Public versus Secret Voting in Committees*

Andrea Mattozzi† Marcos Y. Nakaguma‡
EUI and MOVE University of Sao Paulo

January 2017

Abstract

This paper studies a committee decision-making problem. Committee members are heterogeneous in two private dimensions: their competence about what the correct alternative is, and their bias. Furthermore, they are career oriented and they can abstain. The interaction between career concern and bias affects the voting behavior of members depending on transparency of individual votes. We show that transparency 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 provide experimental evidence supporting our theoretical conclusions.

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

*We thank Veronica Rattini and Riccardo Ghidoni for excellent research assistance. We thank 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, Luigi Proietti, Debraj Ray, Aniol Llorente-Saguer, Arthur Schram, Justin Valasek and seminar audiences at Bologna, CORS-RWIO, 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
‡Department of Economics, USP, Sao Paulo, Brazil. Email: nakaguma@usp.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. 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.[1]

The issues confronted by committees are typically multi-faceted and complex, and may involve a variety of conflicts and personal interests. Consider the case of a company deciding whether to downsize a particular division, legislators voting on a constitutional reform that may be harmful to some of them, members of an academic department hiring committee with different views on hiring priorities or, more generally, committees of experts. For example most of the questions presented to the advisory committees of the U.S. Food and Drug Administration (FDA) related to specific products involve important conflict of interests. In addition, committee members are usually motivated by the desire to advance their own careers and, therefore, care about being perceived as competent decision makers. For example, the reputation for making correct decisions is crucial for the reappointment and career prospects of members of a company's board of directors, or for top bureaucrats.[2] Finally, since different committee members may have distinct abilities and competences, it is not unusual to observe situations where members abstain when unable to provide a firm answer to a particular question.[3]

This paper studies a committee decision-making problem that combines all elements discussed above. Specifically, in our model, committee members 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 the context of this model we study how the degree of transparency of individual votes affects equilibrium

[1] See Gerling et al [17] and Li and Suen [27] for reviews of this literature.

[2] See Wilson [44].

[3] For example, the advisory committees of the FDA adopt a simultaneous public voting system where members can vote "yes", "no" or "abstain". During a committee meeting charged with reviewing a new drug application, three out of thirteen members abstained from answering the following question: "*Do the clinical results of the single historically-controlled study 201/202 provide substantial evidence (...) that eteplirsen is effective for the treatment of Duchenne Muscular Dystrophy?*". One of the abstainers later explained his vote as follows: "*I voted to abstain and the reason is I was basically just torn between my mind and my heart. And I don't want to make type 1 error, and I don't make a type 2 error.*" In another vote, one member mentioned: "*I abstained only because I really just have never seen this product. (...) And I just don't know enough to say yes or no.*" Transcripts can be accessed from http://www.fda.gov/AdvisoryCommittees/

voting behavior and the quality of the decisions.

From a positive point of view, our goal is to understand how the incentives for an agent to abstain, vote for her bias or vote for the correct alternative depend on career concerns and the degree of transparency of individual votes. From a normative point of view, our main question is whether voting should be public or secret.

We consider a simple theoretical environment where agents are heterogenous in two private dimensions, competence and preferences. A decision over a binary agenda is taken by simple majority and committee members can vote for either alternative or abstain. The payoff of a member depends on three components: i) 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 agent.

Our analysis highlights that the interaction between career concerns and transparency leads to qualitatively different implications depending on the agent'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 either for their biases or for the ex-ante more likely alternative, 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 and always more inclined to indiscriminately follow their biases than competent members, since they always believe they might be right with some probability. We know that 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 competent members. This is the well-known swing voter's curse, first studied by Feddersen and Pesendorfer [11]. In the presence of career concerns, however, such behavior affects perceived competence negatively, since abstentions may signal incompetence in equilibrium. This creates an incentive for incompetent members to vote under transparency and, when they do so, they choose to vote either for their biases or for the ex-ante more likely alternative. Hence, while transparency attenuates the pre-existing biases of competent members, it may actually exacerbate the pre-existing biases of incompetent members.

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 mitigating 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 reduces the incentives for incompetent members to "gamble" and vote just in order to avoid revealing their lack of competence.

The distinction between public and secret voting is driven in our model by the fact that the effect of an agent's correct vote on his own reputation is diluted across all committee members under secrecy. Moreover, since this dilution effect is proportional to the size of the committee, it follows from our analysis that the choice between secret and public voting becomes more relevant as the committee increases.

Our analysis has implications for the design of committee decision-making rules. Our model emphasizes the idea that voting should be transparent in committees where members are highly subjected to the influence of ideological or self-interested motives. This is often the case of committees composed by politicians 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 charged with highly technical decisions such as top bureaucrats.[4]

More specifically, our analysis suggests that the level of transparency of a committee could be made contingent on the nature of the particular issue being discussed. For example, in hiring committees of academic departments − where members maybe biased towards hiring in their own areas − the influence of individual biases is likely to be larger when the search involves candidates of several fields rather than when it involves candidates of a one field only. In this respect, our model suggests that voting should be public in the former case and secret in the latter. Similarly, while most of the questions presented to the advisory committees of the FDA involve conflict of interest, there are also a number of meetings on broad scientic issues, not dealing with a specific product or class of branded products, for which a secret vote could lead to better decisions in terms of aggregation of information.[5]

Finally, some recent empirical findings by Mian et al [31] suggest an additional

[4] Alesina and Tabellini [1] and [2] study theoretically the optimal assignment of policy tasks to elected politicians or to non-elected bureaucrats. For an empirical analysis see Iaryczower et al [24].

[5] Distinguishing between different types of meetings maybe difficult but not unfeasible. For example, in order to analyze possible conflicts of interest in FDA Advisory Committees, Pham-Kanter [39] screens all transcripts and distinguishes between voting meetings where a conflict of interest issue was likely or not to arise.

possible implication of our results. The authors provide evidence that politicians and voters become more politically polarized in the aftermath of financial crises. In light of these findings, our results suggest that voting in committees should be transparent in relatively "bad times" when ideological biases tend to be exacerbated. Conversely, secret vote might perform better in relatively "good times" when ideological positions are less polarized.

We test the main theoretical predictions of the model by means of a 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.[6] Furthermore, as will become clear later in the paper, there are regions of the parameters space where our model features multiple equilibria with different properties. From this perspective, a controlled experiment can inform about whether individuals coordinate on certain equilibria and not on others.[7]

Our study contributes to the experimental literature on committee decision-making with career concerns and, to the best of our knowledge, is the first to combine in a single setting common value, biases, heterogeneous information, career concerns and the possibility of abstention. We consider a 2 by 2 design: low versus high bias and secret versus public voting. Consistently with our theoretical predictions, secret vote performs better (worse) than public voting in aggregating information with relatively low (high) bias. While half of the incompetent subjects abstains under secret vote and low bias, this proportion drops dramatically with public vote and it is almost zero in the case of high bias. Furthermore, our results in the secret low-bias treatment are in line with the results of the experimental literature on the swing voter's curse. When there are multiple equilibria, our results suggests that subjects tend to coordinate on the efficient equilibrium.

[6] For an alternative approach, which exploits a natural experiment related to the release of the Federal Open Market Committee (FOMC) transcripts, see Hansen et al [22], Meade and Stasavage [29] and Swank et al [43].

[7] An additional reason for investigating experimentally our theory - where individual behavior changes in response to transparency of actions - is whether our mechanism survives the existences of psychological costs of lying, which have been documented for experimental subjects (Gneezy [18]).

2 Literature Review

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 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 transparency of voting procedures on decision making in committees. To the best of our knowledge, however, none of the existing papers has investigated how competence, individual biases and career concerns interact in shaping individuals' voting behavior in a committee, and how this interaction is affected by transparency.

Gersbach and Hahn [15] and Levy [26] 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 [26] 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 [11].

Gersbach and Hahn [14] and Stasavage [41] 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 [9], Morris [33] and, more recently, Shapiro [40] 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.[8] 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 per se are punished.

In addition to these papers, Gersbach and Hahn [16] show that transparency induces agents to exert more effort in order to improve their chances of reappointment, Dal

[8] For single decision makers with career concerns see also Maskin and Tirole [28] and Prat [38].

Bo [7] and Felgenhauer and Gruner [10] argue that public voting makes the committee more vulnerable to the influence of special interest groups, and Swank and Visser [42] show that career concerns create an incentive for committees to conceal internal disagreements and show a united front in public. Finally, Midjord et al [32] point out that career concerns induces experts to be too conservative in order not to put their reputation at risk, and Gradwohl [19] 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, the most related paper to ours is Fehrler and Hughes [12]. As in our paper, they focus on the effect of transparency on committee decision making where agents are career concerned. Differently from our approach, however, their committee members are unbiased, committees are composed of two individuals, and the experimental focus is mostly on deliberation.[9] Also related is Battaglini et al [3] who provided the first test of the swing voters' curse in a laboratory setting.[10]

3 The Model

We consider a committee of $n \geq 3$ members, with n odd, that must decide between two alternatives, A and B. There are two states of the world, $\omega \in \{A, B\}$, with $\Pr(\omega = A) = q \in (0,1)$. While the true state is a priori unknown, committee members may receive an informative signal about it $s_i \in \{A, \emptyset, B\}$. An agent may be either competent, $\mathtt{c}$, in which case he receives a perfectly informative signal, or incompetent, $\mathtt{nc}$, in which case he receives an uninformative signal. We assume that each member knows his own competence type $\tau_i \in \{\mathtt{c},\mathtt{nc}\}$ and the distribution of other members' competences, which is given by $\Pr(\tau_i = \mathtt{c}) = \sigma \in (0,1)$. After observing their private signals, all members decide simultaneously whether to vote for A or B or to abstain, $v_i \in \{A, \emptyset, B\}$.[11] The final decision, $x \in \{A, B\}$, is determined by simple majority rule

[9] See also Morton and Ou [34] for an empirical investigation of whether secret voting leads to less prosocial voting behavior than public voting.

[10] See also Morton and Tyran [35] and [36] for related experiments and Herrera et al [23] for a theory on strategic abstention in proportional elections.

[11] In the model we abstract from any form of deliberation, that is we assume that agents cannot share their information. In Section 5.2, we discuss whether competent agents would actually have an incentive to reveal their information and how this decision could impact our basic comparative static results.

and ties are broken randomly.

Throughout the analysis we have assumed that the signals received by the members of the committee were private and that competent agents were not allowed to share their information with other players. In this subsection, we discuss whether competent agents would actually have an incentive to reveal their information and how this decision could impact our basic comparative static results.

The committee members care about making correct decisions and receive a common value $\alpha > 0$ whenever the final choice is equal to the state of the world, i.e. $x = \omega$. Additionally, we suppose that every member is biased towards one of the alternatives, i.e. each agent is biased towards either A or B. Every committee member knows his own bias type, $\beta_i \in \{A, B\}$, as well as the distribution of other agents' biases, $\Pr(\beta_i = A) = p \in (0,1)$, which we assume to be common knowledge. An agent biased towards β_i, receives an extra payoff $\gamma > 0$, irrespective of the state of the world, when alternative $x = \beta_i$ is chosen by the committee.[12]

The members of the committee are also concerned with building a reputation for competence and making correct decisions. We assume the existence of an additional agent, the external evaluator, whose task is simply to update his beliefs about the likelihood that each member is competent and voted correctly, conditional on the state of the world plus any other relevant information that might be available to him. We suppose that the state of the world is always revealed ex-post. Furthermore, under public voting, the evaluator is 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 an agent i is competent and voted correctly is, therefore, given by:

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

where ω is the state of the world, $\lambda \in \{\mathtt{p}, \mathtt{s}\}$ denotes whether voting is public or secret, and $\mathcal{I}^\lambda$ represents all relevant information available under λ. As equation (1) shows, we are assuming that a committee member competence is valued only if his vote is correct. This assumption greatly simplifies the analysis and the exposition. Furthermore, it proves particularly useful in the experimental implementation of our

[12] This model extends the setting studied by Nakaguma [37] to an asymmetric environment.

[13] Alternatively, we could have assumed that only the final decision of the committee was observed under secrecy. See the discussion in the online Appendix A about changes in voting rule and degree of transparency.

theory. As we detail in Section 5.1, our results are robust to using a more standard definition of competence, which is based only on the posterior probability that the agent is competent.[14]

Thus, 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:

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

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

4 Equilibrium Analysis

We solve the model for symmetric pure-strategy equilibria, where committee members of the same type (i.e., with the same bias and competence level) choose identical strategies. We also assume that agents 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 players and the external evaluator's beliefs. At the same time, the evaluator's beliefs must be consistent with the agents' 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 strategy, $v_i \in \{A, \emptyset, B\}$, induces a probability distribution over final outcomes, which is represented by the mapping $\rho_{\mu_i}^{\omega} : \{A, \emptyset, B\} \to [0,1]$, where $\rho_{\mu_i}^{\omega}(v_i)$ denotes the

[14] Under our assumption, a committee member receives zero reputation whenever he abstains or votes incorrectly under public voting. Intuitively, this assumes an external evaluator very tough on whoever says "I am not sure what to do" or who expresses blatantly wrong opinions. While it is not always the case that not taking a position is detrimental for expected competence, our assumption seems plausible in a variety of cases. For example, an expert who candidly reveals in public that he does not know what is the right policy to implement would most probably harm his reputation for competency.

probability, as perceived by the agent, that the committee's decision is A when the agent chooses v_i, given μ_i and ω. Observe that the probability $\rho^{\omega}_{\mu_i}(v_i)$ already takes into account all the uncertainty related to the realization of types of other committee members. Furthermore, it must be the case that:

$$\rho^{\omega}_{\mu_i}(B) \leq \rho^{\omega}_{\mu_i}(\emptyset) \leq \rho^{\omega}_{\mu_i}(A), \tag{3}$$

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 individual abstains) and, similarly, a vote for B can never increase the probability that the final outcome is A (relative to the case where he abstains).[15]

Next, let μ_e be the external evaluator's beliefs about the behavior of committee members. Under public voting, all individual votes are observable ex-post, so that career concern reward depend only on each member's own vote according to the following expression:

$$r^{\omega,\mathsf{p}}_{i,\mu_e} = \Pr_{\mu_e}(t = \mathsf{c}|v = \omega)\mathbb{I}_{\{v_i=\omega\}}, \tag{4}$$

where $\Pr_{\mu_e}(t = \mathsf{c}|v = \omega)$, is computed based on the external evaluators' beliefs about the behavior of voters and $\mathbb{I}_{\{v_i=\omega\}}$ is an indicator function that equals one when agent i votes correctly, $v_i = \omega$. Under secret voting, on the other hand, only the aggregate vote is observable ex-post, so that career concern rewards can be made contingent only on the total number of correct votes, $V^c \equiv \sum_i \mathbb{I}_{\{v_i=\omega\}}$, according to the following expression:

$$r^{\omega,\mathsf{s}}_{\mu_e} = \Pr_{\mu_e}(t = \mathsf{c}|v = \omega)\frac{V^c}{n}\ , \tag{5}$$

where V^c/n represents the probability that a particular agent voted correctly. Observe that the evaluator expects that each member is equally likely to have cast one of the V^c correct votes, given that all agents are ex-ante identical. Therefore, in this case, the career concern rewards are the same across all members and equal to the average expected competence in the committee.

In equilibrium, each committee member correctly anticipates the beliefs of the external evaluator and, before casting a vote, forms an expectation about the career concern reward that he will receive as a function of his strategy. Suppose, first, that the state of the world is observed by the agent. Under public voting, each agent can

[15]The inequalities are weak since there may be situations where the committee member is not expected to be pivotal.

perfectly anticipate his career concern reward in equilibrium:

$$\widetilde{r}^{\omega,\mathsf{p}}\left(v_i\right) = \Pr(t = \mathsf{c}|v = \omega)\mathbb{I}_{\{v_i=\omega\}}, \tag{6}$$

where we omit the index for the evaluator's beliefs for simplicity. Under secret voting, expected career concern reward depends also on how each agent expects other members to vote:

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

where $\mathbb{E}(\sum_{j\neq i}\mathbb{I}_{\{v_j=\omega\}})$ is the number of correct votes expected to be cast by the other committee members. Hence, under secret voting, the impact of an agent's correct vote on his own career concern is diluted in proportion to the size of the committee. When the state of the world is not observed as it is the case, each agent must compute his expected reward by averaging his career concern under each state:

$$\widetilde{r}^{\lambda}\left(v_i\right) = q\widetilde{r}^{\omega=A,\lambda}\left(v_i\right) + \left(1-q\right)\widetilde{r}^{\omega=B,\lambda}\left(v_i\right). \tag{8}$$

Based on the elements defined above, and assuming that the state of the world is A, the expected utility of a competent member can be expressed as a function of his vote v_i as follows:

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

and

$$U^{\beta_i=B,\lambda}(v_i, s_i = A) = \phi\widetilde{r}^{\omega=A,\lambda}(v_i) + \rho^{\omega=A}(v_i)\alpha + (1-\rho^{\omega=A}(v_i))\gamma, \tag{10}$$

depending on whether the agent is biased towards A or B, respectively. Similar expressions can be derived for the case where the state of the world is B. The next lemma provides a general characterization of the behavior of competent members.[16]

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

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

[16] All proofs can be found in Online Appendix C.

b. Abstaining is a weakly dominated strategy for a competent member whose bias is different than the 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 [11]), i.e. the risk of unwillingly shifting the committee's decision away from the correct outcome. Therefore, there is no reason for them to abstain, since by voting for either alternative they can push the decision towards a particular outcome and abstentions are associated with lack of competence. Lemma 1 also implies that a competent member who receives a signal equal to his bias, $s_i = \beta_i$, always prefers (weakly) to vote in accordance with the state of the world, given that both common and private interests are aligned in this case, while a competent member who receives a signal different than his bias, $s_i \neq \beta_i$, may either vote for the state of the world or in accordance with his bias. Note that the above result guarantees that, in any equilibrium, every competent members who is biased towards the state of the world votes correctly. Thus, by Bayes rule, the likelihood that an agent is competent given that he voted correctly is strictly positive, $\Pr(t = \mathsf{c} | v = \omega) > 0$. The next lemma follows as a direct implication of this result.

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

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

Furthermore, we have that:

$$\widetilde{r}^{\omega,\mathtt{p}}(v_i = \omega) > \widetilde{r}^{\omega,\mathtt{s}}(v_i = \omega)$$

and

$$\widetilde{r}^{\omega,\mathtt{p}}(v_i \neq \omega) < \widetilde{r}^{\omega,\mathtt{s}}(v_i \neq \omega)$$

Interestingly, conditional on a correct vote, the expected career concern reward is larger under public than under secret voting, whereas the opposite is the case conditional on an incorrect vote or an abstention. Intuitively, this result follows from the fact that under secrecy career concern rewards are distributed equally across members and depend only on the total number of correct votes. The next lemma characterizes the equilibrium behavior of incompetent members relative to competent ones.

Lemma 3. *There exists no equilibrium in which a competent member who receives a signal different than his bias votes against the state of the world and an incompetent member abstains.*

Intuitively, incompetent agents are relatively more inclined to follow their "biases" by either voting for the ex-ante more likely alternative or for the alternative that matches their bias types. When a competent agent decides to vote against his signal, he knows for sure that he is casting an incorrect vote, while an incompetent agent always attributes positive probability to the event that his vote is correct, in which case he obtains larger career concern rewards. That is, incompetent agents are "naively" optimistic that their vote will coincide with the state of the world, which makes them more willing to vote even without having any information.

Finally, based on the above results, it is possible to show that there are only three types of equilibria in the model.

Proposition 1. *The 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 incompetent members abstain;

iii. A biased equilibrium, where at least some competent members vote against their signals and all incompetent members vote either to the ex-ante more likely alternative or in accordance with their biases.

Note that this characterization holds under both public and secret voting, any value of the prior and any distribution of types. However, the region of the parameters where each class of equilibrium can be sustained do depend on the transparency of the voting rule, as we shall discuss in detail in the next subsection.

4.2 Main Comparative Statics Results

In this subsection, we provide a characterization of each type of equilibrium under secrete and public voting. Let the subscript $\mu = \{full, part, bias\}$ denote equilibrium beliefs of all agents. The following proposition summarizes the main properties of the fully competent equilibrium.

Proposition 2. *A fully competent equilibrium can be sustained, if and only if*

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

Furthermore, if a fully competent equilibrium can be supported under public voting, then it can also be supported under secret voting.

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 be supported under secret voting. Intuitively, the interaction between transparency and career concerns creates an incentive for incompetent members to vote, since abstaining perfectly reveals their lack of competence in this case.

The next proposition provides a general characterization of the partially competent equilibrium.

Proposition 3. *A partially competent equilibrium can be sustained, if and only if*

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

where $\underline{\gamma}^{\lambda}_{part}\left(\alpha, \phi, \sigma, n\right) < \alpha$ *and* $\overline{\gamma}^{\lambda}_{part}\left(\alpha, \phi, \sigma, n\right) > \alpha$. *Furthermore, if a partially competent equilibrium can be supported under secret voting, then it can also be supported under public voting.*

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

Finally, the next proposition summarizes the main properties of the biased equilibrium.

Proposition 4. *A biased equilibrium can be sustained, if and only if*

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

Furthermore, if a biased equilibrium can be supported under public voting, then it can also be supported under secret voting.

A biased equilibrium is more likely to be sustained under secret voting since secrecy reduces the career concern reward associated with a correct vote, and makes competent members more willing to disregard their information about the state of the world and vote in accordance with their biases.

Overall, our analysis highlights the fact that transparency affects the behavior of competent and incompetent agents in markedly different ways. On the one hand, transparency *attenuates* the preexisting biases of competent members by inducing them to vote correctly, even when 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 either for the ex-ante more likely alternative or in accordance with their biases to avoid revealing their lack of competence.

4.3 The Symmetric Case

In this subsection, we provide a precise characterization of the equilibria by assuming that both the prior probability and the distribution of biases are symmetric, i.e. $q = p = 1/2$. The symmetric prior assumption implies that, when an incompetent member decides to vote, he will always vote for the alternative towards which he is biased, while the uniform distribution of biases simplifies the analysis by making the equilibrium behavior of incompetent members symmetric between agents of different bias types. Under these assumptions, we derive closed forms for the thresholds defined above. The following proposition characterizes the structure of the equilibria under both public and secret voting.

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

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

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

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

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

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

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

Furthermore, $\overline{\gamma}_{full}^{\lambda}(\alpha, \phi, \sigma, n) < \underline{\gamma}_{bias}^{\lambda}(\alpha, \phi, \sigma, n) < \overline{\gamma}_{part}^{\lambda}(\alpha, \phi, \sigma, n)$, $\overline{\gamma}_{full}^{\mathtt{p}}(\alpha, \phi, \sigma, n) < \overline{\gamma}_{full}^{\mathtt{s}}(\alpha, \phi, \sigma, n)$, $\overline{\gamma}_{part}^{\mathtt{p}}(\alpha, \phi, \sigma, n) > \overline{\gamma}_{part}^{\mathtt{s}}(\alpha, \phi, \sigma, n)$, *and* $\underline{\gamma}_{bias}^{\mathtt{p}}(\alpha, \phi, \sigma, n) > \underline{\gamma}_{bias}^{\mathtt{s}}(\alpha, \phi, \sigma, n)$.

The term $\frac{n-1}{n}\mathbb{I}_{\{\lambda=\mathtt{s}\}}$ which appears inside parenthesis in the above expressions captures the effect of the dilution of career concern under secret voting. Hence, a change from public to secret voting is qualitatively equivalent to a reduction in the weight attached to career concerns. Figure 1 shows the values of the parameters α and γ for which each class of equilibria can be sustained, given a level of transparency λ, and for fixed values of ϕ, σ and n.

Observe that since $\overline{\gamma}_{full}^{\lambda} < \overline{\gamma}_{part}^{\lambda}$, the region of parameters where a fully competent equilibrium exists is contained inside the region where a partially competent equilibrium can be supported. 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 is only optimal for an incompetent member if he expects other incompetent members to abstain as well. If, on the other hand, he expects other incompetent members to vote for their biases, then it becomes optimal for him to also do so.

Similarly, since $\underline{\gamma}_{bias}^{\lambda} < \overline{\gamma}_{part}^{\lambda}$, there exists a region of parameters where both a partially competent and a biased equilibria 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 equilibrium, either all of them vote correctly or all of them vote in accordance with their biases.

Figure 2 summarizes the main comparative static results of the model. Observe that in region I, where $\overline{\gamma}_{part}^{\mathtt{s}} < \gamma < \overline{\gamma}_{part}^{\mathtt{p}}$, a partially competent equilibrium can be sustained under public but not under secret voting; while in region II, where $\overline{\gamma}_{full}^{\mathtt{p}} < \gamma < \overline{\gamma}_{full}^{\mathtt{s}}$, a fully competent equilibrium can be sustained under secret but not under public voting. Intuitively, when the magnitude of the bias is relatively large, like in region I, incompetent members always vote in accordance with their biases, but public voting may actually induce competent members to vote correctly rather than to follow their biases since this increases the career concern gain associated with a correct vote. On the other hand, when the magnitude of the bias is relatively small, like in region II, competent members always vote correctly, but secret voting may help incompetent agents to abstain rather than to vote for their biases by reducing the expected career concern gain associated with voting.

For each class of equilibrium, it can be shown that the probability of a correct decision is given by

$$\Pi_{full} = 1 - \tfrac{1}{2}(1-\sigma)^n \tag{11}$$

$$\Pi_{part} = \textstyle\sum_{i=(n+1)/2}^{n} \binom{n}{i} \left(\sigma + \tfrac{1}{2}(1-\sigma)\right)^i \left(\tfrac{1}{2}(1-\sigma)\right)^{n-i} \tag{12}$$

and

$$\Pi_{bias} = \tfrac{1}{2}, \tag{13}$$

with $\Pi_{full} > \Pi_{part} > \Pi_{bias}$. Observe that the likelihood of a correct decision is lower than one even under a fully competent equilibrium, given that with probability $(1-\sigma)^n$ all committee members are incompetent, in which case the correct alternative would be chosen only half of the time. It is also interesting to note that the expected difference in the quality of decisions between a fully competent and a partially competent equilibrium increases with n, provided that the proportion of competent members is small enough. Intuitively, the theoretical difference between the two classes of equilibria is expected to be particularly pronounced whenever there is a large proportion of incompetent agents in the committee. Given these results, it is possible to rank public and secret voting in terms of the quality of decisions expected under each of them.

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.

Thus, it follows that, when the magnitude of the bias is relatively large, a correct decision is more likely under public voting; while when the magnitude of the bias is relatively small, a correct decision is more likely under secret voting. Note that the possible existence of multiple equilibria in both of the regions considered above prevents us from ordering transparency and secrecy in strict terms, given that it is not possible to guarantee that a change from public to secret voting, or vice-versa, will necessarily lead to a change in the class of equilibrium that ultimately prevails. In light of this, a controlled laboratory experiment is a particularly useful tool that can inform on whether individuals coordinate on certain equilibria.

Finally, in the next proposition, we show that the region of parameters where it is possible to sustain different classes of equilibria under public and secret voting becomes larger as both the relevance of career concerns and the proportion of competent agents increase.

Proposition 7. *Suppose that* $q = p = 1/2$. *Then the distance* $\overline{\gamma}^{\mathsf{p}}_{part}(\alpha, \phi, \sigma, n) - \overline{\gamma}^{\mathsf{s}}_{part}(\alpha, \phi, \sigma, n)$ *and the distance* $\overline{\gamma}^{\mathsf{s}}_{full}(\alpha, \phi, \sigma, n) - \overline{\gamma}^{\mathsf{p}}_{full}(\alpha, \phi, \sigma, n)$ *are increasing in* ϕ *and* σ.

Therefore, the more career oriented are the members of the committee and the larger the proportion of competent agents, the larger is the region of parameters where the choice between secret and public voting is expected to matter.[17] Finally, it is possible to show that these regions become arbitrarily large as the number of committee members goes to infinity, implying that our conclusions become possibly even more relevant for large committees.

5 Discussion

Throughout our analysis we have made a number of simplifying assumptions that deserve to be discussed. In Section 5.1, we present a detailed analysis of the robustness of our findings to our modelling of career concerns. In Section 5.2, we examine the implications of allowing for information sharing prior to the voting stage and we consider the incentives of different types of agents to choose between secret and public

[17] Intuitively, as σ increases, the career concern gains associated with a correct vote under a partially competent equilibrium increase, which generates an even stronger incentive for competent members to vote correctly under public voting relatively to secret voting. At the same time, as σ gets larger, the probability that an uninformed agent is pivotal when he decides to cast a vote under a fully competent equilibrium decreases, which diminishes the risk of the "swing voter's curse", thus increasing even more the incentive for incompetent members to vote under public voting relatively to secret voting.

voting. We cover other generalizations and extensions of the basic model in the Online Appendix A.[18]

5.1 Career Concern Rewards

We have assumed in our basic model that the career concern reward of a committee member is proportional to the conditional probability that the agent is competent *and* voted correctly (see equation [1]). However, our main qualitative results would remain the same even if we allow for the career concern reward to be based only on the posterior probability that agent i is competent, $r_i^{\omega,\lambda} \equiv \Pr(\tau_i = \mathsf{c}|\omega, \mathcal{I}^\lambda)$. In particular, both fully competent and partially competent equilibria would still be characterized by Propositions 2 and 3, respectively, and all comparative static results regarding these two types of equilibria would remain unchanged. The intuition is that in both cases the career concern reward associated with a correct vote is strictly larger than that associated with an abstention or an incorrect vote, since all competent members vote correctly in equilibrium.[19] It is in this sense that we can say that our basic conclusion that transparency attenuates the biases of competent members while it exacerbates the biases of incompetent members is robust to how career concern is defined.

The main implication of relaxing the assumption that career concern materializes only in connection with a correct vote is that it is now possible to sustain a larger set equilibria than those described in Proposition 1. In particular, we may also have equilibria involving the following "new" behaviors: (i) competent members with biases that are consistent with the state of the world voting against the state of the world and (ii) competent members abstaining. The next proposition provides a characterization of some basic aspects of these equilibria.

Proposition 8. *Assume that the career concern rewards depend only on the posterior probability that the agent is competent,* $r_i^{\omega,\lambda} \equiv \Pr(\tau_i = c|\omega, \mathcal{I}^\lambda)$*, then we have:*

i. An equilibrium in which a competent member with bias equal to the state votes

[18] In the Online Appendix A we discuss the assumption that the state of the world is observed ex-post and study the case in which competent and incompetent members receive signals of different precision. We elaborate on changes in the voting rule and in the assumption about what is revealed ex-post under secret voting. Finally, we show that the model can be easily extended to the existence of unbiased agents and to possible correlations between competence and bias.

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

against the state can be sustained only if the career concern reward associated with an incorrect vote is strictly larger than that associated with a correct vote.

ii. An equilibrium in which a competent member with bias equal to the state abstains can be sustained only if the career concern reward associated with an abstention is strictly larger than that associated with a correct vote.

An equilibrium involving a competent member with bias equal to the state either abstaining or voting incorrectly requires a very particular structure of incentives, namely: an agent 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, regardless of the correctness or incorrectness of the vote, may actually happen provided that career concerns are large enough. We believe that this element is not likely to be dominant in most applications of our model and this is one reason why our initial assumption that career concern is related to the joint probability that an agent is competent and voted correctly may be viewed as a reasonable form of refinement.[20] Still, even if we do not take these issues into account, it is possible to show that the equilibria discussed above can only exist in certain specific regions of the parameter space.

Proposition 9. *An equilibrium in which a competent member with bias equal to the state either abstains or votes against the state can be sustained only if the sum of the common value and the bias term, $\alpha + \gamma$, is small enough.*

Intuitively, since in this case the bias and the state of the world are aligned, voting for the state would increase the likelihood that the agent gets a payoff of $\alpha + \gamma$. Therefore, for such agent 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.

Next, we define that beliefs are *monotone* if the evaluator's beliefs are such that $\Pr(t = \mathsf{c} | v = \omega, \omega) \geq \Pr(t = \mathsf{c} | v \neq \omega, \omega)$ for any ω. Note that this condition implies that

[20] Incidentally, in a different model where committee members have incentive to signal both that they are competent and relatively unbiased, it would be reasonable to expect the existence of equilibria where abstentions are associated with relatively large career concern rewards. Note that a situation like that makes less sense in the context of our model, because here career concern depends solely on competence. A formal analysis of this other version of the model is beyond the scope of the present paper and is left for future research.

the career concern reward associated with a correct vote is not strictly smaller than that associated with an abstention or an incorrect vote, *i.e.* $\widetilde{r}_i^{\omega,\lambda}(v_i = \omega) \geq \widetilde{r}_i^{\omega,\lambda}(v_i \neq \omega)$ *for* $\omega \in \{A, B\}$. Proceeding with our analysis, the following proposition provides a characterization of the main properties of the equilibrium where a competent member biased against the state of the world abstains.

Proposition 10. *Assume that* $r_i^{\omega,\lambda} \equiv \Pr(\tau_i = c|\omega, \mathcal{I}^\lambda)$*, then we have:*

i. An equilibrium in which a competent member biased against the state abstains can be sustained only if $\gamma - \alpha$ *is strictly positive and small enough.*

ii. If in equilibrium a competent member biased against the state abstains, then a competent member with bias equal to the state can never vote against the state.

iii. Any equilibrium with monotone beliefs where a competent member biased against the state abstains can be sustained only if:

$$\alpha < \underline{\gamma}^\lambda_{abst}(\alpha, \phi, \sigma, n) \leq \gamma \leq \overline{\gamma}^\lambda_{abst}(\alpha, \phi, \sigma, n)$$

Furthermore, we have that:

$$\underline{\gamma}^{\mathtt{s}}_{abst}(\alpha, \phi, \sigma, n) \leq \underline{\gamma}^{\mathtt{p}}_{abst}(\alpha, \phi, \sigma, n)$$

and

$$\overline{\gamma}^{\mathtt{s}}_{abst}(\alpha, \phi, \sigma, n) \leq \overline{\gamma}^{\mathtt{p}}_{abst}(\alpha, \phi, \sigma, n).$$

There are several interesting facts contained in the above proposition.

First, part (i) emphasizes that equilibria where competent members biased against the state of the world abstain are not pervasive. In particular, they can only exist if agents are somewhat indifferent between voting for the correct alternative and following their biases.

Second, we can provide a sharper characterization of the equilibrium by focusing exclusively on equilibria with monotone beliefs. Indeed, as part (iii) of the proposition shows, an equilibrium where a competent member biased against the state abstains can only be sustained if the bias term is larger than the common value. Hence, combining parts (i) and (iii) of Proposition 10, we have that under the monotone beliefs assumption the bias term must be larger than the common value, but not too large, so that

$\gamma - \alpha$ is small enough. Thus, equilibria where competent members biased against the state abstain exist in a "small" subset of the parameters space case in the sense that γ can neither be too small nor too large relatively to α, for otherwise agents would have an incentive to vote correctly or to follow their biases, respectively. The fact that we must have $\gamma > \alpha$ for the equilibrium to be sustained is important, because it guarantees that the region of parameters where an equilibrium involving a competent member abstaining can never overlap with the region where a fully competent equilibrium exists. Therefore, our result that secret voting leads to better decisions when the magnitude of the bias is small relative to the common value is not affected in any way by the possible existence of multiple equilibria in that region.

Third, while equilibria where competent members abstain can be supported in the same region where a partially competent equilibrium exists, part (*iii*) of Proposition 10 also shows that public voting always leads competent members to behave "better". Specifically, they are both more likely to abstain rather than to vote incorrectly, given that $\overline{\gamma}^{\mathrm{s}}_{abst}(\alpha, \phi, \sigma, n) \leq \overline{\gamma}^{\mathrm{p}}_{abst}(\alpha, \phi, \sigma, n)$, and more likely to vote correctly rather than to abstain, since $\underline{\gamma}^{\mathrm{s}}_{abst}(\alpha, \phi, \sigma, n) \leq \underline{\gamma}^{\mathrm{p}}_{abst}(\alpha, \phi, \sigma, n)$. Note, however, that this result refers only to the behavior of competent agents, as it is not possible to guarantee that incompetent agents will behave better as well. Intuitively, there may now exist a region of parameters with $\gamma > \alpha$, where it is possible to support an equilibrium where both competent members biased against the state of the world and some incompetent members abstain. In this case, a move from secrecy to transparency could lead both competent members to vote correctly and incompetent members to vote for their biases. However, there is a sense in which such equilibria are difficult to be supported in that they require a very particular set of conditions to hold. For example, in the symmetric case discussed in Subsection 4.3, under the same parameter values used to construct Figure 1, one can show that there exists no equilibrium with monotone beliefs where a competent member abstains.[21]

Finally, to complement these results, we can also show that if beliefs are monotone, then a biased equilibrium is still characterized by the same properties stated in Proposition 4, and it is still the case that such equilibrium is less likely to be sustained under public voting.

[21] Specifically, it is possible to show that an equilibrium where competent members biased against the state abstain and, likewise, all incompetent members abstain can only be supported, in the symmetric case, if the proportion of competent members, σ, is very large.

5.2 Information Sharing

Throughout the analysis we have assumed that the signals received by the members of the committee were private and that competent agents were not allowed to share their information with other players. In this subsection, we discuss whether competent agents would actually have an incentive to reveal their information and how this decision could impact our basic comparative static results. In a setting where the members' interests are aligned, Coughlan [6] showed that voters would have strong incentives to share information, since this can only lead to a larger probability that the right decision is taken. However, the direction of incentives in our setting is not so clear-cut given the presence of biases and career concerns. For instance, competent members may prefer not to reveal their private information in order to separate themselves from incompetent agents. Moreover, a competent member may be particularly unwilling to share information if he is biased against the state of the world, since revealing information in this case could lead to the correct decision being taken with higher likelihood.

Let us consider a version of the basic model where we introduce a "mechanism" that collects all private signals and reveals them truthfully to the committee before the voting stage.[22] Note that, in this case, all members become fully informed about the state of the world whenever there is at least one competent agent in the group. Furthermore, it is possible to show that, if all members are informed, then there can be only two symmetric equilibria: one in which all members vote in accordance with the state of the world and another one in which all members vote for their biases. In particular, we can show that the equilibrium where all vote correctly always exists, whereas the equilibrium where all vote for their biases can only be sustained if the size of the bias is large relatively to the common value.[23] Naturally, there is no incentive for anyone to abstain in this case.

Would competent members actually have incentive to voluntarily participate in the mechanism described above? Note that career concern rewards of competent agents are significantly diluted under the mechanism, since information sharing prevents them distinguishing from the incompetent agents. In particular, the external evaluator now applies an extra discount to the career concern reward assigned to any correct vote

[22] For a general model of committee decision making with deliberation, see Gerardi and Yariv [13].

[23] Observe that if an informed agent expects all other members to vote correctly, then he is never pivotal and better off by also voting correctly, since by doing so he guarantees himself larger career concern rewards. Therefore, the equilibrium where all vote in accordance with the state of the world can be sustained for all possible parameter values.

in order to account for the fact that incompetent members may also learn the state of the world. It then follows that the willingness of competent members to take part in the mechanism should be especially low if voting is public, since the losses caused by the dilution effect are larger in this case. Similarly, they are less likely to share information when the size of the committee is large and when the importance attached to career concerns is high. On the other hand, competent members are more likely to participate in the mechanism if the common value is high relatively to the bias, given that information sharing is expected to lead to better decisions in this case.

Thus, from a normative point of view, it follows that if all members are expected to vote correctly after information is collected and shared, then secret voting is more likely to lead to better decisions, since it makes competent agents more willing to participate in the mechanism ex-ante. Alternatively, if the members of the committee are expected to vote in accordance with their biases even after information about the state is revealed, then the quality of the decisions cannot be improved by the mechanism. In fact, under certain conditions, public voting could lead to better decisions in this case by creating incentives for competent members to withhold information and then vote correctly in equilibrium (i.e. partially competent equilibrium). Overall, these results reinforce our previous conclusions and highlight another dimension in which the degree of transparency might be relevant for the quality of decisions.

Which level of transparency would the members of the committee prefer if, prior to voting, they could choose between public and secret voting? Here, we examine the institutional preferences of committee members by competence type. As discussed before, the choice between public and secret voting affects the payoffs of agents both in terms of how the career concern rewards are distributed across agents and 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 interesting exceptions to this general observation. First, if the weight associated with career concerns is small and the common value is high relatively to the bias, then competent members may actually prefer a secret voting rule, since secrecy is more likely to lead to better decisions in this case. Furthermore, whenever a biased equilibrium is expected to prevail anyway, then competent agents who are biased against the state of the world would actually prefer a secret voting rule, since in this case they always receive zero career concern rewards under public voting.

6 Experimental Design

In this section we explore the main theoretical predictions of our model by means of a controlled laboratory experiment. As discussed in the Online Appendix A, 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. This makes particularly difficult to evaluate the impact of transparency on voting outcomes using non-experimental data. A controlled experiment allows us to both collect data on individuals' behavior 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, a controlled experiment can inform on whether subjects eventually coordinate on the efficient equilibrium.

For the experimental implementation, we decided to amend the basic model imposing two simplifying assumptions on the structure of the career concern rewards. First, we assume that the career concern reward associated with a correct vote is exogenous under both public and secret voting. Specifically, before voting, each committee member knows, and is guaranteed to receive, a certain payoff $R^{\lambda} > 0$, $\lambda \in \{\mathtt{p}, \mathtt{s}\}$, whenever his or her vote is correct. Note that this simplified version retains all basic features of the general model, except that now the updating process of the external evaluator is not being explicitly modelled. In this way we can implement the experiment by avoiding the need for an extra subject whose role would be to guess the competence of each committee member, a complex task that would certainly add noise to the experimental results.[24] Second, while it is natural to suppose that $R^{\mathtt{p}} > R^{\mathtt{s}}$, we further assume that $R^{\mathtt{s}} = 0$, i.e. the career concern gain associated with a correct vote is zero under secret voting. We make this assumption in order to sharpen the contrast between the two treatments.

While the assumptions above simplify the model in significant ways, its basic structure remains unchanged. In particular, the same three classes of equilibria still exist, there are multiple equilibria in some regions of the parameter space and all previous comparative static results hold. We focus the experimental analysis on committees of three members with uniform prior $q = 1/2$ and symmetric distribution of both biases

[24]Both Fehrler and Hughes [12] and Meloso and Ottaviani [30] find that experimental subjects have a hard time updating beliefs correctly in the lab. In particular, Meloso and Ottaviani [30] 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.

$p = 1/2$ and competent types $\sigma = 1/2$. Under this parametrization, it is possible to show that the conditions for the existence of a fully competent, partially competent and biased equilibria are, respectively, the given by:

$$\gamma \leq \overline{\gamma}^{\lambda}_{full} \equiv \frac{1}{2}\alpha - 2R^{\lambda} \tag{14}$$

$$\gamma \leq \overline{\gamma}^{\lambda}_{part} \equiv \alpha + \frac{8}{3}R^{\lambda} \tag{15}$$

and

$$\gamma \geq \overline{\gamma}^{\lambda}_{bias} \equiv \alpha + 2R^{\lambda}, \tag{16}$$

where, as before, α is the common value, γ is the bias term and R^{λ} is the career concern reward associated with a correct vote under voting rule λ.[25]

We concentrate our analysis on regions of the parameter 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 summarized in Table 1. The common value was set to $\alpha = 10$ in all treatments, while the magnitude of the bias could be either low, $\gamma = 1$, or high, $\gamma = 14$. Moreover, the career concern rewards were chosen so that the payoff associated with a correct vote was $R^{\mathsf{p}} = 9$ under public voting and $R^{\mathsf{s}} = 0$ under secret voting. Accordingly, the treatments were labelled as: Low/Secret, Low/Public, High/Secret and High/Public.[26]

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. Each treatment was repeated for 32 periods, the first two of which being practice non-paid rounds. In every session, the value of the bias term (low or high) 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 2 summarizes the sequence of treatments and number of participants in each session. In total, 144 different subjects took part in the experiments.

The experiment was implemented via computer terminals and programmed in z-

[25] See Online Appendix D for the derivation of these conditions.

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

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.[27] 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 individual received a message about the group's color that could be either perfectly informative or non-informative with equal probability.[28] Specifically, subjects were told that 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.[29] 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 individuals in a given committee could be anything between zero and three.

Also before voting, 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 as described above: among all subjects present in a given session, half of them was randomly assigned the blue color and the other half was assigned the yellow color. 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 randomly. At the end of each period, subjects were provided with information about their group's color, the decision taken

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

[28] 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".

[29] This distribution procedure was adopted in order to make the experiment as direct and 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 (14)-(16). However, for the number of participants and parameter values used in each session, all of our equilibrium predictions remain unchanged.

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 he or she received 1 extra point under low bias treatments and 14 extra points under high bias treatments. Finally, under public voting treatments, subjects were also given an additional payoff of 9 points if his or her vote was equal to the group's color, while no points were given to a correct vote under secret voting treatments. 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 periods at the end of the experiment. The average earning was around €13.9, including a show-up fee of €2, with each session lasting for approximately 60 minutes.

7 Experimental Results

7.1 Decisions

We begin our analysis of the experimental results by investigating how the degree of transparency affects the quality of the committees' decisions, as measured by the proportion of correct choices made by the committees. Table 3 presents the fraction of correct decisions observed under each treatment, alongside with the fractions predicted by the model. Observe, first, that 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.[30]

7.2 Individual Choices

Table 4 summarizes the aggregate choices of uninformed subjects. Note that, when the magnitude of the bias is low, uninformed voters 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

[30] 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 $p = 0.00$.

(46.20%).[31] 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 public voting (84.26%).[32] These results are all in line with our theoretical comparative statics. 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 B). We also observe between 3% and 16% of uninformed agents 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 individuals do so as an attempt to guess the state of the world. This finding is consistent with experimental results previously obtained by Elbittar et al [8], who argue that a large proportion of uninformed subjects vote based on "hunches" (subjective beliefs).[33]

In Table 5 we summarize the behavior of informed voters who received a signal different than their biases. Among informed agents, these individuals are the ones most interesting to our analysis, since they 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 small, the vast majority of them vote correctly under both secret (95.96%) and public voting (97.71%).[34] The percentage of individuals who vote correctly under High/Secret (21.86%) and the percentage of individuals 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.[35] We also observe a fraction of informed voters who abstain under High/Secret (14.70%). This result is puzzling given that, in theory, abstaining is weakly dominated for agents 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, so that some informed subjects may simply prefer to abstain.

Finally, we complement our analysis of individual choices by classifying subjects in

[31] The χ^2 statistic for the difference in abstention rates is 158.5, with $p = 0.00$, and the χ^2 statistic for the difference in biased voting is 75.7, with $p = 0.00$.

[32] The χ^2 statistic for this difference is 6.1, with $p = 0.01$.

[33] Similar findings are also in Guarnaschelli et al [21] and in Bouton et al [5].

[34] 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 $p = 0.11$.

[35] See Online Appendix B for a detailed discussion.

accordance with their overall behavior during a session. Table 6 presents the distribution of the types of strategies used by uninformed agents in Low/Secret and Low/Public treatments. Note that the majority of individuals (44.44%) vote for their biases more than any other alternative in both treatments, while a substantial proportion of subjects (20.83%) mostly abstain under Low/Secret and vote for their biases under Low/Public. Interestingly, a considerable fraction of individuals (18.06%) abstain more than any other choice in both Low/Secret and Low/Public treatments. Next, table 7 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 individuals (65.38%) mostly vote for their biases under High/Secret and vote for their signals under High/Public.

7.3 Voting Profiles

We start 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 involved at least one uninformed agent and one informed agent who received a signal different than his bias. This restriction is imposed in order to allow us to associate each voting profile to a single class of equilibria. As shown in Table 8, 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.[36] Note that 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.[37] Moreover, the fraction of voting profiles consistent with a biased equilibrium drops significantly from 48.71% under High/Secret to 8.56% under High/Public.[38] 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.[39,40] We also

[36] The χ^2 statistic for this difference is 28.9, with $p = 0.00$.
[37] The χ^2 statistic for this difference is 20.3, with $p = 0.00$.
[38] The χ^2 statistic for this difference is 150.6, with $p = 0.00$.
[39] The χ^2 statistic for this difference is 161.3, with $p = 0.00$.
[40] Note that in all treatments there is a significant percentage of voting profiles that cannot be strictly categorized in one of the three classes of equilibria. Observe, however, that the fact that a voting profile belongs to this residual category, which we denote by "others", does not necessarily mean that individual behavior is incompatible with rationality. In fact, there are other classes of equilibria that may involve either asymmetric and/or mixed strategies, which we have not characterized in our theoretical analysis, but may be played in practice.

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 provides extra indication that subjects were gradually learning to coordinate on the more efficient equilibrium. In fact, 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, 44.33% in periods 21-30.

7.4 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.[41] We start by examining the determinants of a correct vote by informed agents. Table 9 presents the results of linear probability models where the dependent variable is a dummy that equals one if the individual voted correctly in a given period and zero otherwise. The sample is restricted to subject-period observations where the agent 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.[42]

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 agent 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.[43] 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.[44] 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

[41] Our results remain unchanged when we control for random effects instead of fixed effects.

[42] Clustering by session and adjusting the standard errors to account for the small number of clusters using a procedure proposed by Ibragimov and Müller [25] does not change any of our main results.

[43] 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.

[44] 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.

periods 11-20 and 68.3 p.p. in periods 21-30, which corroborates the existence of a strong learning effect for informed voters.[45]

Finally, we create a dummy variable that captures whether a subject performed poorly in the comprehension quizzes administered before the beginning of each treatment.[46] We interpret a low performance in these tests as evidence that either the individual did not fully understand a particular aspect of the experiment 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 voters. Table 10 presents the results of linear probability models where the dependent variable is a dummy that equals one if the agent abstained in a given period and zero otherwise. The sample is restricted to subject-period observations where the agent 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 individual level.[47] 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 agents 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 agent 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

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

[46] 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.

[47] 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.

after controlling for individual fixed effects. Specifically, the impact of a change from public to secret voting on the probability that an uninformed voter 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.[48]

Overall, the above results are consistent with our 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 better understand why uninformed voters do not coordinate more heavily on the Pareto optimal equilibrium, which involves all of them abstaining in order to let the "experts" decide. Our discussion here is related to previous studies by Elbittar et. al [8], and Grosser and Seebauer [20] who found, in a setting with common values, that a substantial proportion of individuals vote even though they have no information about the state of the world.

One possible explanation for this finding could 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 [11]). 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 individuals 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 individuals may have recognized the potential benefits of abstaining, they were discouraged from doing so by the fact that other uninformed agents were not abstaining as well. Indeed, the optimal behavior for an uninformed agent is for him to vote in accordance with his bias if he believes that other uniformed agents are also voting in accordance with their biases. In order to examine whether a negative feedback in one period impacts the subsequent decisions of agents, 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

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

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 an agent 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 voters may have, indeed, significantly limited the convergence of voting behavior towards the Pareto optimal equilibrium.

References

[1] Alesina, A. and G. Tabellini. 2007. "Bureaucrats or Politicians? Part I: A Single Policy Task." *American Economic Review*, 97(1): 169-79.

[2] Alesina, A. and G. Tabellini. 2008. "Bureaucrats or Politicians? Part II: Multiple Policy Tasks." *Journal of Public Economics*, 92(3-4): 426-47.

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

[4] Blinder, A. S. 2004 "*The Quiet Revolution: Central Banking Goes Modern.*" New Haven, Yale University Press.

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

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

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

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

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

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

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

[12] Fehler, S. and N. Hughes. 2016. "How Transparency Kills Information Aggregation: Theory and Experiments." Working Paper.

[13] Gerardi, D and L. Yariv. 2007. "Deliberative Voting." *Journal of Economic Theory*, 134: 317-338.

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

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

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

[17] 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

[18] Gneezy, U. 2005. "Deception: The role of consequences." *American Economic Review*, 95: 384-394.

[19] Gradwohl, R. 2014. “Voting in the Limelight." Mimeo

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

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

[22] Hansen, S., M. McMahon and A. Prat. 2015. “Transparency and Deliberation within the FOMC: A Computational Linguistic Approach.” Working Paper.

[23] Herrera, H., J. McMurray and A. Llorente-Saguer. 2016. “The Marginal Voter’s Curse.” Working Paper.

[24] Iaryczower, M., G. Lewis and M. Shum. 2014. “To Elect or to Appoint? Bias, Information, and Responsiveness of Bureaucrats and Politicians.” *Journal of Public Economics*, 97: 230-44.

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

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

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

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

[29] 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.

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

[31] Mian, A., A. Sufi and F. Trebbi. 2014. “Resolving Debt Overhang: Political Constraints in the Aftermath of Financial Crises.” *American Economic Journal: Macroeconomics*, 6(2): 1-28.

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

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

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

[35] 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.

[36] 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.

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

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

[39] Pham-Kanter, G. 2014. "Revisiting Financial Conflicts of Interest in FDA Advisory Committees." *The Milbank Quarterly*, 92(3): 446-470.

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

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

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

[43] 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.

[44] Wilson, J. Q. 1989. "*Bureaucracy: What Government Agencies Do and Why They Do it.*" Basic Books.

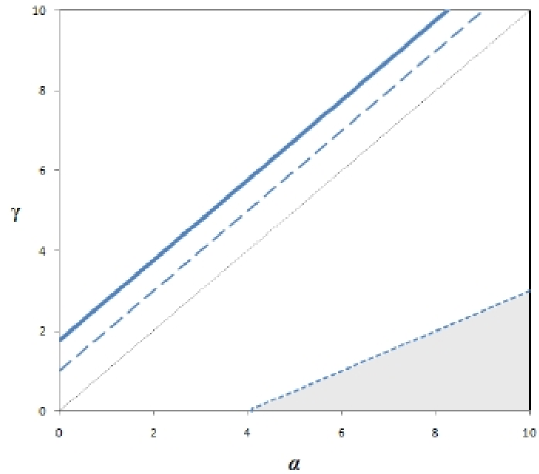

Panel a. Fully Competent Equilibrium

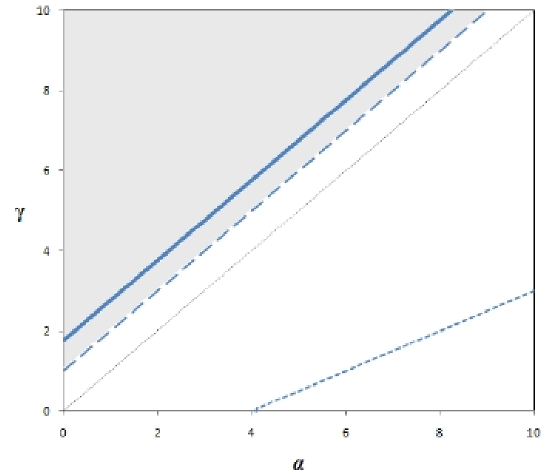

Panel c. Biased Equilibrium

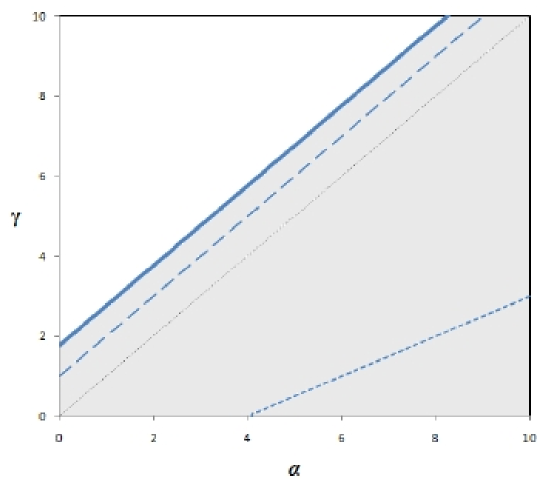

Panel b. Partially Competent Equilibrium

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, ϕ=1 and σ=0.5.

Figure 1. Equilibria: The Symmetric Case

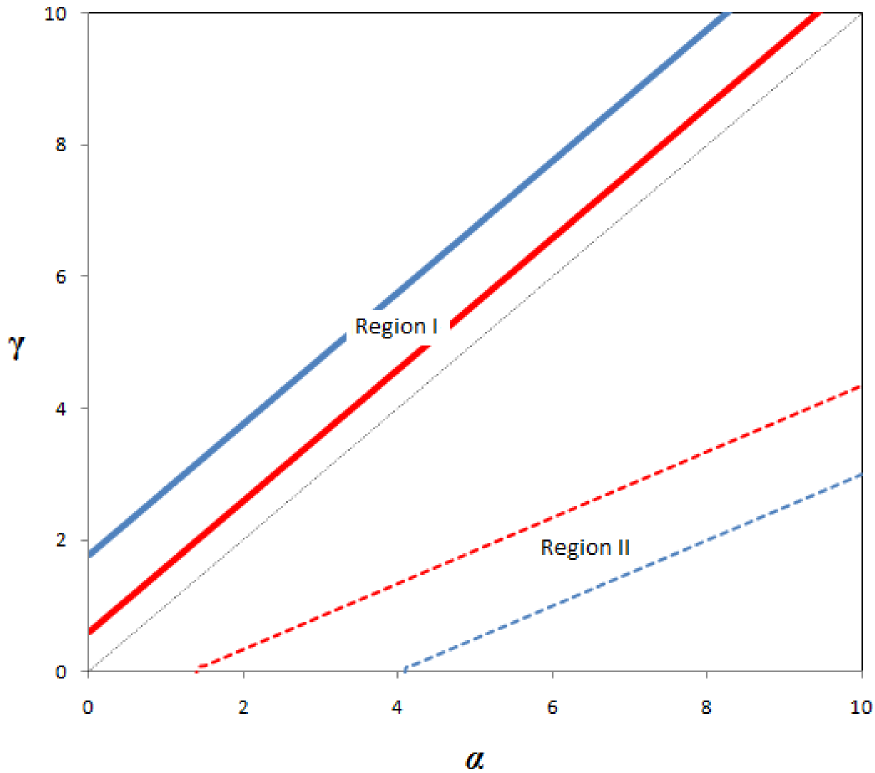

Notes. This figure provides a comparison of the parameter regions where a fully competent and a partially competent equilibrium can be supported under each voting rule. The relevant thresholds for the public and secret voting rules are depicted in blue and red, respectively. Region I represents the set 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 2. Comparative Static Result

Treatment	alpha	gamma	Reputation	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

Table 1. 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

Table 2. 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

Table 3. 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

Table 4. Individual Choices: Uninformed Subjects

Treatment	Obs	Informed Voters with Signal ≠ Bias		
		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

Table 5. Individual Choices: Informed Subjects

		Low/Secret		
		Mostly Vote for Signal (%)	Mostly Vote for Bias (%)	Mostly Abstain (%)
	Mostly Vote for Signal (%)	18.06	0.00	0.00
Low/Public	Mostly Vote for Bias (%)	20.83	44.44	1.39
	Mostly Abstain (%)	6.94	4.17	4.17

Table 6. Types of Strategies: Uninformed Subjects

		High/Secret		
		Mostly Vote for Signal (%)	Mostly Vote for Bias (%)	Mostly Abstain (%)
	Mostly Vote for Signal (%)	19.44	65.28	12.50
High/Public	Mostly Vote for Bias (%)	0.00	0.00	1.39
	Mostly Abstain (%)	0.00	0.00	1.39

Table 7. Types of 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

Table 8. Voting Profiles

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 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 9. 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 10. Regression Analysis: Uninformed Subjects