

IZA DP No. 9027

How Transparency Kills Information Aggregation: Theory and Experiment

Sebastian Fehrler
Niall Hughes

April 2015

How Transparency Kills Information Aggregation: Theory and Experiment

Sebastian Fehrer

*University of Konstanz,
University of Zurich and IZA*

Niall Hughes

University of Warwick

Discussion Paper No. 9027
April 2015

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

How Transparency Kills Information Aggregation: Theory and Experiment^{*}

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 Classification: C92, D71, D83

Keywords: committee decision-making, deliberation, transparency, career concerns, information aggregation, experiments, voting, strategic communication

Corresponding author:

Sebastian Fehrler
University of Konstanz
Box 131
78457 Konstanz
Germany
E-mail: sebastian.fehrler@uni-konstanz.de

^{*} 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 Schneider, Francesco Squintani, Jean-Robert Tyran 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). All remaining errors are of course our own.

“Secrecy is the first essential in affairs of state.” (Cardinal Richelieu)

1. Introduction

Transparency in decision-making is a recurring and controversial topic in public debate. A recent exemplary case is that of FIFA, soccer’s governing body, where its executive committee’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 example is the discussion on how transparent decision-making in monetary policy committees such as the Federal Open Market Committee (FOMC) should be (e.g., Williams (2012)). Here, the argument has been made that too much transparency could be harmful (Meade and Stasavage (2008); Swank et al. (2008); Swank and Visser (2013)). This debate over transparency is not a recent phenomenon. In the early 19th century Jeremy Bentham ([1816] 1999) argued for more transparency in parliamentary decision-making, while reservations had long been voiced by influential thinkers like Thomas Hobbes and John Stuart Mill (see Stasavage (2007)).

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 principals.² The downside of transparency is that if agents care about their reputations, they may pander to the principal, choosing an action which makes them appear smart but which is not necessarily in the principal’s interest (Prat (2005); Fox and Van Weelden (2012)). Indeed, recent theoretical literature has shown that reputational concerns influence committee behaviour differently under secrecy and transparency in the case of political committees (e.g. Stasavage (2007)), corporate boards (Malenko (2013)), and monetary policy committees (Levy (2007); Visser and Swank (2007); Swank and Visser (2013); Swank et al. (2008); Meade and Stasavage (2008); Gersbach and Hahn (2008)).

In the studies that allow for deliberation - an obviously realistic assumption for most

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

²A positive case for transparency is also made in the deliberative democracy literature (e.g., Habermas (1996) and Cohen (1996)) whose proponents view public deliberation as an essential element of legitimate democratic decision-making. See Landa and Meiowitz (2009) for a critical discussion of this literature.

committees - strategic communication is the key driver of these effects. However, recent experimental literature shows for a variety of settings that subjects are too truthful and too trusting as compared to theoretical ‘cheap talk’ predictions (Cai and Wang (2006); Wang et al. (2010); Battaglini and Makarov (2014)) and that a substantial fraction of subjects display psychological costs of lying (Gneezy (2005); Gneezy et al. (2013); Fischbacher and Föllmi-Heusi (2013)). In a study on committee voting, Goeree and Yariv (2011) even find that open deliberation leads to perfect information sharing in a set-up where committee members have private values regarding the committee decision and thus a clear incentive to keep their information private. These findings cast some doubt on the predictions from the theoretical literature on the role of career concerns in committees.

Unfortunately, empirical evidence on the effects of transparency on career concerned committee members is scarce and does not identify the mechanisms through which behavior is affected. Both Meade and Stasavage (2008) and Swank et al. (2008) present evidence of changes in deliberation in the FOMC after the decision to publish minutes of their meetings, but they present the same data as supporting evidence for their different career concerns models.³ The difficulty of identifying mechanisms is not surprising given the shortcomings of using field data to study the role of transparency in committee decision-making. Firstly, it is almost certainly impossible to observe the level of ability, prior information and biases which committee members may have. To the extent that transparency may be interacting with these unobserved variables, it is impossible to identify the mechanism through which transparency affects behavior. Secondly, a controlled comparison of institutions is always very difficult as their different elements, e.g. the committee’s voting rule, deliberation protocol and other factors are endogenous. For these reasons, we believe a laboratory experiment provides the cleanest way of approaching the important and so far unanswered questions of whether and how transparency affects the behavior of career concerned committee members. By carefully controlling the information players have and varying the level of transparency, effects and mechanisms can be clearly identified.

³Another study is presented by Cross (2013), who does not test a particular model but compares behavior of members of the Council of the European Union under different levels of transparency and shows evidence suggesting that more extreme positions are taken under higher levels of transparency.

Instead of testing one of the existing models we construct and test a new model in which career concerns play out very differently under three levels of transparency. This set-up 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 non-binding 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. 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. This corresponds to the practice of the FOMC before 1993 (see Meade (2005)).⁴ We implement the three transparency regimes in a laboratory experiment in which the principal's belief about an agent's type is elicited with a proper scoring rule. This stated belief directly affects the agent's payoff. As free form communication has been shown to matter greatly for the level of strategic deliberation in Goeree and Yariv (2011), we implement the second deliberation stage as an open chat.

The following key insights from the model are corroborated in the experimental data. Firstly, in the most informative equilibrium committee members truthfully share all their information with each other under secrecy, but fail to do so under transparency. When

⁴Mild Transparency also reflects the main structure of a scenario in which deliberation is nominally transparent but committee members arrange secretive pre-meetings for deliberation and then stage a public show meeting without real deliberation and public votes, which arguably describes what happened with decision-making in the FOMC after 1993 (Swank et al. (2008); Swank and Visser (2013)).

the principal is watching, nobody wants to admit to being a low type, so information aggregation is incomplete. Secondly, incorrect group decisions are more prevalent under transparency than secrecy. This occurs for two reasons: (a) members are less truthful under transparency and thus aggregate poorer information, and (b) low ability members have an incentive to vote according to their private information even if they believe it to be wrong. This effect stems from the fact that if members change their position between the deliberation and voting stages, the principal can infer that they are a low ability type. Thirdly, a principal biased in favour of the status quo can actually be better off under transparency than secrecy. The failure of committee members to share information under transparency means they find it difficult to vote unanimously against the status quo. This means more mistakes when they should change the status quo, but less mistakes when they shouldn't. 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. Finally, while the most informative equilibrium under mild transparency predicts behavior as under secrecy, the experiment shows aggregate level results quite similar to the transparency case. However, these result from quite different individual level behavior. Mild transparency leads to a significant level of deception at the first deliberation stage, whereas under transparency information aggregation breaks down at the second deliberation stage and at the voting stage. Under mild transparency, there are equilibria in which high types do not truthfully reveal their signals. This makes them easier to separate from low types in the voting stage, resulting in a more favorable belief of the principal regarding their type.⁵ Our results suggest that the existence of these equilibria leads a number of high types to deliberate non-truthfully and thus to very different aggregate outcomes than predicted by the most informative equilibrium.

We also observe a number of deviations from our theoretical predictions regarding individual level behavior. Principals are on average too optimistic in their assessment of committee members. However, they update their beliefs in the correct direction to all pieces of information they receive. Most committee members play a best response to

⁵A structurally similar situation is described in Sobel (2013), section 2, in which he calls for experiments to study the predictive power of the most informative equilibrium in cheap talk games with additional equilibria which are preferred by at least one type of player.

the stated beliefs of the principals under secrecy and mild transparency, while many of them fail to best respond under transparency, where belief updating is more complicated and therefore also more complicated to anticipate for committee members. Also under transparency, open chat deliberation is more truthful and informative than predicted. This suggests that some subjects do indeed face psychological costs of lying, and confirms results from previous experimental studies. However, despite this truth-telling preference, the level of truth-telling varies greatly between treatments. This suggests that Goeree and Yariv's (2011) result on the power of open chat to moderate institutional differences does not extend to the case of transparency versus secrecy.

In the next section, we present and solve the model. We proceed with the experimental design and theoretical predictions for the chosen parameter values before discussing the aggregate and individual level results. We conclude with a discussion of the main results and their implications for the literature on career concerns in committees and cheap talk games, and for debates about the optimal level of transparency.

2. The Model

A committee of two members must make a decision $D \in \{B(lue), R(ed)\}$ on behalf of a principal.⁶ There are two equally likely states of the world $S \in \{B, R\}$. The utility of the principal in each state is

$$\begin{aligned} U_{principal}(D = B|S = B) &= x \\ U_{principal}(D = R|S = B) = U_{principal}(D = B|S = R) &= 0 \\ U_{principal}(D = R|S = R) &= 1 - x \end{aligned}$$

with $x \in (0.5, 1)$. That is, the principal gets higher utility from a correct group decision when the state is B .⁷ The group decision is made by unanimity rule, whereby option R is

⁶We present a two member model for clarity and simplicity, as we run two member games in the lab. In an n member committee players' behaviour would not change, though the probability of incorrect group decisions would.

⁷Equivalently we could say that the utility from a correct decisions is the same across states but a mistake is far costlier in state B .

implemented only if both members vote for it, otherwise the status quo B is upheld.

Committee members are not perfectly informed of the state of the world. Instead, each member gets an informative signal about the true state, where the level of informativeness depends on his ability. Each member i is either of high or low ability. Accordingly, there are four types of committee member $\{hb, hr, lb, lr\}$, where, for example, hb refers to someone of high ability who receives a signal b . 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 $Pr(t^i = h) = q \in (0, 1)$. The accuracy of the signals $s_t^i \in \{b, r\}$ are given by $Pr(s_h^i = b | S = B) = Pr(s_h^i = r | S = R) = 1$ and $Pr(s_l^i = b | S = B) = Pr(s_l^i = r | S = R) = \sigma \in (0.5, 1)$ respectively.⁸

Abilities and signals are private information. However, once nature has chosen member's ability levels and signals have been received, members can communicate with each other. In modelling communication we want to capture the key features of real world committees while still giving clear predictions to test in the laboratory. First, in a simple straw poll, each member simultaneously announces a message $m^i \in \{m_b, m_r, m_\emptyset\}$, i.e. he can raise his hand in favor of B or R or can abstain. Next, each member simultaneously announces a message $\tau^i \in \{\tau_h, \tau_l, \tau_\emptyset\}$, i.e. he can announce that he is of low type, high type or can remain silent.⁹ After these two stages of communication, committees have 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, after both rounds of communication, each member casts a vote $v^i \in \{v_B, v_R\}$.¹⁰ We model communication in two stages because we want players to be able to communicate in open chat in the experiment, as members in a real committee are. However, we want to avoid any effects of sequencing which may occur. For example, one player copying the announcement of another, or both players rushing to make the first announcement. The

⁸We could allow for imperfect high type signals; the strategy of committee members in the most informative equilibria would remain largely unchanged. The only change is that there could now be a committee of two high types with conflicting signals. Under secrecy or mild transparency such a group would implement each state with probability 0.5.

⁹We could collapse all communication into one simultaneous round or could have the second stage of communication be sequential rather than simultaneous - the theoretical predictions would not change.

¹⁰We could also allow for abstention in the final voting stage but it would not change equilibrium predictions.

problem is solved by splitting the lab communication into two stages: first, players can simultaneously announce their signals m^i , then they can communicate in open chat about their ability levels and can decide on a group strategy if necessary. This second stage of communication differs slightly between the model and the experiment because a message space and a sequence has to be fixed for the model. However, the only sensible way subjects can make use of the open chat is to share info on their ability and to coordinate. To also allow for coordination in the model we add the coordination device.

Committee members' primary concern is not in making the correct decision for the principal but rather in advancing their own individual careers or reputations. The payoff of a committee member is simply the principal's posterior belief that he is of high ability, given by $\hat{q} \in [0, 1]$. This is standard in models of career concerns (see Prat (2005); Fox and Van Weelden (2012); Levy (2007)).¹¹ Before the game starts, the principal's prior for each member is $Pr(t^i = h) = q$.

A committee member's strategy consists of a communication strategy and a voting strategy. A communication strategy is a pair (m^i, τ^i) , where m^i is mapping from a pair (s^i, t^i) into a probability distribution over messages $\{m_b, m_t, m_\emptyset\}$ and τ^i is a mapping from (s^i, t^i) and messages exchanged in the straw poll into a probability distribution over announcements $\{\tau_h, \tau_b, \tau_\emptyset\}$. A voting strategy is a mapping from signal, ability and messages exchanged in both communication rounds into a probability distribution over votes $v^i \in \{v_B, v_R\}$.

Once votes are cast, the true state is revealed and utilities are realized. The principal updates her prior beliefs as best she can, where this ability depends on how much of the decision-making process she observes, i.e. on the level of transparency. We compare three different regimes. Our primary interest is in comparing *secrecy*, where the principal only observes the group decision D , with *transparency*, where she witnesses each member's messages m^i, τ^i , and final vote v^i . Subsequently, we also examine *mild transparency*, where the principal observes only the group decision and how each individual votes.

¹¹In line with the career concerns literature mentioned, we ignore the possible benefits to the principal of learning a member's ability level. One setting in which there would be no benefits to the principal is when there are competitive labor markets and long term contracting is not possible (see discussion in Fox and Van Weelden (2012)).

We study symmetric perfect Bayesian equilibria under the three transparency regimes. As is standard in voting games, we restrict attention to strategies which are not weakly dominated.¹² As committee members' payoffs depend on the principal's beliefs and because talk is cheap, there will be many equilibria in each of the three settings. We employ 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 do exist, we focus on these equilibria as: (a) it seems reasonable that with perfect knowledge h types will vote for the true state (especially as the principal knows they have perfect knowledge); (b) we can test whether focusing on these equilibria is sensible. Indeed, in the laboratory h types vote to signal 98.2% of the time, rising to 100% in the final 5 periods.¹³ Secondly, we follow the cheap talk literature (Crawford and Sobel (1982); Ottaviani and Sørensen (2001); Chen et al. (2008)) in focusing on the most informative equilibrium. That is, in each setting we consider the equilibrium where the greatest possible amount of information is shared across the two communication stages.¹⁴ As is standard, 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".

2.1. Equilibrium

We can now compare the behavior of two person committees under our different transparency regimes. A useful benchmark to consider is what the principal would do if she could observe the two signals and their strengths directly. If either of the signals were from a h type, i.e. fully revealing, she would obviously choose that decision. With two low signals she would choose B whenever there is a tie or both signals favour B . Finally, when she observes two low signals in favour of R she would choose R only if $(\frac{\sigma}{1-\sigma})^2 > \frac{x}{1-x}$, that is, if the evidence in favor of R is strong enough to overturn her bias

¹²In our setting this prevents cases where each player votes v_B simply because he expects the other player to do so, and so no vote is pivotal.

¹³Alternatively, we could restrict attention to responsive and monotone strategies as in Fox and Van Weelden (2012).

¹⁴We discuss the existence of further equilibria, which might preferred by h types, in the appendix.

for B .

Secrecy As the principal can only see the group decision, she must hold the same posterior \hat{q} for each individual. For this reason committee members have a common interest in making the correct group decision.

Proposition 1. *In the most informative equilibrium under secrecy, all members truthfully reveal their signal and ability, and then jointly implement the policy with the highest posterior probability of matching the state. The probability of a group mistake in each state is $(1 - q)^2(1 - \sigma)$. Each member earns an evaluation of zero if the group decision is wrong and an evaluation $\frac{q}{1 - (1 - q)^2(1 - \sigma)}$ if the group decision is correct.*

Proof. See Appendix □

It is not surprising that sharing all their information is an equilibrium for players with a common interest; Coughlan (2000) shows that allowing communication between players with a common interest can lead to full aggregation. There are of course other, less informative, equilibria in which players babble in one or both stages of communication, however Guarnaschelli et al. (2000) and Goeree and Yariv (2011) show that in the laboratory players with a common interest are overwhelmingly truthful. Regardless of the level of communication, two things are worth noting about the most informative equilibrium under secrecy: firstly, the voting rule plays no role. The voting outcome would be the same regardless of whether the rule was stacked in favor of B or R or neutral. Secondly, we saw in the benchmark case that the principal trades off the evidence against her bias. Here, however, committee members face no such trade-off - they simply want to choose the most likely state, regardless of the principal's bias.

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

Proposition 2. *In the most informative equilibrium under transparency, all members truthfully reveal their signal and vote according to their private signal, however no information on abilities can be credibly communicated. 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 an evaluation of zero if their individual vote doesn't match the state, and an evaluation $\frac{q}{q+(1-q)\sigma}$ if it does.

Proof. See Appendix □

Here, in the most informative equilibrium, members reveal their signals truthfully in the straw poll but cannot credibly reveal their ability. This is because l types will mimic the strategy of h types in the open chat in an effort to appear competent in front of the principal. Information aggregation is incomplete when compared to secrecy: while signals are shared, players cannot differentiate the quality of those signals. Interestingly, even though some information is aggregated, players ignore this when deciding how to vote. An lb type who sees a m_r message in the straw poll will believe R to be 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 type who sees a m_b message in the straw poll. Why is this? A 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.¹⁵ This means that l types will stick to their straw poll announcement in the final vote even when they believe it is the wrong state.

Mild Transparency In this section, we look at an intermediate case of transparency where 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, as was the case for the FOMC before 1993. Furthermore, mild transparency reflects the settling where an attempt to introduce transparency leads to the emergence of pre-meetings in which members can discuss what to do before they are under the watchful eye of the principal. Swank and Visser (2013) have shown that pre-meetings can undo the benefits of transparency in their setting, while similarly Alan Greenspan noted that in-

¹⁵This equilibrium is sustained by beliefs that any member who switches in the final vote is of type l . As nobody actually switches in equilibrium, this node is never reached and so we are free to choose any beliefs of the principal off the equilibrium path.

roducing transparency into Federal Open Market Committees would mean “a tendency would arise for one-on-one pre-meeting discussions, with public meetings merely announcing already agreed-upon positions or each participant to enter the meeting with a final position not subject to the views of others” (quoted in Meade and Stasavage (2008)). The following proposition shows that the most informative equilibrium under mild transparency gives the same outcomes as under secrecy.

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

Proof. See Appendix □

The only difference between the most informative equilibrium under secrecy and mild transparency is that under mild transparency each member must vote for the project most likely to match the state, while under secrecy it is only the group implementation that matters. This makes no difference to the probability of mistakes. This implies that, if players do indeed play the most informative equilibrium, the emergence of pre-meetings would actually lead to more information aggregation and less mistakes than would otherwise be the case.¹⁶

Optimal Transparency A natural question is which transparency regime would the principal prefer *ex ante* assuming the most informative equilibria are played. We focus here on secrecy and transparency, as the intermediate case of mild transparency predicts the same outcome as secrecy. A principal who is “not too biased” will prefer secrecy to transparency. The superior information aggregation ensures a lower probability of mistakes. Meanwhile, a principal who is sufficiently biased in favor of B will prefer transparency to secrecy exactly *because* information aggregation breaks down under transparency.

¹⁶There are other equilibria with less information sharing under mild transparency which are preferred by h types. See discussion in the appendix.

Proposition 4. *There always exists an x^* such that if $x < x^*$ the principal is ex ante better off under secrecy than under transparency and if $x > x^*$ then transparency is preferred.*

Proof. See Appendix □

A transparent committee, voting according to their own signals, will make the unanimous decision needed to implement R less often. Such a committee will, thus, make more mistakes in state R but less in B when compared to secrecy. A principal with a high value of x will be more wary of losses in state B and so will prefer transparency to secrecy. Indeed, this is that case for the parameter values we bring to the lab. With $\sigma = 0.55$ and $x = 0.75$ the principal would prefer transparency to secrecy.

3. Experiment

To test the main predictions of the model, we ran a laboratory experiment with three treatments - one for each level of transparency. In a slight departure from the model, in which we had to specify the message space for the second deliberation stage, we allow for free form communication in the second deliberation stage in the experiment. As subjects can now coordinate through communication, we did not implement the correlation device. Note, that the equilibria that we have characterized in the previous section are unaffected by this change. Given the effects of free-form communication in Goeree and Yariv's (2011) study and the fact that open deliberation takes place in most real world committees, we decided that it is better to test the theoretical predictions in this setting rather than in a setting with completely structured communication. To control for possible effects of the sequence of messages we chose to implement the first deliberation stage, in which the sequence would matter under transparency, exactly as it is in the model, i.e. as a simultaneous straw poll. For the second deliberation stage the sequence of messages does not play a role for our theoretical predictions.

3.1. Experimental Design

We ran six sessions, two for each transparency regime (see Table 1). Each session consisted of 20 rounds with random matching of subjects into groups. In the first round of the experiment, subjects were randomly assigned to matching groups of nine people.¹⁷ In every period, new groups of three were randomly formed within the matching groups to avoid the emergence of reciprocal behavior and at the same time provide independent matching groups. In each group and round, one member was randomly assigned the role of the principal (called the “observer” in the instructions) and the other two were assigned the role of committee members (called “voters” in the instructions).¹⁸ With probability $q = 0.25$ a committee member was of type h (“well-informed voter”), with probability $1 - q = 0.75$ of type l (“informed voter”). The task of the committee members was to vote on the true color of a randomly selected jar.¹⁹ The blue jar ($S = B$) contained 11 blue and 9 red balls, the red jar ($S = R$) contained 11 red and 9 blue balls. Jars were chosen with equal probability. Type h committee members received a ball with the true color of the jar (the perfectly accurate signal), type l committee members received a ball that was drawn from the selected jar (a signal of accuracy $\sigma = 0.55$). We chose a low signal accuracy for the low type to make equilibrium predictions sufficiently different across treatments which would allow us to distinguish them statistically in the data in case behavior resembles equilibrium predictions.

On the first screen, the principal learned that she is a principal, while committee members were informed about their type and the color of their ball. Committee members then had to simultaneously send a message {red, blue, not specified} to the other committee member in their group. On the next screen, committee members saw the message from their partner and had the opportunity to chat with him for 90 seconds.²⁰ In the transparency treatment, the principal could see the committee members’ messages and follow

¹⁷In one session we had only 15 subjects and therefore only one matching group.

¹⁸Screenshots and instructions are in the online appendix.

¹⁹This has become the standard task in information aggregation experiments (e.g. Guarnaschelli et al. (2000); Battaglini et al. (2009); Goeree and Yariv (2011); Bhattacharya et al. (2014); Bouton et al. (2014)).

²⁰The timeout was not strictly enforced. When the time was up a message appeared on the screen asking them to finish their sentence and proceed. Most subjects did so immediately and the few others were kindly asked to proceed after 120 seconds by an experimenter.

Table 1: Experimental sessions

	N sessions	N matching groups	N subjects
Secrecy	2	5	45
Transparency	2	4	42
Mild Transparency	2	5	45

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

the chat on her second screen. In the other two treatments the principal could not. On the third screen, the committee members could review the communication and then make their final decision by voting for red or blue. Votes were then aggregated to the group decision. For a decision for red ($D = R$) two votes were required, for a decision for blue ($D = B$) only one. After the voting stage, the principal received information on her next screen that varied between treatments. In the secrecy treatment, she could only see the group decision and the true color of the jar. In the mild transparency treatment, she could, in addition, see the individual votes of the committee members. In the transparency treatment, she could also review the whole communication (messages and chat) between the committee members. On this screen, the principal had to indicate her belief about the probability that the committee members are of type h , by entering this probability in per cent. In the secrecy treatment, in which both committee members are indistinguishable to the principal, she had to evaluate one randomly chosen committee member from her group, in the other two treatments she had to evaluate both committee members in her group. On the final screen of each round, subjects received feedback information regarding the types of the committee members in their group, the group decision, the true color of the jar and their pay-offs.

Subjects earned points in each round. The points of a committee member in one round was twice the probability that he was a high type, as entered in per cent by the principal, e.g., if the principal entered 30% the committee member's payoff was 60 points. The principal's payoff was 3 points for a correct group decision if the true state of the jar was

blue and 1 point for a correct group decision if the true state was red, i.e. the principal's bias is given by $x = 0.75$. In addition, the principal earned a number of points between 0 and 100 for accurate evaluation of the committee members' types. In the transparency and mild transparency treatments, the evaluation of one of the two committee members (committee member i) was randomly selected and a principal j 's earnings were determined by the following quadratic scoring rule (under secrecy the principal evaluated only one member i).

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

where $Pr_j(t_i = h)$ denotes the probability that committee member i is of type h , as entered by principal j , in per cent. This rule makes it optimal for expected pay-off maximizing subjects to truthfully enter their beliefs (see, e.g., Nyarko and Schotter 2003) and subjects were directly told so in the instructions.²¹ To keep potential effects of social preferences of committee members toward the principal limited, we kept the principal's payoff from correct group decisions low relative to the payoff from accurate evaluations.

Four rounds were randomly chosen at the end of the session and the points earned in these rounds 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 spend about 2 hours in the lab and earned, on average, CHF 47 in addition to their show-up fee of CHF 10. Earnings per hour are comparable to an hourly wage for student jobs in Zurich.

3.2. Equilibrium Predictions

Ex ante equilibrium error rates are reported in Table 2. In this section, whenever we say

²¹More complicated belief elicitation procedures have been proposed in the literature 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. In telling subjects directly that truthfully reporting their belief is the action which maximizes their expected payoff we follow Schotter and Trevino (2014) who suggest this approach to help subjects understand the scoring rule.

equilibrium predictions we refer to the most informative equilibrium.

Table 2: *Ex ante* equilibrium error rates (in %)

	$S = B$	$S = R$	overall
Secrecy	25.3	25.3	25.3
Transparency	11.4	56.1	33.8
Mild Transp.	25.3	25.3	25.3

Note: In the experiment $S = B$ ($S = R$) corresponds to the case of the blue (red) jar.

Table 3 reports the principal's predicted evaluation of player types after observing the true state of the world, the group decision and, in case of transparency, the individual decisions. In the same table we also report the *ex ante* expected evaluations for h and l type committee members, i.e. the expected evaluation before the realization of the state of the world and the signals. Type l committee members are predicted to do best under secrecy and mild transparency and type h committee members under transparency. The reason is that principals are predicted to be better able to tell them apart under transparency.

Table 3: Principal's equilibrium beliefs about members' types $Pr(t_i = h)$ in %

	Eq. evaluation after correct and wrong decisions		Eq. expected <i>ex ante</i> evaluation for committee members	
	dec. corr.	dec. wr.	type h	type l
Secrecy (group decision)	33.5	0	33.5	22.2
Transparency (individual decision)	37.7	0	37.7	20.8
Mild Transp. (ind./group decision)	33.5	0	33.5	22.2

Note: In equilibrium under secrecy the principal only takes the group decision into account for her evaluation, under transparency she looks at messages and individual decisions (which are predicted to be the same). Under mild transparency she looks at individual decisions which are the same as the group decision in equilibrium.

Equilibrium predicts that committee members should succeed in aggregating informa-

tion under secrecy and mild transparency but fail to do so under transparency. Under secrecy and mild transparency all committee members are expected to truthfully announce their signal in the straw poll and their type in the chat. Under transparency committee members are also predicted to share information about their signal truthfully in the straw poll but the second deliberation stage is predicted to be uninformative. Moreover, all committee members are predicted to vote according to their signal and straw poll message, which are identical, in the voting stage. Furthermore, our restriction on equilibria in which high types always vote according to their signal is a testable assumption that will also be addressed.

3.3. Experimental Results

Aggregate Behavior

Table 4 summarizes the observed error rates and, for comparison, the *ex post* equilibrium predictions, i.e. the predictions after realization of the true state of the world, the types, and the signals. It also contains a column with the hypothetical full information aggregation benchmark case, i.e. the case where committee members share all of their information and implement the decision that matches the more likely state.

Table 4: Observed, *ex post* equilibrium, and full information aggregation error rates by state and in total

true color	$S = B$			$S = R$			total		
	Observed	m.i.e.	f.i.a.	Observed	m.i.e.	f.i.a.	Observed	m.i.e.	f.i.a.
Secrecy	28.3 (5.5)	27.2	27.2	25.8 (1.8)	24.8	24.8	27 (3.1)	26	26
Transparency	15.8 (1.0)	8.2	21.9	50.7 (6.4)	61.9	26.9	33.3 (3.5)	33.9	24.3
Mild Transp.	15.5 (1.2)	24	24	46.1 (3.6)	25.3	25.3	30.8 (1.8)	24.7	24.7

Note: m.i.e. = most informative equilibrium, f.i.a. = full information aggregation. Equilibrium error rates are *ex post* error rates. Standard errors of the observed error rates (in parentheses) are adjusted for clustering in matching groups.

Equilibrium predictions are very accurate for the secrecy treatment, but less so for the

transparency treatment. Under transparency the error rate is higher in state B and lower in state R than the model predicts. Nonetheless the key predicted difference between transparency and secrecy is borne out in the lab: when 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, $p = 0.024$). In contrast to the theoretical predictions, this is also true for the error rates under mild transparency (Wald-tests, $p < 0.001$ for $S = R$, and $p = 0.023$ for $S = B$) which are statistically not different from those under transparency (Wald-tests, $p = 0.526$ for $S = R$, and $p = 0.894$ for $S = B$). Even though the total error rate is higher under transparency than under secrecy, principals earned minimally more points from the group decisions (0.11 points on average) in the two transparency treatments because, given their bias, correct decisions in state B are more valuable to them. However, this difference is not statistically significant (t -test, $p = 0.235$).

The most likely source of the differences between the treatments are (a) differences in information aggregation and (b) coordination behavior of groups consisting only of low types. Information aggregation helps if there is one type h committee member and one type l committee member in a group, they have conflicting signals and the true state is red. In this scenario all groups in the secrecy treatment aggregated information successfully (see Table 5), while the error rate is substantially and significantly higher in the transparency treatment (Wald-test, $p < 0.001$). However, the error rate is much lower than its predicted value of 100% (Wald-test, $p < 0.001$). There is a significant difference between the error rate for these groups under the two forms of transparency. Under mild transparency the error rate is significantly lower than under transparency (Wald-test, $p = 0.026$) but still significantly higher than under secrecy (Wald-test, $p < 0.001$).

If the true state of the world is blue, groups with conflicting signals and one high type always choose the right decision in all regimes with the exception of one observation under mild transparency. This is not surprising because it must be the high type that received the blue signal and only one vote for blue is required.

The case of committees with two l types is also illuminating. Under transparency the theory predicts that, when such a group has conflicting signals, B will be implemented.

However, we see in Table 5 that this is not the case; mistakes occur in state B 27.5% of the time and only 83.3% of the time in state R , so some coordination on implementing R does take place. The error rates are similar and not statistically different under mild transparency. Under secrecy, we expect such a group to implement each state with equal probability, giving us an error rate of 50%, and indeed the data is relatively close to this.²²

Table 5: Error rates in groups with conflicting signals

true color	$\{h, l\}$ Group		$\{l, l\}$ Group	
	$S = R$	$S = B$	$S = R$	$S = B$
Secrecy	0	0	54.3 (11.0)	38.7 (15.3)
Transparency	44.8 (8.3)	0	83.3 (10.3)	27.5 (8.9)
Mild Transp.	22.6 (5.5)	3.3 (3.4)	75.7 (4.4)	27.9 (6.8)

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

Interestingly, not all groups with two signals in the same direction vote for that alternative because communication is not always truthful and leads some low types to switch in the wrong direction, as we will see in the next section. As a consequence, the error rates of these groups are not 100% if their signals go in the wrong direction (but range from 82.4% in the mild transparency treatment to 95.6% under secrecy, with 93.3% under transparency) but also not 0% if they go in the right direction (and range from 4.7% under secrecy to 13.6% under mild transparency, with 7.4% under transparency, instead).²³

²²Note that none of the differences in error rates between secrecy and the transparency treatments in these groups is statistically significant at the 5% 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% level (Wald-tests). The differences in error rates between the two transparency regimes are not statistically significant even at the 10% level.

²³The error rates after two correct signals are all significant at the 5% level (Wald-tests). However, only the error rate under mild transparency is significantly different from 100% at the 5% level after two false signals (Wald-tests).

Individual Behavior

Deliberation We start our analysis of individual behavior with the deliberation stage. In the straw poll, we observe that under secrecy there is almost completely truthful revelation of signals by both high and low types, consistent with equilibrium predictions (Table 6). Under transparency, high types are also almost always truthful and reveal their signal. However, there are 8.3% low types that lie about their signal and another 10.5% who stay silent, which goes against our predictions.²⁴ Note that lying is not very costly in expectation as σ -signals are not very informative. Staying silent might be motivated by the hope to learn more about the true state and then vote accordingly without being punished for the behavior in the deliberation stage. When we turn to the principals' reactions, we will see whether this strategy pays off or not.

Table 6: (Non-)truthful straw poll messages from different types

	type	lying	silent
Secrecy	<i>h</i> type	0	0
	<i>l</i> type	0.7 (.3)	1.8 (1.2)
Transparency	<i>h</i> type	1.4 (0.9)	0.7 (0.7)
	<i>l</i> type	8.3 (2.1)	10.5 (2.8)
Mild Transp.	<i>h</i> type	19.2 (6.6)	4.8 (2.0)
	<i>l</i> type	3.5 (1.1)	3.3 (1.4)

Note: 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.

Under mild transparency, 19.2% of the high types lie about their signal and another 4.8% stay silent which is against our prediction of truthful communication, while the low types are almost always truthful and reveal their signal.²⁵ The degree of lying from high

²⁴Both percentages are significantly larger than zero (Wald-tests, $p < 0.001$).

²⁵All percentages of incorrect messages or abstentions are significantly larger than zero under mild transparency (Wald-tests, $p < 0.05$).

types is significantly higher than in the other treatments (Wald-tests, $p < 0.001$). Out of 26 subjects who were high type committee members under mild transparency in at least 3 rounds, 15 never lied, 1 always lied, and the remaining 8 lied in some periods and told the truth in others. The share of high type lies is lower (16.9% as compared to 21.7%) in the last ten periods than in the first ten but the difference is not statistically significant (Wald-test, $p = 0.375$). Throughout the 20 rounds, the degree of lying is not high enough to make the straw poll uninformative for low types.

Table 7: Chat messages about type

	report type		claim	lying
Secrecy	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)
Transparency	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)
Mild Transp.	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)

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

Next, we turn to communication in the chat. After exchanging information about signals in the straw poll the only relevant information left to talk about is the type.²⁶ Consequently, we coded whether a type was announced in the chat, and if so, what type was announced. Overall, 77.3% of the committee members announce a type in the chat with substantial differences between the treatments (Table 7). Under transparency many low types stay silent, possibly due to an aversion to lying or because they believe the principal would ignore the chat. Under secrecy and mild transparency most subjects announce a type. While announced types are almost always truthful under secrecy, 51.4% of the

²⁶We give examples of chats under secrecy and transparency in the online appendix.

claims to be a high type are lies under transparency and 7.8% under mild transparency. The level of lying of high types is significantly higher under transparency than under secrecy or mild transparency (Wald-tests, $p < 0.001$).

Voting Now we turn to the voting stage and first study whether high types really vote according to their signal as our refinement assumes. Indeed, high types vote according to their signal 98.2% of the time across all treatments which is not statistically different from 100% (Wald-test, $p = 0.097$). Next, we turn to information aggregation again and study how many of the committee members who might update their beliefs after receiving a low quality blue signal and seeing the other committee member send the message “red” vote against their signal in the final vote (Table 8).

Table 8: Percentage of low types voting against their blue signal when other committee member reports a red signal

	overall	other’s claim: h	other’s claim: l	no claim
Secrecy	56.2 (6.5)	100	42.4 (11.3)	50.0 (36.7 / 2 obs.)
Transparency	41.9 (6.0)	47.2 (5.9)	44.4 (9.4 / 9 obs.)	37.5 (10.6)
Mild Transp.	45.9 (2.9)	91.3 (4.5)	32.9 (1.9)	50.0 (12.2 / 6 obs.)

Note: 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.

Low types always vote against their signal under secrecy when it matters most, i.e. when the other group member is a high type, which is facilitated by the truthfulness of announcements of types in the chat. As announcements of types and signals are also often truthful under mild transparency, low types also switch very often under this regime when the other group member is a high type. However, because of the more frequent lies of high types regarding their signal under mild transparency, aggregate error rates stay much higher than under secrecy as we have seen in the previous section. The number of subjects voting against their signal is surprisingly high for the transparency treatment where voting according to signal is predicted. Still it is substantially and significantly lower for the case where the other committee member claims to be a high type as compared to the other two

regimes (Wald-tests, $p < 0.001$).

Evaluations Next, we study how the principals evaluate. A first look at their evaluations shows that they are well able to distinguish between low types and high types (Table 9).

Table 9: Evaluations

		evaluation		
		avg.	dec. corr.	dec. wr.
Secrecy (group decision)	overall	41.6 (1.5)	54.5 (3.2)	6.9 (1.9)
	<i>h</i> type	53.5 (3.4)	53.5 (3.4)	
	<i>l</i> type	37.4 (1.5)	55.0 (3.3)	6.9 (1.9)
Transparency (individual decision)	overall	37.2 (3.0)	48.2 (1.7)	14.7 (4.0)
	<i>h</i> type	60.8 (3.2)	62.8 (3.8)	
	<i>l</i> type	29.4 (2.9)	40.2 (2.2)	14.6 (4.0)
Mild Transp. (individual decision)	overall	35.7 (1.9)	47.1 (1.9)	11.7 (1.6)
	<i>h</i> type	47.7 (1.2)	48.3 (1.1)	
	<i>l</i> type	31.9 (2.3)	46.4 (2.6)	11.8 (1.6)

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

As predicted, low types do best under secrecy and high types under transparency. However, average evaluations are too high in all treatments and even significantly positive after wrong decisions (t -tests, $p < 0.01$). However, the incentives to make a correct group decision under secrecy and a correct individual decision under (mild) transparency are about as strong as in our theoretical predictions and deciding correctly entails an evaluation that is, on average, at least 33.5 percentage points higher than after a wrong decision in all treatments. Regressing the evaluation on the round number of the experiment, controlling for the treatment and correctness of decisions gives a significant but small negative coefficient for “round” (coeff.=0.36, $p = 0.013$, Table OA1, online appendix),

suggesting that behavior moves slowly in the direction of the theoretical predictions. Under transparency, principals are able to distinguish high and low types very well even if the individual decision is correct. The evaluations do not differ significantly between the two states of the world (t -test, $p > 0.99$).

Regressing the evaluations on the pieces of information that the principal sees before evaluating, shows that making the wrong group decision has a big effect on the evaluation under secrecy but no effect under mild transparency, where she also sees the individual votes and the accuracy of these votes has the biggest influence on the evaluation (Table 10, models M1 and M3). Under transparency, the evaluation is negatively influenced by a wrong individual vote, by a wrong message or staying silent in the straw poll, and if the committee member communicates her type to be l in the chat (Table 10, M2). It is positively influenced if the committee member communicates her type to be h .

(Best) Responses In the final part of the results section we study in how far the individual behavior is a best response to the behavior of the other players.

We start with the principals. We have already seen that the evaluation levels they choose are too high, on average, to maximize their pay-offs and even positive after wrong decisions. This might be due to social preferences toward the committee members, as their payoff directly depends on the evaluation, or due to flawed belief updating. The high evaluations after correct decisions suggests that principals do not correctly take the low prior probability for being a high type into account and the regression results, which show that they lower their evaluations over time (Table OA1, online appendix), suggest that they are learning. However, the positive evaluations after wrong decisions make it also appear likely that social preferences play a role, as the updating problem is simple in this case. More important, however, is the questions whether principals react in the right direction after reviewing the different bits of information in the three treatments. This question can be best answered by comparing models M1-M3 with linear probability models with the same explanatory variables but a binary variable indicating a high type as the dependent variable (models M1c-M3c in Table 10). We see that while the average level of evaluations is too high as compared to the data, which is reflected in the

Table 10: Evaluation responses

	M1	M1c	M2	M2c	M3	M3c
Group decision D wrong	-47.6***	-36.1***	0.7	-1.6	-0.9	-2.3
	(4.2)	(1.7)	(2.9)	(2.9)	(3.3)	(4.4)
Individual vote v wrong					-34.8***	-32.7***
					(2.1)	(3.0)
<i>Combinations of m and v, reference category: message and vote are correct</i>						
v wrong, m wrong			-34.0***	-34.4***		
			(4.9)	(3.5)		
v wrong, m right			-31.3***	-19.3		
			(3.9)	(10.1)		
v right, m wrong			-20.6**	-27.9***		
			(3.8)	(3.0)		
v right, silent in straw poll			-15.5**	-30.3***		
			(4.6)	(3.7)		
v wrong, silent in straw poll			-40.6***	-32.3***		
			(5.2)	(3.0)		
Claimed to be type h in chat			15.3***	29.7***		
			(2.1)	(2.4)		
Claimed to be type l in chat			-17.5*	-7.5		
			(5.9)	(3.8)		
Constant	54.5***	36.1***	46.2***	28.2***	47.1***	35.5***
	(3.4)	(1.7)	(2.2)	(3.6)	(1.9)	(2.5)
N observations	600	600	560	560	600	600
N clusters	5	5	4	4	5	5
R^2	0.49	0.13	0.29	0.33	0.31	0.14

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (in parentheses) are adjusted for clusters in matching groups. The dependent variable in the comparison models M1c-M3c is rescaled to the same range as the dependent variable in M1-M3 and takes the value 100 if the subject is a high type and 0 if not.

differences in the constants, evaluations are influenced by all pieces of information in the right directions and most of the time but not always, as we will see, with quite accurate magnitudes.

Next, we turn to the behavior of committee members. Under secrecy low types clearly best respond by switching between the straw poll message and the final vote when the other committee member reports to be a high type and a conflicting signal. Truthful communication is a best response of high types to this behavior. All committee members act in order to maximize the probability of a correct group decision almost all the time, and a correct committee decision leads to a substantially higher evaluation than a wrong decision which means that subjects are best responding.

Under mild transparency the low types best respond by switching in case of conflicting straw poll messages and a second committee member that claims to be a high type. The level of information sharing by high types is high enough to make switching optimal in order to vote in the right direction, and correct individual votes lead to substantially higher evaluations than wrong individual votes. As the accuracy of the group decision does not influence the evaluation any level of information sharing is a (weakly) best response by high types.

Under transparency, we see a substantial fraction of low types who do not best respond. This is most obvious for those who claim to be a low type in the chat instead of announcing to be a high type. The accuracy of the group decision, which is positively influenced by this honesty, does not affect the evaluation of the principal and evaluations are substantially lower for these subjects. As a consequence, the number of committee members who claim to be a low type goes down from 12.9% in the first ten to 6.8% in the last ten rounds and this difference is statistically significant (Wald-test, $p = 0.031$). However, a substantial fraction of subjects also stays silent about their type in the chat which also leads to a lower evaluation than claiming to be a high type, although not as low as claiming to be a low type. Interestingly, the fraction of subjects staying silent about their type increases significantly from 45.7% in the first ten rounds to 55.3% in the last ten rounds (Wald-test, $p < 0.001$). Truthfully reporting to be a low type or staying silent about the type in the chat might be motivated by an aversion to lying (e.g. Fischbacher

and Föllmi-Heusi (2013); Gneezy et al. (2013)).

Behavior in the straw poll and the voting stage requires a more detailed analysis. The vast majority of high types clearly best responds by announcing their signal and type truthfully and voting for the true state of the world. Low types face a more difficult task. If they send a message, it is clearly better to announce the signal truthfully, as it matches the true state with a probability higher than one half and announcing a wrong message leads to a lower evaluation (Table 10, M2). If the other committee member announces the same message or stays silent it is also obvious that the best choice is to vote according to your signal because this maximizes the probability of being right and all subjects who switch in these cases (Table 8) are not best responding. However, if the other member reports a conflicting signal which is correct, your expected payoff would be higher if you switch than if you do not. If the other member's message is wrong, though, your expected payoff from switching is far lower than from not switching. So, it makes only sense to switch if you know with a high probability that the other member is a high type, as high types know the true state and report their signal truthfully most of the time (Table 6). However, roughly half of the claims to be a high type are lies (Table 7). As a consequence, there is a chance of roughly one quarter that you are wrong if you switch. Comparing the evaluations of low types who switch in this scenario, i.e. after observing the other committee member announce a conflicting signal and claiming to be high type, with those who do not switch, shows that those who switch are slightly better off and only the 47.2% who do switch (Table 8) are thus best responding.²⁷ Note that this results from too optimistic evaluations of switchers by the principals. A look at Table 10 reveals that evaluators give too much credit for correct votes after wrong messages. This is because high types almost always announce their signal truthfully and, as a consequence, switching between the straw poll and the voting stage reveals that you are a low type with a very high probability.

What about the low types that stay silent in the straw poll? To assess whether they are best responding we compare them to low types who announce their signal truthfully

²⁷The difference in evaluations is 6.2 percentage points and is not statistically different from zero (*t*-test, $p = 0.209$).

and do not switch. The latter group receives an average evaluation of 61.5 percentage points if they are right, i.e. with probability 0.55, and a payoff of 27.5 if they are wrong, i.e. with probability 0.45, which gives an expected payoff of 46.2.²⁸ If you stay silent in the straw poll, it is optimal to vote according to your own signal if the other committee member claims to be a low type (or stays silent in the straw poll as well). In this case your payoff would be 46 if you are right and 20.9 if you are wrong with a slightly lower (equal) probability for the good outcome as above. The best you can hope for after staying silent in the straw poll is to be in a group in which the other member has a conflicting signal and claims to be a high type. As we have seen, this claim is true with probability 0.486, which results in a probability of being right after voting against your signal of slightly less than three quarters and an expected payoff of 39.5. So, even in the best case the expected payoff is lower than from announcing your signal truthfully and not switching. Thus the subjects who stay silent in the straw poll do not play a best response.

Career concerns play out in a more subtle way under transparency than in the other two regimes. So, it is not surprising to see more subjects not playing a best response here. Despite these deviations from our theoretical predictions, information aggregation is still far worse than under secrecy, as predicted, which is driven by a substantial fraction of subjects who do not switch if the other member announces to be a high type and a conflicting signal, and by the high degree of lying with respect to types, i.e. signal strengths.

4. Discussion

In this paper we constructed a model of committee decision-making in which members are career concerned. We examined how the incentives of committee members to share their private information varies with the level of transparency i.e. how much of the decision process the principal observes. The mechanisms through which transparency harms efficiency involve strategic communication of committee members, and a number of ex-

²⁸Here, we are using the regression results from model M2 and assume for simplicity that the low type optimally announces to be a high type in the chat.

perimental studies have shown for various settings that subjects communicate far less strategically than theoretically predicted. As such, we were eager to empirically test our model. By bringing it to the laboratory, we gained a number of insights.

First, under transparency committee members are more truthful than the model's most informative equilibrium allows. While it is somewhat startling that players tell the truth when it is optimal to lie, this is in line with previous experimental results which show that people are often overly truthful and apparently face a cost of lying. This suggests that the negative effect of transparency on group decision-making, though important, may not be as severe as theory suggests.

Second, under mild transparency play is significantly less truthful than the most informative equilibrium. This seems to stand in sharp contrast to our previous point. However, as we show in the appendix, under secrecy and transparency the most informative equilibrium coincides with the best equilibrium for high ability players (those with the most information to share). Under mild transparency, however, the most informative 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 large degree of deception from high ability types under mild transparency. As a result, group errors and player evaluations here are much closer to the transparency case than to secrecy, casting doubt on the predictive power of the most informative equilibrium. This is a finding which is relevant beyond this particular context. Sobel (2013) called for experimental tests of the predictive power of the most informative equilibrium in cheap talk games in which other equilibria exist which are preferred by at least one type of player. What does this tell us about the desirability of pre-meetings? Swank and Visser (2013) show in their model that, when players do not know their own abilities, pre-meetings can undo the benefit of transparency. In our setting, where players know their abilities, pre-meetings can reduce the probability of mistakes if players share information while if no information is credibly shared pre-meetings will simply have no effect.

Finally, despite some deviations from theoretical predictions, we find large differences in behavior between the three regimes: information aggregation is successful under secrecy while it breaks down, for different reasons, under transparency and mild trans-

parency. This results in fewer mistakes under secrecy than the two other treatments and suggests that the level of transparency is a highly important element of institutional design - setting it wrong might indeed have considerable negative consequences for a principal. The difference in the level of truth-telling between treatments contrasts sharply with the results of Goeree and Yariv (2011). They find that free-form communication greatly diminishes institutional differences; players are overwhelmingly truthful regardless of the voting rule. Our results show that the degree of transparency in a committee strongly affects members' behavior even with open chat. This suggests that the level of openness may be a more important institutional choice than that of the voting rule.

References

- BATTAGLINI, M. AND U. MAKAROV (2014): "Cheap talk with multiple audiences: An experimental analysis," *Games and Economic Behavior*, 83, 147–164.
- BATTAGLINI, M., R. B. MORTON, AND T. R. PALFREY (2009): "The Swing Voter's Curse in the Laboratory," *Review of Economic Studies*, 77, 61–89.
- BENTHAM, J. ([1816] 1999): *Political Tactics*, Oxford: Oxford University Press.
- BHATTACHARYA, S., J. DUFFY, AND S.-T. KIM (2014): "Compulsory versus voluntary voting: An experimental study," *Games and Economic Behavior*, 84, 111–131.
- BOUTON, L., M. CASTANHEIRA, AND A. LLORENTE-SAGUER (2014): "Divided Majority and Information Aggregation: Theory and Experiment," *Mimeo*.
- CAI, H. AND J. T.-Y. WANG (2006): "Overcommunication in strategic information transmission games," *Games and Economic Behavior*, 56, 7–36.
- CHEN, Y., N. KARTIK, AND J. SOBEL (2008): "Selecting cheap-talk equilibria," *Econometrica*, 76, 117–136.

- COHEN, J. (1996): "Deliberation and Democratic Legitimacy," in *Deliberative Democracy: Essays on Reason and Politics*, ed. by J. Bohman and W. Rehg, Cambridge, MA: MIT Press, 67–92.
- COUGHLAN, P. (2000): "In Defense of Unanimous Jury Verdicts : Mistrials , Communication , and Strategic Voting," *American Political Science Review*, 94, 375–393.
- CRAWFORD, V. AND J. SOBEL (1982): "Strategic information transmission," *Econometrica*, 50, 1431–1451.
- CROSS, J. P. (2013): "Striking a pose: Transparency and position taking in the Council of the European Union," *European Journal of Political Research*, 52, 291–315.
- FISCHBACHER, U. AND F. FÖLLMI-HEUSI (2013): "Lies in Disguise-an Experimental Study on Cheating," *Journal of the European Economic Association*, 11, 525–547.
- FOX, J. AND R. VAN WEELDEN (2012): "Costly Transparency," *Journal of Public Economics*, 96, 142–150.
- GERSBACH, H. AND V. HAHN (2008): "Should the Individual Voting Records of Central Bankers be Published?" *Social Choice and Welfare*, 30, 655–683.
- GNEEZY, U. (2005): "Deception: The role of consequences," *The American Economic Review*, 95, 384–394.
- GNEEZY, U., B. ROCKENBACH, AND M. SERRA-GARCIA (2013): "Measuring lying aversion," *Journal of Economic Behavior & Organization*, 93, 293–300.
- GOEREE, J. K. AND L. YARIV (2011): "An Experimental Study of Collective Deliberation," *Econometrica*, 79, 893–921.
- GUARNASCHELLI, S., R. D. MCKELVEY, AND T. R. PALFREY (2000): "An Experimental Study of Jury Decision Rules," *American Political Science Review*, 94, 407.
- HABERMAS, J. (1996): *Between Facts and Norms*, Cambridge, MA: MIT Press.

- LANDA, D. AND A. MEIROWITZ (2009): “Game Theory, Information, and Deliberative Democracy,” *American Journal of Political Science*, 53, 427–444.
- LEVY, G. (2007): “Decision Making in Committees: Transparency, Reputation, and Voting Rules,” *American Economic Review*, 97, 150–168.
- MALENKO, N. (2013): “Communication and Decision-Making in Corporate Boards,” *Review of Financial Studies*, 27, 1486–1532.
- MEADE, E. (2005): “The FOMC: preferences, voting, and consensus,” *Federal Reserve Bank of St. Louis Review*, 93–102.
- MEADE, E. E. AND D. STASAVAGE (2008): “Publicity of Debate and the Incentive to Dissent: Evidence from the US Federal Reserve,” *The Economic Journal*, 118, 695–717.
- NYARKO, Y. AND A. SCHOTTER (2003): “An experimental study of belief learning using elicited beliefs,” *Econometrica*, 70, 971–1005.
- OFFERMAN, T., J. SONNEMANS, G. VAN DE KUILEN, AND P. P. WAKKER (2009): “A Truth Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes,” *Review of Economic Studies*, 76, 1461–1489.
- OTTAVIANI, M. AND P. SØRENSEN (2001): “Information aggregation in debate: who should speak first?” *Journal of Public Economics*, 81, 393–421.
- PRAT, A. (2005): “The Wrong Kind of Transparency,” *American Economic Review*, 95, 862–877.
- SCHOTTER, A. AND I. TREVINO (2014): “Belief Elicitation in the Laboratory,” *Annual Review of Economics*, 6, 103–128.
- SOBEL, J. (2013): “Ten possible experiments on communication and deception,” *Journal of Economic Behavior & Organization*, 93, 408–413.

- STASAVAGE, D. (2007): “Polarization and publicity: rethinking the benefits of deliberative democracy,” *Journal of Politics*, 69, 59–72.
- 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, 478–486.
- SWANK, O. H. AND B. VISSER (2013): “Is Transparency to No Avail?” *Scandinavian Journal of Economics*, 115, 967–994.
- VISSER, B. AND O. SWANK (2007): “On committees of experts,” *The Quarterly Journal of Economics*, 122, 337–372.
- WANG, J. T.-Y., M. SPEZIO, AND C. 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, 984–1007.
- WILLIAMS, J. C. (2012): “Opening the Temple,” *Federal Reserve Bank of San Francisco 2011 Annual Report*.

Appendix

Proof of Proposition 1

First we show that (1) truth-telling and implementing the posterior most likely state is an equilibrium; then (2) we calculate the principal's evaluations and the probability of mistakes. The most informative outcome of the two communication stage stages is, by definition, where each player reveals his signal and ability. Suppose this is the case in equilibrium and all players truthfully reveal their signals and abilities. Both committee members will agree on which state is most likely and will wish to implement the policy which maximises their expected evaluation. As h types vote to signal, it must be that the principal's posterior beliefs satisfy $\hat{q}(D \neq S) < \hat{q}(D = S)$. Given the principal's beliefs, the optimal strategy is for the committee to implement the policy with the highest posterior probability of matching the state. In the case of a balanced posterior, i.e. a group (lr, lb) , the committee would choose each project with probability 0.5. This can be achieved by making use of the coordination device (e.g both voting v_B if the draw is below 0.5 and voting v_R if it is above). If such committees chose one state with a higher probability than the other, the principal would incorporate this in his evaluations, and the resulting payoffs would be lower. So, in equilibrium a committee with a balanced posterior must choose each state with equal probability. Given full information sharing in communication followed by the group implementing the policy most likely to match the state, is there an incentive to deviate from full information sharing? As players have a common interest in matching the policy to the state of the world, a deviation from truth-telling can only decrease expected utility. For example, if a hr player announced m^b, τ^l and was matched with a lb player, the former would vote v_B , and both players would receive a low evaluation $\hat{q}(D \neq S)$. As players can only do worse by deviating from full information sharing, it is indeed an equilibrium.

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 option R in state B is $(1-q)^2(1-\sigma)$. The case of a mistake in state R is symmetric.

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:

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

Proof of Proposition 2

First we show that (1) there is no equilibrium with full information revelation; then (2) the most informative equilibrium involves all players truthfully announcing the signal and voting to signal; finally (3) we calculate the probability of mistakes and the principal's evaluations.

Suppose there existed an equilibrium in which each player truthfully reveals his ability and type. Here 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 pool with hb types by announcing τ^h rather than τ^l and thus gain an evaluation $\hat{q} = 1$ with probability σ . An lr type has the same incentive to pool with hr types by announcing τ^h . Therefore, truthful revelation of types cannot be an equilibrium. In fact, any strategy a h type follows in announcing an ability level can be mimicked perfectly by l types. As these low types have an incentive to pool, it follows that no information can be communicated about ability levels in equilibrium. Therefore the most informative level of communication we can hope for is where players each reveal their signals truthfully but pool on the same strategy in announcing abilities.

In the voting stage, h types vote to signal. Therefore if players reveal their signal in the straw poll the principal's posterior can only be positive in two cases: $\hat{q}(m^B, v_B | S = B)$ and $\hat{q}(m^R, v_R | S = R)$ and possibly for out-of-equilibrium actions. If we set the out-of-

equilibrium belief for announcing an $m \neq v$ to zero the l types must follow the same strategy as h types. If an l type voted differently to his announced signal, it would be clear to the principal that he is an l type and so would receive an evaluation of zero. As σ signals are informative, the optimal strategy for an l type is to announce his signal truthfully, mimic the strategy of the h type in the second communication stage, and then vote to signal.

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 $1 - (q + (1 - q)\sigma)^2$.

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 $\frac{q}{q+(1-q)\sigma}$ if a member announces the correct signal and also votes for that signal.

Proof of Proposition 3

If players are sharing all information and h types are voting to signal, then $\hat{q}(v_B|S = B) > \hat{q}(v_R|S = B)$ and $\hat{q}(v_R|S = R) > \hat{q}(v_B|S = R)$. The best response for a l type is to vote for the policy with the higher posterior probability of matching the state. In the case of a balanced posterior, i.e. a group (lr, lb) , the committee would vote unanimously for B or R each with probability 0.5. This 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). There are two reasons for doing so. First, as under secrecy, if such committees chose one state more than the other, the principal would incorporate this in his evaluations, and the resulting payoffs would be lower. Second, if the group did not come to a unanimous decision, the principal would know there are no h types. In equilibrium there will be no non-unanimous decisions, so the incentive for h types to tell the truth is pinned down by off path beliefs that $\hat{q}_i(v_B^i, v_B^j|S = B) > \hat{q}_i(v_B^i, v_R^j|S = B)$. With such beliefs, no player will have an incentive to deviate from truth-telling. The calculation of evaluations and probability of mistakes are found in the proof of the secrecy proposition.

Proof of Proposition 4

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

$$\begin{aligned} x[Pr_{tran}(D = B|B)] + (1-x)[Pr_{tran}(D = R|R)] > \\ x[Pr_{sec}(D = B|B)] + (1-x)[Pr_{sec}(D = R|R)] \end{aligned}$$

which can be rearranged to

$$(1-x)[Pr_{tran}(D = B|R) - Pr_{sec}(D = B|R)] < x[Pr_{sec}(D = R|B) - Pr_{tran}(D = R|B)]$$

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

$$\begin{aligned} \frac{(1-q)^2(1-\sigma)\sigma + 2q(1-q)(1-\sigma)}{2(1-q)^2(1-\sigma)\sigma + 2q(1-q)(1-\sigma)} < x \\ \frac{(1-q)\sigma + 2q}{2(1-q)\sigma + 2q} \equiv x^* < x \end{aligned} \quad (2)$$

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

When is the most informative equilibrium payoff dominated?

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

$$\frac{q}{1 - (1-q)(1-\sigma)} > \frac{q}{1 - (1-q)^2(1-\sigma)} \quad (3)$$

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.

$$\frac{\sigma q}{1 - (1-q)(1-\sigma)} < \frac{((1-q)\sigma + q)q}{1 - (1-q)^2(1-\sigma)} \quad (4)$$

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. As individual votes are observed by the principal and h types can learn nothing from l types (they already know the state of the world), the highest evaluation a h type can get, is when all players vote to signal. If some credible information could be communicated such that l types sometimes do not vote to signal, this could only decrease the expected evaluation of h types.

Mild Transparency Here the most informative equilibrium coincides with that of secrecy. However, as the principal now observes individual votes, a 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) which 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, $\frac{q}{1-(1-q)(1-\sigma)}$. There are a series of equilibria between these two poles which 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 a 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 10 shows, this is what we find in our laboratory setting.

For Online Publication

Online Appendix

OA1 Evaluations over time

Table OA1: Evaluations over rounds

MOA1		
<i>dependent variable: evaluation</i>		
Round number	-0.36**	(0.12)
<i>Secrecy</i>		
decision correct	58.21***	(3.75)
decision wrong	10.62***	(1.78)
<i>Transparency</i>		
decision correct	51.89***	(2.05)
decision wrong	18.58***	(4.14)
<i>Mild Transparency</i>		
decision correct	50.86***	(2.60)
decision wrong	15.36***	(2.39)
<i>N</i> observations	1760	
<i>N</i> clusters	14	
R^2	0.31	

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses are adjusted for clusters in matching groups. The model was estimated without a constant. The R^2 statistic was computed with a constant, leaving out one group as a reference group.

OA2 Sample chat conversations (translated from German)

Further chat protocols are available from the authors upon request.

Conversation 1 (Secrecy):

hey (voter 1)

type h (voter 2)

hello (voter 2)

I am type l (voter 1)

Red! (voter 2)

ok (voter 1)

Conversation 2 (Transparency):

May I guess, you will also again say you are of type h? (voter 2)

Cannot be the case that both are always h. (voter 2)

You can really give up on the chat. (voter 1)

OA3 Instructions to the experiment and screenshots

Instructions and screenshots for the transparency treatment (translation; original in German). Instructions and screens for the other treatments where very similar and are therefore omitted here. The original instructions can be obtained from the authors upon request.

Overview

Welcome to this experiment. We kindly ask you not to communicate with other participants during the experiment and to switch off your phones and other mobile devices.

At the end of the experiment you will be paid out in cash for your participation in today's session. The amount of your pay-off depends in parts on your decisions, on the decisions of other participants and on chance. For this reason it is important that you read the instructions carefully and understand them before the start of the experiment.

In this experiment all interactions between participants take place via the computers that you are sitting in front of. You will interact anonymously and your decisions will only be stored together with your random ID number. Neither your name nor names of other participants will be made public, not today and not in future written analyses.

Today's session consists of several rounds. **At the end, 4 rounds will be randomly selected and paid out. The rounds that are not chosen will not be paid out.** Your pay-off results from the points that you earn in the selected rounds, converted to Swiss Francs, plus your show-up fee of CHF 10. The conversion of points to Swiss Francs happens as follows: Every point is worth 15 cents, which means that

20 points = CHF 3.00.

Every participant will be paid out in private at the payment counter, so that no other participant can see how much you have earned.

Experiment

This experiment consists of 20 procedurally identical rounds. In each round a group decision has to be made that can be correct or wrong.

Two members in each group of three make the group decision (henceforth we will call them the voters). There are well and less well informed voters and the task of the third group member is to observe the decision process of the other two members and then to indicate the probability with which he thinks that the other group members are well or less well informed (henceforth, we will call this member the observer).

The higher the evaluation of the observer with respect to the level of information of a voter is, the higher is the pay-off to that voter in the round. The more accurate the evaluation of the observer with respect to the level of information of the voters is, the higher is his or her pay-off in the round. In addition, the observer receives a pay-off for correct group decisions.

The Group

In the **first** round you will be assigned a meta-group of 9 members. In the beginning of **every** round you will be randomly assigned to a **new group** which consists of randomly selected members of your meta-group. Every group has three members: 2 voters and 1 observer.

Whether you will be assigned the role of a voter or an observer, is randomly determined each round. The voters receive, again randomly, the labels “voter 1” and “voter 2”.

All interactions in a round take place within your group of three.

The Voters

There are two types of voters, well informed (type G) and (less well) informed (type I) voters. Of which type the group members are, is again determined randomly. With probability $\frac{1}{4}$ (or 25%) a voter receives good information which means he is of type G; with probability $\frac{3}{4}$ (or 75%) he receives less good information which means he is of type I.

Because the assignment of types to the voters is independent of the assignment to other voters, there can be two voters of type G, two voters of type I, or one of each type in a group.

The voters learn their type on the first screen of a round but not the type of the other voter in their group. The observer learns that he is an observer on the first screen but not the types of the voters in his group.

Later, after observing the behavior of the voters, it will be the task of the observer to estimate the probabilities that voter 1 and voter 2 are of type G.

The Jar

There are two jars: one red jar and one blue jar. The red jar contains 11 red and 9 blue balls, the blue jar 11 blue and 9 red balls. Each round one jar will be randomly selected.

The task of the voters is to vote on the color of the jar. Each jar has an equal probability of being selected, that is, it will be selected with 50% probability.

The Ball


The well informed voters (type G) receive a ball with the actual color of the jar, that is they are directly informed about the color of the jar.

The informed voters (type I) receive a randomly drawn ball from the selected jar. They are not told the color of the jar. If there are two type I voters in a group, each of them receives a ball from the jar. Every ball in the jar has the same selection probability for the type I voters, that is for each voter of type I a ball is drawn from a jar containing 20 balls (11 with the color of the jar, 9 with the other color).

The voters learn the color of their ball on the first screen. Every voter only sees the color of his ball, not the color of the other voter's ball.

Communication

After learning their type and the color of their ball, the voters can communicate the color of their ball to the other voter in their group. They can also communicate the color that their ball did not have or stay silent. The communication is made through the following entry mask.



The image shows a user interface for communication. At the top, there are three radio buttons labeled "BLAU", "ROT", and "Keine Angabe". Below these, there is a red-bordered input field containing the text "blue red abstain".

On the following screen the voters learn the message of the other voter in their group and have the option to chat with him. The chat happens via the following entry mask.

You can now chat with the other voter in your group. You will have the ID voter 1 and the other voter the ID voter 2 for this chat. Every entry has to be confirmed by pressing the enter button.

Sie können nun mit dem anderen Wähler in ihrer Gruppe chatten. Im Chat bekommen Sie die ID "Wähler 1" und der andere Wähler die ID "Wähler 2". Jede Texteingabe muss mit der Enter-Taste bestätigt werden.

Wähler 2: Hallo

You can enter arbitrary text messages into the blue entry field. Pay attention to confirm every entry by pressing the enter button to make it visible for the other voter. It will then appear in the grey field above.

The observer cannot participate in the communication but sees the messages of the two voters regarding the color of their ball as well as the chat.

Group Decision

After the communication stage the voters make their decision in a group vote.

So, if you are a voter, you have to vote either for blue or for red.

Once both voters have made their decision, the votes for blue and red are counted and the group decision results from the following rule:

- **If the color RED receives 2 votes, the group decision is RED**
- **If the color BLUE receives 1 or 2 votes, the group decision is BLUE**

That is for a group decision for blue only one vote is necessary while a group decision for red requires two votes.

Evaluation of the Observer

After the voters have cast their vote and the group decision is determined, the evaluator learns the group decision as well as the decisions of the individual voters in his group.

Moreover, he learns the true color of the jar, that is, whether the group decision and the individual decisions were correct or wrong.

On the same screen the observer can review the entire communication between the voters in his group once again.

If you are an observer, you now have to enter for each of the two voters the probability with which you believe that this voter is of type G.

To do so you enter a number between 0 and 100 which expresses your evaluation in percentage points. The entry mask looks as follows.

Geben Sie die Wahrscheinlichkeit in Prozent an, mit der Sie glauben, dass **Wähler 1** ein Typ G ist (eine Zahl zwischen 0 und 100):

Enter the probability in per cent, with which you believe that voter 1 is a high type.

The complete screen of the observer looks as follows (example screen).

Die Dose hat die Farbe **ROT** The color of the jar is red.

Die Wähler haben wie folgt abgestimmt The voters have voted as follows

Wähler 1 hat für folgende Farbe abgestimmt: **ROT** Voter 1 voted for RED. Wähler 2 hat für folgende Farbe abgestimmt: **ROT** Voter 2 voted for RED.

Die Gruppenentscheidung ist somit: **ROT** The resulting group decision is: RED.

Hier sehen Sie die Kommunikation zwischen den Wählern noch einmal Here you can once again review the communication between the voters.

Die Kugel, die Wähler 1 angab, ist: **ROT** Voter 1 indicated a red ball. Die Kugel, die Wähler 2 angab, ist: **ROT** Voter 2 indicated a red ball.

Wähler 2: Hallo

Geben Sie nun Ihre Einschätzung zu Wähler 1 und Wähler 2 an. Now, please, indicate your evaluation of voter 1 and voter 2.

Geben Sie die Wahrscheinlichkeit in Prozent an, mit der Sie glauben, dass **Wähler 1** ein Typ G ist (eine Zahl zwischen 0 und 100): Enter the probability in per cent, with which you believe that voter 1 is a high type.

Geben Sie die Wahrscheinlichkeit in Prozent an, mit der Sie glauben, dass **Wähler 2** ein Typ G ist (eine Zahl zwischen 0 und 100): Enter the probability in per cent, with which you believe that voter 2 is a high type.

Weiter

Pay-off in each Round

If you are a **voter** your pay-off is determined by the evaluation of the observer. If the observer believes that you are of type G with **X%** probability, you receive a pay-off of **2*X** points in

this round. This means that your pay-off directly depends on the probability with which the observer believes you are a **well-informed voter (type G)**.

If the observer has entered the probability 25%, for example, your pay-off is 50 points, if he has entered 50%, it is 100 points.

If you are an **observer** you receive a pay-off for correct group decisions and a pay-off for the accuracy of your evaluations of the types of the voters.

- If the group decision is RED and the jar is indeed RED, you as an observer receive **1 point**.
- If the group decision is BLUE and the jar is indeed BLUE, you as an observer receive **3 points**.
- If the group decision is wrong, you receive **0 points**, independently of the true color of the jar.

For your evaluation regarding the types of the voters you receive a pay-off between 0 and 100 points. It will be randomly determined whether you will be paid out for the evaluation of voter 1 or voter 2.

If you have evaluated both voters correctly with certainty (that is with 0 or 100%) (if you entered the probability 0 for both voters, for example, and both are indeed not of type G but of type I), you receive 100 points. If you are completely wrong (if both are of type G in the example) you receive 0 points.

The formula that determines your pay-off is a little complicated.

Put simple the formula assures **that it is best for you (gives you the highest expected pay-off) if you truthfully indicate the probability with which you believe that a voter is indeed of type G**. Every other evaluation lowers your expected pay-off.

If you believe, for example, that voter 1 in your group is of type G with 30% probability and voter 2 with 60% probability, it is best for you to enter exactly these values.

In case you want to know in more detail how your payoff is determined: for the evaluation of the randomly selected voter you receive:

$$100 - \frac{1}{100} (100 - \text{prob}(\text{voter is of type G}))^2, \text{ if this voter is of type G and}$$
$$100 - \frac{1}{100} (\text{prob}(\text{voter is of type G}))^2, \text{ if this voter is of type I,}$$

where $(\text{prob}(\text{voter is of type G}))$ is your indication of probability in percentage points that that voter is of type G. The resulting number is rounded up to a whole number and gives, together with your pay-off in case of a correct group decision, your pay-off in the round.

Remember: At the end of the experiment 4 rounds are randomly selected, the point incomes converted to Swiss Francs and paid out in private. The rounds that were not selected will not be paid out.

Questions?

Take your time to read the instructions carefully. If you have any questions, raise your hand. An experimental administrator will then come to your seat.

Screenshots *(not part of the instructions)*

First screen of a committee member

(The observer's first screen only informed the subject that he is an observer in that round.)

Sie sind ein Wähler vom
Typ I

Die Kugel, die Sie erhalten haben ist
ROT

You are a voter of type L

The ball that you have received is RED

You can now tell the second voter in your group what color your ball had. You can also stay silent or indicate the color that your ball did not have.

Sie können nun dem zweiten Wähler in ihrer Gruppe mitteilen, welche Farbe ihre Kugel hatte. Sie können auch schweigen, oder die Farbe, die ihre Kugel nicht hatte, als die Farbe ihrer Kugel angeben:

BLAU ROT Keine Angabe

blue red abstain

Weiter

Second screen of a committee member

Verbleibende Zeit [sec]: 72

You are a voter of type L

Sie sind ein Wähler vom
Typ I

The ball you received is red.

Die Kugel, die Sie erhalten haben ist
ROT

Die Kugel, die der andere Wähler in ihrer Gruppe angibt bekommen zu haben ist
ROT

The color of the ball of the other member of your group as indicated by him is red.

You can now chat with the other voter in your group. You will have the ID voter 1 and the other voter the ID voter 2 for this chat. Every entry has to be confirmed by pressing the enter button.

Sie können nun mit dem anderen Wähler in ihrer Gruppe chatten. Im Chat bekommen Sie die ID "Wähler 1" und der andere Wähler die ID "Wähler 2". Jede Texteingabe muss mit der Enter-Taste bestätigt werden.

Wähler 2: Hallo

Weiter

In the transparency treatment the principal could follow the chat in real time on a screen with a very similar layout. Under mild transparency and secrecy the principal just saw a waiting screen during communication.

Third screen of a committee member

You are a voter of type L
Sie sind ein Wähler vom
Typ I

The color of the ball of the other member of your group as indicated by him is red.
Die Kugel, die der andere Wähler in ihrer Gruppe angab bekommen zu haben ist
ROT

Die Kugel, die Sie erhalten haben ist
ROT

Here you can review the chat. Your ID was voter 1 and the other voter's ID was voter 2 for this chat.

Hier können Sie den Chat noch einmal nachlesen. Ihre ID war "Wähler 1" und die des anderen Wählers "Wähler 2".

Wähler 2: Hallo

Now make your final voting decision:
Treffen Sie nun Ihre Entscheidung für die Abstimmung:

BLAU ROT
blue red

Weiter

Evaluation screen of a principal in transparency treatment

Die Dose hat die Farbe **ROT** The color of the jar is red.

Die Wähler haben wie folgt abgestimmt The voters have voted as follows

Wähler 1 hat für folgende Farbe abgestimmt: **ROT** Voter 1 voted for RED. Wähler 2 hat für folgende Farbe abgestimmt: **ROT** Voter 2 voted for RED.

Die Gruppenentscheidung ist somit: **ROT** The resulting group decision is: RED.

Hier sehen Sie die Kommunikation zwischen den Wählern noch einmal
Here you can once again review the communication between the voters.

Die Kugel, die Wähler 1 angab, ist: **ROT** Voter 1 indicated a red ball. Die Kugel, die Wähler 2 angab, ist: **ROT** Voter 2 indicated a red ball.

Wähler 2: Hallo

Geben Sie nun Ihre Einschätzung zu Wähler 1 und Wähler 2 an.
Now, please, indicate your evaluation of voter 1 and voter 2.

Geben Sie die Wahrscheinlichkeit in Prozent an, mit der Sie glauben, dass **Wähler 1** ein Typ G ist (eine Zahl zwischen 0 und 100):

Enter the probability in per cent, with which you believe that voter 1 is a high type.

Geben Sie die Wahrscheinlichkeit in Prozent an, mit der Sie glauben, dass **Wähler 2** ein Typ G ist (eine Zahl zwischen 0 und 100):

Enter the probability in per cent, with which you believe that voter 2 is a high type.

Weiter

The evaluation screen had the same lay-out in the other two treatment but with several elements left out. Under mild transparency the communication part was left out. Under secrecy the individual votes and communication were left out and only one randomly selected committee member had to be evaluated.

Feedback screen for a voter at the end of a round

The screenshot shows a feedback screen for a voter at the end of a round. The screen is divided into several sections:

- Top section:** Two columns of information. The left column shows "Type of voter 1" (Typ von Wähler 1) and "Type L" (Typ I). The right column shows "Type of voter 2" (Typ von Wähler 2) and "Type H" (Typ G).
- Middle section:** A horizontal line separates the voter types from the voting information. Below the line, it states "The other voter in your group voted for: RED" (Der andere Wähler in ihrer Gruppe hat für folgende Farbe abgestimmt: Rot).
- Bottom section:** It displays the group decision: "Die Gruppenentscheidung ist Rot" (The group decision is: RED) and the jar color: "Die Dose hat die Farbe Rot" (The color of the jar is: RED).
- Payoff section:** At the bottom, it states "Ihre Rundenauszahlung (d.h. die Einschätzung des Beobachters verdoppelt) beträgt 60 Punkte" (Your payoff from this round (i.e. double the observer's evaluation) is 60 points).
- Navigation:** A red button labeled "Weiter" (Next) is located in the bottom right corner.

The feedback screens looked very similar for principals. At the end of the last round subjects saw a final screen which reported the rounds which were randomly selected to be paid out and the total earnings in Swiss Francs.