



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2018

How transparency kills information aggregation: theory and experiment

Fehrler, Sebastian ; Hughes, Niall

DOI: <https://doi.org/10.1257/mic.20160046>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-158357>

Journal Article

Published Version



The following work is licensed under a Creative Commons: Attribution-NoDerivs-NonCommercial 1.0 Generic (CC BY-ND-NC 1.0) License.

Originally published at:

Fehrler, Sebastian; Hughes, Niall (2018). How transparency kills information aggregation: theory and experiment. *American Economic Journal: Microeconomics*, 10(1):181-209.

DOI: <https://doi.org/10.1257/mic.20160046>

How Transparency Kills Information Aggregation: Theory and Experiment[†]

By SEBASTIAN FEHRLER AND NIALL HUGHES*

We investigate the potential of transparency to influence committee decision-making. We present a model in which career concerned committee members receive private information of different type-dependent accuracy, deliberate, and vote. We study three levels of transparency under which career concerns are predicted to affect behavior differently and test the model's key predictions in a laboratory experiment. The model's predictions are largely borne out—transparency negatively affects information aggregation at the deliberation and voting stages, leading to sharply different committee error rates than under secrecy. This occurs despite subjects revealing more information under transparency than theory predicts. (JEL C92, D72, D82, D83)

Transparency in decision-making is a recurring and controversial topic in public debate. A recent example comes from the world of sport, where FIFA's decision to hand the 2022 World Cup to Qatar—amid allegations of bribing—spurred calls for more transparency from soccer fans worldwide.¹ Another prominent debate is whether monetary policy committees should be secretive or transparent. Decision-making in the the Federal Open Market Committee has become more transparent over time (Williams 2012), yet many—including former Federal Reserve Chairman Alan Greenspan—view it as a step in the wrong direction.²

The supposed boon of transparency is accountability; by making the decision-making process more transparent, this should align the incentives of the agents with those of the principal. The downside of transparency is that if agents care about their reputations, they may pander to the principal, choosing an action that makes them appear smart but is not necessarily in the principal's interest (Prat 2005; Fox and Van Weelden 2012). This is the case both for single agents and for committees.

*Fehrler: Department of Economics, University of Konstanz, Box 131, D-78457 Konstanz, Germany (email: sebastian.fehrler@uni-konstanz.de); Hughes: King's Business School, King's College London, Bush House, London WC28 4BG, UK (email: niall.hughes@kcl.ac.uk). We are especially grateful to Alessandra Casella for her support and many great discussions. We also wish to thank Peter Buisseret, Aniol Llorente-Saguer, Katja Michaelowa, Massimo Morelli, Becky Morton, Maik T. Schneider, Francesco Squintani, Jean-Robert Tyran, Bauke Visser, and participants at several conferences and seminars for their valuable comments. This project was financially supported by the Swiss National Science Foundation (grant 100017_150260/1).

[†]Go to <https://doi.org/10.1257/mic.20160046> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

¹<http://www.theguardian.com/football/2014/sep/24/michael-garcia-fifa-investigation-public-world-cup>.

²Greenspan's concern was that forced transparency would lead to the emergence of premeetings where honest exchanges would take place. For a discussion of this, see Meade and Stasavage (2008) and Swank and Visser (2013).

The recent theoretical literature on transparency in committees falls into two strands. In one strand, committee members do not know their own ability but can deliberate on which decision is best. Here, whether transparency is a blessing or a curse depends crucially on whether communication is simultaneous (Visser and Swank 2007) or sequential (Meade and Stasavage 2008). In the other strand, each member has private information on his ability but cannot communicate with others (Stasavage 2007; Levy 2007; Mattozzi and Nakaguma 2017). Here, transparency can mitigate the influence of biased members, but with a loss of information.³

We believe that both communication *and* private information on abilities are important features of real world committees. Our theoretical analysis reveals important interactions between these two features, which determine if and when transparency is optimal. Excessive transparency prevents committee members from revealing their abilities to one another and, moreover, it forces members to ignore any useful information they get in deliberation. We derive these results from a model in which career concerns play out very differently under three levels of transparency. This setup allows us to study how subjects react to the changing opportunities to act smart both at the deliberation and the voting stages of the decision process. As in most career concerns models, there are two types of committee members: high and low ability. First, committee members receive either a fully informative or a noisy signal about the true state of the world, depending on their ability. Then they have the opportunity to deliberate and, finally, vote for or against changing the status quo in favor of an alternative option. A change of the status quo requires a unanimous vote of the committee. The first deliberation stage consists of a nonbinding straw poll in which subjects can exchange information about their signals. This is followed by a second stage in which they can exchange information about their type, i.e., the signal strength. A committee member's utility depends on the principal's belief that he is of high ability.⁴ We study three different transparency regimes: *secrecy*, where votes and communication are secret and the principal only learns the committee decision; *transparency*, where both communication and individual votes are public; and the intermediate case of *mild transparency*, where individual votes are made public but deliberation is secret. Mild transparency thus reflects a committee practice of publishing individual voting records but not transcripts of the deliberation.

The model provides several testable predictions. It predicts that committee members will truthfully share all their information with each other under secrecy and mild transparency, but will fail to do so under transparency. When the principal is watching nobody will admit to being a low type, so information aggregation is incomplete. Our model further predicts that committee members will implement the action favored by the group posterior under secrecy and mild transparency, but each member will vote according to his own signal under transparency. Low-ability members will do so even if they believe their signal to be wrong. This occurs because if members change their position between the deliberation and voting stages, the

³Recent papers that study the trade-off between transparency and information aggregation in a signaling framework include Frankel and Kartik (2017) and Ali and Bénabou (2016). Closely related, Hahn (2017) studies self-selection of career concerned experts into committees under different levels of transparency (see also Fehrer, Fischbacher, and Schneider 2017).

⁴For simplicity, the committee members in our text are male while their principal is female.

principal can infer that they are a low type.⁵ Based on these predictions on individual behavior, we can state the first key prediction regarding aggregate effects: incorrect group decisions will be more prevalent under transparency than secrecy. However, if the state of the world favors retaining the status quo, transparency will dominate secrecy. This stems from the fact that a unanimous decision is needed to overturn the status quo. As committee members vote according to their private signals under transparency, they find it difficult to vote unanimously against the status quo. This means more mistakes when they should change the status quo but fewer mistakes when they should not. This asymmetry of errors appeals to a principal who is very concerned about wrongly changing the status quo. Here, transparency aids the principal *because* it hinders information aggregation in the committee. Another key aggregate level prediction, resulting from our predicted individual-level behavior, is that the decision accuracy will be the same under secrecy and mild transparency.

These predictions are based on equilibrium refinements, and assumptions of fully strategic behavior of committee members and flawless belief updating of principals. To deal with the multiplicity of equilibria, we follow the cheap talk literature in focusing on the most informative equilibrium. In Section IIIA, we discuss reasons why this equilibrium selection and our assumption of fully strategic behavior might fail to predict behavior well in all treatments. In this context, mild transparency is of particular interest; here, high ability types prefer babbling equilibria to the most informative equilibrium. We then take our hypotheses, which we explicitly state in Section III, to the laboratory in what is the first experimental test of theoretical predictions on committee deliberation and decision-making with career concerns.⁶

We find strong evidence in favor of our hypotheses regarding differences in deliberation and voting between secrecy and transparency. The resulting differences in committee error rates confirm our first key hypothesis regarding the aggregate effect of transparency—there *are* indeed more incorrect group decisions under transparency than under secrecy. We further find that our principals in the lab are not perfect Bayesians—they are too optimistic in their assessment of committee members. However, they update their beliefs in the correct direction to all pieces of information they receive. Finally, we roundly reject the hypotheses stating that behavior and error rates will be the same under secrecy and mild transparency. The experiment shows instead that individual behavior is quite different under mild transparency and aggregate level error rates are very similar to the transparency case. Overall, our results show that transparency indeed has a strong potential to influence committee behavior and decision-making accuracy. They further show that experimental tests can be useful complements to committee voting theory as many of its predictions

⁵In this respect, our model has some similarities with that of Gersbach and Hahn (2008)—they too find that low ability types will pose as high types under transparency, leading to worse information aggregation. However, while committee members are not allowed to change their position between the deliberation and the voting stages in their model, the “sticking to signal” result arises endogenously in our model.

⁶Mattozzi and Nakaguma (2017) experimentally examine the role of transparency on career concerned committees where members are biased and there is no deliberation. Using a completely different (machine-learning) approach, Hansen, McMahon, and Prat (2017) analyze transcripts of Federal Open Market Committee meetings and find that debates became more scripted under transparency. Koch, Morgenstern, and Raab (2009) and Irlenbusch and Sliwka (2006) conduct experiments based on Holmström’s (1999) career concerns model.

necessarily rely on strong behavioral assumptions and equilibrium refinements—and it is not clear *ex ante* that they will predict behavior accurately.

In the next section, we present and solve the model. We proceed with the experimental design before discussing the aggregate and individual level results. We conclude with a discussion of the main theoretical and empirical results and their implications for the literature on career concerns in committees and on cheap talk.

I. The Model

A committee of two members must make a decision $D \in \{B(lue), R(ed)\}$ on behalf of a principal. The voting rule is unanimity: option R is implemented only if both members vote for it, otherwise the status quo B is upheld. There are two equally likely states of the world $S \in \{B, R\}$ and each member i gets an informative signal about the true state $s^i \in \{b, r\}$. Each member can be of low or high ability. A high-ability player receives a perfectly informative signal while a low-ability player receives a correct signal with probability $\sigma \in (0.5, 1)$. Thus, there are four possible types of committee member $\{hb, hr, lb, lr\}$. With a slight abuse of terminology we will refer to $\{hb, hr\}$ as high types and $\{lb, lr\}$ as low types; $t^i \in \{h, l\}$, where $\Pr(t^i = h) = q \in (0, 1)$. Abilities and signals are private information.

The principal gets a positive utility if the committee decision matches the state and zero otherwise. In Section IID, we examine the case where she cares more about one state than the other. The payoff of a committee member is simply the principal's posterior belief that he is of high ability, given by $\hat{q}^i \in [0, 1]$ —he gains no utility from the group decision *per se*. This is standard in models of career concerns (see Prat 2005; Fox and Van Weelden 2012; Levy 2007).

The timing of the game is as follows. Each committee member learns his ability and signal. Members can then communicate in two steps. First, in a simple straw poll, each member simultaneously announces a message $m^i \in \{m_b, m_r, m_\emptyset\}$, i.e., raises his hand in favor of B or R or abstains. Next, each member simultaneously announces a message $\tau^i \in \{\tau_h, \tau_l, \tau_\emptyset\}$, i.e., says he is a high type, low type, or remains silent. After these two stages of communication, the committee has access to a coordination device—a publicly observable random draw from a uniform distribution $u[0, 1]$ which allows them to coordinate on a group decision. Finally, each member casts a vote $v^i \in \{v_B, v_R\}$, the group decision is taken and the true state is revealed to everyone.

Once the group decision and state are revealed, the principal updates her beliefs on member abilities using all available information. Of course, the available information depends on the level of transparency. We compare three different regimes: secrecy, transparency, and mild transparency. Under *secrecy*, the principal only observes the group decision D . Under *transparency*, she witnesses each member's messages m^i, τ^i , and final vote v^i . Under *mild transparency*, she observes only the group decision and how each individual votes.

A committee member's strategy consists of a communication strategy and a voting strategy. A communication strategy is a pair (\mathbf{m}^i, τ^i) , where \mathbf{m}^i is a mapping from (s^i, t^i) into a probability distribution over messages $\{m_b, m_r, m_\emptyset\}$ and τ^i is a mapping from (s^i, t^i) and messages exchanged in the straw poll into a probability

distribution over announcements $\{\tau_h, \tau_l, \tau_\emptyset\}$. A voting strategy \mathbf{v}^i is a mapping from signal, ability, and messages exchanged in both communication rounds into a probability distribution over votes $\{v_B, v_R\}$.

We study symmetric perfect Bayesian equilibria under the three transparency regimes. As is standard in voting games, we restrict attention to strategies that are not weakly dominated. As talk is cheap and payoffs depend on the principal's beliefs, there will be many equilibria in each of the three settings. We apply some restrictions to reduce the set of equilibria. Firstly, we restrict attention to equilibria in which h types vote in line with their private signal. Though other equilibria exist, it seems reasonable that h types will vote for the true state (especially as the principal knows they are perfectly informed). Our lab results support this restriction— h types vote to signal 98.2 percent of the time, rising to 100 percent in the final 5 periods. Secondly, we follow the cheap talk literature (Crawford and Sobel 1982; Ottaviani and Sørensen 2001; Chen, Kartik, and Sobel 2008) in focusing on the most informative equilibrium, i.e., in each setting we look for equilibria where the maximal amount of information is shared across the two communication stages. Finally, we ignore equilibria with inverted language; for example, where a message m_r is interpreted as “I have a signal in favor of state B ” or an announcement τ_h is interpreted as “I am of low ability.” In Section IIIA, we discuss our focus on the most informative equilibrium in more detail.

II. Equilibrium

A. Secrecy

As the principal can only see the group decision, she must hold the same posterior \hat{q} for each individual. Committee members therefore have a common interest.

PROPOSITION 1: *In the most informative equilibrium under secrecy, all members truthfully reveal their signal and ability, then jointly implement the decision with the highest posterior probability of matching the state. In case of a balanced posterior after two conflicting signals from low ability types, the committee coordinates on implementing each decision with probability 0.5. The probability of a group mistake in each state is $(1 - q)^2(1 - \sigma)$. Each member earns a payoff of $q/(1 - (1 - q)^2(1 - \sigma))$ if the group decision matches the state, and zero otherwise.*

The proofs of all propositions in this paper appear in the Appendix. It is hardly surprising that players with a common interest share their information. Coughlan (2000) shows that allowing communication between players with a common interest can lead to full information aggregation. There are other, less informative, equilibria in which players babble in one or both stages of communication. In the laboratory, however, Guarnaschelli, McKelvey, and Palfrey (2000) and Goeree and Yariv (2011) show that players with a common interest are overwhelmingly truthful.

The proposition states that a committee picks the decision most likely to match the state, and mixes with probability 0.5 in case of a balanced posterior belief. To see why this is the case, let all (lb, lr) committees implement B with (possibly

degenerate) probability π . The expected utility of each committee member in equilibrium is then

$$\frac{0.5\pi q}{1 - (1 - q)^2(1 - \sigma)^2 - 2(1 - \pi)(1 - q)^2\sigma(1 - \sigma)} + \frac{0.5(1 - \pi)q}{1 - (1 - q)^2(1 - \sigma)^2 - 2\pi(1 - q)^2\sigma(1 - \sigma)},$$

where the first term corresponds to the expected evaluation in state B and the second to state R . The committees have to be indifferent between the two decisions in order to mix, i.e., both terms have to be equal. This is only the case if $\pi = 0.5$, when all (lb, lr) committees mix equally between implementing B and R . If (lb, lr) committees made one decision more often than another, the principal would take this into account, lowering his posterior belief following a correct decision in this state vis-à-vis the other state. Thus, the fact that committee mistakes are equally likely in both states stems entirely from the reputational concerns of l types.

B. Transparency

Under this regime, the principal sees all stages of communication and observes each individual's vote.

PROPOSITION 2: *In the most informative equilibrium under transparency, all members truthfully reveal their signal, information on abilities cannot be credibly communicated, and members vote according to their signal. The probability of a group mistake in state B is $(1 - q)^2(1 - \sigma)^2$, while the (much larger) probability of a mistake in state R is $1 - (q + (1 - q)\sigma)^2$. Each member earns a payoff of $q/(q + (1 - q)\sigma)$ if their individual vote matches the state, and zero otherwise.*

Here we cannot have full information sharing. If players were truthful about their ability with the principal watching, then l type players would receive a utility of zero. Instead, in an effort to appear competent, l types will mimic the strategy of h types so that in equilibrium no information on ability is revealed. This means information aggregation is incomplete when compared to secrecy: while signals are shared truthfully, players cannot differentiate the quality of those signals.

Even though some valuable information is shared, players actually ignore this when deciding how to vote. An lb player who sees an m_r message in the straw poll believes R is the most likely state. However, this lb player will be better off voting v_B in the final vote. The same holds for an lr player who sees an m_b message in the straw poll. Why is this? An h type would never vote against his signal, therefore any l type who switches his choice between the straw poll and the final vote can be

identified by the principal.⁷ So, l types will stick to their straw poll announcement in the final vote even when they believe it is wrong. This effect is similar to that found in Prendergast and Stole (1996) and Majumdar and Mukand (2004) in single agent models of managerial decision-making and political policy experimentation, respectively.

The combination of less information sharing and players “sticking to their signals” gives a higher aggregate probability of committee mistakes under transparency than under secrecy. However, there are differences across states. Transparency leads to far more mistakes in state R but somewhat fewer mistakes in state B . The reason is the different behavior of (lb, lr) committees under each regime. Under secrecy, the committee will mix equally between implementing each decision, under transparency everyone votes according to their signal—meaning the status quo will always be upheld. If the true state is B , an (lb, lr) committee will never make a mistake under transparency, while it will half of the time under secrecy.

C. Mild Transparency

Here, the principal cannot observe any communication, but she does observe the individual votes of committee members as well as the final group decision. This corresponds to many real world cases where voting records are released but discussions are kept secret. It also reflects the setting where the introduction of transparency leads to the emergence of secret pre-meetings.

PROPOSITION 3: *In the most informative equilibrium under mild transparency, communication, payoffs, and the probability of mistakes are the same as under secrecy. In the voting stage, each member votes for the decision with the highest posterior probability of matching the state.*

The proposition implies that, if players do indeed play the most informative equilibrium, the emergence of premeetings would actually lead to more information aggregation and fewer mistakes than would otherwise be the case under transparency. This contrasts with the result of Swank and Visser (2013)—pre-meetings can undo the benefits of transparency in their setting. However, in our setting there are also other equilibria with less information sharing under mild transparency, which are preferred by h types. We discuss this aspect further in Section IIIA.

There is a small difference between the most informative equilibria under secrecy and mild transparency—though it makes no difference to the probability of mistakes. Under secrecy, only the group decision matters for belief updating; decisions in favor of B need not be unanimous. Under mild transparency, however, any member whose vote doesn’t match the state will be revealed as an l type. Moreover, the principal knows that committee members perfectly share their information—so in an (h, l) committee both will vote correctly. The knock-on effect is that (lb, lr)

⁷This equilibrium is sustained by beliefs that any member who switches in the final vote is an l type. As nobody actually switches in equilibrium, this information set is never reached and so we are free to choose any beliefs of the principal off the equilibrium path.

committees must also vote unanimously. Otherwise, both will be revealed as l types, even if one player votes correctly.

D. Optimal Level of Transparency

A natural question is which transparency regime the principal would prefer ex ante, assuming the most informative equilibrium is played. As theory predicts the same outcome under secrecy and mild transparency, we compare the principal's utility under secrecy and transparency. In the analysis thus far, the principal's utility was irrelevant—committee members act to maximize the principal's beliefs of their ability, not to maximize her utility. To analyze the principal's welfare we must be more specific about her utility function than before. Let the principal gain a utility of x for a correct decision in state B and a utility of $(1 - x)$ in state R , where $x \in [0, 1]$, i.e., she may care more about correct decisions in one state than the other. The following proposition shows that the optimal level of transparency depends on x .

PROPOSITION 4: *There always exists an $x^* > 0.5$, such that if $x < x^*$, the principal prefers secrecy, while if $x > x^*$, she prefers transparency.*

Why might transparency be preferred despite leading to more committee mistakes? By voting according to their signals, a transparent committee will seldom make the unanimous decision needed to implement R . As a result, they will make more mistakes than a secret committee in state R but less in state B . A principal with a high value of x is more wary of losses in state B , so will prefer transparency to secrecy.

III. The Experiment

The model gives us several key predictions that we can formulate into testable hypotheses:

- H1:** *Decision accuracy will be higher under secrecy than transparency if the true state of the world is R but lower if the state is B .*
- H2:** *Mild transparency will produce the same decision accuracy as secrecy.*
- H3:** *Deliberation will be fully truthful under secrecy and mild transparency but not under transparency.*
- H4:** *Voting will involve all members “sticking to signal” under transparency, even if that doesn't match the member's posterior belief. Under secrecy and mild transparency, members will implement the decision with the highest posterior probability of matching the state. If the posterior is balanced, the committee will vote so as to implement each decision with equal probability.*

H5: *Principals will use all available information in updating their posterior. Members who are observed voting against the true state or whose message and vote do not agree will receive an evaluation of zero.*

To test these hypotheses, we ran a laboratory experiment with three treatments—one for each level of transparency. Before describing the experimental setup in detail, we first discuss possible reasons our hypotheses may fail and key design choices in this context.

A. Why Might H1–H5 Fail?

Subtle Strategies.—Several of the hypotheses rely on quite subtle strategic mechanisms. H1 relies on the fact that under secrecy an (lb, lr) committee should realize that they must implement each decision with equal probability and that this involves coordination among group members. H4 relies on committee members believing that if the principal sees them report one message but voting differently, he will infer they are an l type. Furthermore, H3 relies on members believing that all communication about ability is useless under transparency yet still playing the strategy ascribed to them in this babbling equilibrium. This seems to be a strong assumption as deviating unilaterally to truth-telling would not change the evaluation if the principal indeed disregarded messages. A number of experimental studies on strategic communication have repeatedly found that many subjects often tell the truth even when they would benefit from lying strategically (e.g., Cai and Wang 2006; Wang, Spezio, and Camerer 2010; Goeree and Yariv 2011; Battaglini and Makarov 2014; Le Quement and Marcin 2016). Thus, we wanted to see how strategically players behave in our environment, and if possible deviations affect aggregate outcomes.

Equilibrium Selection.—Our predictions rely on the assumption that members play the most informative equilibrium. This selection is common in the cheap talk literature (Crawford and Sobel 1982; Ottaviani and Sørensen 2001; Chen, Kartik, and Sobel 2008). A reason given for the focus on the most informative equilibrium in the canonical sender-receiver game is that it is *ex ante* optimal. In our setup, the most informative equilibrium is optimal for h types under secrecy and transparency but not under mild transparency, where a babbling equilibrium is optimal. Nonetheless, H2–H4 suppose that the most informative equilibrium is played under mild transparency. In this equilibrium, both h and l types share all of their information and vote unanimously for the decision most likely to match the state. However, the equilibrium most preferred by h types involves high types babbling about their signal (e.g., by mixing between announcing m_b and m_r). Each member then votes according to his own signal. In fact, there is a continuum of equilibria that are less informative than the most informative, all of which are preferred by h types (see Appendix B). Apart from the level of truth-telling by h types, another difference between the most informative equilibria and those preferred by h types is that in the latter the payoff of each committee member will be independent of the other member's vote. In a setting where h types prefer equilibria in which less information is shared, it is not obvious that the most informative equilibrium will be played and,

thus, to what degree H2–H4 will hold. Structurally similar situations are described in Farrell and Gibbons (1989, section 3) and Sobel (2013, section 2). The latter calls for experiments to study the predictive power of the most informative equilibrium in cheap talk games that feature additional equilibria, which are preferred by one type of player.

Belief Updating.—In theory, the principal’s beliefs should be correct and the strategies of members should be optimal given the principal’s beliefs. These beliefs are especially important in this setting as they entirely determine the payoffs of committee members. If the principal has incorrect beliefs or does not update correctly, a failure of H5, this may have strong knock-on effects on the behavior of committee members and aggregate error rates (H1–H4).

We believe that this interaction effect may be important in real world committees, so we chose to have subjects play the role of the principal and elicit beliefs using a proper scoring rule. Another option would have been to have a computer take the role of the principal. But how a principal responds to each piece of information depends on which equilibrium is played. As we saw, it is not obvious which equilibrium will be played under every transparency regime. Using a computer with a fixed updating rule would likely push members to play one particular equilibrium, as we would have to tell players how the computer updates. Thus, we would not be able to learn about which equilibria are more likely to be played. Our approach also has a clear advantage over more implicit incentives (such as a rehiring rule) because it allows us to study how the different pieces of information are processed and whether committee members best-respond to this belief updating.

Open Chat.—Real world deliberation does not tend to be as structured as in our model. Rather, committee members have open discussions before moving to a formal vote. Yet, there is evidence that people behave differently depending on whether their communication is structured or unstructured. Goeree and Yariv (2011) find that open deliberation in the lab leads to perfect information sharing in a setup where committee members have a clear incentive to keep their information private.

We were keen to see if our various hypotheses are supported in the more realistic setting of open chat. Therefore, we chose to allow for free-form communication in the second deliberation stage rather than simultaneous announcements and a public coordination device. The equilibria characterized in the previous section are unaffected by this change—the only sensible way subjects could make use of the open chat is to share information on their ability and to coordinate. We decided against open chat in the first stage of deliberation as this might lead to strange behavior under transparency.⁸

⁸There is an incentive for l types to learn from the announcement of their partner. This might mean that under transparency a player who is slower in revealing his signal is interpreted as an l type by the principal. This, in turn, would lead to a “rush to announce”—where both players announce their signals immediately. This rushed open chat would then produce outcomes equivalent to the simultaneous announcements we use in the lab.

TABLE 1—EXPERIMENTAL SESSIONS

	Sessions	Matching groups	Subjects
<i>Secrecy</i>	2	5	45
<i>Transparency</i>	2	4	42
<i>Mild transparency</i>	2	5	45

Notes: All sessions were run at the DeSciL Lab at the ETH Zurich in May 2013 with 132 students (48 percent female) studying various majors at the ETH or the University of Zurich. Psychology students were not recruited.

B. Experimental Design

We ran six sessions, two for each transparency regime (see Table 1). Each session consisted of 20 rounds. In the first round, subjects were randomly assigned to matching groups of nine people, which remained fixed for all 20 rounds.⁹ Within each matching group, groups of three were randomly formed in every round. This was done to avoid the emergence of reciprocal behavior and at the same time provide a larger number of independent matching groups. In each group, one person was randomly assigned the role of “observer” (the principal) and the other two the role of “voters” (committee members). Screenshots and instructions are in the online Appendix. We set $\sigma = 0.55$ and $q = 0.25$ in all treatments.

There was a blue jar ($S = B$) and a red jar ($S = R$). The blue jar contained 11 blue and 9 red balls, the red jar contained 11 red and 9 blue balls. In each round, one jar was selected by a fair coin toss.¹⁰ A “well-informed voter” (h type) received a ball matching the color of the jar, an “informed voter” (l type) received a ball that was drawn from the jar.

Committee Members.—On the first screen, a member learns his type and the color of his ball, and then sends a message {red, blue, not specified} to the other member. On the next screen, he sees the message from his partner and has the opportunity to chat with him for 90 seconds. On the third screen, he can review the communication and make his final decision by voting for red or blue. On the final screen of each round, he learns the other member’s type, how the other member voted, the group decision, the true color of the jar, and his own payoff.

The Principal.—On the first screen, the player learns that she is a principal. Under secrecy, she sees nothing else until voting is concluded. Once the committee has voted, she sees the true color of the jar and which group decision was taken. On this screen, she has to indicate her belief about the ability of a randomly chosen committee member, by entering a probability of him being an h type. Under mild transparency, she also observes individual votes and has to evaluate each member separately. Under transparency, she can see committee members’ messages and follow their chat in real time. After the vote, she sees the same information as under

⁹In one session, we had only 15 subjects and therefore one matching group.

¹⁰This has become the standard task in information aggregation experiments (e.g., Guarnaschelli, McKelvey, and Palfrey 2000; Battaglini, Morton, and Palfrey 2009; Goeree and Yariv 2011; Bhattacharya, Duffy, and Kim 2014; Bouton, Castanheira, and Llorente-Saguer 2016; Bouton, Llorente-Saguer, and Malherbe 2017).

mild transparency, but can also review the communication (messages and chat) before evaluating each member. In all three treatments, on a final (feedback) screen she learns each member's true type, their votes, the resulting group decision, and her own payoff.

Payoffs.—The principal's payoff for a correct group decision was 3 points if the jar was blue and 1 point if the jar was red. This corresponds to $x = 0.75$.¹¹ In addition, the principal earned a number of points between 0 and 100 for accurate evaluation of the committee members' types. Her points were determined by the following quadratic scoring rule:

$$\text{Points} = \begin{cases} 100 - \frac{1}{100}(100 - \text{Pr}_j(t_i = h))^2 & \text{if committee member } i \text{ is an } h \text{ type} \\ 100 - \frac{1}{100}(\text{Pr}_j(t_i = h))^2 & \text{if committee member } i \text{ is an } l \text{ type,} \end{cases}$$

where $\text{Pr}_j(t_i = h)$ denotes principal j 's submitted probability (in percent) that committee member i is an h type. This rule makes it optimal for expected payoff maximizing subjects to truthfully enter their beliefs (see, e.g., Nyarko and Schotter 2003)—participants were told as much in the instructions.¹² Under secrecy, the principal earned points from her single evaluation, while under the other treatments one of her two evaluations was randomly chosen. We kept the principal's payoff from correct group decisions low relative to the payoff from accurate evaluations so as to limit the potential effects of social preferences. The concern here is that committee members would strive to make more correct decisions in state B as these are more valuable to the principal. Each committee member's payoff was determined by the principal. If she judged that member i had a y percent chance of being an h type, that committee member would gain $2y$ points.

Four rounds were randomly chosen at the end of the session and the points earned in these rounds were converted to Swiss Francs at a rate of 1 point = CHF 0.15 (at the time of the experiment CHF 1 was worth USD 1.04). Subjects spent about 2 hours in the lab and earned CHF 47 on average, in addition to their show-up fee of CHF 10. Earnings per hour are comparable to an hourly wage for student jobs in Zurich.

IV. Experimental Results

A. Aggregate Committee Behavior

In the model, we spoke of probabilities of a committee mistake; the analog in the data are observed error rates, i.e., the share of committees that implement the wrong

¹¹ We chose $x = 0.75$ to see if a principal really would be better off under transparency, as Proposition 4 suggests. The behavior of all players should be the same for any level of x , it only alters the final payoff of the principal.

¹² More complicated belief elicitation procedures have been proposed for risk-averse subjects (e.g., Offerman et al. 2009). However, to avoid making the instructions overly complicated (and thus distracting subjects from the game) we chose to implement a standard quadratic scoring rule. We follow Schotter and Trevino (2014) in telling subjects that truthfully reporting their belief maximizes their expected payoff.

TABLE 2—OBSERVED AND EQUILIBRIUM ERROR RATES

True color	$S = B$		$S = R$		Total	
	Observed	Equilibrium	Observed	Equilibrium	Observed	Equilibrium
<i>Secrecy</i>	28.3 (5.5)	25.3 [27.2]	25.8 (1.8)	25.3 [24.8]	27.0 (3.1)	25.3 [26.0]
<i>Transparency</i>	15.8 (1.0)	11.4 [8.2]	50.7 (6.4)	56.1 [61.9]	33.3 (3.5)	33.8 [33.9]
<i>Mild transparency</i>	15.5 (1.2)	25.3 [24.0]	46.1 (3.6)	25.3 [25.3]	30.8 (1.8)	25.3 [24.7]

Notes: Equilibrium = ex ante expected error rate in the most informative equilibrium. Ex post expected equilibrium error rates, i.e., theoretical predictions after the realization of types and signals, are reported in brackets. Standard errors of the observed error rates (in parentheses) are adjusted for clustering in matching groups.

decision. In Table 2, we compare the observed error rates with equilibrium predictions. We see immediately that the most informative equilibrium almost perfectly predicts error rates under secrecy. In the model, transparency generates more errors than secrecy in state R but less in state B . Indeed, this is borne out in the data.^{13,14}

RESULT 1: *As predicted by H1, error rates are very different between secrecy and transparency and the differences go in the predicted direction.*

Principals earned more points from group decisions under transparency than under secrecy—as the model with $x = 0.75$ predicts—though this difference was not statistically significant. The most informative equilibrium predicts the same play under mild transparency and secrecy. Strikingly, however, error rates under mild transparency are statistically indistinguishable from those of transparency.

RESULT 2: *Contrary to H2, decision accuracy is very different under secrecy and mild transparency.*

What explains the differences in error rates? To investigate, in Table 3, we analyze groups that “should” have differences across treatments, i.e., those with conflicting signals. In (h, l) groups, we expect no errors in state B under any treatment—the h type can unilaterally implement the correct decision. This is what we see in the data (apart from a single observation under mild transparency). In state R , we expect no errors from conflicted (h, l) groups under secrecy, but a large number under transparency. We do find a significant difference between the two treatments, though the 44.8 percent error rate in transparency falls well short of the predicted 100 percent. This suggests that, under transparency, a significant share of l types are not voting according to their signal. In mild transparency, the error rate in state R

¹³When compared to secrecy, the transparent committees performed significantly and substantially worse in state R at conventional levels (Wald-test, $p < 0.001$) but better in state B (Wald-test, $p = 0.024$).

¹⁴For every parametric test of treatment differences presented in the paper, we also conducted a nonparametric (rank sum) test, treating each matching group as one observation. In the following, we will only report these test results if they lead to different test decisions than the parametric tests at the 5 percent significance level. Also in the following, if we speak of statistically (in)significant differences without reporting the p -values we mean statistically different at the 5 percent level.

TABLE 3—ERROR RATES IN GROUPS WITH CONFLICTING SIGNALS

		$S = B$		$S = R$	
		Observed	Equilibrium	Observed	Equilibrium
$\{h, l\}$ groups	<i>Secrecy</i>	0	0	0	0
	<i>Transparency</i>	0	0	44.8 (8.3)	100
	<i>Mild transparency</i>	3.3 (3.4)	0	22.6 (5.5)	0
$\{lb, lr\}$ groups	<i>Secrecy</i>	38.7 (15.3)	50	54.3 (11.0)	50
	<i>Transparency</i>	27.5 (8.9)	0	83.3 (10.3)	100
	<i>Mild transparency</i>	27.9 (6.8)	50	75.7 (4.4)	50

Note: Standard errors of the observed error rates (in parentheses) are adjusted for clustering in matching groups.

is significantly higher than the predicted value of zero, yet significantly lower than under transparency.¹⁵

Clearly, information is not being fully aggregated here. Next, we turn to (lb, lr) groups. Under secrecy, we expect each decision to be implemented half of the time, otherwise committee evaluations would be lower in one state than the other. Indeed, the data is relatively close to an even split. In the model, an (lb, lr) group under transparency would always implement B . Table 3 does show much higher error rates in state R than B , but the difference between states is not as extreme as in theory. Some coordination on implementing R must be taking place here, though the accuracy is significantly lower than in (h, l) groups. Finally, observed error rates under mild transparency are virtually the same as those under transparency.¹⁶

B. Individual Committee Member Behavior

Deliberation.—We begin our analysis of individual behavior with the first deliberation stage, the straw poll. Recall that in the most informative equilibrium all players truthfully reveal their signals, regardless of the transparency regime.

Table 4 shows two cases that clearly violate this: l types under transparency, and h types under mild transparency. In the first case, 8.3 percent of low types lie about their signal and another 10.5 percent stay silent, and both percentages are significantly different from zero. If an l type lies and then sticks with this lie, this would not be very costly, after all, σ -signals are not very informative. Why do some l types remain silent? Perhaps they hope to learn more about the true state and then vote accordingly. When we turn to the principals' reactions, we will see whether this strategy pays off. In the second case, 19.2 percent of high types lie about their signal and another 4.8 percent stay silent.¹⁷ This behavior of h types is clearly inconsistent

¹⁵The p -value of a Wilcoxon rank-sum test comparing mild transparency to transparency is 0.065 while that of a Wald test is $p = 0.026$.

¹⁶Note that none of the differences in error rates between secrecy and the transparency treatments in these groups is statistically significant at the 5 percent level. In state R , however, the differences between secrecy and transparency ($p = 0.084$) and between secrecy and mild transparency ($p = 0.06$) are significant at the 10 percent level (Wald-tests). The differences in error rates between the two transparency regimes are not statistically significant even at the 10 percent level. The overall error rate under transparency is significantly higher in (lb, lr) than in (h, l) groups (Wald test $p < 0.001$).

¹⁷Both percentages are significantly larger than zero (Wald-tests, $p < 0.001$).

TABLE 4—(NON)-TRUTHFUL STRAW POLL MESSAGES FROM DIFFERENT TYPES

	Type	Lying	Silent
<i>Secrecy</i>	<i>h</i> type	0	0
	<i>l</i> type	0.7 (0.3)	1.8 (1.2)
<i>Transparency</i>	<i>h</i> type	1.4 (0.9)	0.7 (0.7)
	<i>l</i> type	8.3 (2.1)	10.5 (2.8)
<i>Mild transparency</i>	<i>h</i> type	19.2 (6.6)	4.8 (2.0)
	<i>l</i> type	3.5 (1.1)	3.3 (1.4)

Notes: Percentage of non-truthful messages (lying) and “not specified” messages (silent). Standard errors of the observed error rates (in parentheses) are adjusted for clustering in matching groups.

with the most informative equilibrium. It seems that some *h* types are hoarding their information in an attempt to separate from *l* types. Indeed, in Appendix Section B, we show that there are equilibria of this type that yield higher payoffs for *h* types than full information sharing. Overall, however, the significant degree of lying is not enough to make the straw poll uninformative for *l* types.

Next, we turn to communication in open chat. After exchanging information about signals in the straw poll, the only relevant information left to discuss is member types.¹⁸ We coded whether a type was announced in the chat, and if so, what type was announced. Under secrecy and mild transparency, we expect truthful revelation of types. This is, by and large, what we observe in Table 5, though there is a small amount of lying under mild transparency. This suggests that the deviation of mild transparency from theory is explained by the frequent straw poll lies of *h* types rather than anything in the open chat.

We have no theoretical prediction for chat behavior in the transparency treatment, only that chat should be uninformative in equilibrium. In Table 5, we see that almost half of subjects announce a type, with 80.1 percent of those claiming to be an *h* type. This does not make communication informative per se, but a look at the truthfulness of these announcements shows that this stage is, in fact, informative. Of the 19.9 percent who announce an *l* type almost none are lying. In contrast, 51.4 percent of those claiming to be an *h* type are lying. This rate of lying is actually lower than what we would expect if chat was pure babbling: given that *h* types only make up 25 percent of the population, 75 percent of *h* claims should be lies for the chat to be uninformative. What about those who don't announce a type? Eighty-nine percent of them are *l* types. In sum, a member's announcement (or lack thereof) is actually informative about his type.

RESULT 3: *As predicted by H3, deliberation is truthful under secrecy, while it is not under transparency. Contrary to H3, deliberation is not fully truthful under mild transparency.*

Voting.—Do *h* types really vote as we assumed in our refinement? Yes, they vote according to their signal 98.2 percent of the time across all treatments. Next, in

¹⁸We give examples of chats under secrecy and transparency in the online Appendix.

TABLE 5—CHAT MESSAGES ABOUT TYPE

	Report type		Claim	Lying
<i>Secrecy</i>	91.5 (1.5)	<i>h</i> type	25.9 (1.1)	0.7 (0.6)
		<i>l</i> type	74.1 (1.1)	0.5 (0.3)
<i>Transparency</i>	49.5 (8.7)	<i>h</i> type	80.1 (3.9)	51.4 (5.8)
		<i>l</i> type	19.9 (3.9)	1.8 (1.1)
<i>Mild transparency</i>	89.0 (2.1)	<i>h</i> type	21.5 (1.0)	7.8 (5.3)
		<i>l</i> type	78.5 (1.0)	3.8 (1.4)

Notes: The first column reports the fraction of committee members who report a type, the second reports the fraction of low- and high-type claims out of those reports, and the third reports the fraction of lies out of these claims. Standard errors (in parentheses) are adjusted for clustering in matching groups.

TABLE 6—PERCENTAGE OF *l* TYPES VOTING AGAINST THEIR SIGNAL WHEN PARTNER REPORTS A CONFLICTING SIGNAL

		Other's claim: <i>h</i>	Other's claim: <i>l</i>	No claim
Own signal blue				
<i>Secrecy</i>	Observed	100	42.4 (11.3)	50.0 (36.7) = 2 obs.
	Equilibrium	100	≥ 50	
<i>Transparency</i>	Observed	47.2 (5.9)	44.4 (9.4) = 9 obs.	37.5 (10.6)
	Equilibrium	0	0	
<i>Mild transparency</i>	Observed	91.3 (4.5)	32.9 (1.9)	50.0 (12.2) = 6 obs.
	Equilibrium	100	50	
Own signal red				
<i>Secrecy</i>	Observed	100	45.9 (12.6)	40.0 (17.6) = 2 obs.
	Equilibrium	100	≤ 50	
<i>Transparency</i>	Observed	47.4 (5.1)	28.6 (7.9) = 2 obs.	57.1 (13.1) = 4 obs.
	Equilibrium	0	0	
<i>Mild transparency</i>	Observed	95.23 (4.2)	32.4 (6.1)	34.8 (5.1)
	Equilibrium	100	50	

Notes: Standard errors of the observed error rates (in parentheses) are adjusted for clustering in matching groups. Number of observations are reported if less than 10.

Table 6, we look at how many *l* types vote against their signal when the other member announces a conflicting signal.

Under secrecy and mild transparency, each member should announce their type to aid information aggregation. Indeed, we saw overwhelmingly truthful revelation of types in Table 5, meaning that *l* types should believe those who claim to be *h* types. If one member claims to be an *hr* (*hb*) type, an *l* member should always vote v_R (v_B). We see that this is always the case under secrecy, and almost always under mild transparency. This is despite the moderate amount of deception under mild transparency. Under these two treatments, we would expect (*lb*, *lr*) committees to implement each decision half of the time. Under mild transparency, this involves players unanimously voting for each decision 50 percent of the time. Under secrecy, players have a wider set of strategies, though all involve *lb* members voting against their signal at least 50 percent of the time and *lr* members doing so at most 50 percent of the time. In Table 6, we see that subjects get very close to 50 percent under secrecy and a little lower under mild transparency.

TABLE 7—EVALUATIONS

		Evaluation		
		Average	Decision correct	Decision wrong
<i>Secrecy</i> (group decision)	Equilibrium	25	33.5	0
	Observed	41.6 (1.5)	54.5 (3.2)	6.9 (1.9)
	Observed for h types	53.5 (3.4)	53.5 (3.4)	
	Observed for l types	37.4 (1.5)	55.0 (3.3)	6.9 (1.9)
<i>Transparency</i> (individual decision)	Equilibrium	25	37.7	0
	Observed	37.2 (3.0)	48.2 (1.7)	14.7 (4.0)
	Observed for h types	60.8 (3.2)	62.8 (3.8)	
	Observed for l types	29.4 (2.9)	40.2 (2.2)	14.6 (4.0)
<i>Mild transparency</i> (individual decision)	Equilibrium	25	33.5	0
	Observed	35.7 (1.9)	47.1 (1.9)	11.7 (1.6)
	Observed for h types	47.7 (1.2)	48.3 (1.1)	
	Observed for l types	31.9 (2.3)	46.4 (2.6)	11.8 (1.6)

Note: Standard errors (in parentheses) are adjusted for clustering in matching groups.

The really interesting case is transparency. The theory predicts that no player should ever switch their final vote from their straw poll announcement. In the lab, however, we see 30.8 percent of l types switching from their announcement¹⁹. This is puzzling; switching from your straw poll announcement should alert the principal that you are an l type, resulting in an evaluation of zero. In the next section, we will investigate whether such switching was indeed punished.

RESULT 4: *As predicted by H4, under secrecy and mild transparency committee members implement the decision most likely to match the state given the posterior. Also as predicted by H4, groups with a balanced posterior implement R roughly half of the time under secrecy. However, they implement R less often under mild transparency, and under transparency many l types do not vote according to their announced message, which contradicts H4.*

C. Evaluations

Next, we study how the principals evaluate. Table 7 shows that average evaluations are too high in all treatments and even significantly positive after wrong decisions.

Nonetheless, evaluations are much higher for correct decisions than incorrect ones, meaning the incentives to make a correct decision are about as strong as in theory.²⁰ In each treatment, h types receive higher evaluations on average than l types. Under transparency, this is even the case when both types vote correctly. This shows that principals can distinguish the two types well.

¹⁹Note, that this number is slightly different from the percentage of l types voting against their signal as some of them stay silent in the straw poll.

²⁰Evaluations look very similar if split by the state of the world (not reported here). There is no evidence that principals reward correct decisions in state B more than in state R , which might have been the case if reciprocity had played a role.

TABLE 8—EVALUATION RESPONSES

	Sec. (1)	Sec-Lin. (2)	Trans. (3)	Trans-Lin. (4)	Mild (5)	Mild-Lin. (6)
Group decision D wrong	-47.6 (4.2)	-36.1 (1.7)	0.7 (2.9)	-1.6 (2.9)	-0.9 (3.3)	-2.3 (4.4)
Individual vote v wrong					-34.8 (2.1)	-32.7 (3.0)
<i>Combinations of m and v, reference category: Message and vote are correct</i>						
v wrong, m wrong			-34.0 (4.9)	-34.4 (3.5)		
v wrong, m right			-31.3 (3.9)	-19.3 (10.1)		
v right, m wrong			-20.6 (3.8)	-27.9 (3.0)		
v right, silent in straw poll			-15.5 (4.6)	-30.3 (3.7)		
v wrong, silent in straw poll			-40.6 (5.2)	-32.3 (3.0)		
Claimed to be type h in chat			15.3 (2.1)	29.7 (2.4)		
Claimed to be type l in chat			-17.5 (5.9)	-7.5 (3.8)		
Constant	54.5 (3.4)	36.1 (1.7)	46.2 (2.2)	28.2 (3.6)	47.1 (1.9)	35.5 (2.5)
Observations	600	600	560	560	600	600
Clusters	5	5	4	4	5	5
R^2	0.49	0.13	0.29	0.33	0.31	0.14

Notes: Standard errors (in parentheses) are adjusted for clusters in matching groups. The dependent variable in the comparison columns 2, 4, and 6 is rescaled to the same range as the dependent variable in 1, 3, and 5, and takes the value 100 if the subject is an h type and 0 if not.

In Table 8 (models 1, 3, and 5), we regress the principal's evaluations on the information she sees before evaluating. Making the wrong group decision greatly reduces a member's evaluation under secrecy, but not under the other two treatments (where individual voting is observed). If the most informative equilibrium was played under mild transparency, then the group decision should matter (as the principal expects members to share information). The fact that only individual decisions matter here means that h types face no risk in lying and leading l types astray. Under transparency, anything apart from a correct message and correct individual vote results in a lower evaluation—so switchers are actually punished. Interestingly, claiming to be an h type leads to a higher evaluation than remaining silent, while claiming to be an l type gives a lower evaluation. It seems the principals were aware of the relatively high level of truth-telling seen in Table 5.

D. (*Best*) Responses

In the previous sections, we saw that the observed behavior of both principals and committee members deviated somewhat from theoretical predictions. But perhaps committee members take the principals deviation into account when deciding on

their own strategy. Similarly, principals may anticipate that committee members will not behave in line with theory. Therefore, we now examine whether individuals are best responding to the behavior of other players.

Principals.—The evaluations that principals give are too high to maximize their payoffs. This may be because they find it difficult to update beliefs correctly.²¹ They do perform better over time. Table A1 in the online Appendix shows that principals lower their evaluations in later rounds, which suggests learning. However, committee members who vote incorrectly often receive positive evaluations, even in the final rounds of play. These positive evaluations occur when updating is extremely simple and an evaluation of zero would maximize the principal's payoff. This suggests that principals have social preferences toward committee members as their evaluations directly affect members' payoffs. Another factor that might play a role is that the scoring rule does not correct for risk-aversion. Entering 50 percent as an evaluation gives you a certain payoff of 75, whereas any other evaluation leaves you with a lottery if you attach positive probabilities to either type. Risk-averse subjects would, therefore, be drawn toward 50 percent, which is double the ex ante probability of a high type q .

However, we are primarily interested in how evaluations differ across regimes and player behavior. Does the principal react in the right direction to each piece of information? To answer this, we can compare columns 1, 3, and 5 to columns 2, 4, and 6 in Table 8. The odd columns have the principals' actual evaluations as the dependent variable, while the even columns are linear probability models where the dependent variable is binary, taking a value of 100 (percent) if the member is an h type and 0 if not. If principals were best responding, they would behave as in the even columns. We see that, in fact, they come very close. Principals use all pieces of information in the right direction, often with accurate magnitudes.

RESULT 5: *Principals' beliefs about committee members' types are too high. However, they interpret the available information correctly and update their beliefs in the correct directions providing partial support for H5. This creates clear incentives for committee members to behave differently under transparency than under secrecy.*

Committee Members.—Under secrecy, committee members best respond to the principal's behavior, as well as the behavior of other committee members. They first share their private information in the group and then implement the decision most likely to match the state. We know that behavior under mild transparency deviates from theory, but are committee members best responding given these deviations? High types vote to signal, as they should. Some tell the truth while others lie about their signal. Are some of these players doing better than others? Actually, as realized evaluations do not depend on group decisions, any level of information sharing by h

²¹This is consistent with the existing experimental literature. In a strategic communication experiment with computer and human evaluators, Meloso, Nunnari, and Ottaviani (2017) find that humans have difficulty in assessing the truthfulness of reports and in learning which strategies are played.

types is a best response. Most h types do reveal their information. This makes it optimal for l types to switch their vote when they encounter an h type with a conflicting signal. This is exactly what we observe in Table 6, l types are best responding under secrecy and mild transparency.

Under transparency, the vast majority of h types best respond by announcing their signal and type truthfully and then voting for the correct decision. However, a substantial fraction of low types do not best respond. In the most informative equilibrium, an l type should truthfully announce his signal and stick to this signal in the final vote. Whichever type he claims to be, this should be uninformative in equilibrium. The following analysis shows that it is in fact best for l types in the lab to truthfully announce their signal, claim to be an h type, and vote to signal.

In the straw poll, 8.3 percent of l types lie, while 10.5 percent remain silent. Neither is a best response. Staying silent here may be part of a learning strategy whereby an l type waits to see the signal of his partner. If the principal believed the straw poll was idle babbling, then this might pay off. However, from Table 8, we see that this strategy results in lower evaluations than announcing the signal truthfully.

In the open chat, many l types either truthfully reveal their type or simply avoid declaring a type. This behavior makes the chat informative and leads to lower evaluations than those who claim to be h types. The number of committee members who claim to be an l type goes down from 12.9 percent in the first 10 to 6.8 percent in the last 10 rounds, while the number remaining silent climbs from 45.7 to 55.3 percent.²² There are several explanations for the behavior of these l types. They may believe that the principal will ignore announcements on type as pure babbling. It could be that some l types are trying to achieve the correct group decision, so don't want to fool their partner by pretending to be an h type. Or, it may be that some players have an aversion to lying (Gneezy, Rockenbach, and Serra-Garcia 2013 and Fischbacher and Föllmi-Heusi 2013) which might also explain excessive truth-telling in the other experiments that we discussed in Section IIIA. In the other treatments, truth-telling is predicted anyway. However, we also see a significant number of subjects who lie about their signal under mild transparency and about their type in the chat under transparency.

In theory, we would never expect an l type to switch between his straw poll announcement and his final vote. Doing so would perfectly reveal his type. In the lab we actually see a large amount of switching from l types. The fact that evaluators give too much credit for correct votes after incorrect messages, and around half the claims of being an h type are legitimate, makes it not very costly for an l type to switch when his partner claims to be an h type with a conflicting signal.²³ Those who switch do not play a best response but it only costs them CHF 0.65, on average. This might explain why we see such a high level of switching, contrary to H4.

²²Both of these differences are statistically significant with $p = 0.031$ for the decrease in l claims and $p < 0.001$ for the increase in no claims.

²³The difference in evaluations is 4.3 percentage points and is not statistically different from zero (t -test, $p = 0.399$).

Switching if the other member does not claim to be an h type is more costly and, as a consequence, occurs less often (see Table 6).²⁴

RESULT 6: Most committee members best-respond to the actual belief updating of principals under secrecy and mild transparency. However, under transparency, a substantial number of subjects do not best-respond to the actual belief updating of principals.

V. Discussion

In this paper, we set out to investigate the potential of transparency to affect committee decision-making when committee members are career concerned. For this purpose, we constructed a model to study how the incentives of committee members to share their private information and vote strategically varies with the level of transparency, i.e., how much of the decision process the principal observes. We then formulated several hypotheses that we tested experimentally to see how well the model predicts behavior.

We found that behavior does not fully correspond to that predicted by the model. Principals tend to reward committee members even when they are clearly low-ability types, and members tend to be overly truthful to their own detriment (Results 3, 5, and 6). This second effect is found in several other experimental papers on strategic communication (cited above), especially when players chat openly as they do in our setting. Nonetheless, key predictions of the model are borne out in the lab; there is less information sharing and more committee mistakes under transparency than under secrecy (Results 1 and 3). The difference in outcomes is in fact driven by players understanding the subtle incentives to maximize their expected payoffs. Under secrecy, committees with a balanced posterior appear to realize that they should coordinate to implement each decision with equal probability. Under transparency, 87.2 percent of low types understand that they should not reveal their true type as this will hurt their evaluation (Result 3). Also under transparency, a majority of low types vote in line with their private signal even when the evidence points in favor of the other decision (Table 6). They do this as voting against their signal announcement would out them as a low type. Furthermore, principals use all the available information to update their belief in the correct direction (Result 5), though they fail to update perfectly.

One of our hypotheses is roundly rejected. In H2, we posited that error rates would be the same under secrecy and mild transparency. The data instead show error rates under mild transparency more in line with transparency (Result 2). In our theoretical analysis, we focused on the most informative equilibrium. However, under mild transparency this equilibrium is payoff-dominated for high ability types; they prefer equilibria in which no information is shared. It is perhaps unsurprising, therefore, that there is a substantial degree of deception from high ability types under mild transparency. As a result, group errors and player evaluations here are

²⁴The difference in evaluations is 9.9 percentage points (1.48 CHF) and is not statistically different from zero (t -test, $p = 0.409$). The number of observations is low for this scenario.

much closer to the transparency case than to secrecy, casting doubt on the predictive power of the most informative equilibrium for situations in which some players prefer a different equilibrium.

Our setting combined communication and private information on abilities. It is worth exploring how this relates to the other theoretical literature. If we had no communication as in Levy (2007), then players in the model would vote according to signal under transparency, but would vote strategically to overcome the voting rule under secrecy.²⁵ By adding communication we do even better under secrecy and do no worse under the other two regimes; communication is useless in avoiding committee errors under transparency, but may reduce errors under mild transparency if informative equilibria are played. If members were unaware of their ability level but could communicate as in Visser and Swank (2007), then the downside of transparency is curtailed. Once each member has announced a message, his reputational payoff is determined and so he no longer has the incentive to stick to his signal when voting. In their setting, transparency performs at least as well as secrecy and mild transparency. If instead, members know their own type, secrecy produces far fewer committee errors than transparency. This is because better information is aggregated under secrecy once players can reveal their abilities, and there is now an incentive to stick to your signal under transparency. To summarize, the optimal transparency regime according to theory will depend on how the principal weighs success in each state of the world (as we show in Section IID), and on whether committee members have private information on their ability or not. However, whether or not these theoretical results are good predictions for actual behavior are empirical questions and further experimental studies could be very useful to complement the existing body of theoretical results and our study's findings.

Despite the deviations from theory that we found in the lab, we can conclude, in line with the theory, that transparency indeed has a strong potential to influence behavior in committees and therefore appears to be an important element of institutional design.

APPENDIX A: PROOFS

A. Proof of Proposition 1

The most informative outcome of the two communication stages is, by definition, where each player reveals his signal and ability. Let a candidate for equilibrium be where members truthfully announce their signal and ability, then implement the decision with the highest posterior probability of matching the state. In case of a balanced posterior (after two conflicting signals from low ability types), the committee coordinates on implementing each decision with probability 0.5. Assuming this is an equilibrium, the probability of mistakes and member evaluations are as follows.

²⁵Transparency and mild transparency coincide here as the only thing to be observed is individual votes.

Mistakes.—A mistake occurs in state B when we have a pair (lr, lr) (with probability $(1 - q)^2(1 - \sigma)^2$) or a pair (lb, lr) who implement R (with probability $(1 - q)^2(1 - \sigma)\sigma$). Thus, the total probability of implementing decision R in state B is $(1 - q)^2(1 - \sigma)$. The case of a mistake in state R is symmetric.

Evaluations.—Given truthful communication, an incorrect group decision reveals that there are no h types on the committee; thus, each player gets an evaluation of zero. Instead, when a committee makes the correct decision the principal updates her beliefs in the following way:

$$\begin{aligned}\hat{q}_{\text{sec}}(D = S) &= \frac{\sum_{k=0}^2 \binom{k}{2} (1 - \Pr_{\text{sec}}(D \neq S | k \# \text{ of } h \text{ types})) \binom{2}{k} q^k (1 - q)^{2-k}}{1 - \Pr_{\text{sec}}(D \neq S)} \\ &= \frac{q^2 + 0.5q(1 - q)2}{1 - (1 - q)^2(1 - \sigma)} \\ &= \frac{q}{1 - (1 - q)^2(1 - \sigma)}.\end{aligned}$$

We will show this is indeed an equilibrium and discuss its uniqueness.

Existence.—As h types always vote according to their signal, we have that $\hat{q}_{\text{sec}}(D = S) > \hat{q}_{\text{sec}}(D \neq S) = 0$. Thus, each member strictly prefers the committee decision to match the state. As the principal only observes the group outcome, a player would only have an incentive to deviate if it increases the probability of a correct group decision. Any group with an h member will implement the correct decision, so in these groups no player has an incentive to deviate. It remains to show that there is no incentive for any member to deviate in (l, l) groups. In an (lr, lr) group implementing R would give an expected evaluation of $\hat{q}_{\text{sec}}(D = S)\sigma^2/(\sigma^2 + (1 - \sigma)^2)$, while implementing B would lead to a lower expected evaluation of $\hat{q}_{\text{sec}}(D = S)(1 - \sigma)^2/(\sigma^2 + (1 - \sigma)^2)$. Sending a non-truthful message can only lower the expected evaluation (for example by inducing the other lr member to vote v_B). Hence, there is no incentive for any member here to deviate. Similarly, there is no profitable deviation in an (lb, lb) group. In an (lb, lr) group, the proposed equilibrium has members implementing each decision with probability 0.5. This yields an expected payoff for each member of $0.5q/(1 - (1 - q)^2(1 - \sigma)^2 - (1 - q)^2\sigma(1 - \sigma))$. As the group posterior is balanced and all (lb, lr) groups are mixing equally between each decision, the expected evaluation from implementing B or R is the same. Therefore there is no strict incentive to deviate from this mixing strategy. As there is no benefit in implementing one decision more than another, there is also no incentive to deviate from truth-telling in communication. Thus, the candidate equilibrium is in fact an equilibrium.

Uniqueness.—Restricting attention to full information sharing, there is still more than one equilibrium. For example, when the group's posterior favors B , it could

be that one or both members vote v_B . Furthermore, in an (lb, lr) group, there are many voting strategies which will implement each group decision with probability 0.5. Nonetheless, in any equilibrium with full information sharing the group will implement the decision with the highest posterior probability of matching the state, and will implement each decision with probability 0.5 if they have a balanced posterior. To see why mixing with probability 0.5 is unique, let all (lb, lr) committees implement B with (possibly degenerate) probability π . The expected utility of each committee member in equilibrium is then

$$\frac{0.5\pi q}{1 - (1 - q)^2(1 - \sigma)^2 - 2(1 - \pi)(1 - q)^2\sigma(1 - \sigma)} + \frac{0.5(1 - \pi)q}{1 - (1 - q)^2(1 - \sigma)^2 - 2\pi(1 - q)^2\sigma(1 - \sigma)},$$

where the first term corresponds to the expected evaluation in state B and the second to state R . To be willing to mix, the committees must be indifferent between the two decisions, i.e., both terms have to be equal. This is only the case if $\pi = 0.5$. ■

B. Proof of Proposition 2

First, we show that there can be no information on ability revealed in communication. Suppose there existed an equilibrium in which each player truthfully reveals his signal and ability. As the principal sees these announcements, l types would get an evaluation $\hat{q} = 0$, while h types would get an evaluation $\hat{q} = 1$. An lb type would have an incentive to deviate, announcing τ_h rather than τ_l and thus pooling with hb types. This deviation would earn him an evaluation $\hat{q} = 1$ if the state of the world is B , a clear improvement on zero. An lr type has the same incentive to pool with hr types by announcing τ_h . Therefore, full truth-telling cannot be an equilibrium under transparency. Can any information on ability be revealed in communication? No. Fix arbitrary communication and voting strategies for hr and hb types. As the principal observes all individual behavior under transparency, only actions consistent with the strategies of h types will receive positive evaluations. As such, lb and lr types will have an incentive to perfectly mimic the strategies of hb and hr types. By matching high types' distribution over ability announcements, low types succeed in signal jamming—no information on ability is revealed.

The most informative outcome of the two communication stages is, therefore, where each player reveals his signal truthfully and where no meaningful information on ability is communicated. Let a candidate for equilibrium be where members truthfully announce their signal, hb and lb members follow the same strategy in announcing ability, τ^b , hr and lr members follow the same strategy in announcing ability, τ^r , and each member votes according to his signal. Assuming this is an equilibrium, the probability of mistakes and member evaluations are as follows.

Mistakes.—In state B , a mistake only occurs when we have an (lr, lr) committee, as R is wrongly implemented. This occurs with probability $(1 - q)^2(1 - \sigma)^2$. In

state R , a correct decision will be made by a committee composed of hr or lr members. Thus, the probability of a mistake is $1 - (q + (1 - q)\sigma)^2$.

Evaluations.—Only members who receive a signal in line with the state will be given positive evaluations. All h types will receive the correct signal as will a share σ of l types. The principal will thus give an evaluation $q/(q + (1 - q)\sigma)$ if a member announces the correct signal and also votes for that signal.

We will show this is indeed an equilibrium and discuss its uniqueness.

Existence.—In this candidate for equilibrium, each member's message m^i and vote v^i always agree. Furthermore, hb and lb players only make ability announcements permitted in the strategy τ^b , while hr and lr players only make ability announcements permitted in the strategy τ^r . To check the existence of this equilibrium we need to fix the principal's off-path beliefs. We set off-path beliefs to be that the only positive evaluations are $\hat{q}(m_b, \tau^i, v_B | S = B, \tau^i \in \tau^b)$ and $\hat{q}(m_r, \tau^i, v_R | S = R, \tau^i \in \tau^r)$. Given these beliefs, it is immediate that hb and hr members have no incentive to deviate. An lb member knows that he must announce either m_b or m_r to gain a positive evaluation. As his prior is informative, his best response is to announce m_b . Given the beliefs of the principal, the only course of action open to this lb member is then to announce an ability consistent with τ^b , and then vote v_B . Any other course of action will lead to an evaluation of zero for sure.

Uniqueness.—The strategies τ^b and τ^r are arbitrary. All that an equilibrium requires is that hb and lb follow the same communication strategy, as do hr and lr . For example, it could be that all members announce τ_h or all remain silent or perhaps $\tau^b = \tau_h$ while $\tau^r = \tau_l$. Nonetheless, in each equilibrium there is no information on ability levels revealed in communication. In the set of most informative equilibria, only equilibria in which all players truthfully announce their signal and then vote according to their signal exist. To see this, note that h types will vote according to their signal (by assumption), so the only possible deviation in this set is for l types not to vote according to their signal. As the messages and votes of h types always agree, an l type (who truthfully reports his signal) would reveal himself as an l type if he voted against his signal, getting an evaluation of zero. ■

C. Proof of Proposition 3

The most informative outcome of the two communication stages is, by definition, where each player reveals his signal and ability. Let a candidate for equilibrium be where members truthfully announce their signal and ability, then unanimously vote for the decision with the highest posterior probability of matching the state. In case of a balanced posterior (after two conflicting signals from low ability types), the committee unanimously votes for each decision with probability 0.5. Assuming this is an equilibrium, the probability of mistakes and member evaluations are as in Proposition 1. We will show this is indeed an equilibrium and discuss its uniqueness.

Existence.—In this candidate for equilibrium, all group decisions are unanimous. To check its existence we must set off-path beliefs for nonunanimous decisions. We impose that $\hat{q}_i(v_B^i, v_B^j | S = B) > \hat{q}_i(v_B^i, v_R^j | S = B)$, $\hat{q}_i(v_B^i, v_R^j | S = R)$. Given these beliefs, it is immediate that a h type will never have an incentive to deviate from the proposed strategy. Similarly, the off-path beliefs make it optimal for those in an (lr, lr) or (lb, lb) group to tell the truth, while the fact that $\hat{q}_i(v_B^i, v_B^j | S = B) = \hat{q}_i(v_R^i, v_R^j | S = R) > \hat{q}_i(v_B^i, v_B^j | S = R) = \hat{q}_i(v_R^i, v_R^j | S = B) = 0$ means that there is no incentive to deviate from a strategy that unanimously implements the decision with the highest posterior probability of matching the state. In an (lb, lr) group, the proposed equilibrium has members unanimously implementing each decision with probability 0.5.²⁶ As shown in the proof of Proposition 1, there is no incentive for any member to deviate from this equal mixing strategy. The requirement that (lb, lr) groups vote unanimously is sustained by the off-path beliefs. Thus, the candidate equilibrium is in fact an equilibrium.

Uniqueness.—This is the only equilibrium under mild transparency in which all members fully share their information, and h types vote according to their signal. This uniqueness is guaranteed by the principal’s off-path beliefs, which lead l types to vote unanimously for the decision most likely to match the state. If the principal’s off-path beliefs were instead $\hat{q}_i(v_B^i, v_B^j | S = B) \leq \hat{q}_i(v_B^i, v_R^j | S = B)$, then truthful communication would break down. For more on this see below. ■

D. Proof of Proposition 4

The principal will be better off under transparency than under secrecy when

$$x[\text{Pr}_{tran}(D = B|B)] + (1 - x)[\text{Pr}_{tran}(D = R|R)] > x[\text{Pr}_{sec}(D = B|B)] + (1 - x)[\text{Pr}_{sec}(D = R|R)],$$

which can be rearranged to

$$(1 - x)[\text{Pr}_{tran}(D = B|R) - \text{Pr}_{sec}(D = B|R)] < x[\text{Pr}_{sec}(D = R|B) - \text{Pr}_{tran}(D = R|B)];$$

substituting in with the values from Proposition 1 and 2 and rearranging, we get

$$\frac{(1 - q)^2(1 - \sigma)\sigma + 2q(1 - q)(1 - \sigma)}{2(1 - q)^2(1 - \sigma)\sigma + 2q(1 - q)(1 - \sigma)} < x,$$

$$\frac{(1 - q)\sigma + 2q}{2(1 - q)\sigma + 2q} \equiv x^* < x.$$

²⁶This is achieved using the coordination device (e.g., both voting v_B if the draw is below 0.5 and voting v_R if it is above).

Thus, secrecy is preferred if $x < x^*$ while transparency is preferred if $x > x^*$. ■

APPENDIX B: WHEN IS THE MOST INFORMATIVE EQUILIBRIUM PAYOFF DOMINATED?

In the most informative equilibrium, an h player will have a higher expected utility under transparency than secrecy or mild transparency:

$$\frac{q}{1 - (1 - q)(1 - \sigma)} > \frac{q}{1 - (1 - q)^2(1 - \sigma)}.$$

An l type player must weigh these expected evaluations by the probability of voting for the correct state. Unsurprisingly, l types have a higher expected utility when they can pool with h types—they prefer secrecy to transparency:

$$\frac{\sigma q}{1 - (1 - q)(1 - \sigma)} < \frac{((1 - q)\sigma + q)q}{1 - (1 - q)^2(1 - \sigma)}.$$

As h types are those with the bulk of information to share, we examine if the most informative equilibrium in each setting is the equilibrium with the highest payoff for h types. Here, we only relax the assumption that the most informative equilibrium is played. We maintain our refinement assumptions regarding voting behavior.

Secrecy.—Here players face a common-value problem. The most informative equilibrium allows players to aggregate their private information and then make a decision which maximizes the group (and each player's) expected evaluation. No player can be better off by withholding information.

Transparency.—Here the most informative equilibrium involves all players voting to signal. The combination of voting to signal and individual votes being observed means the principal can distinguish h and l types very well. In fact, when all players vote to signal, h types get the highest evaluation they can get in any equilibrium.

Mild Transparency.—Here the most informative equilibrium coincides with that of secrecy. However, as the principal now observes individual votes, an h type can achieve a higher payoff in another equilibrium in which he separates from l types. That is, there are equilibria with less information sharing (or none) that payoff dominate the most informative equilibrium for h types. One such case is where no information is credibly communicated; for example, h types may mix between announcing m_b and m_r with equal probability. With no information communicated, the best response of l types is to vote to signal. In this polar opposite to the most informative equilibrium we see that h type players gain a higher payoff, $q/(1 - (1 - q)(1 - \sigma))$. There are a series of equilibria between these two poles that are preferred by h types to full truth-telling. In these equilibria, all players reveal their ability, however h types mix between truthfully revealing their signal and remaining silent while l types vote against their private signal when they see a conflicting message from an

h type. Unlike the truth-telling case, in all these “preferred equilibria” the payoff of each committee member is independent of his partner’s action. Indeed, as Table 8 shows, this is what we find in our laboratory setting.

REFERENCES

- Ali, S. Nageeb, and Roland Bénabou. 2016. “Image Versus Information: Changing Societal Norms and Optimal Privacy.” National Bureau of Economic Research (NBER) Working Paper 22203.
- Battaglini, Marco, and Uliana Makarov. 2014. “Cheap talk with multiple audiences: An experimental analysis.” *Games and Economic Behavior* 83: 147–64.
- Battaglini, Marco, Rebecca B. Morton, and Thomas R. Palfrey. 2009. “The Swing Voter’s Curse in the Laboratory.” *Review of Economic Studies* 77 (1): 61–89.
- Bhattacharya, Sourav, John Duffy, and Sun-Tak Kim. 2014. “Compulsory versus voluntary voting: An experimental study.” *Games and Economic Behavior* 84: 111–31.
- Bouton, Laurent, Micael Castanheira, and Aniol Llorente-Saguer. 2016. “Divided majority and information aggregation: Theory and experiment.” *Journal of Public Economics* 134: 114–28.
- Bouton, Laurent, Aniol Llorente-Saguer, and Frédéric Malherbe. 2017. “Unanimous rules in the laboratory.” *Games and Economic Behavior* 102: 179–98.
- Cai, Hongbin, and Joseph Tao-Yi Wang. 2006. “Overcommunication in strategic information transmission games.” *Games and Economic Behavior* 56: 7–36.
- Chen, Ying, Navin Kartik, and Joel Sobel. 2008. “Selecting Cheap-Talk Equilibria.” *Econometrica* 76 (1): 117–36.
- Coughlan, Peter J. 2000. “In Defense of Unanimous Jury Verdicts: Mistrials, Communication, and Strategic Voting.” *American Political Science Review* 94 (2): 375–93.
- Crawford, Vincent P., and Joel Sobel. 1982. “Strategic Information Transmission.” *Econometrica* 50 (6): 1431–51.
- Farrell, Joseph, and Robert Gibbons. 1989. “Cheap Talk with Two Audiences.” *American Economic Review* 79 (5): 1214–23.
- Fehrler, Sebastian, Urs Fischbacher, and Maik T. Schneider. 2017. “Honesty and Self-Selection into Cheap Talk.” <http://sebastianfehrler.bplaced.net/wordpress/wp-content/uploads/Honesty-and-Self-Selection-08-17.pdf>.
- Fehrler, Sebastian, and Niall Hughes. 2018. “How Transparency Kills Information Aggregation: Theory and Experiment: Dataset.” *American Economic Journal: Microeconomics*. <https://doi.org/10.1257/mic.20160046>.
- Fischbacher, Urs, and Franziska Föllmi-Heusi. 2013. “Lies in Disguise: An Experimental Study on Cheating.” *Journal of the European Economic Association* 11 (3): 525–47.
- Fox, Justin, and Richard Van Weelden. 2012. “Costly transparency.” *Journal of Public Economics* 96 (1–2): 142–50.
- Frankel, Alex, and Navin Kartik. Forthcoming. “What Kind of Central Bank Competence?” *Theoretical Economics*.
- Gersbach, Hans, and Volker Hahn. 2008. “Should the individual voting records of central bankers be published?” *Social Choice and Welfare* 30 (4): 655–83.
- Gneezy, Uri, Bettina Rockenbach, and Marta Serra-Garcia. 2013. “Measuring lying aversion.” *Journal of Economic Behavior and Organization* 93: 293–300.
- Goeree, Jacob K., and Leeat Yariv. 2011. “An Experimental Study of Collective Deliberation.” *Econometrica* 79 (3): 893–921.
- Guarnaschelli, Serena, Richard D. McKelvey, and Thomas R. Palfrey. 2000. “An Experimental Study of Jury Decision Rules.” *American Political Science Review* 94 (2): 407–23.
- Hahn, Volker. 2017. “Committee design with endogenous participation.” *Games and Economic Behavior* 102: 388–408.
- Hansen, Stephen, Michael McMahon, and Andrea Prat. Forthcoming. “Transparency and Deliberation within the FOMC: A Computational Linguistics Approach.” *Quarterly Journal of Economics*.
- Holmström, Bengt. 1999. “Managerial Incentive Problems: A Dynamic Perspective.” *Review of Economic Studies* 66 (1): 169–82.
- Irlenbusch, Bernd, and Dirk Sliwka. 2006. “Career concerns in a simple experimental labour market.” *European Economic Review* 50 (1): 147–70.
- Koch, Alexander K., Albrecht Morgenstern, and Philippe Raab. 2009. “Career concerns incentives: An experimental test.” *Journal of Economic Behavior and Organization* 72 (1): 571–88.

- Le Quement, Mark T., and Isabel Marcin.** 2016. "Communication and voting in heterogeneous committees: An experimental study." Max Planck Institute for Research on Collective Goods Bonn Paper 2016/5.
- Levy, Gilat.** 2007. "Decision-Making Procedures for Committees of Careerist Experts." *American Economic Review* 97 (2): 306–10.
- Majumdar, Sumon, and Sharun W. Mukand.** 2004. "Policy Gambles." *American Economic Review* 94 (4): 1207–22.
- Mattozzi, Andrea, and Marcos Y. Nakaguma.** 2017. "Public versus Secret Voting in Committees." <http://apps.eui.eu/Personal/Amattozzi/Papers/Committees.pdf>.
- Meade, Ellen E., and David Stasavage.** 2008. "Publicity of Debate and the Incentive to Dissent: Evidence from the US Federal Reserve." *Economic Journal* 118 (528): 695–717.
- Meloso, Debrah, Salvatore Nunnari, and Marco Ottaviani.** 2017. "Looking into Crystal Balls: A Laboratory Experiment on Reputational Cheap Talk." http://www.salvatorenunnari.eu/mno_crystalballs.pdf.
- Nyarko, Yaw, and Andrew Schotter.** 2003. "An Experimental Study of Belief Learning Using Elicited Beliefs." *Econometrica* 70 (3): 971–1005.
- Offerman, Theo, Joep Sonnemans, Gijs Van De Kuilen, and Peter P. Wakker.** 2009. "A Truth Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes." *Review of Economic Studies* 76 (4): 1461–89.
- Ottaviani, Marco, and Peter Sørensen.** 2001. "Information aggregation in debate: Who should speak first?" *Journal of Public Economics* 81 (3): 393–421.
- Prat, Andrea.** 2005. "The Wrong Kind of Transparency." *American Economic Review* 95 (3): 862–77.
- Prendergast, Canice, and Lars Stole.** 1996. "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning." *Journal of Political Economy* 104 (6): 1105–34.
- Schotter, Andrew, and Isabel Trevino.** 2014. "Belief Elicitation in the Laboratory." *Annual Review of Economics* 6: 103–28.
- Sobel, Joel.** 2013. "Ten possible experiments on communication and deception." *Journal of Economic Behavior and Organization* 93: 408–13.
- Stasavage, David.** 2007. "Polarization and Publicity: Rethinking the Benefits of Deliberative Democracy." *Journal of Politics* 69 (1): 59–72.
- Swank, Otto H., and Bauke Visser.** 2013. "Is Transparency To No Avail?" *Scandinavian Journal of Economics* 115 (4): 967–94.
- Visser, Bauke, and Otto H. Swank.** 2007. "On Committees of Experts." *Quarterly Journal of Economics* 122 (1): 337–72.
- Wang, Joseph Tao-yi, Michael Spezio, and Colin F. Camerer.** 2010. "Pinocchio's Pupil: Using Eyetracking and Pupil Dilation to Understand Truth Telling and Deception in Sender-Receiver Games." *American Economic Review* 100 (3): 984–1007.
- Williams, John C.** 2012. "Opening the Temple." In *Federal Reserve Bank of San Francisco 2011 Annual Report*, edited by John C. Williams, 7–14. San Francisco: Federal Reserve Bank of San Francisco.