma}_{bias}^{\mathtt{s}}(\cdot)=\alpha$. ■

D.13 Proposition 9

See subsection B.2 of Online Appendix B.

D.14 Proposition 10

See subsection B.2 of Online Appendix B.

Appendix E. Model for Lab Experiment

Consider a committee of three members, $n=3$, with uniform prior, $q=\frac{1}{2}$, and symmetric distribution of both bias, $p=\frac{1}{2}$, and competence types, $\sigma=\frac{1}{2}$. Assume that the career-concern reward associated with a correct vote is exogenous and given by R^{λ} for $\lambda\in\{\mathtt{p},\mathtt{s}\}$.

E.1 Fully-Competent Equilibrium

Suppose that all committee members act in accordance with a fully-competent equilibrium and consider the behavior of a competent member biased against the state of the world. Note that for members of this type the expected utility of voting in accordance with the state of the world is:

$$U_{full}^{\beta,\lambda}(v_i=s_i,s_i\neq\beta)=\alpha+R^{\lambda},$$

while the expected utility of voting in accordance with his or her bias is:

$$U_{full}^{\beta,\lambda}(v_i=\beta,s_i\neq\beta)=\frac{1}{2}\alpha+\frac{1}{2}\gamma$$

Therefore, the condition for a competent member to always prefer to vote correctly in equilibrium is:

$$\alpha + R^{\lambda} \geq \frac{1}{2}\alpha + \frac{1}{2}\gamma \;\Rightarrow\; \gamma \leq \alpha + 2R^{\lambda} \tag{E.1}$$

Now, consider the behavior of an incompetent member. Observe that for members of this type the expected utility of abstaining is:

$$U^{\beta,\lambda}_{full}\left(v_i = \emptyset, s_i = \emptyset\right) = \frac{7}{8}\alpha + \frac{1}{2}\gamma,$$

while the expected utility of voting in accordance with his or her bias is:

$$U^{\beta,\lambda}_{full}\left(v_i = \beta, s_i = \emptyset\right) = \frac{3}{4}\alpha + \frac{3}{4}\gamma + \frac{1}{2}R^{\lambda}$$

Thus, the condition for an incompetent member to always prefer to abstain in equilibrium is:

$$\frac{7}{8}\alpha + \frac{1}{2}\gamma \geq \frac{3}{4}\alpha + \frac{3}{4}\gamma + \frac{1}{2}R^{\lambda} \;\Rightarrow\; \gamma \leq \frac{1}{2}\alpha - 2R^{\lambda} \tag{E.2}$$

Finally, note that the condition on incompetent members (E.2) is always harder to satisfy than condition on competent members (E.1), so that a fully-competent equilibrium can be sustained if, and only if:

$$\gamma \leq \frac{1}{2}\alpha - 2R^{\lambda} \tag{E.3}$$

E.2 Partially-Competent Equilibrium

Next, suppose that all committee members act in accordance with a partially-competent equilibrium and consider the behavior of a competent member biased against the state of the world. Note that for members of this type the expected utility of voting in accordance with the state of the world is:

$$U^{\beta,\lambda}_{part}\left(v_i = s_i, s_i \neq \beta\right) = \frac{15}{16}\alpha + \frac{1}{16}\gamma + R^{\lambda},$$

while the expected utility of voting in accordance with his or her bias is:

$$U^{\beta,\lambda}_{part}\left(v_i = \beta, s_i \neq \beta\right) = \frac{9}{16}\alpha + \frac{7}{16}\gamma$$

Therefore, the condition for a competent member to always prefer to vote correctly in equilibrium is:

$$\frac{15}{16}\alpha + \frac{1}{16}\gamma + R^{\lambda} \geq \frac{9}{16}\alpha + \frac{7}{16}\gamma \;\Rightarrow\; \gamma \leq \alpha + \frac{8}{3}R^{\lambda} \tag{E.4}$$

Now, consider the behavior of an incompetent member. Observe that for members of this type the expected utility of abstaining is:

$$U^{\beta,\lambda}_{part}\left(v_i = \emptyset, s_i = \emptyset\right) = \frac{3}{4}\alpha + \frac{1}{2}\gamma,$$

while the expected utility of voting in accordance with his or her bias is:

$$U_{part}^{\beta,\lambda}\left(v_i=\beta,s_i=\emptyset\right)=\frac{3}{4}\alpha+\frac{11}{16}\gamma+\frac{1}{2}R^{\lambda}$$

Thus, the condition for an incompetent member to prefer to vote in accordance with his bias rather than to abstain is given by:

$$\frac{3}{4}\alpha+\frac{11}{16}\gamma+\frac{1}{2}R^{\lambda}\geq\frac{3}{4}\alpha+\frac{1}{2}\gamma\ \Rightarrow\ \frac{3}{16}\gamma+\frac{1}{2}R^{\lambda}\geq 0 \tag{E.5}$$

Note that this condition is always satisfied, so that we can guarantee that incompetent members do not have any incentive to deviate from the equilibrium.

Therefore, it follows that a partially-competent equilibrium can be sustained if, and only if:

$$\gamma\leq\alpha+\frac{8}{3}R^{\lambda} \tag{E.6}$$

E.3 Biased Equilibrium

Finally, suppose that all committee members act in accordance with a biased equilibrium and consider the behavior of a competent member biased against the state of the world. Note that for members of this type the expected utility of voting in accordance with the state of the world is:

$$U_{bias}^{\beta,\lambda}\left(v_i=s_i,s_i\neq\beta\right)=\frac{3}{4}\alpha+\frac{1}{4}\gamma+R^{\lambda},$$

while the expected utility of voting in accordance with his or her bias is:

$$U_{bias}^{\beta,\lambda}\left(v_i=\beta,s_i\neq\beta\right)=\frac{1}{4}\alpha+\frac{3}{4}\gamma$$

Therefore, the condition for a competent member to always prefer to vote for his or her bias in equilibrium is:

$$\frac{3}{4}\alpha+\frac{1}{4}\gamma+R^{\lambda}\leq\frac{1}{4}\alpha+\frac{3}{4}\gamma\ \Rightarrow\ \gamma\geq\alpha+2R^{\lambda} \tag{E.7}$$

Next, consider the behavior of an incompetent member. Observe that for members of this type the expected utility of abstaining is:

$$U_{bias}^{\beta,\lambda}\left(v_i=\emptyset,s_i=\emptyset\right)=\frac{1}{2}\alpha+\frac{1}{2}\gamma,$$

while the expected utility of voting in accordance with his bias is:

$$U_{bias}^{\beta,\lambda}\left(v_i=\beta,s_i=\emptyset\right)=\frac{1}{2}\alpha+\frac{3}{4}\gamma+\frac{1}{2}R^{\lambda}$$

Thus, the condition for an incompetent member to prefer to vote in accordance with his bias rather than to abstain is given by:

$$\frac{1}{2}\alpha + \frac{3}{4}\gamma + \frac{1}{2}R^{\lambda} \geq \frac{1}{2}\alpha + \frac{1}{2}\gamma \;\Rightarrow\; \frac{1}{4}\gamma + \frac{1}{2}R^{\lambda} \geq 0 \tag{E.8}$$

Note that this condition is always satisfied, so that, consistently with Lemma 3, incompetent members do not have any incentive to deviate from the equilibrium in this case.

Therefore, it follows that a biased equilibrium can be sustained if, and only if:

$$\gamma \geq \alpha + 2R^{\lambda} \;\blacksquare \tag{E.9}$$

Appendix F. Experiment Instructions

This section presents the English version of the experiment instructions for treatments Low/Secret and Low/Public.[24] See Figures F.1 and F.2 for a depiction of the two main screens of the experiment.

Instructions

Thank you for your participation! The goal of this study is to investigates how people make decisions in group. You will be paid 2 euros for your presence. Your total earnings will depend partly on your decisions, partly on the decisions of other participants, and partly on chance. Your gains will be calculated in points and will be converted in euros at the rate of 1 euro per 80 points. You will be paid in cash at the end of the experiment.

During the experiment, you are not allowed to communicate with anyone. Please turn off your cell phone. If you have any question, please raise your hand.

This study is divided in 2 parts. We will begin by reading the instructions for the first part. Please, pay careful attention. After the instructions are read, there will be a short comprehension quiz.

First Part

This part consists of 32 rounds. The first two rounds are practice rounds and will not be paid. All other rounds are paid.

Groups. We begin every round by randomly dividing you into groups of three people. Every group receives one color: Blue or Yellow. In each round, your group's color may be Blue or Yellow with equal probability. The color of your group may be different from the colors of other groups and may change from one round to another. The computer will randomly choose your group's color in every round. Some people observe their group's color, while others do not.

[24] The full set of instructions in Italian is available upon request.

Votes. In each round, your group will choose one color by voting. Each member of the group may vote for Blue, vote for Yellow or abstain. Whichever color receives more votes is the group's choice. Ties are broken randomly by the computer. Examples:

i. If the number of votes for Blue is 2, the number of votes for Yellow is 1 and the number of abstentions is 0, then Blue is the group's choice;

ii. If the number of votes for Blue is 0, the number of votes for Yellow is 2 and the number of abstentions is 1, then Yellow is the group's choice;

iii. If the number of votes for Blue is 1, the number of votes for Yellow is 1 and the number of abstentions is 1, then we have a tie and the group's choice will be Blue or Yellow with equal probability;

iv. If all members of the group abstain, then we have a tie and the group's choice will be Blue or Yellow with equal probability.

Messages. Before voting, each of you will receive a message that may reveal the color of your group. There are three types of message.

1. The message says: "*The color of your group is Blue.*" In this case, you know for sure that your group's color is Blue.
2. The message says: "*The color of your group is Yellow.*" In this case, you know for sure that your group's color is Yellow.
3. The message says: "*The color of your group is Blue or Yellow with equal probability.*" In this case, the message does not provide any additional information with respect to what was already known.

Messages 1 and 2 are informative messages, while the third one is an uninformative message. In every round, half of the people in this room will receive an uninformative message, while the other half will receive an informative message and, therefore, will know exactly what is the color of their groups. For every group, there are four possible cases.

1. All members of the group know the group's color;
2. Two members of the group know the group's color while one member does not know;
3. One member of the group knows the group's color while two members do not know;
4. No member of the group knows the group's color.

Why does your vote matter? Your payoff in a given round depends on the choice made by your group, which is the color that receives the largest number of votes. If your group chooses the alternative that matches your group's color, then all members of the group receive 10 points; otherwise, everyone receives zero points.

Roles. Your payoff also involves an additional component that depends on your "role". In every round, the computer will randomly assign a role to each of you, which can be either Blue or Yellow. In every round, half of the people in this room will receive the Blue role and the other half will receive the Yellow role. Your role in a given round does not depend on the role of other members of your group nor on your role in previous rounds. For a given group, the number of members with the Blue role can be 3, 2, 1 or none. Your role is not known by anyone except you. If your group's choice is equal to your role, then you receive 1 extra point; otherwise, you receive no extra point.

Examples. Suppose that your role is Blue. The following table summarizes all possible payoffs in this case:

	Group's Color	Group's Choice	Group's Choice = Group's Color		Group's Choice = Role	Total Payoff
i	Blue	Blue	10	+	1	11
ii	Yellow	Yellow	10	+	0	10
iii	Yellow	Blue	0	+	1	1
iv	Blue	Yellow	0	+	0	0

The first line corresponds to the case where your group's color is Blue and your group's choice is Blue. In this case, your total payoff is 11 points: 10 points because your group's choice is equal to your group's color plus 1 extra point because your group's choice is equal to your role. In the second line, we have, instead, the case where your group's color is Yellow and your group's choice is Yellow. In this case, your total payoff is 10 points because your group's choice is equal to your group's color but not equal to your role. Next, in the third line, your group's color is Yellow and your group's choice is Blue. In this case, your total payoff is 1 point because your group's choice is equal to your role, but not equal to your group's color. Finally, in the fourth line, your payoff is zero, because your group's choice is neither equal to your group's color nor to your role.

Similarly, suppose that your role is Yellow. The following table summarizes all possible payoffs in this case:

	Group's Color	Group's Choice	Group's Choice = Group's Color		Group's Choice = Role	Total Payoff
i	Yellow	Yellow	10	+	1	11
ii	Blue	Blue	10	+	0	10
iii	Blue	Yellow	0	+	1	1
iv	Yellow	Blue	0	+	0	0

The first line corresponds to the case where your group's color is Yellow and your group's choice is Yellow. In this case, your total payoff is 11 points: 10 points because your group's choice is equal to your group's color plus 1 extra point because your group's choice is equal to your role. In the second line, we have, instead, the case where your group's color is Blue and your group's choice is Blue. In this case, your total payoff is 10 points because your group's choice is equal to your group's color but not equal to your role. Next, in the third line, your group's color is Blue and your group's choice is Yellow. In this case, your total payoff is 1 point because your group's choice is equal to your role, but not equal to your group's color. Finally, in the fourth line, your payoff is zero, because your group's choice is neither equal to your group's color nor to your role.

Summary. To conclude, please remember the following information.

- At the beginning of each round, you will see a screen with information about your message and your role.
- In every round, the number of members of your group who know the group's color can be 3, 2, 1 or none.
- In every round, the number of members of your group with the Blue role can be 3, 2, 1 or none.
- You can vote for Blue, vote for Yellow or abstain. Remember that the group's choice is taken by majority and that ties are broken randomly by the computer.
- After every round, you will be able to see what were your group's color and choice in that round. You will also receive information about your payoff and how many members of your group voted for Blue, voted Yellow and abstained.
- Your payoff in every round is determined by the sum of two components:

If your group's choice is equal to your group's color, then all members of the group earn 10 points. Otherwise, everyone gets zero points.

+

If your group's choice is equal to your role, then you earn 1 extra point. Otherwise, you get zero extra points.

- Remember that the decision of each group is independent of the decisions of other groups and that new groups are formed randomly in every round.

Second Part

The second part of the experiment is almost exactly the same as the first part, with a single difference. In the first part, your payoff depended on your group's choice, your group's color and your role. In this part of the experiment, your payoff will depend on your group's choice, your group's color, your role and on *how you vote*. In particular, if you vote for your group's color, you will now earn 9 extra points. Otherwise, if you vote for a color that is different than your group's color or if you abstain, you will earn zero extra points. For example, if you vote for Yellow and your group's color is Yellow, then you receive 9 extra points independently of what your group chooses. Remember that you still earn 10 points if your group's choice is equal to your group's color and 1 extra point if your group's choice is equal to your role.

Examples. Suppose that your role is Blue and that you voted for Blue. The following table summarizes all possible payoffs in this case:

	Group's Color	Group's Choice	Group's Choice = Group's Color		Group's Choice = Role		Vote = Group's Color	Total Payoff
i	Blue	Blue	10	+	1	+	9	20
ii	Yellow	Yellow	10	+	0	+	0	10
iii	Yellow	Blue	0	+	1	+	0	1
iv	Blue	Yellow	0	+	0	+	9	9

The first line corresponds to the case where your group's color is Blue and your group's choice is Blue. In this case, your total payoff is 20 points. You earn 10 points because your group's choice is equal to your group's color plus 1 extra point because your group's choice is equal to your role. These two components of your payoff are exactly the same as in the first part of the experiment, but now you also earn 9 extra points because you voted for your group's color. In the second line, we have, instead, the case where your group's color is Yellow and your group's choice is Yellow. In this case, your total payoff is 10 points because your group's choice is equal to your group's color, but not equal to your role, and you did not vote for your group's color. Next, in the third line, your group's color is Yellow and your group's choice is Blue. In this case, your total payoff is 1 point because your group's choice is equal to your role, but not equal to your group's color, and you did not vote for your group's color. Finally, in the fourth line, your payoff is 9, because you voted for your group's color, but your group's choice is neither equal to your group's color nor to your role.

Similarly, suppose that your role is Blue and that you voted for Yellow. The following table summarizes all possible payoffs in this case:

	Group's Color	Group's Choice	Group's Choice = Group's Color		Group's Choice = Role		Vote = Group's Color	Total Payoff
i	Blue	Blue	10	+	1	+	0	11
ii	Yellow	Yellow	10	+	0	+	9	19
iii	Yellow	Blue	0	+	1	+	9	10
iv	Blue	Yellow	0	+	0	+	0	0

The first line corresponds to the case where your group's color is Blue and your group's choice is Blue. In this case, your total payoff is 11 points, because your group's choice is equal to your group's color and to your role, but you did not vote for your group's color. In the second line, we have, instead, the case where your group's color is Yellow and your group's choice is Yellow. In this case, your total payoff is 19 points; $10 + 0$ points because your group's choice is equal to your group's color, but not equal to your role, plus 9 extra points because you voted for your group's color. Next, in the third line, your group's color is Yellow and your group's choice is Blue. In this case, your total payoff is 10 point because your group's choice is equal to your role, but not equal to your group's color, and you voted for your group's color. Finally, in the fourth line, your payoff is zero, because you did not vote for your group's color and your group's choice is neither equal to your group's color nor to your role.

Finally, suppose that your role is Blue and that you abstained. The following table summarizes all possible payoffs in this case:

	Group's Color	Group's Choice	Group's Choice = Group's Color		Group's Choice = Role		Vote = Group's Color	Total Payoff
i	Blue	Blue	10	+	1	+	0	11
ii	Yellow	Yellow	10	+	0	+	0	10
iii	Yellow	Blue	0	+	1	+	0	1
iv	Blue	Yellow	0	+	0	+	0	0

The first line corresponds to the case where your group's color is Blue and your group's choice is Blue. In this case, your total payoff is 11 points, because your group's choice is equal to your group's color and to your role, but you did not vote for your group's color. In the second line, we have, instead, the case where your group's color is Yellow and your group's choice is Yellow. In this case, your total payoff is 10 points, because your group's choice is equal to your group's color, but not equal to your role, and you did not vote for your group's color. Next, in the third line, your group's color is Yellow and your group's choice is Blue. In this case, your total payoff is 1 point because your group's choice is equal to your role, but not equal to your group's color, and you did not vote for your group's color. Finally, in the fourth line, your payoff is zero, because you did not vote for your group's color and your group's choice is neither equal to your group's color nor to your role.

In a similar way, you can calculate your payoffs in case your role is Yellow.

Summary. To conclude, please remember the following information.

- At the beginning of each round, you will see a screen with information about your message and your role.
- In every round, the number of members of your group who know the group's color can be 3, 2, 1 or none.
- In every round, the number of members of your group with the Blue role can be 3, 2, 1 or none.
- You can vote for Blue, vote for Yellow or abstain. Remember that the group's choice is taken by majority and that ties are broken randomly by the computer.
- After every round, you will be able to see what were your group's color and choice in that round. You will also receive information about your payoff and how many members of your group voted for Blue, voted Yellow and abstained.
- Your payoff in every round is determined by the sum of three components:

If your group's choice is equal to your group's color, then all members of the group earn 10 points. Otherwise, everyone gets zero points.

+

If your group's choice is equal to your role, then you earn 1 extra point. Otherwise, you get zero extra points.

+

If your vote is equal to your group's color, then you earn 9 extra points. Otherwise, you get zero extra points.

- Remember that the decision of each group is independent of the decisions of other groups and that new groups are formed randomly in every round.

References

[1] Battaglini, M., R. Morton and T. Palfrey. 2010. "The Swing Voter's Curse in the Laboratory." *Review of Economic Studies*, 77(1): 61-89.

[2] Canes-Wrone, B., M. Herron and K. Shotts. 2001. "Leadership and Pandering: A Theory of Executive Policymaking." *American Journal of Political Science*, 45(3): 532-50.

[3] Elbittar, A., A. Gomberg, C. Martinelli and T. Palfrey. 2014. "Ignorance and Bias in Collective Decisions." Working Paper, California Institute of Technology.

[4] Feddersen, T. and W. Pesendorfer. 1996. "The Swing Voter's Curse." *American Economic Review*, 86(3): 408-24.

[5] Grosser, J. and M. Seebauer. 2015. "The Curse of Uninformed Voting: An Experimental Study." Working Paper.

[6] Ibragimov, R. and U. Müller. 2010. "t-Statistic Based Correlation and Heterogeneity Robust Inference." *Journal of Business and Economic Statistics*, 28: 453-468.

[7] Levy, G. 2007. "Decision Making in Committees: Transparency, Reputation and Voting Rules." *American Economic Review*, 97(1): 150-68.

[8] Maskin, E. and J. Tirole. 2004. "The Politician and the Judge." *American Economic Review*, 94(4): 1034-54.

[9] Morris, S. 2001. "Political Correctness." *Journal of Political Economy*, 109(2): 231-65.

[10] Stasavage, D. 2007. "Polarization and Publicity: Rethinking the Benefits of Deliberative Democracy." *Journal of Politics*, 69(1): 59-72.

[11] Swank, O. H. and B. Visser. 2007. "On Committees of Experts." *Quarterly Journal of Economics*, 122(1): 337-72.

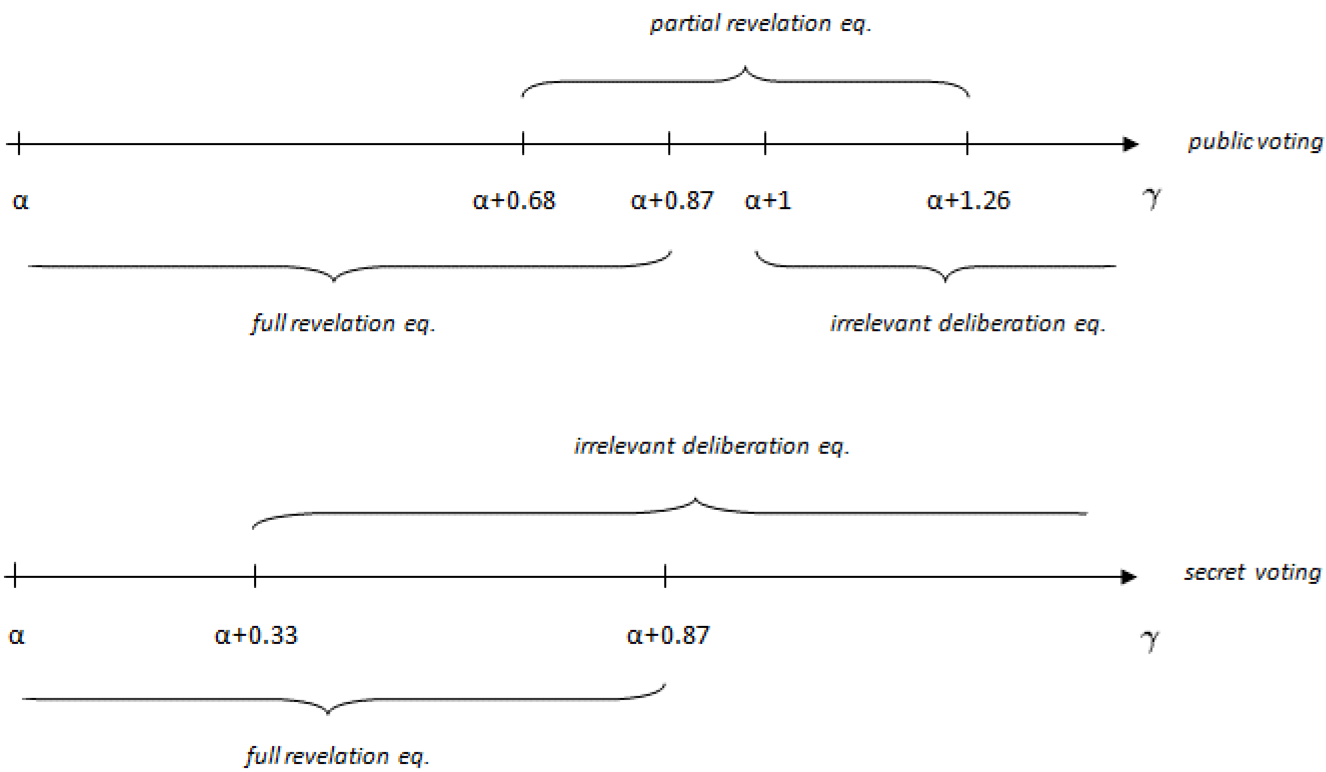

Notes. This figure illustrates the parameter regions where each class of equilibrium can be sustained under public and secret voting in a model with deliberation. The horizontal axis represents values of the parameter γ. The parameter values assumed for the construction of this graph were n=3, ϕ=1 and σ=0.5.

Figure B.1. Equilibria: Model with Deliberation

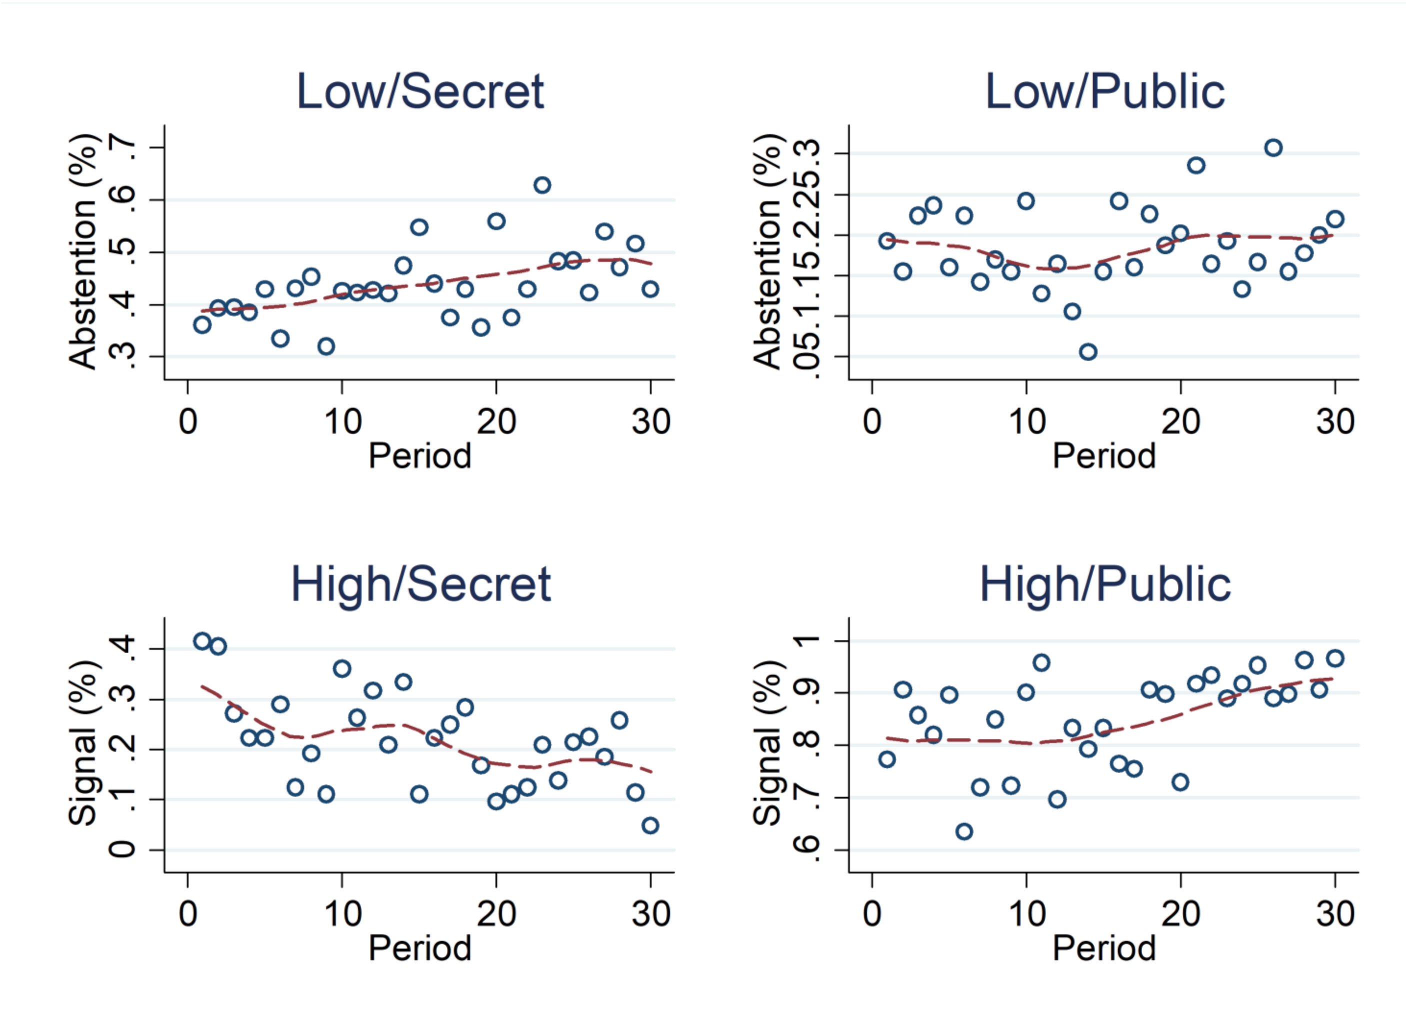

Notes. This figure plots the percentage of abstentions in Low/Secret and Low/Public treatments and the percentage of votes in accordance with the signal in High/Secret and High/Public treatments across all periods. The dashed lines represent local polynomial smooths.

Figure C.1. Learning across Treatments

Figure F.1. Voting Screen

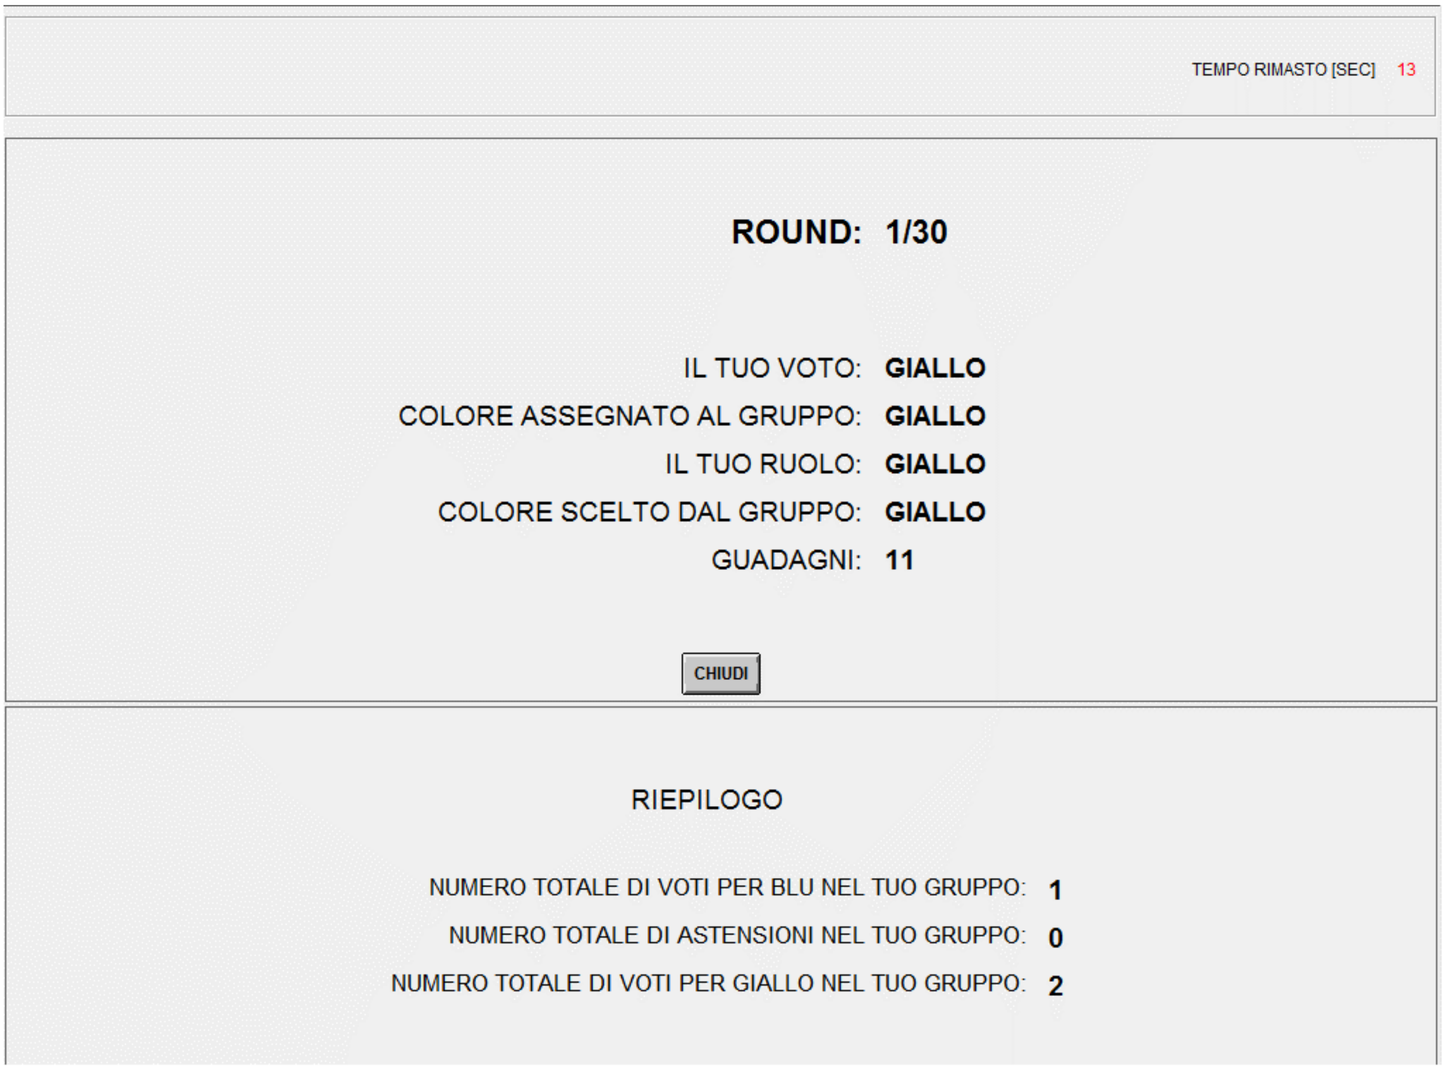

Figure F.2. Feedback Screen

Treatment	Periods	Obs	Uninformed Voters		
			Abstention (%)	Bias (%)	Against-Bias (%)
Low/Secret	1 – 10	360	39.17	50.56	10.28
	11 – 20	360	45.00	44.72	10.28
	21 – 30	360	48.33	43.33	8.33
Low/Public	1 – 10	360	19.44	65.28	15.28
	11 – 20	360	16.67	66.94	16.39
	21 – 30	360	20.83	62.22	16.94
High/Secret	1 – 10	360	11.11	85.83	3.06
	11 – 20	360	10.00	87.22	2.78
	21 – 30	360	6.94	90.83	2.22
High/Public	1 – 10	360	6.39	86.39	7.22
	11 – 20	360	6.11	82.50	11.39
	21 – 30	360	5.00	83.89	11.11

Table C.1. Learning Effects: Uninformed Subjects

Treatment	Periods	Obs	Informed Voters with Signal ≠ Bias		
			Signal (%)	Bias (%)	Abstention (%)
Low/Secret	1 – 10	180	94.44	1.67	3.89
	11 – 20	171	96.49	0.58	2.92
	21 – 30	169	97.04	2.37	0.59
Low/Public	1 – 10	178	96.07	3.93	0.00
	11 – 20	168	97.62	2.38	0.00
	21 – 30	178	99.44	0.56	0.00
High/Secret	1 – 10	170	25.88	61.18	12.94
	11 – 20	172	23.26	64.53	12.21
	21 – 30	175	16.57	64.57	18.86
High/Public	1 – 10	186	80.11	17.20	2.69
	11 – 20	192	81.25	14.06	4.69
	21 – 30	200	92.00	5.00	3.00

Table C.2. Learning Effects: Informed Subjects

Sequence	Treatment	Obs	Uninformed Voters		
			Abstention (%)	Bias (%)	Against-Bias (%)
Low/Secret – Low/Public	Low/Secret	810	47.65	43.46	8.89
	Low/Public	810	22.59	61.85	15.56
Low/Public – Low/Secret	Low/Secret	270	33.70	54.44	11.85
	Low/Public	270	8.15	73.70	18.15
High/Secret – High/Public	High/Secret	720	11.39	85.69	2.92
	High/Public	720	5.28	83.33	11.39
High/Public – High/Secret	High/Secret	360	5.28	92.50	2.22
	High/Public	360	6.94	86.11	6.94

Table C.3. Sequencing Effects: Uninformed Subjects

Sequence	Treatment	Obs	Informed Voters with Signal ≠ Bias		
			Signal (%)	Bias (%)	Abstention (%)
Low/Secret – Low/Public	Low/Secret	394	95.69	1.27	3.05
	Low/Public	390	97.69	2.31	0.00
Low/Public – Low/Secret	Low/Secret	126	96.83	2.38	0.79
	Low/Public	134	97.76	2.24	0.00
High/Secret – High/Public	High/Secret	341	13.49	70.38	16.13
	High/Public	395	82.28	14.43	3.29
High/Public – High/Secret	High/Secret	176	38.07	50.00	11.93
	High/Public	183	89.62	6.56	3.83

Table C.4. Sequencing Effects: Informed Subjects

Dependent Variable: Correct Vote	[1]	[2]	[3]	[4]
High/Secret	-0.627 ***	-0.618 ***		-0.731 ***
	[0.040]	[0.043]		[0.055]
High/Secret × Periods 1-10			-0.565 ***	
			[0.051]	
High/Secret × Periods 11-20			-0.604 ***	
			[0.053]	
High/Secret × Periods 21-30			-0.683 ***	
			[0.048]	
High/Secret × Low Performance in Comprehension Quiz				0.264 ***
				[0.080]
Individual Fixed-Effects	N	Y	Y	Y
Observations	1095	1095	1095	1095
R^2	0.39	0.55	0.55	0.56

Notes. This table reports OLS regressions in which the dependent variable is a dummy indicating whether the subject voted correctly. The sample is restricted to include only High/Secret and High/Public treatments and subject-period observations where the individual received a signal different than her bias. All standard errors are clustered at the individual level. ***, ** and * denote significance at 1%, 5% and 10%, respectively.

Table C.5. Regression Analysis: Informed Subjects

Dependent Variable: Abstention	[1]	[2]	[3]	[4]	[5]
Low/Secret	0.251 ***	0.243 ***		0.347 ***	0.096 *
	[0.040]	[0.041]		[0.073]	[0.041]
Low/Secret × Periods 1-10			0.205 ***		
			[0.041]		
Low/Secret × Periods 11-20			0.247 ***		
			[0.043]		
Low/Secret × Periods 21-30			0.276 ***		
			[0.046]		
Low/Secret × Low Performance in Comprehension Quiz				-0.164 *	
				[0.087]	
Low/Secret × N° of Abstentions in Periods 1-10					0.106 **
					[0.022]
Low/Secret × N° of Bad Abstentions in Periods 1-10					-0.139 *
					[0.081]
Individual Fixed-Effects	N	Y	Y	Y	Y
Observations	2160	2160	2160	2160	1440
R^2	0.07	0.55	0.55	0.56	0.63

Notes. This table reports OLS regressions in which the dependent variable is a dummy indicating whether the subject abstained. The sample is restricted to include only Low/Secret and Low/Public treatments and subject-period observations where the individual did not receive any information about the state of the world. The regression reported in column [5] further restricts the sample to include only observations from periods 11-30 of Low/Secret and Low/Public treatments. All standard errors are clustered at the individual level. ***, ** and * denote significance at 1%, 5% and 10%, respectively.

Table C.6. Regression Analysis: Uninformed Subjects