Public versus Secret Voting in Committees*

Andrea Mattozzi† EUI and MOVE

Marcos Y. Nakaguma‡ EESP-FGV

Abstract

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

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

*We thank Veronica Rattini, Hanno Kase, Chiara Santantonio and Riccardo Ghidoni for excellent research assistance. We also 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, Emanuel Ornelas, Luigi Proietti, Debraj Ray, Aniol Llorente-Saguer, Arthur Schram, Justin Valasek and seminar audiences at Bologna, CORS-RWIO, EESP-FGV, Erasmus University, EUI, Insper, King's College, Leicester, Ridge Forum, Royal Holloway, SBE, St.Andrews, University of Oslo, USP, Warwick,Venice 4th Political Economy conference for useful conversations, comments, and encouragement.

†Department of Economics, EUI, Florence, Italy. Email: andrea.mattozzi@eui.eu

‡Sao Paulo School of Economics, FGV, Sao Paulo, Brazil. Email: marcos.nakaguma@fgv.br

1 Introduction

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

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

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

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

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

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

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

We extend our basic model to allow for a behind closed-doors deliberation stage prior to voting, where committee members may choose to share their private informa-

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

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

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

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

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

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

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

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

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

2 Stylized Facts

In this section we present novel stylized facts about decision-making in committees. We exploit data from the national legislatures in Brazil and Italy and from the FDA's

advisory committees to illustrate the relevance of our theory in a number of real-world settings.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

3 The Model

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

result of the shift from sequential to simultaneous voting.

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

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

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

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

nc, in which case he receives an uninformative signal $s_i = \emptyset$. We assume that each member knows her own competence type $\tau_i \in \{\mathtt{c},\mathtt{nc}\}$ 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, $v_i \in \{A, \emptyset, B\}$, where abusing notation we denote abstention by $v_i = \emptyset$. The final decision $x \in \{A, B\}$ is determined by simple majority rule and ties are broken randomly.

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

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

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

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 assume that a committee member's competence is rewarded only if his vote is correct.[14] This assumption simplifies both the analysis and presentation of the results.

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

[14] Under our assumption, a committee member receives zero reputation whenever he abstains or votes incorrectly and his vote is observed. 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

Furthermore, it proves particularly useful for the experimental implementation of our theory. In Online Appendix A (see Subsection A.5) we provide a detailed analysis of a version of the model where we use a more standard definition of career-concerns based only on the posterior probability that the agent is competent. We show that our main results are robust to this alternative assumption.

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

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

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.

4 Equilibrium Analysis

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

4.1 Basic Properties

We begin our analysis by providing a general characterization of the basic properties of the equilibria. Let μ_i denote the conjecture held by a committee member i about the behavior of other members and the beliefs of the external evaluator. Suppose first that member i observes the state of the world prior to voting, i.e. he receives a perfectly informative signal. Given the conjecture μ_i and the state of the world ω, player i's action v_i induces a probability distribution over final outcomes, which is represented by the mapping $\rho^{\omega}_{\mu_i} : \{A, \emptyset, B\} \to [0, 1]$, where $\rho^{\omega}_{\mu_i}(v_i)$ denotes the probability (as perceived by the member) that the committee's decision is A when his choice is v_i, given μ_i and ω. Note that we must have $\rho^{\omega}_{\mu_i}(B) \leq \rho^{\omega}_{\mu_i}(\emptyset) \leq \rho^{\omega}_{\mu_i}(A)$ since a vote for A can never lead to a lower probability that the committee's decision is A relative to the case where the member abstains or votes for B.[15]

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

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

where $\Pr_{\mu_e}(\tau_i = \mathtt{c}|v_i = \omega)$ is computed based on the external evaluators' beliefs about the behavior of members. Under secret voting, on the other hand, only the aggregate voting outcome is observed ex-post, so that career-concern rewards depend on the total

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

number of correct votes, $V^c \equiv \sum_i \mathbb{I}_{\{v_i=\omega\}}$, in accordance with the following expression

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

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

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

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

where we omit the index for the evaluator's beliefs to simplify the notation. Under secret voting, on the other hand, the expected career-concern reward is given by

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

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

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

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

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

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

and

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

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

The next lemma provides a general characterization of the behavior of competent members.[16]

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

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

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

Intuitively, competent members observe the state of the world and, as a consequence, are not subject to the "swing voter's curse" (Feddersen and Pesendorfer [10]). Therefore, there is no reason for them to abstain. Lemma 1 implies that a competent member who receives a signal equal to her bias always (weakly) prefers to vote for the state. Instead, a competent member who receives a signal different from her bias may either (weakly) prefer to vote for the state or for her bias. Note also that Lemma 1 guarantees that, in any equilibrium, every competent members biased towards the state of the world votes correctly. Thus, by Bayes rule, the likelihood that a member is competent given that she voted correctly is strictly positive in any equilibrium, $\Pr(t = \mathsf{c} | v = \omega) > 0$. The next lemma follows as an implication of this result.

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

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

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

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

Intuitively, incompetent members are always more inclined to follow their biases relative to competent members. Note that when a competent individual decides to vote against the signal, she is certain to be casting an incorrect vote, while an incompetent member always attributes positive probability to the event that her vote is correct. Specifically, Lemma 3 guarantees that if, for instance, in equilibrium, competent members biased towards A vote for their bias when the state is B, then all incompetent members biased towards A must vote for their bias. As for the behavior of incompetent members with bias B in this case, the lemma just says that they will never abstain − they might vote either for their bias or for the ex-ante more likely alternative (against their bias).[17]

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

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

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

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

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

A fully-competent equilibrium completely pins down the behavior of all committee members, while both a partially-competent and a biased equilibrium allow for a variety of different behaviors within each class. A partially-competent equilibrium pins down only the behavior of competent members and is consistent with either both types of incompetent members voting − either for the ex-ante more likely alternative or for their bias − or with one type voting (e.g., A biased) and the other abstaining (e.g., B biased). A biased equilibrium is consistent with either both types of competent members always voting for their biases or with one type always voting for the state and the other always voting for the bias. As for the incompetent members, they never

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

abstain in a biased equilibrium and, as prescribed by Lemma 3, they vote for their bias if the competent members of the same type are doing so as well.

Proposition 1 helps organizing the set of all possible equilibria by grouping them in terms of key qualitative features of voters' behavior.[18] Importantly, our characterization holds under both public and secret voting, although the region of parameters where each class of equilibrium exists does depend on the level of transparency of the voting rule.

4.2 Main Comparative Statics Results

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

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

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

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

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

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

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

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

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

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

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

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

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

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

4.3 The Symmetric Case

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

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

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

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

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

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

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

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

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

The term $((n-1)/n)\cdot\mathbb{I}_{\{\lambda=\mathtt{s}\}}$ appearing in the expressions above captures the impact of the dilution of career-concern rewards under secret voting.[19] Note that a change from public to secret voting is qualitatively equivalent to a reduction in the weight on reputation ϕ. Figure 2 shows the values of the parameters α and γ for which each class of equilibrium can be sustained, for given transparency λ and for fixed values of ϕ, σ and n.

Observe that since $\overline{\gamma}^{\lambda}_{full} < \overline{\gamma}^{\lambda}_{part}$, the region of parameters where a fully-competent equilibrium exists is contained inside the region where a partially-competent equilibrium exists. Recall that the main reason for an incompetent member to abstain is to avoid adding "noise" to the decision process. However, a coordination issue arises in the region where the two equilibria overlap in that abstaining can only be optimal for an incompetent member if she expects other incompetent members to abstain as well. Similarly, since $\underline{\gamma}^{\lambda}_{bias} < \overline{\gamma}^{\lambda}_{part}$, there exists a region of parameters where both a partially-competent and a biased equilibrium can be sustained simultaneously. The multiplicity of equilibria arises in this case due to the existence of a coordination issue among competent members who are biased against the state of the world. In the re-

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

gion where the two equilibria overlap, voting in accordance with one's bias can only be optimal if the member expects other competent members of the same type to do the same. The reason is that an individual is less likely to be pivotal when she is the only competent member voting against the state, in which case she would prefer to vote correctly in order to guarantee a larger career-concern reward for herself.

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

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

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

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

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

with $\Pi_{full}(\sigma, n) > \Pi_{part}(\sigma, n) > \Pi_{bias}(\sigma, n)$ for any $0 < \sigma < 1$ and $n \geq 3$.[20] We are, therefore, able to rank public and secret voting in welfare terms, based on the expected quality of the decisions.

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

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

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

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

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

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

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

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

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

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

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

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

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

here is that an increase in σ raises the opportunity cost of voting against the state, given that the career-concern rewards associated with a correct vote are increasing in the fraction of competent members.

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

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

i. Under public voting:

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

ii. Under secret voting:

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

As n gets arbitrarily large a fully-competent equilibrium can never be supported. Indeed, the probability that an incompetent member is pivotal in a fully-competent equilibrium converges to zero as $n \to \infty$, so that incompetent members have a large incentive to vote. Thus, contrarily to Feddersen and Pesendorfer (1996), information is never fully aggregated in large elections.[21] Furthermore, note that a partially-competent equilibrium exists everywhere in the parameters' space under both public and secret voting for large n. Finally, a biased equilibrium can only exist under secrecy. Thus, overall, our analysis shows that in large elections with career-concerns, transparency is expected to lead to (weakly) better decisions

5 Extensions

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

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

5.1 Model with Deliberation

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

We analyze an extension of the symmetric version of the model with two stages: (i) a deliberation stage and (ii) a voting stage. At the beginning of the game, each member receives a signal $s_i \in \{A, \emptyset, B\}$ about the state, as in the benchmark model. We focus our analysis on the case where, at the deliberation stage, members are able to exchange hard information: each member sends a message $m_i \in \{s_i, \emptyset\}$ to the committee, which may reveal or not his private signal truthfully.[22] At the voting stage, after having observed all messages $m = (m_1, ..., m_n)$, the members of the committee vote and the decision is taken by majority. As before, under public voting, all votes are observed by the external evaluator, while under secret voting only the vote tally is observed. We assume that in both cases the messages exchanged at the deliberation stage are never observed by the evaluator.

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

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

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

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

model.[25]

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

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

Notice that behavior at the voting stage is straightforward. If information about the state is revealed during the deliberation process, then at the voting stage no member ever wants to vote incorrectly, since the probability of being pivotal is zero. On the other hand, if information is not revealed on the equilibrium path, then it becomes common knowledge among committee members that everyone is incompetent and it is, therefore, optimal for all members to simply vote for their biases.

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

Partial Revelation Equilibrium. In a partial revelation equilibrium, only the competent members with bias equal to the state reveal their signals, yet everyone votes correctly when information is shared. Note that, conditional on the state not being revealed during deliberation, the equilibrium is consistent with a variety of voting behaviors. We focus on the more interesting case: a partial revelation equilibrium with "partially-competent voting" when no information is revealed.[26]

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

[26] A partial revelation equilibrium with "fully-competent voting" does not exist under public voting. Indeed, a competent member with bias equal to the state would also have an incentive to withhold

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

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

Comparative Results. While it is difficult to derive general comparative static results in the model with deliberation, our analysis highlights the fact that the level of transparency still matters, although the nature of the trade-off is slightly different in this case.[27] First, in line with Coughlan (2000), we conclude that when $\alpha > \gamma$

information and vote correctly in an attempt to separate himself from incompetent members. The case of a partial revelation equilibrium with "biased voting" has qualitative features that make it very similar to an irrelevant deliberation equilibrium, which we analyze below in detail.

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

a full deliberation equilibrium always exists under both public and secret voting. In this case the state is revealed at the deliberation stage − provided that there is at least one competent member in the committee − and the correct decision is taken by unanimity. Interestingly, the same force that induced incompetent members to abstain in a model without deliberation now generates an incentive for competent members to share information in equilibrium.

Conversely, when the bias term is relatively large, $\alpha < \gamma$, transparency might involve a trade-off between the quality of information aggregation at the deliberation stage and the quality of the decisions at the voting stage. Under certain conditions, secrecy may actually make it more likely that information about the state is revealed at the deliberation stage, while transparency creates an incentive for informed members to vote correctly. In particular, note that while public voting always makes the irrelevant deliberation equilibrium less likely, there are some regions of the parameters where a change from public to secret voting might lead to a shift from a partial deliberation to a full deliberation equilibrium.[28] Altogether, these results reinforce our previous conclusions, highlighting another dimension in which the level of transparency might be relevant for the quality of the decisions of the committee.[29]

5.2 Dynamic Model

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

We focus our analysis on the characterization of the conditions under which a

and $\phi = 1$ in the Online Appendix B.1.

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

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

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

5.2.1 The IID Case

We begin our analysis by considering the case where competence is iid across periods, with $\Pr(\tau = \mathsf{c}) = \sigma$.[30] Intuitively, this case can be viewed as capturing situations where a committee member may be an expert in some issues but not in others. Since both competence and bias are iid over time, the structure of the game in any period is very similar to that of the static model, except that the external evaluator is now able to condition the career-concern rewards on the entire history of observed votes. In this case a *dynamic fully-competent equilibrium* is defined as an equilibrium where committee members vote correctly in all periods in which they are competent and abstain in all periods in which they are incompetent. Thus, on the equilibrium path, no member ever votes incorrectly and a correct vote in a given period perfectly reveals competence in that particular period. Outside of the equilibrium path, after an incorrect vote is observed, we assume that committee members switch to voting in accordance with a partially-competent equilibrium.[31]

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

Under secret voting, on the other hand, the career-concern rewards in any period are such that, on the equilibrium path, the reputation of the committee is equal to the total number of correct votes in that period divided by the size of the committee. We assume that when an incorrect vote is observed, and in every period thereafter,

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

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

the external evaluator simply attributes a reward that is equal to the prior probability that a member is competent σ multiplied by a factor $(n-1)/n$. The scaling captures the idea that the external evaluator "randomly" punishes one individual with zero reputation. Thus, as before, punishment is diluted among committee members.[32] The main results of our analysis are summarized in the following proposition:

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

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

Furthermore, we have that:

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

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

Our analysis shows that it is easier to sustain behavior consistent with a fully-competent equilibrium when the game is played repeatedly and competence is iid across periods. Formally, we can show that $\overline{\gamma}^{\lambda}_{full}(\alpha, \phi, \sigma, n) < \overline{\gamma}^{\lambda}_{dfull-i}(\alpha, \phi, \sigma, n, \delta)$ for any $\delta \in (0,1)$, so that the condition for the existence of a dynamic fully-competent equilibrium is easier to satisfy than the condition for a fully-competent equilibrium in the static model. Indeed, the fact that an incompetent member may turn out to be competent in the future helps to discipline her behavior in the current period.[33]

Regarding the comparison between public and secret voting (point *ii.* above), a dynamic fully-competent equilibrium is more likely to be sustained under public voting if and only if the discount factor is large enough. Note that a novel trade-off between public and secret voting emerges in the dynamic version of the model. Besides the benefit associated with voting in the current period, there now exists the expectation of punishments in future periods. Transparency increases both the incentive for an

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

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

incompetent member to vote today and the expected punishment associated with an incorrect vote, so that the overall effect depends on the discount factor. When the discount factor is large enough, committee members place a higher weight on the future stream of payoffs, so that a harsher expected punishment makes it easier to sustain a dynamic fully-competent equilibrium. In this case, public voting is expected to lead to better voting outcomes in terms of information aggregation. On the other hand, when the discount factor is small, committee members place a higher weight on the current payoff, so that lower career-concern rewards associated with voting in the current period make it easier to sustain a dynamic fully-competent equilibrium. In this case, secret voting is expected to lead to better outcomes.

5.2.2 Persistent Competence

We now consider the case where the competence of committee members is realized at the beginning of the game, with $\Pr(\tau = \mathsf{c}) = \sigma$, and remains fixed over time. As before, we maintain the assumption that both the state of the world and the individual biases are iid across periods. A dynamic fully-competent equilibrium in this case is defined as an equilibrium where competent members vote correctly and incompetent members abstain in all periods − unless it becomes common knowledge that all members are incompetent, in which case everyone votes for their biases.[34]

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

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

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

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

Furthermore, we have that:

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

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

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

6 Experimental Design

In this section we test the main theoretical predictions of our basic model by means of a controlled laboratory experiment. Since the choice of adopting secret or public voting may be endogenous to the composition of the committee as well as to the types of decisions being made, it is particularly difficult to evaluate the impact of transparency on voting outcomes using non-experimental data.[35] A lab experiment allows us to both collect data on individuals' behaviors and compare the quality of the decisions under public and secret voting, while controlling for the degree of information and biases of committee members. Furthermore, since our model features multiple equilibria with

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

different information aggregation properties, the experimental results can inform on whether subjects select a particular equilibrium.

For the experimental implementation, we amend the basic model imposing two simplifying assumptions on the structure of the career-concern rewards. First, we assume that the career-concern rewards associated with a correct vote are exogenous under both public and secret voting. Specifically, before voting, each committee member knows, and is guaranteed to receive, a certain payoff $R^{\lambda} > 0$ when she votes correctly. Note that this simplification maintains all basic features of the original model, except that now the updating process of the external evaluator is not being explicitly modelled.[36] 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 reward associated with a correct vote is zero under secret voting. We make this assumption in order to sharpen the contrast between the two treatments. The rest of the model remains unchanged. In particular, the same three classes of equilibria exist, there are multiple equilibria in some regions of the parameters' space and all previous comparative static results hold.

We consider committees of three members with uniform prior ($q = 1/2$) and symmetric distribution of both biases ($p = 1/2$) and competence ($\sigma = 1/2$). Recall that in this case there exists a unique equilibrium in each class.[37] Under this parametrization, the conditions for the existence of each class of equilibria are given by

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

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

and

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

with the usual notation.[38] We concentrate our analysis on regions of the parameters'

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

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

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

space where a change in the degree of transparency is expected to lead to a change in observed behavior. The choice of parameters as well as the equilibrium predictions associated with each of the four treatments considered in the experiments are reported in Table 2. The common value is set to $\alpha = 10$ in all treatments, while the magnitude of the bias can be either low, $\gamma = 1$, or high, $\gamma = 14$. Furthermore, the career-concern reward is $R^{\mathtt{p}} = 9$ under public voting and $R^{\mathtt{s}} = 0$ under secret voting. Accordingly, there are four treatments labelled Low/Secret, Low/Public, High/Secret and High/Public, with the Low/Secret treatment being consistent with both a fully-competent and a partially-competent equilibrium.[39]

The experiments were conducted at the Bologna Laboratory for Experiments in Social Science (BLESS) with registered undergraduates from the University of Bologna. We run the experiments in 6 sessions, each consisting of 2 parts with a different treatment being tested in each part (within-subject design). Each treatment was repeated for 32 rounds, the first two being practice non-paid rounds. In every session, the value of the bias was held fixed and only the parameter corresponding to the career-concern reward (public or secret voting) changed from one part to the other. Table 3 reports the sequence of treatments and number of participants in each session. In total, 144 different subjects took part in the experiments.

The experiments were implemented via computer terminals and programmed in z-Tree. In every session, instructions were read aloud at the beginning of each part, after which a short comprehension quiz was administered in order to check basic understanding of the rules.[40] Subjects were randomly divided into groups of three members and were re-assigned, in every period, to different groups using a random matching procedure. The task of each group was to choose between two colors, blue or yellow. The "group's color" (i.e. the state of the world) was ex-ante unknown and could be either one of the two colors with equal probability.

Before voting, each subject received a message about the group's color that could be either perfectly informative or non-informative with equal probability.[41] Specifically, subjects were told that the messages would be randomly assigned so that, among

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

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

[41] 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".

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.[42] At this point, we were explicit in emphasizing that this procedure did not guarantee that there would always be an informed member in every group and that, in fact, the number of informed members in a given committee could be anything between zero and three. Then, each subject was informed about his or her "role" (i.e. bias), which could be either blue or yellow with equal probability. The procedure used to assign individual colors was the same used to assign group colors. After observing their messages and roles, each subject had to choose whether to vote for blue or yellow or to abstain. The "group's decision" was taken by majority rule and ties were broken randomly. At the end of each period, subjects were provided with information about their group's color, the decision taken and the number of members of the group that voted for blue, yellow or abstained.

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

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

7 Experimental Results

7.1 Committee Decisions

Table 4 reports the fraction of correct decisions observed under each treatment, alongside with the model predictions. The quality of the decisions is slightly higher under Low/Secret (85.56%) than Low/Public (84.31%), whereas the fraction of correct decisions under High/Secret (59.58%) is significantly lower than under High/Public (81.53%), as expected.[43]

7.2 Individual Choices

Table 5 summarizes the aggregate choices of uninformed subjects. When the magnitude of the bias is low, uninformed subjects are much more likely to abstain under secret (44.17%) than public voting (18.98%), while being significantly more likely to vote in accordance with their biases under public (64.81%) than secret voting (46.20%). On the other hand, when the magnitude of the bias is high, the vast majority of uninformed subjects vote in accordance with their biases under both secret (87.96%) and public voting (84.26%).[44] These results are all in line with our theoretical comparative static predictions. It should be noted that while 18.98% of subjects abstain under Low/Public, this number decreases substantially when we account for sequencing effects (see the Online Appendix C). We also observe between 3% and 16% of uninformed subjects voting *against* their biases depending on the treatment. Interestingly, the incentive to vote against the bias seems to be larger under public voting, which may be interpreted as evidence that some subjects do so in an attempt to guess the state of the world.[45]

In Table 6 we summarize the behavior of informed subjects who received a signal different than their biases and, therefore, face a trade-off between voting correctly and voting for their biases. Observe that, as predicted by the theory, when the magnitude of the bias is high, these subjects are much more inclined to vote correctly under public (84.60%) than secret voting (21.86%), while when the magnitude of the bias is

[43]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 $p = 0.00$.

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

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

low, the vast majority of them vote correctly under both secret (95.96%) and public voting (97.71%).[46] The percentage of subjects who vote correctly under High/Secret (21.86%) and the percentage of subjects who vote in accordance with their biases under High/Public (11.94%) are larger than expected. We note, however, that these proportions tend to decrease when we account for learning and sequencing effects.[47] We also observe that a fraction of informed subjects abstain under High/Secret (14.70%). This result is puzzling given that, in theory, abstaining is weakly dominated for members of this type. A possible explanation for this result could be attributed to the fact that both the common value (10 points) and the bias (14 points) are relatively close in magnitude in this case, which may lead some informed subjects to simply prefer to abstain.

Finally, we complement our analysis of individual choices by classifying subjects in accordance with their overall behavior during a session. Table 7 reports the distribution of the types of strategies used by uninformed subjects in Low/Secret and Low/Public treatments. Consistently with our theoretical predictions, the majority of subjects (44.44%) vote for their biases more than any other alternative in both treatments and a substantial proportion of subjects (20.83%) mostly abstain under Low/Secret and vote for their biases under Low/Public. We also observe a considerable fraction of subjects (18.06%) abstaining more than any other choice in both Low/Secret and Low/Public treatments. Next, Table 8 reports the most frequent strategies adopted by informed subjects in High/Secret and High/Public treatments when they receive a signal different than their biases. As expected, we find that the vast majority of subjects (65.38%) mostly vote for their biases under High/Secret and vote for their signals under High/Public.[48]

While our results are consistent with the main comparative static predictions about the behavior of uninformed voters, still the fraction of subjects who change from voting to abstaining as a result of a change from public to secret voting is significantly below one. Given that there are multiple equilibria under Low/Secret, it would be interesting to investigate why uninformed voters do not select the efficient equilibrium, although,

[46] 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 $p = 0.11$.

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

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

from a theoretical point of view, it is unclear which equilibrium players should coordinate on.[49] A potential explanation for the relatively low levels of abstention is that, while some subjects may have recognized the potential benefits of abstaining, they were discouraged from doing so by the fact that other uninformed subjects were not abstaining as well.[50] We explore this possibility in the Online Appendix C and provide empirical evidence that a negative feedback in one period (i.e. realizing that at least one other committee member "distorted" the decision by voting for the wrong alternative) significantly impacts the subsequent decisions to abstain in future periods and limits the convergence of voting behavior towards the efficient equilibrium.

7.3 Voting Profiles

We conclude our analysis by examining the frequency with which the observed voting profiles are exactly in accordance with one of the three classes of theoretical equilibria. In order to do so, we restrict the sample to include only decisions that involve at least one uninformed subject and one informed subject who received a signal different than her bias. This restriction allows us to associate each voting profile with a single class of equilibria. As shown in Table 9, the proportion of voting profiles that are consistent with a fully-competent equilibrium decreases, as expected, from 33.23% under Low/Secret to 15.73% under Low/Public. This reduction is accompanied by a proportional increase in the profiles compatible with a partially-competent equilibrium from 35.00% under Low/Secret to 51.96% under Low/Public. Furthermore, the fraction of voting profiles consistent with a biased equilibrium drops significantly from 48.71% under High/Secret to 8.56% under High/Public. Again, this reduction is accompanied by an increase in the profiles compatible with a partially-competent equilibrium from 17.47% under High/Secret to 63.47% under High/Public.[51]

Finally, we find evidence (not reported in Table 8) that the percentage of voting profiles consistent with a fully-competent equilibrium under Low/Secret, a treatment in which there are multiple equilibria, increases substantially within the treatment. This result suggests that subjects were gradually learning to coordinate on the fully-

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

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

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

competent equilibrium. Indeed, the percentage of voting profiles that are exactly in line with a fully-competent equilibrium increases from 27.11% in periods 1-10 to 29.31% in periods 11-20 to finally reach 44.33% in periods 21-30.

8 Concluding Remarks

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

A number of papers in the literature have shown that transparency in decision-making is not always advisable since it creates incentives for agents to distort their behavior in order to convey information about their types. This has been investigated for single decision-makers and only recently − and partly following a trend towards increased procedural transparency in central banking − the literature has started focusing on the effects of the transparency of voting procedures on decision-making in committees. None of the existing papers has investigated how competence, individual biases and career-concerns interact in shaping individuals' voting behaviors in a committee, and how this interaction is affected by transparency.

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

Feddersen and Pesendorfer [10].[52]

Gersbach and Hahn [12] and Stasavage [38] analyze a setting where committee members may be misaligned with the interests of society, but also care about being perceived as "unbiased" to the extent that this enhances their reelection prospects. They show that transparency induces biased agents to act in accordance with the public interest. Conversely, in single decision-makers models, Ely and Välimäki [8], Morris [29] and, more recently, Shapiro [37] argue that transparency and career-concerns create an incentive for an unbiased agent to ignore her private information and choose the alternative that makes her look impartial.[53] Our model can help reconcile these seemingly opposing results: Transparency leads to better decisions when the biases are large, and secrecy leads to better decisions when the biases are small. Furthermore, our model does not assume that individual biases are per se punished.

In addition to these papers, Gersbach and Hahn [14] show that transparency induces agents to exert more effort in order to improve their chances of reappointment, Dal Bo [6] and Felgenhauer and Gruner [9] argue that public voting makes the committee more vulnerable to the influence of special interest groups, and Swank and Visser [39] show that career-concerns create an incentive for committees to conceal internal disagreements and show a united front in public.[54] Finally, Midjord et al [28] point out that career-concerns induce experts to be too conservative in order not to put their reputation at risk, and Gradwohl [16] shows that transparency leads to a trade-off between the accuracy of the decisions and the welfare of agents in a model where committee members have privacy concerns.

As for the experimental literature on committee decision making, a closely related paper to ours is Fehrler and Hughes [11]. As in our paper, they focus on the effect of transparency on committee decision making where agents are career-concerned. Differently from us, they concentrate on committees of two individuals, members are unbiased, and the experimental focus is on deliberation.[55] Also related is Battaglini et al [1] who provided the first test of the swing voters' curse in a laboratory setting.[56]

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

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

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

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

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

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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