HOW DO GROUPS SPEAK AND HOW ARE THEY UNDERSTOOD?*

Ala Avoyan
Indiana University

Paula Onuchic
London School of Economics

August, 2025

Abstract

We experimentally study an environment where a group of senders communicates with a receiver through evidence disclosure. Disclosure decisions aggregate group members' preferences through a given procedure, which we vary across treatments. In line with theoretical results, our experimental evidence establishes a relationship between the aggregation procedure and the receiver's interpretation of "no disclosure:" the receiver's beliefs are more skeptical about group members who have more power to enforce disclosure. In turn, these beliefs are justified by differences in groups' empirical disclosure strategies in each treatment. At the individual level, we show that deviations from the theory are related to the decision-making procedure: senders err towards under- or over-communication, depending on their pivotality, and receivers' deviations from Bayesian updating are related to the complexity of the procedure.

1 Introduction

This paper experimentally studies a group communication game, in which a *group of senders* with distinct interests collectively communicates with a receiver through the disclosure of verifiable information. Disclosure of verifiable evidence by a single sender is a

*We thank João Ramos, Jeanne Hagenbach, and Andrew Schotter for their help and encouragement, and Tim Besley, Aislinn Bohren, Katarina Brütt, Agata Farina, Amanda Friedenberg, Eduardo Perez-Richet, Chloe Tergiman, Peter Schwardmann, Alistair Wilson, Leeat Yariv, and Sevgi Yuksel for engaging discussions and helpful suggestions.

touchstone model of communication, introduced by Grossman (1981) and Milgrom (1981), in which the key *unravelling result* establishes that a sender discloses all available evidence (even if unfavorable), and that the absence of evidence disclosure is regarded with skepticism by the receiver. The experimental literature on disclosure games mainly investigates whether receivers indeed form such skeptical beliefs that incentivize full disclosure.

We instead study the formation of receivers' beliefs — and its impact on senders' strategic disclosure choices — in a group disclosure setting, in which the unravelling result does not apply. Our experimental design is based on an environment of group disclosure introduced by Onuchic and Ramos (2025), which extends the disclosure models of Grossman (1981) and Milgrom (1981) to a context in which a group of agents with conflicting preferences makes a collective decision regarding the disclosure of a piece of evidence. If the evidence is disclosed, the receiver understands it and makes decisions based on that information. If the evidence is not disclosed, the receiver does not learn anything directly, but infers some information indirectly from the fact that not disclosing was a decision made strategically by the group of senders.

The group's disclosure decision aggregates group members' disclosure recommendations (each based on their own preferences) through some pre-determined procedure, which we denote the *deliberation procedure*. In classic disclosure games, where disclosure decisions are made by an individual sender, the receiver interprets the absence of disclosure in equilibrium with skepticism about the value of the state to the sender: the individual's strategic choice not to disclose is a signal that the realized evidence is not favorable. Likewise in the group disclosure setting, the receiver understands "no disclosure" as a strategic choice made by the group, but in this case they must take into account the *process* through which the group reached that decision in order to understand who is to blame for the no disclosure decision, and therefore to which member (or members) of the group the realized evidence is not favorable.

Two theoretical results follow from the attribution problem inherent to the receiver's updating procedure: First, the classic unravelling result does not apply, and equilibria of the group disclosure game typically do not feature full disclosure. Second, the observer attributes more blame for the group's decision to not disclose to group members who have more power, where the balance of power is expressed in the group's deliberation procedure (which is a primitive of the model). Consequently, "no disclosure" is interpreted in equilibrium with more skepticism about the value of the state to group members who have more power over the group's decision. Our experimental inquiry focuses on the relation between the power structure in a group and outcomes of equilibrium communication established by this comparative statics result.

The experimental design closely parallels the group disclosure model. In the lab, subjects are grouped into units of three, two group members (A and B) and an evaluator (the observer). In each round, the group draws a pair of cards (the evidence), which describes the value of that evidence to each of the group members. After seeing their respective values, group members make recommendations, suggesting to report or not to report the drawn cards to the evaluator. These recommendations are aggregated into a group decision to report or not report via a pre-determined deliberation procedure that is known to all players. After seeing/not seeing the cards, the evaluator is asked to guess each group member's value. The evaluator is incentivized to make accurate guesses, and each group member's payoff is increasing in the evaluator's guess of their own value.

In the three main treatments, we vary the deliberation procedure through which the group members' recommendations are aggregated into the group's disclosure decisions, so as to vary the balance of power within the group. The *unilateral* procedure is such that disclosure happens if at least one of the group members recommends it; the *consensus* procedure is such that the group discloses the outcome only if both group members recommend the outcome's disclosure; and in the *leader* procedure the group's decision almost always equals group member A's recommended action (so that group member A is the "leader" and group member B is the "non-leader"). As a baseline, we also consider a comparable *individual* treatment, in which a single individual makes disclosure decisions (regarding an outcome conveying only their own individual value).

Under the unilateral procedure, each group member can individually enforce the disclosure of the evidence; this is the case also for group member A under the leader procedure and for the individual in the non-group benchmark. The theory predicts that the evaluator should regard the absence of disclosure with maximal skepticism about each of these players, given their full power to choose disclosure on behalf of their group. Under the consensus procedure, disclosure only arises after the agreement of both parties, so neither group member has the power to enforce disclosure alone. Consequently, the evaluator should be less than "maximally" skeptical about these players after seeing no disclosure. Finally, the interpretation of no disclosure should be even less skeptical about the value of the realized evidence to group member B in the leader treatment, who is the non-leader and has no power over the group's disclosure decision.

We find in the lab that the interpretation of group communication varies with the group's deliberation procedure as predicted by our theory: after the group chooses not to report, the evaluator's skepticism, gleaned from their guesses of group member A and group member B's values, is larger for subjects who have more power over the group's disclosure decision in their treatment. While our hypothesized ordering on the evaluator's skepticism across treatments is confirmed in the experiment, we find that the exact skepticism levels differ significantly in theory and in practice. In line with previous experimental literature on games of individual disclosure, evaluators in our setting tend to be less skeptical than predicted by the theory, both in treatments where theory predicts full unravelling and in treatments in which equilibria do not feature "maximal skepticism." However, in all but the unilateral treatment, we cannot reject that evaluators use empirical best responses when submitting their "no disclosure" guesses.

By comparing the experimental data in the various group treatments to the baseline individual disclosure treatment, we show that communication coming from a group is interpreted fundamentally differently from communication coming from an individual, even in circumstances where theory predicts group and individual communication equilibria to be comparable. For example, we find that the observer is less skeptical about each group member in the unilateral treatment than about the individual in the individual treatment. In contrast, our theory predicts "maximal skepticism" in both these cases. We interpret this observation in light of the phenomenon of *diffusion of responsibility*,[1] indicating that the evaluator's perception of "collective blame" in the unilateral treatment erodes each group member's "individual blame" for the collective decision to not disclose, compared to the individual treatment in which there is no collective.

With respect to the disclosure choices made by groups, we find that the set of evidence realizations that groups choose to conceal vary significantly between treatments. The comparison of group's communication choices across treatments mirrors the ordering of evaluator skepticism: our data shows that groups reveal evidence realizations that are less positive about the value to group members who have more power. Compared to the empirical best response in each treatment, we see that groups' communication decisions match in around 80% of our observations; and that groups err towards under-communication (too

[1]The term was initially introduced by Darley and Latané (1968) in the psychology literature to denote the idea that individuals may change their own behavior when acting in a group, perhaps towards more "antisocial behavior," due to not perceiving themselves as fully to blame for the group's ultimate decisions. In an experimental economics context, Behnk et al. (2022) documents not only that individuals perceive themselves as less at fault for antisocial behavior done by their group, compared to actions taken individually, but also that outsiders to the group (similarly to evaluators in our group treatments) perceive each individual group member as less responsible for their group's action.

little disclosure) in the individual, leader, and consensus treatments, and towards over-communication in the unilateral treatment.

While our observations about the aggregate behavior of group members and evaluators strongly support our theoretical hypotheses on the relation between the power structure and group communication, the analysis of individual subjects in the lab paints a more nuanced picture. Across all treatments, a Bayesian evaluator's beliefs about each group member's value should be updated negatively after seeing that the group chose not to disclose (to a larger or smaller extent, depending on the treatment). In contrast, we find that a significant portion of evaluators in our group treatments tend to update positively after seeing no disclosure, and to update asymmetrically about the value to group members A and B in symmetric treatments (unilateral and consensus). We show that these deviations from Bayesian updating are related to the complexity of the belief-updating task that follows no disclosure, which we argue differs across treatments.

The individual analysis of subjects who played group member roles shows that, in line with our theory, most group members use threshold disclosure recommendation strategies, suggesting that the group member recommends to disclose the realized evidence if and only if their own value is sufficiently large. However, we find that the disclosure thresholds used by these subjects do not consistently vary across treatments in the way described by our theory. Further, individuals' thresholds deviate from their perceived best responses in patterns that are consistent with behavioral biases such as lying aversion (where we interpret not disclosing as akin to lying) and aversion to making pivotal decisions in groups.

The paper is organized as follows: our theory of group disclosure, based on Onuchic and Ramos (2025), is presented in section 2; section 3 lays out the experimental design and main hypotheses; experimental results regarding the behavior of evaluators are presented in section 4, and those regarding the disclosure behavior of groups are in 5; finally, section 6 documents the dynamics of play over rounds.

1.1 Related Literature

The theoretical portion of our paper contributes to the large literature on disclosure games. For a review on the connection between our stated results and previous theoretical literature, please refer to Onuchic and Ramos (2025).

The main focus of this paper is experimental. To the best of our knowledge, this is the first paper to experimentally study a group communication game, in which a group of senders with distinct interests collectively communicate with a receiver through the disclosure of verifiable information. This focus mainly connects our work to the experimental literature on disclosure games (and communication experiments more broadly) and to the experimental literature on games played by groups or games of collective decision. We comment on each of these connections below.

Communication Experiments. There is a sizable literature that experimentally test predictions of theoretical models of disclosure in the tradition of Grossman (1981) and Milgrom (1981). Some experiments, including those in Forsythe et al. (1989), Li and Schipper (2020), Jin et al. (2021), and Deversi et al. (2021), consider environments in which theoretical analysis predicts that skepticism on the part of the receiver leads to the "unravelling" of equilibria in which not all information is disclosed by the sender. These experiments find evidence of unravelling to different degrees, and scrutinize the mechanisms behind the discrepancy between theoretical and experimental findings. For instance, Li and Schipper (2020) use an iterated admissibility criterion to generate theoretical predictions for finite levels of reasoning about rationality. And Jin et al. (2021) show that the degree of skepticism exhibited in receiver beliefs after no disclosure approach theoretical predictions in treatments that permit more learning (in later rounds, or in their feedback treatment).

There is also a set of experimental papers that consider individual disclosure environments in which full disclosure is not a necessary prediction. For example, King and Wallin (1991) runs an experiment akin to the model in Dye (1985), according to which the sender with some probability does not have access to verifiable information; Dickhaut et al. (2003) consider the possibility that disclosure is costly; Hagenbach and Perez-Richet (2018) study an environment in which the sender does not have monotonic preferences; and Hagenbach and Saucet (2025) run an experiment in which receivers have preferences over the information they learn, as in the literature on motivated beliefs.

In our group disclosure setting, theoretical predictions on the degree of equilibrium skepticism about the value to each group member, reflected in the receiver's beliefs after seeing no disclosure, depends on the deliberation procedure used by the group to make collective disclosure decisions.[2] Our main experimental hypotheses thus regard not the

[2] Hagenbach and Saucet (2025) similarly propose a theoretical environment in which the degree of predicted skepticism varies across the different treatments considered; like in our case, their predicted ordering on skepticism is confirmed experimentally. In their setting, treatments vary whether the state that senders communicate about is ego-relevant or neutral for receivers, and whether skeptical beliefs are aligned or not with what Receivers prefer believing. Compared to neutral settings, they find that the receiver's skepticism is significantly lower when it is self-threatening, and not enhanced when it is self-serving. See also Farina et al. (2024), who vary the evidence structure across treatments and also test a rich set of comparative static

presence or absence of "unravelling," but rather differences in skepticism across treatments with different deliberation procedures. We find that, across treatments, evaluators' beliefs are not as skeptical as equilibrium predictions; however, the relation between skepticism and the structure of power in the group matches that in our theory.

There are a few experimental papers that study communication with multiple senders. For example, Lai et al. (2015) and Vespa and Wilson (2016) consider experiments in which multiple senders communicate with a single receiver via cheap talk, and Sheth (2021) investigates a disclosure setting with competing senders. Our setup differs from the former two papers because the group communicates through a different protocol (information disclosure rather than cheap talk), and from all three papers in that we consider communication by a group as a single coordinated entity, rather than independent communication from multiple competing sources.

Our paper also relates to Farina and Leccese (2025), which considers an experiment in which a single sender makes disclosure decisions regarding a multi-dimensional state. Similarly, in our setting, the state is multi-dimensional in that it reflects the value to each of the group members. Farina and Leccese (2025) document how the presence of multiple dimensions of disclosure can affect the difficulty of the contingent thinking inherent to the updating task performed by receivers. In our context, we argue (in section 4.4) that the complexity of the updating task varies across treatments depending on how group members' decisions are aggregated into the group's disclosure decision, and relate this to receivers' observed updating mistakes.

Experiments Played by Groups. There are two main types of experiments that consider games played by groups of subjects. The first set of papers includes studies that compare group and individual behavior in various games and individual decision problems where all members of the group share the same payoffs. These studies typically find that team play more closely resembles the standard predictions of game theory. Charness and Sutter (2012) and Kugler et al. (2012) are surveys that cover that experimental literature.

The second type of group experiments consider games of collective decision in which group members have different and private information about a state that is relevant to determine the group's ideal action. This literature, surveyed by Martinelli and Palfrey (2018), includes experiments on voting games, on information aggregation in committees, and on legislative bargaining. Our paper resembles some of this work — for example, Goeree and

predictions on receivers' beliefs after communication.

Yariv (2011) — in that our different treatments vary the institutions by which decisions are reached by the group. However, unlike Goeree and Yariv (2011), our main experimental hypotheses regard how these institutions affect the formation of beliefs by a receiver with whom the group of senders communicates.

More generally, our paper distinguishes itself from both these strands of the literature in that we consider a game of group communication. In our game, group members have distinct preferences over communication decisions (unlike in the first literature strand), and have access to all the information relevant to make their own optimal disclosure recommendation (unlike the second strand of the literature, in which information aggregation plays a big role).

Our work is particularly related to an experimental literature that studies the attribution of blame for collective decisions, and how this attribution depends on the procedure through which the group makes decisions. Most closely related is Behnk et al. (2022), who study a sender-receiver game in which a sender (or a group of two senders) can choose to tell the truth (a prosocial action) or profitably lie (an anti-social action) to a receiver. They show that receivers attribute less blame for antisocial actions to senders who make that decision as a group.[3] In our group disclosure context, we consider how the structure of group decision making affects receivers' interpretation of group communication, and specifically of the absence of disclosure.

2 Model and Theoretical Results

There is a group, composed of two group members $i = A, B$. The group draws an observable outcome ω, described by its value to each group member i, $\omega_i \in \{0, 1, ..., 10\}$. The outcome values ω_A and ω_B are independently drawn, each distributed according to the uniform distribution over the set $\{0, 1, ..., 10\}$. The group makes a single decision, of whether to disclose the realized outcome, thereby revealing it to some outside third-party, or to conceal it. Before providing further details on the group's decision making, we discuss possible interpretations of this simple environment.

Interpretation. A possible scenario is one of a team in a tech company that is assigned the project of designing a new tool. The team is made up of various professionals, including an engineer and a marketer. After working on this project for a while, the team produces an

[3] See also, for example, Bartling et al. (2015), Bartling et al. (2014), and Engl (2022), who consider the relation between group members' pivotality and responsibility attribution in other contexts.

initial prototype (the observable outcome), which is very well done in terms of its technical aspects, but poorly "packaged." At this point, the team is approached by a higher-up manager (the outside third-party) who asks them to report on their progress. The team must decide whether to reveal the prototype to the manager or not to do so (maybe claiming that they need more time, or that no prototype has yet been produced). If the team reveals the prototype, the manager will be positively impressed by the engineer, who contributed the technical aspects, but negatively impressed by the marketer, who is responsible for the below-par packaging. In this case, even though the team produced a single observed outcome, its disclosure yields a different value to each team member — a high $\omega_{engineer}$ and a low $\omega_{marketer}$.

Alternatively, think of a meeting of the editorial board of a magazine, where various editors need to decide whether to include an inflammatory piece (the observable outcome) in the upcoming publication (in which case the outcome will be seen by the outside third-party, the potential readers of the magazine). The editors have different views on the ideal editorial leaning for the magazine, maybe relating to their own political views, and therefore assign different value to the inclusion of this piece in the magazine's new issue. Again, even though there is a single observable outcome in hand, the publishable piece, its publication yields a different value to each member of the editorial board — so that $\omega_{editor A} \neq \omega_{editor B}$.

Group Decision-Making. Each group member i sees the realization of their own outcome value ω_i before the group decides on the outcome's disclosure.[4] To reach a group decision, each individual makes a disclosure recommendation $x_i(\omega_i) \in \{0, 1\}$ — $x_i = 1$ indicates that i favors that the group disclose the outcome and $x_1 = 0$ that i favors that the outcome be concealed by the group.[5] Individual recommendations are then aggregated into a group disclosure decision according to some deliberation procedure $D : \{0, 1\}^2 \to [0, 1]$, so that

$$d(\omega) = D\left(x_A(\omega_A), x_B(\omega_B)\right)$$

is the probability that the group discloses outcome ω to the outside third-party.

[4]In our model, each group member's possible payoffs are entirely determined by their own outcome value ω_i and the observer's equilibrium "beliefs of no disclosure." This implies that there is no additional information relevant to group member i that is conveyed by group member j's outcome value; and our assumption that each group member sees only their own value is of very little consequence.

[5]Note that the group's decision is binary, to either reveal fully the realized outcome or to entirely conceal it. This structure is referred to as "simple evidence" in the strategic disclosure literature.

The aggregator function D provides a reduced-form description of the "deliberation procedure" used by the team to reach a collective decision: it describes the disclosure decision that is reached after each possible combination of individual disclosure recommendations made by the group members. We are broadly interested in aggregation procedures that satisfy two properties: *unanimity*, whereby the group follows recommendations that are unanimous across the two group members, and *monotonicity*, so that the probability of disclosure increases when more group members recommend disclosure.[6] In our experimental setting, we consider three deliberation procedures that satisfy these conditions (these procedures correspond to the different treatments in our experiment):

1. The *unilateral procedure*, according to which each group member can unilaterally enforce the disclosure of the outcome. Formally, $D(1,1) = 1$, $D(0,0) = 0$, and $D(1,0) = D(0,1) = 1$.

2. The *consensus procedure*, according to which disclosure must be a consensual decision among group members, or in other words each group member can veto disclosure. Formally, $D(1,1) = 1$, $D(0,0) = 0$, and $D(1,0) = D(0,1) = 0$.

3. The *leader procedure*, according to which group member A is most likely a dictator. Formally, $D(1,1) = 1$, $D(0,0) = 0$, $D(1,0) = 1 - \epsilon$, and $D(0,1) = \epsilon$, for some small $\epsilon > 0$.

Payoffs. If the group chooses to disclose the outcome ω, the outside third-party perfectly observes it, and each group member i receives a payoff equal to their own respective value of the outcome, ω_i. If instead the group chooses to not disclose the outcome, then the outside observer does not see the outcome, but sees that the group chose "no disclosure." In that case, the observer forms a Bayesian posterior belief about the value of ω_i for each group member i, given by $\omega_i^{ND} = \mathbb{E}(\omega_i |\text{ no disclosure})$. Group member i's payoff is then equal to the observer's posterior belief about their own outcome.

Equilibrium. Our equilibrium notion is sequential equilibrium in undominated strategies. Given a deliberation procedure D, individual disclosure strategies x_i for $i \in \{A, B\}$, the group's disclosure decision d, and no-disclosure posteriors ω_i^{ND} for $i \in \{A, B\}$ constitute an equilibrium if

[6] These are the general assumptions made on deliberation procedures by Onuchic and Ramos (2025), where the group may be made up of more than two individuals. With two individuals, monotonicity is implied by unanimity.

1. Group members use undominated recommendation strategies, given ω^{ND}:

$$\omega_i > \omega_i^{ND} \Rightarrow x_i(\omega) = 1 \text{ and } \omega_i < \omega_i^{ND} \Rightarrow x_i(\omega) = 0.$$

2. Given x, no disclosure posteriors ω_A^{ND} and ω_B^{ND} satisfy sequential consistency.

2.1 Equilibrium Group Disclosure

For any deliberation procedure, equilibrium behavior is pinned down by the observer's beliefs of no disclosure. Given beliefs ω_A^{ND} and ω_B^{ND}, we can back out the equilibrium strategy for both group members: each group member recommends disclosure if and only if their own drawn outcome value is larger than the observer's no disclosure belief about their value. In turn, recommendation strategies are aggregated through the deliberation procedure, generating the group's disclosure strategy. Because no disclosure beliefs provide full descriptions of equilibrium behavior, these are essential objects in our analysis. Their equilibrium values are characterized in Theorem 1.

Theorem 1.

1. *If the group uses the* ***unilateral*** *deliberation procedure, the unique equilibrium outcome is full disclosure, with*

$$\omega_i^{ND} = 0 \textit{ for at least one of } i \in \{A, B\}. \tag{1}$$

2. *If the group uses the* ***consensus*** *deliberation procedure, some outcome realizations are not disclosed in the unique equilibrium. No disclosure beliefs are:*

$$\omega_A^{ND} = \omega_B^{ND} = 3.639. \tag{2}$$

3. *If the group uses the* ***leader*** *deliberation procedure — in which $D(1,0) = 1 - \epsilon$ and $D(0,1) = \epsilon$ for some small $\epsilon > 0$ — some outcome realizations are not disclosed in the unique equilibrium. No disclosure beliefs satisfy:*

$$\lim_{\epsilon \to 0} \omega_A^{ND} = 0 \textit{ and } \lim_{\epsilon \to 0} \omega_B^{ND} = 5. \tag{3}$$

To understand this result, let's consider the construction of an equilibrium under each possible deliberation procedure. To that end, first conjecture an equilibrium in which the

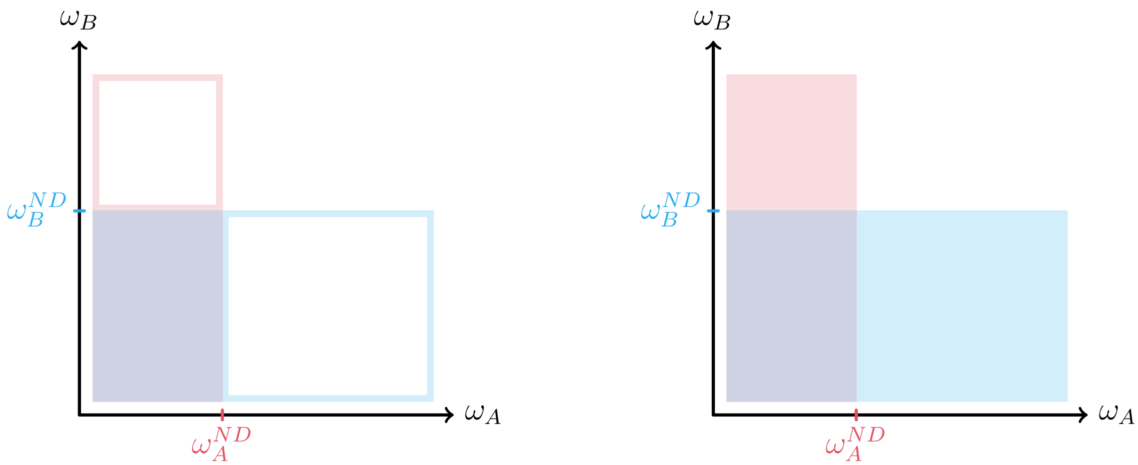

Figure 1: No disclosure regions under unilateral (left) and consensus (right) procedures.

observer's beliefs of no disclosure are $\omega_A^{ND} > 0$ and $\omega_B^{ND} > 0$, that is, in which upon seeing no disclosure, the observer thinks there is some probability that each group member drew an outcome value larger than the lowest possible value, 0. Given these conjectures, each group member recommends to disclose after seeing a realization of their own value that is larger than the observer's no disclosure belief about their value, and to not disclose otherwise. These recommendation strategies are depicted in pink and blue (for group members A and B, respectively) in the two panels of Figure 1. These individual recommendations must then be aggregated through the group's deliberation procedure to attain the group's disclosure strategy.

Under the unilateral procedure, disclosure ensues after at least one group member recommends it, and therefore the "no disclosure region" induced by the individual recommendations is the purple rectangle highlighted in the left-hand panel of Figure 1. Given this group strategy, after seeing no disclosure, the observer understands that the realized values must belong to a point in the purple rectangle; but note that every realization in that region has a value of ω_A lower than the initially conjectured ω_A^{ND} and a value of ω_B lower than the initially conjectured ω_B^{ND}. This implies that the initial conjectures cannot be consistent with Bayesian updating, and therefore do not correspond to an equilibrium. This reasoning applies to any initial conjectures with $\omega_A^{ND} > 0$ and $\omega_B^{ND} > 0$, and we therefore conclude that the only possible equilibrium outcome is full disclosure, supported by beliefs of no disclosure of $\omega_i^{ND} = 0$ for at least one of $i \in \{A, B\}$. This reasoning is exactly the "unravelling logic" from individual disclosure games, where in our setting unravelling happens simultaneously on the dimension of group member A's value and that of B's value.

Consider instead the aggregation of individual recommendations through the consensus procedure. The implied "no disclosure region" is the entire L-shaped region highlighted on the right-hand panel of Figure 1. Given this group disclosure strategy, after seeing no disclosure, the observer understands that group member A might have drawn a very good value, but had its disclosure vetoed by group member B, or vice-versa. As a consequence, their Bayesian-consistent belief of no disclosure might be larger or equal to the initial conjecture ω_i^{ND} for each group member $i \in \{A, B\}$. Our result states that there is exactly one pair of conjectured no disclosure beliefs ($\omega_A^{ND} = \omega_B^{ND} = 3.639$) such that the observer's Bayes-consistent beliefs of no disclosure exactly correspond to the initial conjecture. Under the consensus procedure, the unravelling logic is broken because blame for no disclosure is dissipated among the group members: after seeing no disclosure, the observer understands that at least one group member must have recommended that action, but cannot entirely attribute the group's decision to one of the two individuals.

Finally, under the leader treatment, we show that (for small values of $\epsilon > 0$) there exists a unique equilibrium without full disclosure, supported by beliefs $\omega_A^{ND} > 0$ and $\omega_B^{ND} > 0$. Moreover, as the procedure approaches a dictator rule ($\epsilon \to 0$), we have $\lim_{\epsilon \to 0} \omega_A^{ND} = 0$ and $\lim_{\epsilon \to 0} \omega_B^{ND} = 5$. Under both the consensus and leader procedures, full disclosure is also supported as an outcome in a weak Perfect Bayesian Equilibrium (with undominated strategies). Theorem 3 in Onuchic and Ramos (2025) shows that this outcome is ruled out by our requirement that off-path beliefs satisfy sequential consistency.

Disclosure Power and Skepticism. In each equilibrium described in Theorem 1, the observer's beliefs of no disclosure, ω_A^{ND} and ω_B^{ND} expresses the degree to which the observer is skeptical about each group member's value after seeing no disclosure. Specifically, from an ex-ante perspective, with no additional information, the observer understands that the expected value to a given group member is 5 (remember that values are distributed uniformly between 0 and 10). After the observer sees that the group chose to not disclose, they update this prior belief, accounting for the fact that non-disclosure was a strategic decision. Broadly speaking, the observer becomes more skeptical about each group member's value in that case, updating their average belief of each individual's value to something weakly below 5. To describe the degree of skepticism, we introduce the measure

$$\sigma_i = \frac{5 - \omega_i^{ND}}{5}, \text{ for each } i \in \{A, B\},$$

which reflects how much more skeptical the observer is about i's value after seeing no disclosure than if they had seen no information at all. Using this definition, we can restate conditions (1)-(3) in Theorem 1 respectively as

- Unilateral procedure: $\sigma_i = 1$ for at least one of $i \in \{A, B\}$.
- Consensus procedure: $\sigma_A = \sigma_B = 0.27$.
- Leader procedure: $\lim_{\epsilon \to 0} \sigma_A = 1$ and $\lim_{\epsilon \to 0} \sigma_B = 0$.

These statements comparing skepticism levels across different procedures serve as basis for our main experimental hypotheses.

Skepticism as Experimental Hypotheses. Previous experimental work on disclosure games (where disclosure decisions are made by a single individual) have tested the hypothesis that the unique equilibrium outcome is the "unravelled" full disclosure, which corresponds to an observer who is maximally skeptical about the individual's value after seeing no disclosure. Typically, these papers have found that full disclosure typically does not arise in the lab (see, for example, Jin et al. (2021)), and that information receivers in the lab form beliefs that are insufficiently skeptical about nondisclosed information, and information senders react to these not-so-skeptical beliefs by concealing some unfavorable outcome realizations.

Motivated by these experimental findings, we posit our main hypotheses not in terms of the cardinal predictions of skepticism levels implied by Theorem 1, but rather in terms of ordinal comparisons of skepticism across treatments. Our interpretation of these orderings is that they establish a relationship between a group member's power to enforce disclosure as the group's decision and the observer's skepticism about that group member's value after seeing no disclosure. In particular, the observer is more skeptical about a group member when they have more power to enforce disclosure: full skepticism in the unilateral procedure and about the leader in the leader procedure (in which there is full power to enforce disclosure), intermediate skeptcism in the consensus procedure (in which each individual has some power to enforce disclosure), and no skepticism about the non-leader in the leader procedure (who has no power to enforce disclosure).

A More General Principle. The relation between an individual's power to enforce disclosure and the observer's no-disclosure skepticism about that individual's value is more general than the comparisons established by Theorem 1. The environment in Onuchic and Ramos (2025) has two or more group members, possibly correlated values across group members, and a large set of possible deliberation procedures. In this more general environment, their main comparative statics result formalizes the idea that, when comparing equilibria under two deliberation procedures, if group member i is *relatively more powerful* under procedure D than under procedure D', then the observer is more skeptical about i's value upon seeing no disclosure under procedure D than under procedure D'. Our experimental design considers only three deliberation procedures, but we interpret our results as speaking to this more broad relation between power and skepticism.

3 Experimental Design and Hypotheses

3.1 Basic Experimental Design

We first describe one round of the basic game played in the lab, which is designed to match the environment described in the group disclosure model. The game involves 3 players: group member A, group member B, and an evaluator. The 3 players constitute a unit and play a game consisting of 4 stages: information, reporting, guessing, and feedback.

In the **information stage**, the computer program randomly and uniformly chooses one card from each of two decks, deck A and deck B. Each deck has 11 cards, with numbers $0, 1, ..., 9, 10$. The pair of cards constitutes the group's hand; the value on card A denotes the value of the group's hand for group member A, and the value on card B denotes the value of the group's hand for group member B. At this stage, each group member sees the card representing their respective value, but not the card referring to their partner's value.[7] Additionally, neither card is seen by the evaluator.

The next stage is the **reporting stage**, in which the group chooses whether to disclose the group hand to the evaluator. Towards reaching a decision, each group member makes a recommendation, by clicking one of two buttons: "report" or "not report." The two group members' recommendations are then aggregated into a group disclosure decision through a deliberation procedure.

The deliberation procedure is the object we vary in the different treatments in our experiment. We consider 3 deliberation procedures, following the variations introduced in

[7] As mentioned in the previous section, our theoretical results are not affected by whether group members see each other's card values or not. When bringing our environment to the lab setting, we decided to show each group member only their own value so as to prevent any distortions in group members' behavior caused by social preferences regarding their partners' payoffs, which are not accounted for in our theory.

our theory section: *consensus* procedure, *unilateral* procedure, and *leader* procedure.

- In the consensus treatment, the group hand is reported to the evaluator if and only if both group members recommend reporting it. That is, if both group members recommend reporting, both cards in the group hand are revealed to the evaluator. Otherwise, neither card is revealed to the evaluator.

- In unilateral treatment, if group member A, group member B, or both group members recommend reporting, then both cards in the group hand are revealed to the evaluator. If instead neither group member recommends reporting, the group hand is not revealed to the evaluator.

- In the leader treatment, group member A is the "leader," and the group's reporting decision follows A's recommendation with high probability. Specifically, with 99% chance, group member A's recommendation is followed by the group and with 1% chance, group member B's recommendation is followed by the group.[8]

After the group makes their reporting decision, the evaluator is informed whether the group hand was reported or not. If the group hand is reported, the evaluator sees both cards in the group hand. If the group hand is not reported, the evaluator does not see the group hand, and is alerted of the fact that the group chose not to report the hand.

If the group chooses not to report its hand, we remind the evaluator of the procedure used by the group to reach that decision. For example, in the unilateral treatment, the evaluator sees a message saying "The group hand in this round was NOT reported. That is, both group members recommended not to report the group hand."[9] We choose to remind the evaluator of the procedure so as to replicate the theoretical environment in which the observer knows the group's deliberation procedure, but does not know the recommendations made by each of the group members.

After the evaluator sees the reported/not reported group hand, the game moves to the **guessing stage**, in which the evaluator is asked to make two guesses: to guess group member A's value, and to guess group member B's value. Each of the evaluator's guesses is a

[8]Our theory considers the limit as the probability ϵ that the group follows B's recommendation goes to 0. In the lab, we set $\epsilon = 1\%$. Our main consideration was to maintain group member B's incentives to play truthfully, by keeping a non-zero probability that their recommendation is pivotal.

[9]In the consensus treatment, the message reads "The group hand in this round was NOT reported. That is, group member A, group member B, or both group members recommended not to report the group hand." Finally, in the leader treatment, the message reads "The group hand in this round was NOT reported. That almost certainly means that group member A recommended not to report the group hand."

number between 0 and 10. We allow for guess increases of 0.5. The final stage in a round is the **feedback stage**, in which every participant is shown a screen containing the group hand, whether the group hand was reported or not, and the pair of evaluator guesses made in the current round.[10]

Incentive Implementation. The evaluator is paid for the accuracy of one of the guesses, getting the points earned from either guess A or from guess B, with equal probability. Specifically, the evaluator earns either $110 - 4.4(|\text{Value A} - \text{Guess A}|)^{1.4}$ points or $110 - 4.4(|\text{Value B} - \text{Guess B}|)^{1.4}$ points. As for group members, their payment is increasing in the evaluator's guess of their own value: group member A earns $110 + 4.4(\text{Guess A} - 10)^{1.4}$ points, and group member B earns $110 + 4.4(\text{Guess B} - 10)^{1.4}$ points. The payment scheme is presented to the subjects using a table describing the amount of points to be received for each possible combination of the drawn value i and the evaluator's guess of i's value.[11]

3.2 Questionnaire

In addition to the main experimental setting described above, we ask that the participants complete a short questionnaire. The same questionnaire is presented to every participant, regardless of the role they played during the main portion of the experiment. Participants' answers to parts 1 and 2 of the questionnaire are incentivized.[12]

In the first part of the questionnaire, we elicit each participant's "belief of no disclosure," irrespective of their role. To that end, we ask: " Suppose you are an evaluator, and the group hand is not reported to you by the group. What would be your guess A and guess B for group member A's value and group member B's value, respectively?" In the second part, looking to elicit subjects' reporting strategies, we ask: "Suppose you are group member A, and a group hand is drawn in which value A is x. Would you recommend to report that group hand?" We ask this question 11 times, one for each value of $x \in \{0, 1, ..., 10\}$. In

[10]We provide full feedback, following Jin et al. (2021) feedback treatment, which they found to best replicate the theoretical predictions in their disclosure setting.

[11]The payoff functions are similar to those in Jin et al. (2021) and Deversi et al. (2021), with an adjustment of a constant to ensure all possible payoffs are positive. We also follow their methods ijn presenting the payment scheme using a table.

[12]We communicate to participants at the beginning of the questionnaire, saying: "Please answer the following questions. You can earn additional money with your answers. Your responses will be compared to a randomly chosen participant's behavior in this experiment. To assess your answers, we randomly choose a participant from the main part of this study, and compare your answer to their behavior in one round. Specifically, we will select one of your answers below, and you will receive a \$3 bonus for correctly predicting the answer of the randomly chosen participant." This method of incentivizing subjects is similar to the one introduced by Ba et al. (2022).

the leader treatment, which is asymmetric, we also ask 11 analogous questions, regarding how the subject would report if they were group member B.[13]

We complete the questionnaire with a standard set of questions regarding the participants demographics. These include the participant's major, their gender, their GPA, and whether they have taken a game theory class.

3.3 Individual Disclosure Treatment

As a benchmark, we also run a version of our experiment in which a single individual (rather than a group) makes reporting decisions. Theoretically, individual behavior and observer skepticism in this individual disclosure treatment should be akin to the behavior and skepticism observed in the unilateral treatment and that of the leader in the leader treatment. By making that comparison, we aim to assess whether there is any effect inherent to the fact that our game is played by a group rather than by an individual, even if the strategic interactions are unchanged.

The individual treatment is made up of the same stages as described in section 3.1, with the following changes. First, in the information stage, only one card is drawn (rather than one per group member), and the single individual sees the drawn card. Second, at the reporting stage, the single individual makes a reporting recommendation, and their recommendation is followed. Third, at the guessing stage, the evaluator makes a single guess about the individual's value (rather than two guesses, one per group member).

3.4 Implementation

Our experiment was preregistered using the AEA RCT Registry, under ID 0013276. IRB approval was granted by Indiana University's Institutional Review Board (ID 19303). The experiment was conducted at Interdisciplinary Experimental Laboratory (IELab) at Indiana University (IU) during the spring of 2024 and academic year 2025, using software z-Tree (Fischbacher (2007)). Subjects were recruited from the general student population via ORSEE recruitment system (Greiner 2015). We conducted 8 sessions for each treatment (unilateral, consensus, leader, and individual treatments). Most sessions had 5 units (one unit is made up of 3 subjects for the unilateral, consensus, and leader treatments, and of 2 subjects in the individual treatment); one session of the individual treatment had 7 units. In total, there were 444 subjects.

[13] This question elicits participants' behavior "as if" they were group members using the strategy method, in contrast with the elicitation throughout the rounds. We use this alternative method as robustness to our tests using main data.

The instructions were read aloud, with paper copies distributed to all subjects (see Appendix E for instructions). After reading the instructions, the subjects first engaged in 2 practice rounds before moving onto 30 actual rounds. The experiment lasted around 60 minutes, and subjects earned an average payoff of \$20, which included a \$8 show-up fee. In the experiment, the payoffs in the game were denominated in points. Each point was converted to US dollars at the rate of 10 points to \$1.

We implement a between subjects design for the consensus, unilateral, leader, and individual treatments. In each treatment, each subject is assigned one of the three roles. They keep their role for 30 rounds, but units are re-matched every round to avoid reputation building issues or reciprocity between group members. For instance, group member A stays as group member A in the next round, but is randomly re-matched with another pair of participants playing the roles of group member B and evaluator.

3.5 Hypotheses Regarding Evaluators

These hypotheses follow directly from Theorem 1. In our empirical environment, we measure the evaluator's skepticism in each round in which the group's decision is to *not report* their group hand. We measure the evaluator's skepticism about group member i as i-skepticism $= \sigma_i = (5 - \text{Guess } i)/5$.[14] We refer to the average of A-skepticism and B-skepticism as aggregate skepticism.

Hypothesis 1. *Aggregate skepticism in the consensus treatment is smaller than aggregate skepticism in the unilateral treatment.*

Hypothesis 2. *Aggregate skepticism in the consensus treatment is smaller than A-skepticism in the leader treatment, in which group member A is the leader.*

Hypothesis 3. *Aggregate skepticism in the consensus treatment is larger than B-skepticism in the leader treatment, in which group member B is the non-leader.*

3.6 Hypotheses Regarding Group Members

For group members, individually rational disclosure recommendation strategies are threshold strategies, according to which the group member recommends that the group not report

[14] Deversi et al. (2021) and Hagenbach and Saucet (2025) define skepticism differently from us, as $\hat{\sigma}_i = 1 - \frac{\text{Guess } i - \min(\Omega)}{\max(\Omega) - \min(\Omega)}$. Their measure compares the observer's guess to the support of possible guesses that can be made. Our measure compares the observer's guess to the unconditional mean of 5 (which is the Bayesian mean the observer should hold, were they not to see any additional signal, such as the one conveyed by "no disclosure"). The two measures are linear transformations of each other, and our hypotheses (and validation of hypotheses) would remain unchanged under the alternative measure $\hat{\sigma}_i$.

when their own value is low and that the group report otherwise. The threshold in i's recommendation strategy should coincide with the evaluator's guess of i's value in case the group's hand is not reported. Given this connection between "no disclosure beliefs" and equilibrium recommendation strategies, Hypotheses 1-3 about skepticism imply three parallel hypotheses regarding groups' disclosure behavior (which is determined by the aggregation of recommendations through the procedure in each treatment). We state these hypotheses as follows:

Hypothesis 4. *The distribution of group member i values that are concealed in the consensus treatment dominates that in the unilateral treatment in the sense of first order stochastic dominance.*

Hypothesis 5. *The distribution of group member A values that are concealed in the consensus treatment dominates that in the leader treatment — in which group member A is the leader — in the sense of first order stochastic dominance.*

Hypothesis 6. *The distribution of group member B values that are concealed in the consensus treatment is dominated by that in the leader treatment — in which group member B is the non-leader — in the sense of first order stochastic dominance.*

4 Results: Interpreting Group Communication

4.1 Skepticism across Treatments, in Theory and in Practice

Treatment	Average	SE	N
Individual	0.355	0.015	573
Unilateral	0.240	0.016	582
Leader	0.324	0.018	533
Non-Leader	-0.004	0.016	533
Consensus	0.131	0.009	1804

Table 1: Mean skepticism and standard error by treatment

Our first set of results concerns the evaluator's skepticism about group members' values in the different treatments. Table 1 provides skepticism across all four treatments. For the unilateral and consensus treatments, which are symmetric, we display the aggregate skepticism numbers, whereas the leader treatment is split into "Leader" and "Non-Leader,"

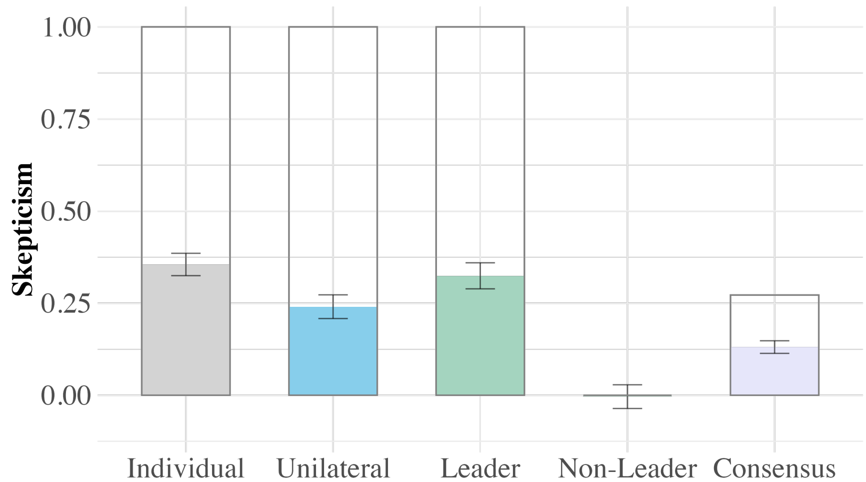

Figure 2: Skepticism in theory and in the lab.

referring to skepticism about group member A's value and group member B's value, respectively. These numbers are averages across all participants and all rounds played in each of these treatments. The number of observations differs across treatments, because skepticism is calculated only in instances in which no disclosure occurred. The same information is displayed in Figure 2, which also highlights (in black outline bars) the skepticism values predicted by our theory: skepticism equal to 1 for both group members in the unilateral treatment, as well as for the individual treatment, and for group member A (the leader) in the leader treatment; skepticism equal to 0.27 for both group members in the consensus treatment; and equal to 0 for group member B (the non-leader) in the leader treatment.

In each treatment, the skepticism values found in practice are smaller than those predicted by our theory. These findings add to a wealth of observations in individual disclosure experiments — for example, Li and Schipper (2020), Jin et al. (2021), and Hagenbach and Saucet (2025) — describing that observers' "no disclosure beliefs" are typically less skeptical than the maximal skepticism/ full unravelling predicted theoretically. Our findings expand on that statement by documenting that "too little skepticism" relative to theoretical predictions arises also in group disclosure contexts, even in contexts where the theory does not predict full unravelling.

In line with our hypotheses 1-3, skepticism (about each of the group members) is smaller in the consensus treatment than in the unilateral treatment, skepticism in the consensus treatment is smaller than skepticism about group member A (the leader) in the leader treatment, and skepticism in the consensus treatment is larger than skepticism about group member B's value (the non-leader) in the leader treatment. Tests are reported in Table 2:[15]

Comparison	Avg. Skepticism	Avg. Skepticism	Adj. p-value
Consensus vs. Unilateral	0.131	0.240	< 0.01
Consensus vs. Leader	0.131	0.324	< 0.01
Consensus vs. Non-Leader	0.131	-0.004	< 0.01
Individual vs. Unilateral	0.355	0.240	< 0.01
Individual vs. Leader	0.355	0.324	1.000
Unilateral vs. Leader	0.240	0.324	< 0.01

Note: Pairwise Wilcoxon rank-sum (Mann-Whitney U) tests with Bonferroni adjustment.

Table 2: Summary of hypotheses testing results.

These tests establish the empirical validity of the main mechanism proposed in our model of group communication: the power structure used to make communication decisions in a group (as given by the deliberation procedure) significantly determines the interpretation of the group's equilibrium messages. In our disclosure setting, "no disclosure" is the only message to be interpreted by the evaluator, and we see that it is interpreted as a less favorable indication of an individual's value whenever that individual has more power to enforce disclosure as the group's decision.

Table 2 also displays comparisons between skepticism values in the unilateral and leader treatments to that in the individual disclosure baseline. According to our theory, the observer should be "maximally skeptical" after seeing no disclosure in each of these three circumstances. Indeed, we observe that A-skepticism in the leader treatment is not significantly different from skepticism in the individual treatment. Contrastingly, we find that skepticism in the unilateral treatment is significantly lower than that in the individual treatment (and than skepticism about the leader in the leader treatment).[16] This observation highlights an empirical distinction (even when there are no theoretical differences) between decisions that are truly collective — as in the unilateral treatment, where the deliberation

[15]In Table 8 in Appendix B, we also verify the significance of the first three hypotheses (our main skepticism hypotheses) using standard errors clustered at the session level.

[16]The difference between skepticism in the unilateral and individual treatments, and between unilateral skepticism and skepticism about the leader in the leader treatment do not remain significant in all of our robustness checks, which are presented in Appendix B.

procedure is symmetric across both partners — and decisions that can be fully (or almost fully) attributable to one individual, as in the individual and leader treatments.

In the unilateral treatment, upon observing "no disclosure," the observer's understanding should be that "both group member A and group member B recommended no disclosure," and therefore both group members should be equally and fully held to blame for that decision. In practice, in this group treatment, each individual is not fully held responsible for the group's decision not to disclose: collective decision making leads to an erosion of "individual blame," perhaps in favor of "social blame" attributed to the group as the true decision-making unit.

This observation connects to the psychology literature on *diffusion of responsibility* — following Darley and Latané (1968) — and to more recent documentation of this phenomenon in experimental economics, for example by Dana et al. (2007), Choo et al. (2019), Brütt et al. (2020), and Behnk et al. (2022). The term was initially introduced to denote the idea that individuals may change their own behavior when acting in a group, perhaps towards more "antisocial behavior," due to not perceiving themselves as fully to blame for the group's ultimate decisions. In an experimental context, Behnk et al. (2022) documents not only that individuals perceive themselves as less at fault for antisocial behavior done by their group, compared to actions taken individually, but also that outsiders to the group (similarly to evaluators in our group treatments) perceive each individual group member as less responsible for their group's action.[17]

Another way of visualizing the differences in skepticism across different treatments is through the comparison of the respective empirical CDFs of skepticism. These are displayed in Figure 9 in Appendix B. Each of the significant differences displayed in Table 2 correspond to ECDFs that are ordered in the first order stochastic sense. We also note that the ECDFs of skepticism in the individual treatment and of A-skepticism in the leader treatment are effectively indistinguishable, whereas the skepticism ECDF in the unilateral treatment has qualitative features more similar to that of the consensus treatment (for example, a significant mass at 0). This observation further highlights differences between treatments in which blame can (almost) perfectly be attributed to one individual and those in which decisions are seen as properly collective.

[17]Compared to the contexts in these experimental settings, in our environment, there is no clear "prosocial" or "anti-social" action. The interpretation of the "no disclosure" message therefore does not reflect whether the evaluator perceives each of the group members to have made a "morally wrong" action recommendation, but simply their learning about each group member's value through the group's strategic choice.

4.2 Robustness of Main Hypotheses Tests

In Appendix B, we report a variety of alternative tests of our main hypotheses. Our initial tests are displayed in the first three lines of Table 2, and are conducted using data from all 30 rounds in all sessions conducted for each of our treatments.

In Tables 9 and 10, we perform the same tests using data either only from early rounds (first 15 rounds) and only from late rounds (last 15 rounds), separately. Tables 11 and 12 test the same hypotheses using the data from our post-experiment questionnaire. Beyond the different elicitation method, this data differs from our 30 rounds of play also because no disclosure beliefs are reported not only by evaluators but also by subjects who played the roles of group members. We perform the main hypotheses tests separately for subjects who played different roles. Table 13 replicates the hypotheses tests using data only from the first 4 sessions of each of our treatments.[18] Finally, for the analysis in Table 14, we define skepticism at the evaluator level: for each subject who played the role of the evaluator, we calculate their *evaluator-level* skepticism by averaging across all instances in which they saw no disclosure.

In each of these alternative specifications, the exact average skepticism values differ from the baseline values reported in Table 1, but our main predictions of how the interpretation of "no disclosure" varies across treatments — those stated in Hypotheses 1-3 — remain valid in all these variations. For the evaluator-level skepticism analysis, we regressed skepticism on an indicator of each of our possible treatments (using consensus as a the reference category), also including demographics as controls. We additionally observe that skepticism is significantly larger for female subjects who played the evaluator role and for subjects with higher GPA.[19]

4.3 Are Evaluators' Empirically Best-Responding?

4.3.1 Guesses after "no disclosure"

We also assess evaluators' guesses after no disclosure as compared to the actual values that were not disclosed in each of our treatments. In a Bayesian benchmark, we should find that, across all treatments, evaluators' elicited beliefs of no disclosure correspond to the actual average values (to each of the group members) that were not disclosed. Table 3 displays average values that were not disclosed and average no disclosure guesses across treatments,

[18] Our preregistration initially stated that we would run 4 sessions of each treatment. Please see details of the original and reviewed preregistration in the AEA RCT Registry, under ID 0013276.

[19] Significance in these regressions is assessed using session-level clustered standard errors.

as well as p-values from testing the difference between values and guesses.

Treatment	Avg. Not Disclosed Value	Avg. No Disclosure Guess	p-value
Individual	2.93	3.23	0.330
Unilateral	3.16	3.80	0.019
Leader	2.98	3.38	0.067
Non-Leader	5.09	5.02	0.795
Consensus	4.50	4.35	0.308

Note: To test for differences between values and guesses in each treatment, we estimate a linear regression in stacked format (with both observed values and guesses as outcomes), including a dummy variable for "guess" versus "value." P-values are calculated from this regression, with standard errors clustered at the session level.

Table 3: Evaluators' no disclosure guesses and values that were not disclosed.

In most treatments — except the consensus treatment — we find that the observer's no disclosure guess is on average higher than the actual average value that was not disclosed by the group. Across treatments, with the exception of the unilateral treatment, we cannot reject the hypothesis that the evaluators' no disclosure guesses are equal to the average value that is not disclosed. We conclude that evaluators on average accurately assess the values that are not disclosed by the group, and do so across treatments. In comparison, previous literature has documented (in individual disclosure experiments) that observers' guesses are larger than the actual non-disclosed values, and often significantly so.

4.3.2 Guesses after Disclosed Values

The focus of our analysis is on how evaluators interpret the "no disclosure message," when it is sent by the group. When the group chooses instead to disclose, the evaluator can perfectly see the drawn values and is able to therefore maximize the accuracy of their guess by submitting a pair of guesses equal to the observed values. Indeed, across treatments, we find that 93% of guesses after the group hand is disclosed correspond to the exact observed values (and over 97% of guesses are within 1 of the observed values).

4.3.3 Evaluators' Payoffs

In theory, the payoffs to evaluators should be larger under deliberation procedures that induce more disclosure in equilibrium. The "amount of disclosure" is ordered in equilibrium as follows: the unilateral procedure induces full disclosure, and so does the individual procedure; the leader procedure induces the disclosure of all but the worst value realizations of the leader's value; and the consensus procedure induces the least disclosure. In our data from the main part of the experiment, the evaluators' average payoff was 105 points in the unilateral treatment, 102 points in the individual treatment, 100 points in the leader treatment, and 92 points in the consensus treatment.

4.4 Individual Analysis: Modal Updating after No Disclosure

Beyond evaluating our main skepticism hypotheses, our data documents other features of the interpretation of communication in a group setting. From the joint distribution of "no disclosure guesses" — the distribution of the pair of guesses made by an evaluator after seeing no disclosure, which we depict for each treatment in Figure 10 in Appendix B — we can make out that evaluators often make no disclosure guesses that are above 5 for either or both group members' values (reflecting "negative" skepticism, relative to the unconditional value mean of 5), and that the pair of guesses are often not symmetric across group members, even in symmetric treatments (unilateral and consensus).

To further investigate these features, we analyze evaluators individually, in terms of their modal pair of guessed values after no disclosure. For each evaluator, if their modal no disclosure guesses are values below 5 for both group member A and group member B, we say the evaluator updates negatively after seeing no disclosure. If their modal guesses are above 5 for both group member A and group member B's values, we say the evaluator updates positively after seeing no disclosure. If modal pairs are such that both guesses are equal to 5, then we say the evaluator does not update after seeing no disclosure. Finally, if the evaluator's modal pair of guesses after no disclosure differs across the two group members (in terms of whether it is below, above, or equal to 5), then we say the evaluator updates asymmetrically after no disclosure.

Figure 3 displays the distribution of evaluators over these modal updating categories across the four treatments. As a benchmark, we observe that in the individual treatment, the vast majority of evaluators (97.6%) see no disclosure as a negative signal of the sender's value, in line with the updating based on the strategic decision to conceal values that are below some threshold. Likewise in the leader treatment, evaluators update in a direction that is largely consistent with the aggregation of group members' strategies of recommending the concealment of low drawn values: 80% of evaluators have modal guesses that interpret "no disclosure" as negative news about the leader's value and not so about the non-leader's value (whose recommendations are almost never followed by the group).

In the symmetric group treatments (unilateral and consensus) we find instead that evaluators' updating after "no disclosure" predominantly differ in direction from our theoretical

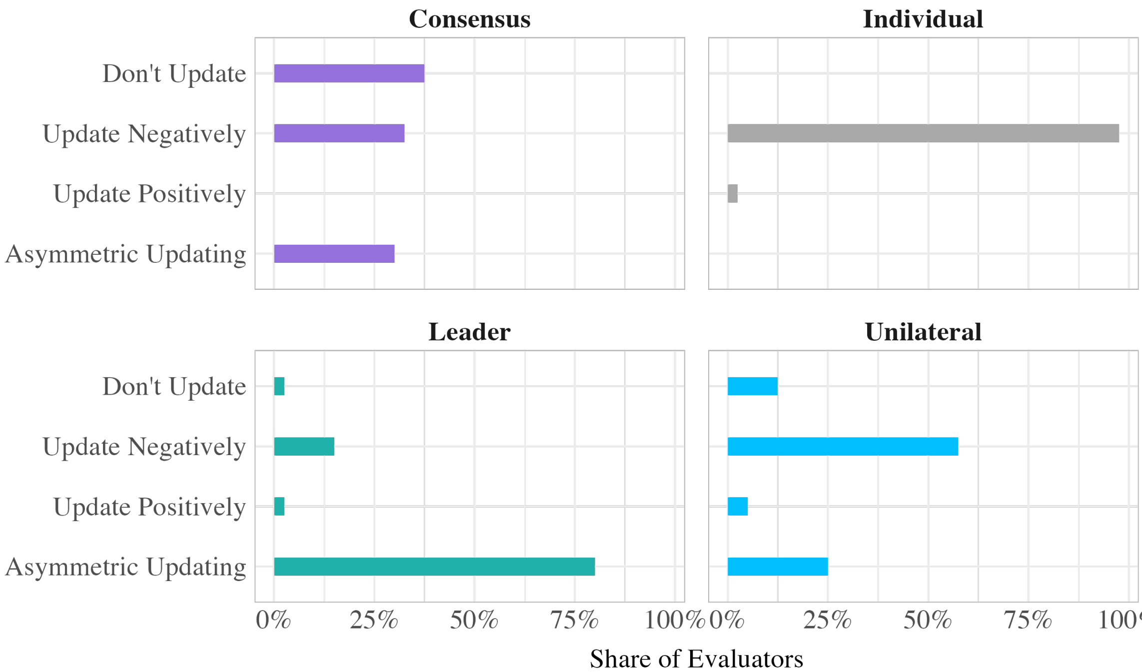

Figure 3: Dominant evaluator types. Categories defined as: “Don’t Update” = modal guesses A and B are both exactly 5; “Update Negatively” = modal guesses A and B are both strictly below 5; “Update Positively” = modal guesses A and B are both strictly above 5; “Asymmetric Updating” = all other modal guess configurations.

predictions. In theory, under both procedures, no disclosure should be interpreted as bad news about both group members’ values, relative to the unconditional mean of 5. In our data, 12.5% of evaluators in the unilateral treatment and 37.5% in the consensus treatment do not significantly update after seeing no disclosure, relative to the unconditional mean. Additionally, 25% of evaluators in the unilateral treatment and 30% in the consensus treatment update asymmetrically after seeing no disclosure (increasing skepticism about one group member’s value, but not the other), despite both treatments being symmetric.

We relate these results to the cognitive difficulty of the strategic learning performed by the evaluators in each of our treatments. We posit that our four treatments are ordered in terms of the complexity of the updating task induced by no disclosure. The individual treatment is the simplest, as a strategic choice to not disclose can only be attributed to the individual sender. The leader treatment is slightly more complex, as in it strategic no disclosure is overwhelmingly more likely to be due to such a recommendation by the leader (group member A) than by the non-leader (group member B). Compared to these two treatments, the updating task is more difficult in the unilateral and consensus treatments, because the group's disclosure decision depends on the non-trivial interaction of the two group members' recommendations. Between the unilateral and consensus treatments, the interpretation of "no disclosure" is less complex in the former than the latter. Under the unilateral procedure, there is a single strategic explanation for a group's decision to not disclose, that both group members recommended that action. In contrast, under the consensus procedure, there are three possibilities: no disclosure can ensue after (a) only group member A recommended no disclosure; (b) only group member B recommended no disclosure; and (c) both group members recommended no disclosure.

There are two layers of complexity to the updating task: strategic complexity (the understanding that group members strategically recommend the concealment of bad value realizations) and procedure complexity (the understanding that group members' recommendations can be non-trivially aggregated into the group's disclosure decision). As summarized by Oprea (2025) in a recent survey, one of the effects of complexity is that it can cause *procedural distortion*, whereby agents respond optimally to a distorted and cognitively simpler version of the task at hand. We see this in the unilateral and consensus procedures, where we can understand the individuals who do not update after seeing no disclosure as ignoring the strategic aspect of the updating task and therefore guessing that each group member's value is equal to the unconditional mean value 5. Another effect of complexity highlighted by Oprea (2025) is that it can lead to *mistakes*. In our context, we can understand evaluators' guesses above 5 after seeing no disclosure as being mistakes, as they cannot be justified by a simple distortion such as ignoring a layer of complexity. The extent to which we observe these effects of complexity is proportional to the complexity of treatments posited above, varying from simplest (individual) to most complex (consensus).

In a recent paper, Aina and Schneider (2025) also propose and document an updating distortion that ensues in situations where there are different models that could explain the observed data. In such a circumstance, a Bayesian updater would weight these different models accordingly when making an inference. In contrast, Aina and Schneider (2025) show that individuals in the lab significantly distort the weights they place on each possible model, relative to the Bayesian benchmark. In our consensus treatment, no disclosure can be explained by three different "models" — (a), (b), or (c) described above, depending on which group member recommended no disclosure. Bayesian reasoning would mean that an evaluator places accurate probability on each of these three events, and accordingly weight the expected value to each group member across them, yielding a guess of no disclosure that is symmetric across the two group members. The prevalent asymmetric guesses we observe can be justified by a non-Bayesian reasoning that outweighs the models (a) or (b), according to which only one group member recommended no disclosure.

5 Experimental Results: Group Disclosure Strategies

5.1 Disclosed and Undisclosed Values across Treatments

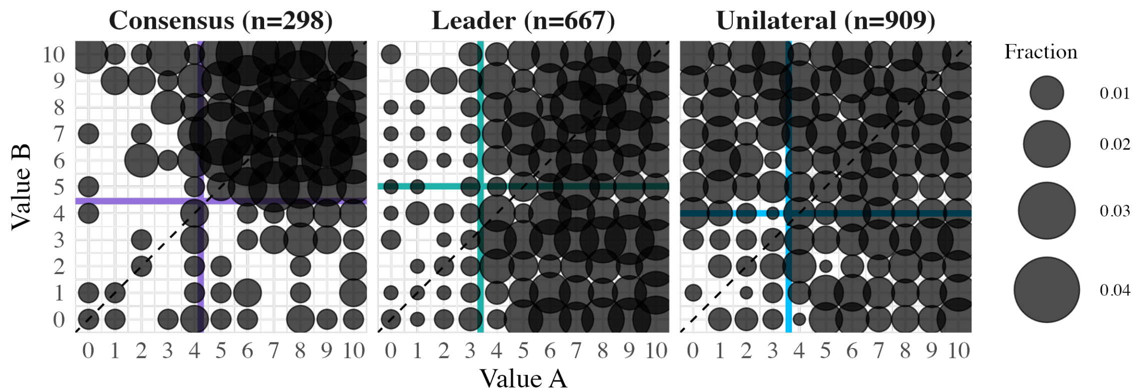

Figure 4: Joint distribution of disclosed values: the black bubbles represent pairs of values that were drawn for which the group chose to disclose. The colored thresholds indicate the average guesses made by the evaluator in each of the treatments.

Figure 4 depicts the joint distribution of values that were disclosed to the evaluator in each of our group treatments. These empirical distributions reflect each instance in which a group's decision was to disclose the group hand, across all participants and rounds played. At a first glance, the figures qualitatively resemble the behavior predicted by our theory, suggesting that group members generally use threshold recommendation strategies (recommending that the hand be disclosed if and only if their own drawn value is sufficiently large), which are then aggregated according to the deliberation procedure in each treatment. In Figure 4, we also highlight in each panel the evaluators' average no disclosure guesses of A and B's values, in the respective treatments.

Closely paralleling our hypothesis 1-3 on skepticism, hypotheses 4-6 regard the comparison of distributions of disclosed/undisclosed values in each of the group treatments. To verify each of these statements, we first used the Kolmogorov-Smirnov test to see whether the undisclosed values in pairs of treatments were drawn from the same distribution. Next, we followed the approach for consistent tests of first order stochastic dominance in Barrett and Donald (2003),[20] conducting tests for first-order stochastic dominance of the distribution of undisclosed values in each treatment pair. For the symmetric group treatments,

[20]Each of these tests utilized 500 bootstrapped resamplings (default option) to calculate p-values. For each pair, two direction null hypotheses are tested and p-values are compared to .005, corresponding to a p-value

consensus and unilateral, we combined the data for both group member roles (A and B), whereas for the asymmetric leader treatment, we split these into "leader" and "non-leader." The results for each of these tests are presented in Tables 4 and 5, and all pairs of relevant ECDFs are contrasted in Figure 11 in Appendix C.

Table 4: Pairwise Kolmogorov–Smirnov test results.

Comparison	KS D-statistic	Adjusted p-value
Consensus vs. Unilateral	0.200	0.000
Consensus vs. Leader	0.230	0.000
Consensus vs. Non-Leader	0.105	0.000
Individual vs. Unilateral	0.043	1.000
Individual vs. Leader	0.012	1.000
Unilateral vs Leader	0.052	0.163

Table 5: Pairwise FOSD results following Barrett and Donald (2003).

Comparison	Dominance
Consensus vs. Unilateral	Consensus FOSD Unilateral
Consensus vs. Leader	Consensus FOSD Leader
Consensus vs. Non-Leader	Non-Leader FOSD Consensus
Individual vs. Unilateral	No dominance
Individual vs. Leader	No dominance
Unilateral vs Leader	No dominance

The tests confirm the direction and significance of hypotheses 4-6: the distribution of undisclosed values in the consensus treatment FOS-dominates that in the unilateral treatment; the distribution of undisclosed values in the consensus treatment FOS-dominates that of undisclosed group member A (leader) values in the leader treatment; and the distribution of undisclosed group member B (non-leader) values in the leader treatment FOS-dominates that of undisclosed values in the consensus treatment. Additionally, we highlight that no dominance is established in comparing the distribution of undisclosed values in the unilateral treatment or those of the leader in the leader treatment to the distribution of undisclosed values in the benchmark individual treatment.

of 0.05 adjusted with Bonferroni correction of 10 comparison pairs. In each case, if the FOSD tests fail to reject one direction of the hypothesis ($p > .005$), while the opposite directional test rejects the hypothesis ($p < .005$), we interpret that as evidence of FOSD dominance, as proposed in Barrett and Donald (2003).

5.2 Are Disclosure Decisions Empirical Best Responses?

Fixing a pair ω_A^{ND} and ω_B^{ND} of "no disclosure guesses" that the evaluator makes about A and B's values, respectively, the best response disclosure recommendation strategy for each group member i is to recommend disclosure if and only if their drawn value is larger than ω_i^{ND}. For each treatment, we define the empirical best response disclosure strategy for the group to be the aggregation (through the deliberation procedure in that treatment) of threshold recommendation strategies where the threshold corresponds to the empirical average guess of A and B's values after the evaluator sees no disclosure — these are the colorful thresholds highlighted in Figure 4.

We compare group disclosure decisions played in the lab to this empirical best response benchmark. In Table 6, we describe the match between empirical best responses and the observed data: across treatments, we find the match rate to be around 80%. We decompose the matches (and mismatches) between group decisions and empirical best responses. The first two columns in Table 6 refer to realizations of the group hand for which the empirical best response is for the group to disclose ("Predicted = D"), and the second two columns refer to realizations for which the empirical best response is for the group to not disclose ("Predicted=ND").

	Best Response = D		Best Response = ND		
Treatment	Observed = D	Observed = ND	Observed = D	Observed = ND	Overall Match Rate
Individual	76.1%	23.9%	11.5%	88.5%	80.3%
Unilateral	81.8%	18.2%	30.0%	70.0%	80.4%
Leader	76.7%	23.3%	15.5%	84.5%	79.4%
Consensus	57.9%	42.1%	10.5%	89.5%	79.9%

Table 6: Match rates between observed group decisions and empirical best-responses.

In the individual, leader, and consensus treatments, we observe that "mistakes" are more often due to too little disclosure than to too much disclosure: no disclosure is chosen close in close to 90% of the instances in which it is a best response, and disclosure is chosen in less than 77% of the instances in which it is a best response. The opposite happens in the unilateral treatment, where "mistakes" due to too much disclosure are more prevalent. This is consistent with the fact that "no disclosure" decisions are harder to reach in that treatment, where they require both group members' support.

5.3 Individual Analysis: Disclosure Recommendation Strategies

We now investigate whether, at an individual level, subjects who played group member roles (a) used threshold recommendation strategies, (b) whether these thresholds correspond to their beliefs about their expected payoff of no disclosure, and (c) whether thresholds vary significantly across treatments as predicted by our theory.

5.3.1 Defining Empirical Threshold Strategies

To evaluate whether a subject played according to a threshold strategy, we take the following steps. For each subject s, we consider their individual recommendations in the last 20 rounds of play. Suppose subject s recommended that the group outcome be concealed from the evaluator in rounds $\{c_1, c_2, ..., c_k\} \subseteq \{11, 12, ..., 30\}$ and that the group outcome be disclosed to the observer in rounds $\{d_1, d_2, ..., d_{k'}\} \subseteq \{11, 12, ..., 30\}$. We create the set $\hat{\Phi}_0^s = \{v_{c_1}^s, v_{c_2}^s, ..., v_{c_k}^s\}$, which records every realization of subject s's own value for which they recommended that the outcome not be disclosed. Analogously, the set $\hat{\Phi}_1^s = \{v_{d_1}^s, v_{d_2}^s, ..., v_{d_{k'}}^s\}$ records every realization of subject s's own value for which they recommended that the outcome be disclosed. (Note that if there were two instances in which subject s drew value 7 and recommended disclosure, then both those instances are separately recorded in set $\hat{\Phi}_1^s$.)

We say there is overlap between sets $\hat{\Phi}_0^s$ and $\hat{\Phi}_1^s$ if their intersection is nonempty; and the *size of the overlap* is equal to $|\hat{\Phi}_0^s \cap \hat{\Phi}_1^s|$.[21] We say subject s uses a threshold strategy if the size of the overlap for subject s is at most 2, and if, after removing the overlaps, we find that the maximal element in the "no reporting" set is lower than the minimal element in the "reporting" set.

Table 7 displays statistics on the size of overlaps in recommendation strategies used by subjects in group member roles in each of our treatments. In all treatments, a significant portion of subjects are classified as having used threshold recommendation strategies.

For subjects who are classified as using threshold strategies, we ascertain their used threshold as follows. First, we remove any overlap from their sets $\hat{\Phi}_0^s$ and $\hat{\Phi}_1^s$, generating sets $\Phi_s^0 = \hat{\Phi}_s^0 \setminus \hat{\Phi}_s^1$ and $\Phi_s^1 = \hat{\Phi}_s^1 \setminus \hat{\Phi}_s^0$. Next, we define the threshold used by subject s as $t_s = \max \Phi_s^0$. That is, t_s is the largest own-value realization for which subject s recommended that the group's outcome be concealed from the evaluator.[22]

[21] If $\hat{\Phi}_0^s = \{0, 1, 1, 1, 2, 2, 3, 3, 3, 4, 4, 5, 7\}$ and $\hat{\Phi}_1^s = \{4, 7, 7, 7, 8, 9, 10\}$, then $\hat{\Phi}_0^s \cap \hat{\Phi}_1^s = \{4, 7\}$ and the size of the overlap for subject s is 2.

[22] As an alternative, we can infer the use of threshold strategies and define individual subject thresholds using the data from our post-play incentivized questionnaire (instead of main data). In the questionnaire, each

Table 7: Statistics on "overlap sizes" and use of threshold strategies.

Treatment	Average Overlap	Median Overlap	% Threshold Strategy
Individual	1.19	1	76.19
Unilateral	1.01	1	83.75
Leader	0.68	0	82.5
Non-Leader	0.98	0	77.5
Consensus	1.05	1	80.0

5.3.2 Individual Recommendation Strategies across Treatments

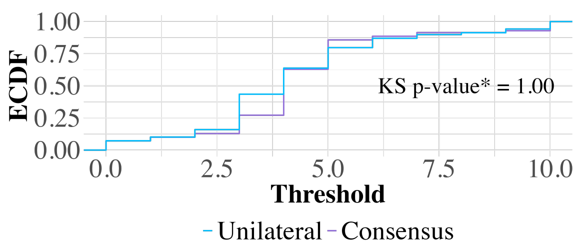

(a) Consensus vs Unilateral

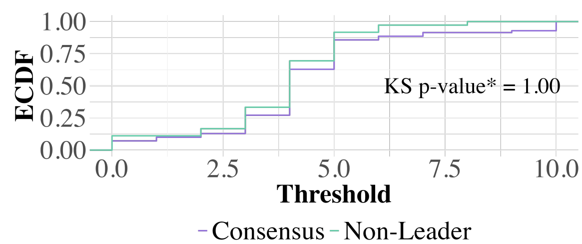

(b) Consensus vs Non-leader

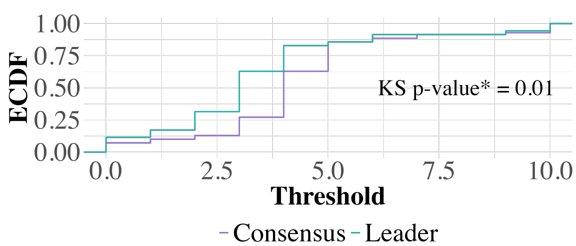

(c) Consensus vs Leader.

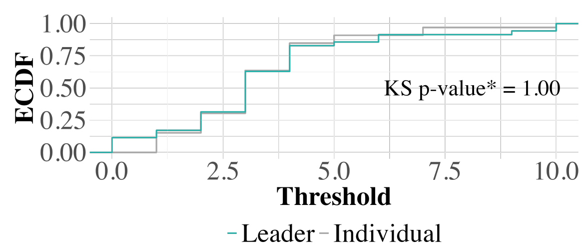

(d) Leader vs Individual

Figure 5: Comparison of subjects' thresholds across treatments.

In Figure 11, we comparatively display the distributions of thresholds across subjects in different treatments, along with p-values of the first order stochastic dominance tests for these distributions. The comparison in panel 5c are in line with our theoretical predictions: group members in the consensus treatment use disclosure recommendation thresholds that are larger than those who play the leader role in the leader treatment.

In contrast, panels 5a and 5b show that the differences between thresholds used in the consensus and unilateral treatments and consensus and non-leaders in the leader treatment,

subject reports their entire strategy, detailing what their recommendation would be after seeing each possible own-value realization. We find that thresholds elicited from the questionnaire data are closely related to those elicited from the main part of the experiment: for 76.8% of our subjects in group member roles, the difference between their thresholds elicited with the two methods is at most 1, and for 40.8%, the thresholds coincide.

which are predicted by our theory, do not exist in our data. This suggests that the differences in the distribution of undisclosed values between the unilateral and consensus treatments, documented in section 5.1, are mostly due to the mechanical differences in aggregation of recommendations, rather than due to individuals using distinct recommendation strategies.

Panel 5d also displays the distribution of thresholds for leaders in the leader treatment and for senders in the individual treatment. We find that these distributions do not differ significantly, as predicted by the theory.

5.3.3 Pivotality and Individual Strategies

Having documented that individuals' disclosure thresholds do not systematically vary across treatments in the way predicted by our theory, we posit that some of these differences between theory and observation correlate with the pivotality structures implied by each deliberation procedure. As a benchmark for what individuals should perceive as their individually optimal thresholds, we elicit their personal "guesses of no disclosure" from the post-play incentivized questionnaire. Remember that all players, including those playing group-member roles, were prompted to respond to the question "if you were the evaluator and saw that the group chose not to disclose their outcome, what would be your guess of group member i's value?"

We interpret this answer as the subject's perception of the evaluator's "no disclosure belief," and therefore of what their payoff would be if the group were to choose not to disclose. As a consequence, each individual's best response to these beliefs is to recommend that the group hand be disclosed if and only if their own drawn value is larger than the belief. In Figure 6, we display the average difference between estimated thresholds and elicited questionnaire beliefs across treatments.[23]

A first observation is that subjects who played the role of the non-leaders in the leader treatment use significantly lower disclosure thresholds (relative to their own elicited beliefs), compared to the individuals, leaders, and group members in the unilateral treatment. Non-leaders deviate from their empirical best responses more often towards using thresholds that are "too small," indicating that they err mostly by recommending too much disclosure, when compared to these other treatments. One interpretation is that, when subjects understand that their recommendations will most likely not be pivotal (because they are non-leaders), they tend to overly recommend disclosure, even at the (possible) expense of

[23]We perform all pairwise Wilcoxon rank-sum tests and find that the average threshold-belief for non-leaders is significantly different from those for individual, leader, and unilateral (Bonferroni-adjusted p-values are .041, .047, and .0009, respectively). None of the other pairwise differences are significant.

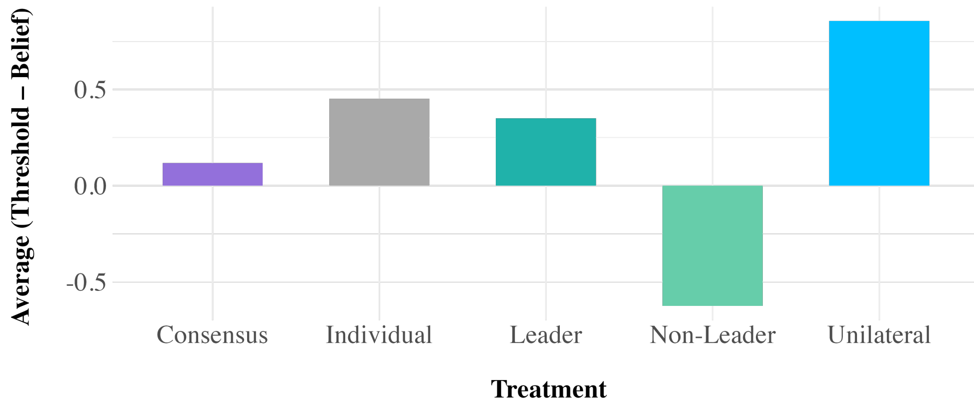

Figure 6: Difference between played threshold and elicited no disclosure belief.

their individual payoff, perhaps indicating subjects' "lying aversion" (where we interpret not disclosing as akin to lying).[24]

We also observe that subjects in the unilateral treatment use thresholds that are larger (relative to their elicited beliefs), compared to the consensus treatment (although the difference between thresholds-beliefs across these two treatments are not significant, when we consider p-values that are adjusted for multiple-hypothesis testing). For interpretation, consider the pivotality dynamics in these two symmetric group treatments.

In the unilateral treatment, if a group member recommends that the hand be disclosed, that recommendation is necessarily heeded by the group (and can therefore be pivotal), whereas a recommendation not to disclose is never pivotal, since in that case the other group member can still recommend disclosure and have that be the chosen group action. Recommending "no disclosure" is therefore an action that delegates all the power to one's partner. Conversely, in the consensus treatment, a no disclosure recommendation by one of the group members is always followed by the group (and may therefore be a pivotal recommendation), and recommending disclosure is never pivotal, and instead delegates the group's decision to the other group member. The difference in individual recommendation behavior across these two treatments may be thus correlated with the pivotality of their actions: subjects in the unilateral treatment "err" more towards choosing their non-pivotal action (which differs from the non-pivotal action in the consensus treatment).

A preference against taking pivotal actions has also been documented in different ex-

[24]Lying aversion has been documented in other communication experiments. For instance, Cai and Wang (2006) suggest this behavioral bias as a justification for "overcommunication" in a cheap talk setting. Jin et al. (2021) also posit that senders' aversion to sending a "no disclosure" message is connected to their social preferences towards the receivers' payoffs.

periments on group decision-making. For example, Bartling et al. (2015) shows that when a group implements an unpopular choice (unpopular from the perspective of an outsider) through a sequential voting process, individuals in the group cast their votes strategically in order to delegate the pivotal choice to their fellow group members (supposedly to avoid being perceived as responsible for the unpopular decision).[25] In our context, the fact that individuals avoid taking pivotal actions can also be understood through the lens of social preferences: by taking a non-pivotal action (for example, recommending no disclosure in the unilateral procedure), an agent delegates the decision to their partner, thereby providing them option value. If a group member values the payoff to their fellow group member, then the non-pivotal action is their optimal action when their own drawn value is close to the belief of no disclosure.[26]

6 Dynamics of Play

Figures 7 and 8 depict the dynamics of skepticism and disclosure recommendations, respectively, over the 30 rounds in each of the treatments. Note that, in Figure 8, we depict the proportion of group hands for which group members recommended disclosure, without controlling for the value that was drawn; a consequence is that a lot of the variability over time is due to changes in the distribution of values that were drawn in different rounds.

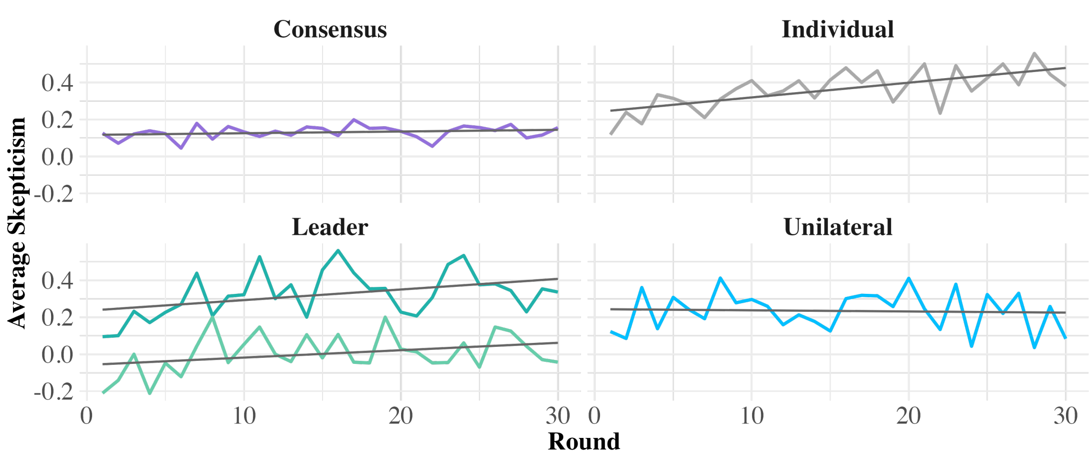

Figure 7: Skepticism dynamics over rounds.

[25]This phenomenon is similarly documented by Engl (2022), who contributes also by formalizing a measure of the pivotality of an individual's vote for a group's decision. In contrast, a different strand of the experimental literature — for example, Fehr et al. (2013), Bartling et al. (2014), and — document that, from an ex-ante stand, individuals prefer to retain control over actions that impact their payoffs.

[26]We formalize this argument in a previous version of this paper, by considering an instance of the group disclosure model in which group members have social preferences a la Fehr and Schmidt (1999).

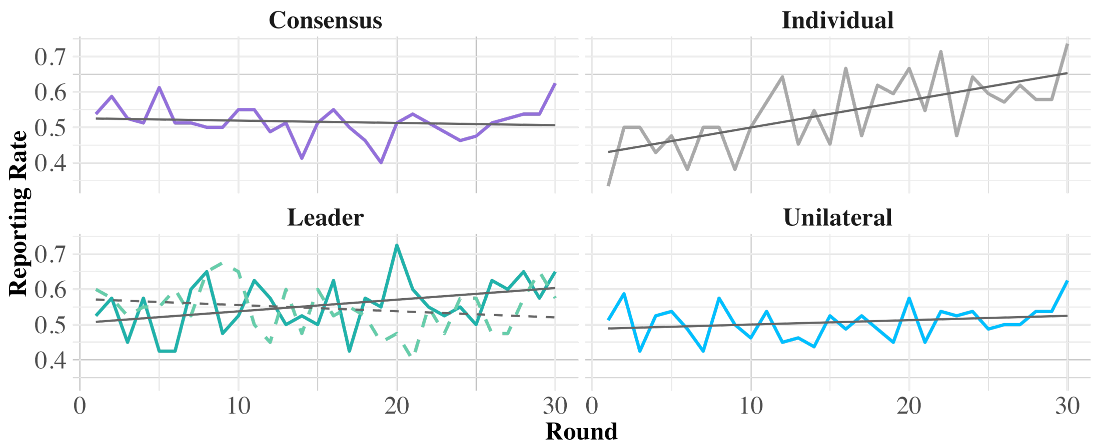

Figure 8: Disclosure recommendation dynamics over rounds.

In the individual treatment, we see that behavior of individuals in both evaluator and group member roles get closer to equilibrium (full disclosure, full skepticism) as rounds progress. This is consistent with observations by Li and Schipper (2020) and Jin et al. (2021), who find that play in later rounds better approaches the theoretically predicted unravelling. The same dynamics arise, though to a smaller extent, in our leader treatment, in which skepticism about the leader increases over time. As for the unilateral and consensus treatment, these features are relatively stable over the 30 rounds.

Despite responses over rounds being very stable in the group treatments, we find that individuals do respond to feedback from previous rounds. Figures 12 and 13 in Appendix D display regression results describing how subjects' decisions depend on previous play: in the first figure, the outcome variable is evaluators' guess of no disclosure, and in the second, the outcome variable is group members' disclosure recommendations. No disclosure guesses vary significantly with the value of previous cards that were hidden by the group, and especially so in early rounds of play.[27] In turn, in early rounds of play, disclosure recommendations made by group members vary significantly with the magnitude of previous guesses made by the evaluator after seeing no disclosure.[28]

[27] We find that a one-unit increase in the most recent hidden value is associated with a 0.155-point rise in the guess ($p < .001$), but this effect is significantly attenuated in the late phase ($p < .05$).

[28] A higher guess reduces the log-odds of recommending disclosure by 0.092 ($p = .006$).

References

Aina, Chiara and Florian H. Schneider, "Weighting Competing Models," *working paper*, 2025.

Ba, Cuimin, J Aislinn Bohren, and Alex Imas, "Over- and Under-Reaction to Information," *working paper*, 2022.

Barrett, Garry F and Stephen G Donald, "Consistent tests for stochastic dominance," *Econometrica*, 2003, *71* (1), 71–104.

Bartling, Björn, Ernst Fehr, and Holger Herz, "The Intrinsic Value of Decision Rights," *Econometrica*, 2014, *82*, 2005–2039.

_ , Urs Fischbacher, and Simeon Schudy, "Pivotality and Responsibility Attribution in Sequential Voting," *Journal of Public Economics*, 2015, *128*, 133–139.

Behnk, Sascha, Li Hao, and Ernesto Reuben, "Shifting Normative Beliefs: on Why Groups Behave More Antisocially than Individuals," *European Economic Review*, 2022, *145*, 104–116.

Brütt, Katharina, Arthur Schram, and Joep Sonnemans, "Endogenous Group Formation and Responsibility Diffusion: an Experimental Study," *Games and Economic Behavior*, 2020, *121*, 1–31.

Cai, Hongbin and Joseph Tao-Yi Wang, "Overcommunication in Strategic Information Transmission Games," *Games and Economic Behavior*, 2006, *16*, 7–36.

Charness, Gary and Matthias Sutter, "Groups Make Better Self-Interested Decisions," *Journal of Economic Perspectives*, 2012, *26* (3), 157–176.

Choo, Lawrence, Veronika Grimm, Gergely Horváth, and Kohei Nitta, "Whistleblowing and Diffusion of Responsibility: an Experiment," *European Economic Review*, 2019, *119*, 287–301.

Dana, Jason, Roberto A. Weber, and Jason Xi Kuang, "Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness," *Economic Theory*, 2007, *33* (1), 67–80.

Darley, John M. and Bibb Latané, "Bystander Intervention in Emergencies: Diffusion of Responsibility," *Journal of Personality and Social Psychology*, 1968, *8.*

Deversi, Marvin, Alessandro Ispano, and Peter Schwardmann, "Spin Doctors: an Experiment on Vague Disclosure," *European Economic Review*, 2021, *139*, 1038–72.

Dickhaut, John, Margaret Ledyard, Arijit Mukherji, and Haresh Sapra, "Information Management and Valuation: an Experimental Investigation," *Games and Economic Behavior*, 2003, *44* (1), 26–53.

Dye, Ronald A., "Disclosure of Nonproprietary Information," *Journal of Accounting Research*, 1985, pp. 123–145.

Engl, Florian, "A Theory of Causal Responsibility Attribution," *working paper*, 2022.

Farina, Agata and Mario Leccese, "Hiding a Flaw?Experimental Evidence on Multi-Dimensional Information Disclosure," *working paper*, 2025.

_, Guillaume Frechette, Alessandro Ispano, Alessandro Lizzeri, and Jacopo Perego, "The Selective Disclosure of Evidence: An Experiment," *working paper*, 2024.

Fehr, E. and K.M. Schmidt, "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 1999, *114*, 817–868.

Fehr, Ernst, Holger Herz, and Tom Wilkening, "The Lure of Authority: Motivation and Incentive Effects of Power," *American Economic Review*, 2013, *103*, 1325–1359.

Fischbacher, Urs, "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments," *Experimental Economics*, 2007, *10* (2), 171–178.

Forsythe, Robert, R. Mark Isaac, and Thomas R. Palfrey, "Theories and tests of "blind bidding" in sealed-bid auctions," *RAND Journal of Economics*, 1989, pp. 214–238.

Goeree, Jacob K. and Leeat Yariv, "An Experimental Study of Collective Deliberation," *Econometrica*, 2011, *79* (3), 893–921.

Greiner, Ben, "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE," *Journal of the Economic Science Association*, 2015, *1* (1), 114–125.

Grossman, Sanford J., "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 1981, *24* (3), 461–483.

Hagenbach, Jeanne and Charlotte Saucet, "Motivated skepticism," *Review of Economic Studies*, 2025, *92* (3), 1882–1919.

_ **and Eduardo Perez-Richet**, "Communication with Evidence in the Lab," *Games and Economic Behavior*, 2018, *112*, 139–165.

Jin, Ginger Zhe, Michael Luca, and Daniel Martin, "Is No News (Perceived as) Bad News? An Experimental Investigation of Information Disclosure," *American Economic Journal: Microeconomics*, 2021, *13* (2), 141–173.

King, Ronald R and David E Wallin, "Voluntary Disclosures When Seller's Level of Information is Unknown," *Journal of Accounting Research*, 1991, *29* (1), 96–108.

Kugler, Tamar, Edgar E. Kausel, and Martin G. Kocher, "Are Groups More Rational than Individuals? A Review of Interactive Decision Making in Groups," *WIREs Cognitive Science*, 2012, *3* (4), 471–482.

Lai, Ernest K., Wooyoung Lim, and Joseph Tao yi Wang, "An Experimental Analysis of Multidimensional Cheap Talk," *Games and Economic Behavior*, 2015, *91*, 114–144.

Li, Ying Xue and Burkhard C Schipper, "Strategic Reasoning in Persuasion Games: an Experiment," *Games and Economic Behavior*, 2020, *121*, 329–367.

Martinelli, Cesar and Thomas R. Palfrey, "Communication and Information in Games of Collective Decision: A Survey of Experimental," *working paper*, 2018.

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 1981, pp. 380–391.

Onuchic, Paula and Joao Ramos, "Disclosure by Groups," *working paper*, 2025.

Oprea, Ryan, "Complexity and its Measurement," *prepared for the Handbook of Experimental Methods in the Social Sciences*, 2025.

Sheth, Jesal D., "Disclosure of Information under Competition: an Experimental Study," *Games and Economic Behavior*, 2021, *129*, 158–180.

Vespa, Emanuel and Alistair J. Wilson, "Communication with Multiple Senders: an Experiment," *Quantitative Economics*, 2016, *7* (1), 1–36.

A Online Appendix: Proof of Theorem 1

Statement 1. The proof of this statement is a direct application of Theorems 1 and 3 in Onuchic and Ramos (2025).

Statement 2. Under the consensus procedure, full disclosure is not a sequential equilibrium (this follows from Theorem 3 in Onuchic and Ramos (2025)). Considering outcomes without full disclosure, ω_A^{ND} and ω_B^{ND} constitute an equilibrium if and only if

$$\omega_A^{ND} = \frac{\sum_{w \leqslant \omega_A^{ND}} \frac{1}{11} w + \sum_{w > \omega_A^{ND}} \left(\sum_{v \leqslant \omega_B^{ND}} \frac{1}{121} \right) w}{\sum_{w \leqslant \omega_A^{ND}} \frac{1}{11} + \sum_{w > \omega_A^{ND}} \left(\sum_{v \leqslant \omega_B^{ND}} \frac{1}{121} \right)}$$

$$\text{and } \omega_B^{ND} = \frac{\sum_{w \leqslant \omega_B^{ND}} \frac{1}{11} w + \sum_{w > \omega_B^{ND}} \left(\sum_{v \leqslant \omega_A^{ND}} \frac{1}{121} \right) w}{\sum_{w \leqslant \omega_B^{ND}} \frac{1}{11} + \sum_{w > \omega_B^{ND}} \left(\sum_{v \leqslant \omega_A^{ND}} \frac{1}{121} \right)}.$$

We verify numerically that the unique solution to this system is $\omega_A^{ND} = \omega_B^{ND} = 3.639$.

Statement 3. Under the leader procedure, for any $\epsilon > 0$, full disclosure is not a sequential equilibrium (this follows from Theorem 3 in Onuchic and Ramos (2025)). Considering outcomes without full disclosure, ω_A^{ND} and ω_B^{ND} constitute an equilibrium if and only if

$$\omega_A^{ND} = \frac{\sum_{w \leqslant \omega_A^{ND}} \left[\sum_{v \leqslant \omega_B^{ND}} \frac{1}{121} + \sum_{v > \omega_B^{ND}} \frac{1-\epsilon}{121} \right] w + \sum_{w > \omega_A^{ND}} \left(\sum_{v \leqslant \omega_B^{ND}} \frac{\epsilon}{121} \right) w}{\sum_{w \leqslant \omega_A^{ND}} \left[\sum_{v \leqslant \omega_B^{ND}} \frac{1}{121} + \sum_{v > \omega_B^{ND}} \frac{1-\epsilon}{121} \right] + \sum_{w > \omega_A^{ND}} \left(\sum_{v \leqslant \omega_B^{ND}} \frac{\epsilon}{121} \right)}$$

$$\text{and } \omega_B^{ND} = \frac{\sum_{w \leqslant \omega_B^{ND}} \left[\sum_{v \leqslant \omega_A^{ND}} \frac{1}{121} + \sum_{v > \omega_A^{ND}} \frac{\epsilon}{121} \right] w + \sum_{w > \omega_B^{ND}} \left(\sum_{v \leqslant \omega_A^{ND}} \frac{1-\epsilon}{121} \right) w}{\sum_{w \leqslant \omega_B^{ND}} \left[\sum_{v \leqslant \omega_A^{ND}} \frac{1}{121} + \sum_{v > \omega_A^{ND}} \frac{\epsilon}{121} \right] + \sum_{w > \omega_B^{ND}} \left(\sum_{v \leqslant \omega_A^{ND}} \frac{1-\epsilon}{121} \right)}.$$

As $\epsilon \to 0$, the unique solution to this system converges to $\omega_A^{ND} = 0$ and $\omega_B^{ND} = 5$. $\square$

B Online Appendix: Other Figures and Tables (Section 4)

Table 8: Skepticism regression results (errors clustered at the session level)

Group	Estimate	Robust SE	p-value
Unilateral - Consensus	0.110	0.052	0.033
Leader - Consensus	0.193	0.047	0.000
Non-Leader - Consensus	-0.135	0.054	0.014

Note: Estimates are from an ordinary least squares regression of skepticism on treatment indicators. Standard errors are cluster-robust, allowing for arbitrary correlation within sessions, and are computed using the Huber–White sandwich estimator. The reference category is Consensus; reported coefficients represent differences in average skepticism relative to the reference category.

Table 9: Skepticism Hypotheses Test Results (Only Rounds 1-15).

Comparison	Avg. Skepticism	Avg. Skepticism	Adj. p-value
Consensus vs. Unilateral	0.125	0.228	< 0.01
Consensus vs. Leader	0.125	0.280	< 0.01
Consensus vs. Non-Leader	0.125	-0.027	< 0.01
Individual vs. Unilateral	0.302	0.228	0.021
Individual vs. Leader	0.302	0.280	1.000
Unilateral vs. Leader	0.228	0.280	0.194

Note: Pairwise Wilcoxon rank-sum (Mann-Whitney U) tests with Bonferroni adjustment.

Table 10: Skepticism hypotheses test results (only rounds 16-30).

Comparison	Avg. Skepticism	Avg. Skepticism	Adj. p-value
Consensus vs. Unilateral	0.137	0.255	< 0.01
Consensus vs. Leader	0.137	0.374	< 0.01
Consensus vs. Non-Leader	0.137	0.022	< 0.01
Individual vs. Unilateral	0.427	0.255	< 0.01
Individual vs. Leader	0.427	0.374	1.000
Unilateral vs. Leader	0.255	0.374	< 0.01

Note: Pairwise Wilcoxon rank-sum (Mann-Whitney U) tests with Bonferroni adjustment.

Table 11: Skepticism hypotheses test results (questionnaire data – evaluators).

Comparison	Avg. Skepticism	Avg. Skepticism	Adj. p-value
Consensus vs. Unilateral	0.121	0.277	< 0.01
Consensus vs. Leader	0.121	0.398	< 0.01
Consensus vs. Non-Leader	0.121	-0.085	< 0.01
Individual vs. Unilateral	0.460	0.277	0.065
Individual vs. Leader	0.460	0.398	1.000
Unilateral vs. Leader	0.277	0.398	0.469

Note: Pairwise Wilcoxon rank-sum (Mann-Whitney U) tests with Bonferroni adjustment.

Table 12: Skepticism hypotheses test results (questionnaire data – group members).

Comparison	Avg. Skepticism	Avg. Skepticism	Adj. p-value
Consensus vs. Unilateral	0.100	0.244	< 0.01
Consensus vs. Leader	0.100	0.314	< 0.01
Consensus vs. Non-Leader	0.100	-0.003	0.028
Individual vs. Unilateral	0.310	0.244	1.000
Individual vs. Leader	0.310	0.314	1.000
Unilateral vs. Leader	0.244	0.314	1.000

Note: Pairwise Wilcoxon rank-sum (Mann-Whitney U) tests with Bonferroni adjustment.

Figure 9: All pairwise ECDF comparisons of skepticism with Bonferroni-adjusted p-values (each p-value is multiplied by 10, leading to a fairly conservative set of p-values).

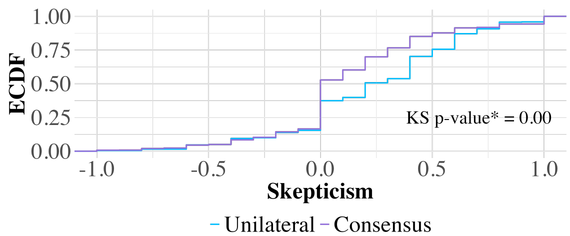

(a) Consensus vs. Unilateral

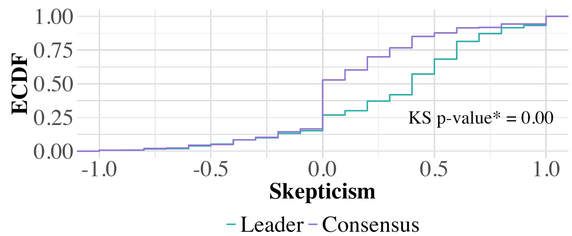

(b) Consensus vs Leader

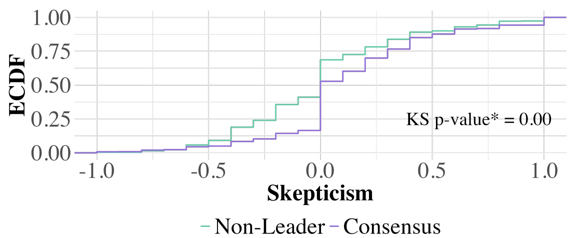

(c) Consensus vs Non-Leader

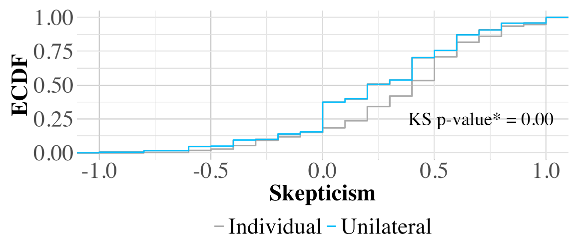

(d) Individual vs. Unilateral

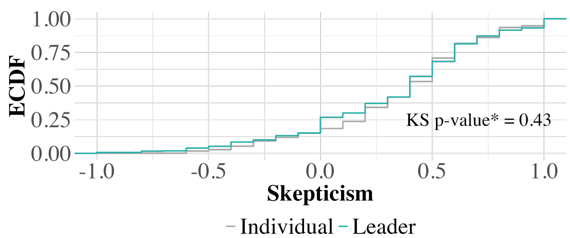

(e) Individual vs. Leader

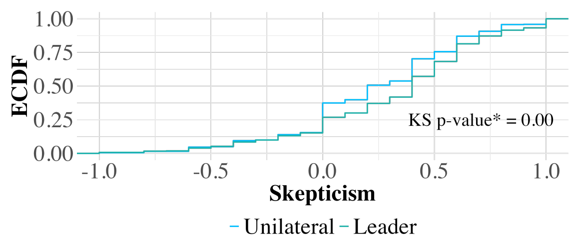

(f) Unilateral vs Leader

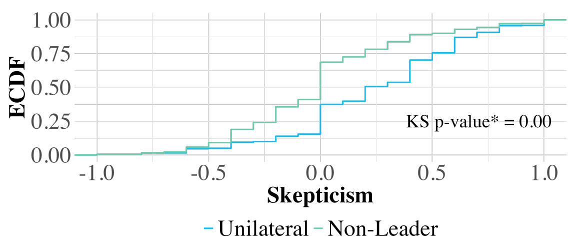

(g) Unilateral vs Non-Leader

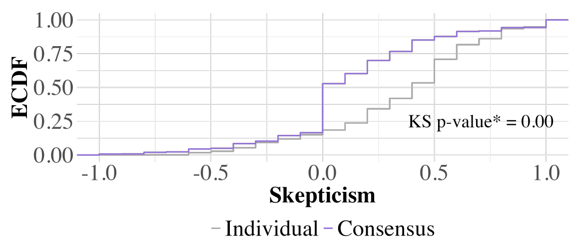

(h) Consensus vs Individual

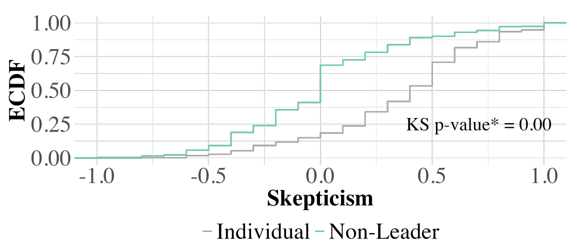

(i) Individual vs Non-Leader

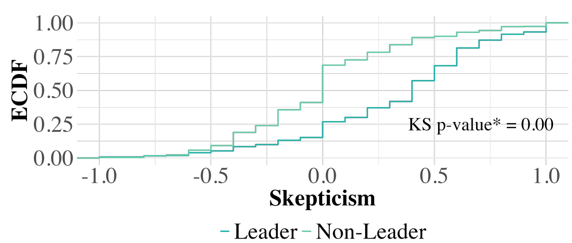

(j) Leader vs Non-Leader

Table 13: Skepticism hypotheses test results (wave 1 data only).

Comparison	Avg. Skepticism	Avg. Skepticism	Adj. p-value
Consensus vs. Unilateral	0.141	0.263	< 0.01
Consensus vs. Leader	0.141	0.286	< 0.01
Consensus vs. Non-Leader	0.141	-0.041	< 0.01
Individual vs. Unilateral	0.295	0.263	0.414
Individual vs. Leader	0.295	0.286	1.000
Unilateral vs. Leader	0.263	0.286	0.319

Note: Pairwise Wilcoxon rank-sum (Mann-Whitney U) tests with Bonferroni adjustment.

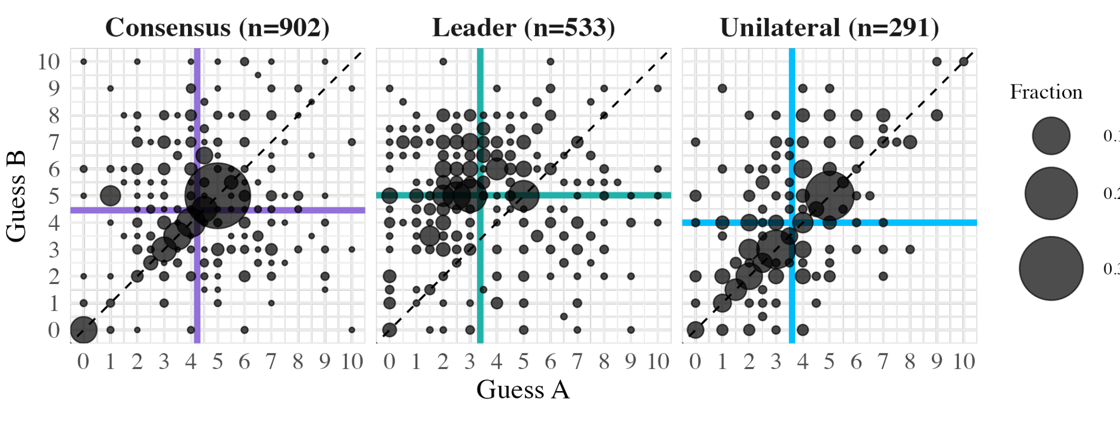

Figure 10: Joint distribution of guesses.

Table 14: Evaluator-level skepticism regressions.

	Evaluator's Average Skepticism	
	Treatment Only	With Demographics and Controls
	(1)	(2)
Unilateral vs Consensus	14.461^{***}	18.520^{***}
	(5.353)	(6.909)
Leader vs Consensus	20.228^{***}	20.523^{***}
	(4.835)	(5.232)
Non-Leader vs Consensus	-13.525^{***}	-13.210^{*}
	(4.755)	(6.957)
Individual vs Consensus	25.242^{***}	21.323^{**}
	(7.463)	(8.583)
GPA		7.671^{**}
		(3.751)
Gender (Female)		7.128^{*}
		(3.890)
Game Theory (Yes)		-2.917
		(4.206)
Comment Left (Yes)		2.303
		(3.774)
Econ/Business Major (Yes)		7.498
		(5.138)
Constant	12.796^{***}	-24.247
	(2.755)	(15.545)
Additional comparative tests		
Unilateral = Leader (p-value)	0.999	
Unilateral = Non-Leader (p-value)	0.002	
Unilateral = Individual (p-value)	0.999	
Leader = Non-Leader (p-value)	0.000	

Notes: Skepticism numbers were scaled from 0 to 100 for interpretation. Coefficients should be interpreted as percentage point changes. Session-level cluster robust SEs in parentheses. For each subject who played the role of the evaluator, we calculate their *evaluator-level* skepticism by averaging across all instances in which they saw no disclosure. For these regressions, the unit of observation is each subject who played the role of the evaluator.

C Online Appendix: Other Figures and Tables (Section 5)

Table 15: Logistic regressions on reporting decision

	Dependent variable:	
	Recommend Report	
	(1)	(2)
Value	0.577^{***} (0.015)	
Value> 4		3.060^{***} (0.071)
Individual	0.401^{***} (0.112)	0.440^{***} (0.113)
Leader	0.309^{**} (0.113)	0.350^{**} (0.114)
Non-Leader	0.226^{*} (0.114)	0.202 (0.113)
Unilateral	0.042 (0.092)	0.043 (0.092)
Constant	-2.876^{***} (0.098)	-1.724^{***} (0.078)
Observations	5,628	5,628
Note:		*p<0.05; **p<0.01; ***p<0.001

Table 15, both regression specifications present evidence that higher underlying values significantly increase the probability of recommending to report. In Model (1), each one-unit increment in the *Value* variable raises the odds of reporting by approximately 78% ($p < .001$), whereas in Model (2) respondents with *Value*> 4 are about 21 times more likely to recommend reporting than those with lower values ($p < .001$). Across both models, the Individual and Leader treatments significantly increase baseline reporting relative to the Consensus, which is the reference category: in Model (1), odds are 49% higher in the Individual condition ($p < .001$) and 36% higher in the Leader condition ($p < .01$), and in Model (2) these treatments increase odds by 55% ($p < .001$) and 42% ($p < .01$), respectively. The Non-Leader condition produces a modest positive effect, attaining significance only in the dummy-based specification, whereas the Unilateral treatment does not differ significantly from Consensus in either model.

Figure 11: All pairwise ECDF comparisons of undisclosed values (Bonferroni-adjusted p-values).

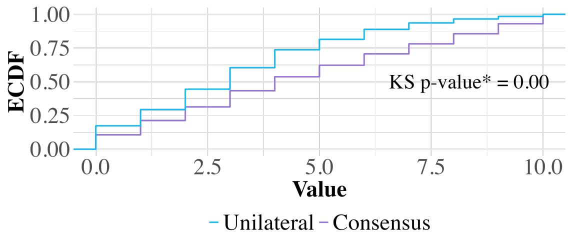

(a) Consensus vs. Unilateral

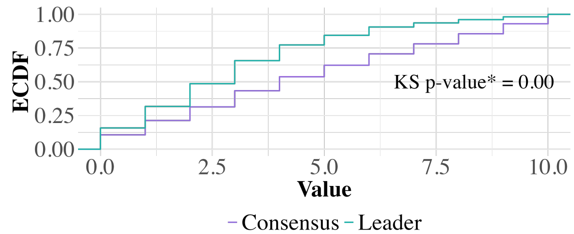

(b) Consensus vs Leader

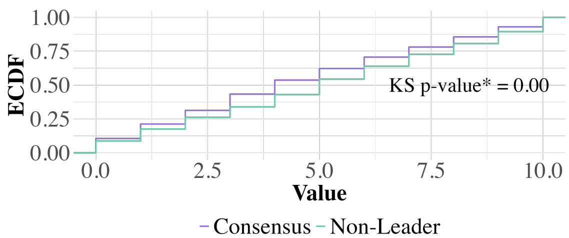

(c) Consensus vs Non-Leader

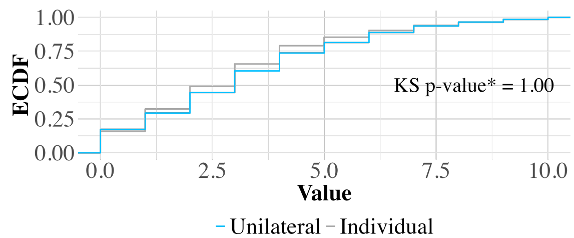

(d) Individual vs. Unilateral

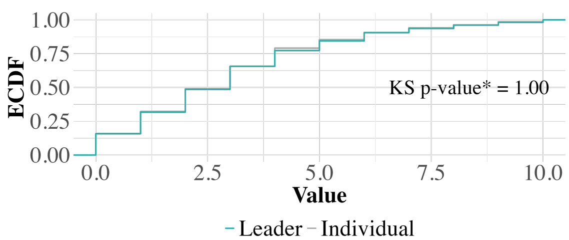

(e) Individual vs. Leader

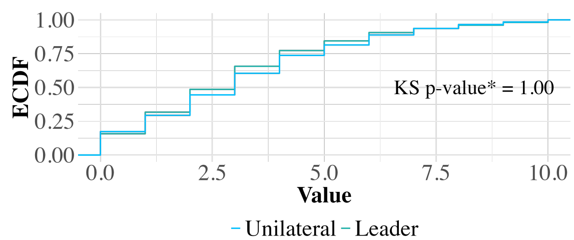

(f) Unilateral vs Leader

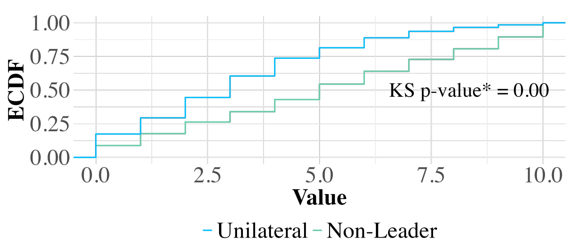

(g) Unilateral vs Non-Leader

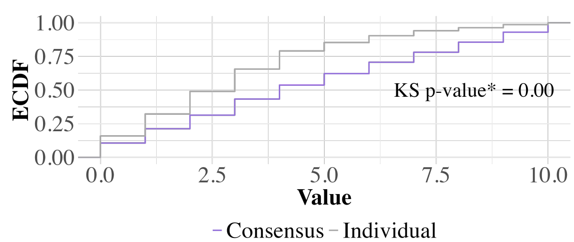

(h) Consensus vs Individual

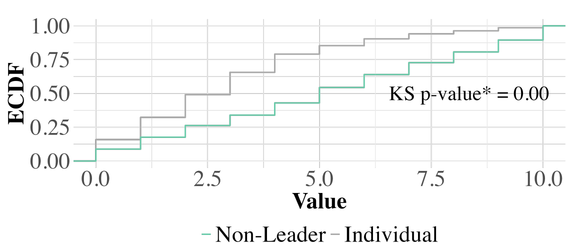

(i) Individual vs Non-Leader

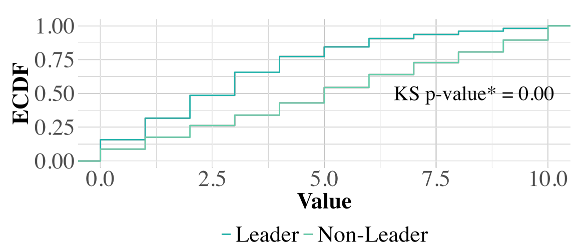

(j) Leader vs Non-Leader

D Online Appendix: Other Figures and Tables (Section 6)

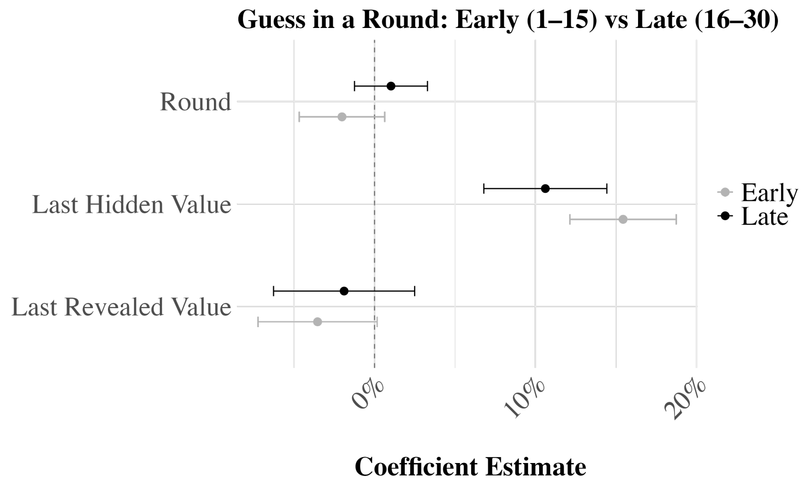

Figure 12: Linear regression of an evaluator's guess of A or B's values on feedback information from previous rounds (errors are clustered at the subject level).

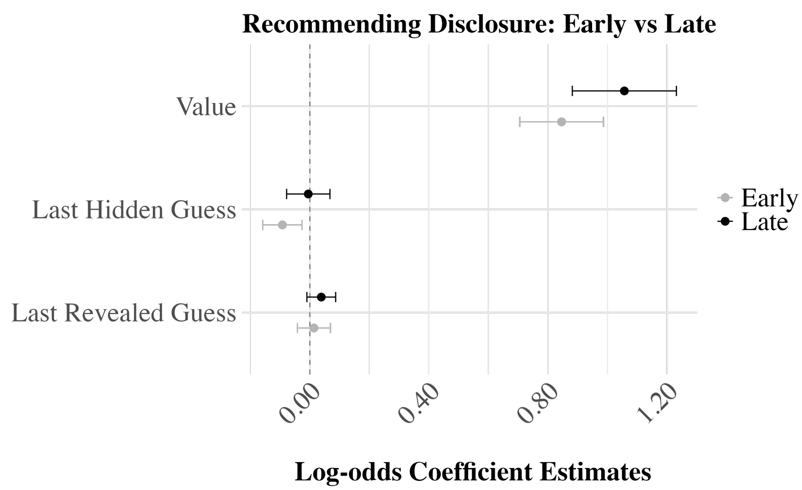

Figure 13: Two fixed-effects logistic regressions of group members' disclosure recommendations on the current drawn value and on feedback from previous rounds, for early rounds and late rounds of play (standard errors are clustered at the subject level in each case).

E Online Appendix: Instructions

The following instructions correspond to the *consensus* treatment.

Instructions

This is an experiment in economic decision making. What you earn in this experiment depends partly on your decisions, partly on the decisions of others and partly on chance. The amount of money you earn will be paid to you privately, in cash, at the end of the experiment.

The entire session will take place through computer terminals, and all interactions between you and the other participants will be done through the computers. Please do not talk, communicate in any way, or use your electronic devices during the session. If you have any questions during the entire session, raise your hand and your question will be answered privately.

Role assignment

You will be randomly assigned to one of three possible roles: you could be a *group member A*, a *group member B* or an *evaluator*. You will keep the same role for all 30 rounds. In each round, you will be randomly matched with two other participants in this room who have been assigned the other two roles. There will be a new random matching at the beginning of each round, so it is unlikely that you will be matched with the same two participants in consecutive rounds.

Round description

Each round consists of four stages:

1. Card-drawing stage; 2. Reporting stage; 3. Guessing stage; 4. Feedback stage.

1. ***Card-drawing stage (only group members participate)*** There are two decks of cards, deck A and deck B. Each deck has 11 cards labeled 0,1,2,...,9, and 10. The computer program will randomly pick one card from deck A and one card from deck B. The pair of cards drawn by the computer is referred to as the *group's hand*.

The number on the card drawn from deck A is called Value A; it represents the value of the group's hand to group member A. Similarly, the number on the card drawn from deck B is called Value B, representing the value of the group's hand to group member B. Within each deck a card is picked at random with equal chance. Note that the computer separately picks the card from each deck, so that group member A's value is not related to group member B's value.

At the card-drawing stage, each group member sees the card representing their respective value, but not the card representing the value of the group's hand to the other group member. Moreover, at this stage, neither card is seen by the evaluator.

2. *Reporting stage (only group members participate)* After observing their respective values, group member A and group member B decide whether to report the group hand to the evaluator. Group members can choose to report the entire group hand, or to not report it; reporting each card separately is not possible.

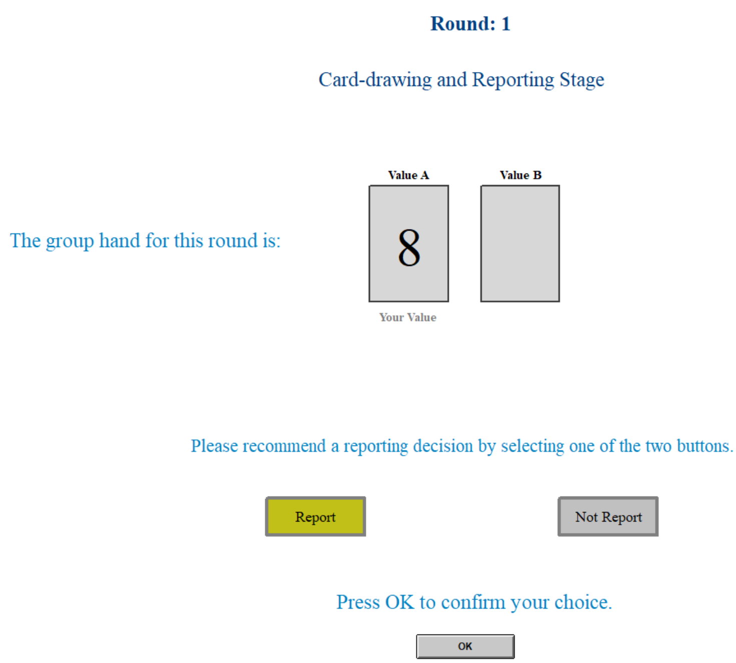

Figure 14: Sample Screen - Group Member A

Figure 14 presents a sample screen for a group member A. There are two buttons on the screen labeled 'Report' and 'Not Report' corresponding to two choices. The group member can move the cursor over one of these buttons and that button will light up, as button 'Report' is in Figure 14. After deciding on the selection, group member presses the "OK" button to confirm the recommendation. Each group member makes their own recommendation without observing that of the other group member.

- If *both* group members choose the 'Report' button then the evaluator **will see** both cards in the group hand, thereby revealing Value A and Value B.

- If neither group member A nor group member B choose the ‘Report’ button or if only one group member chooses the ‘Report’ button, the evaluator **will not see** the group hand. In that case, the evaluator will be informed that the group chose not to report the group’s hand.

3. *Guessing stage (only evaluator participates)* The evaluator is informed whether the group hand was reported or not.

- If the group hand was reported, the evaluator sees both cards in the group hand, thereby revealing Value A and Value B (as in Figure 15).

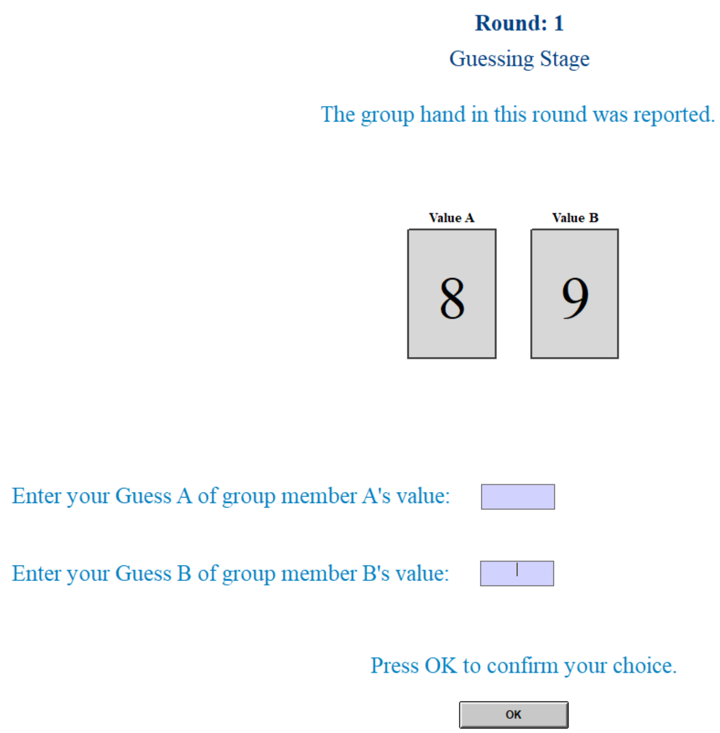

Figure 15: Sample Screen - Reported

- If the group hand is **not** reported, the evaluator does not see the group hand, and is instead informed that the group chose not to report the group’s hand (as in Figure 16).

After seeing the reported/not reported group hand, the evaluator is asked to make two guesses: to guess group member A’s value (Guess A), and to guess group member B’s value

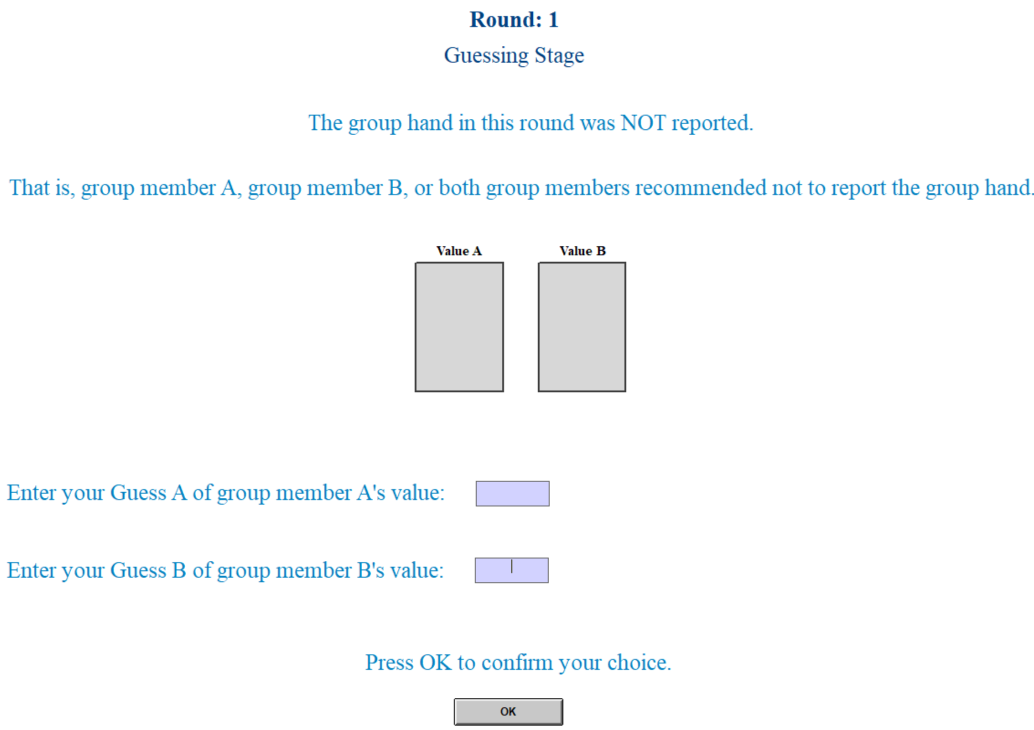

Figure 16: Sample Screen - Not Reported

(Guess B). Each guess is entered as a number between 0 and 10, and increments of 0.5 are allowed. For instance, if the evaluator thinks 3 and 4 are equally likely, they can insert a guess of 3.5. Instead, an evaluator who, for instance, would like to make a guess of 6.7 needs to "settle" for a guess of 6.5 or 7.

Once the evaluator makes their guesses (Guess A and Guess B) and confirms, every participant in the unit moves to the feedback stage.

4. ***Feedback stage (everyone participates)*** Every participant will see a feedback screen. The screen will show both cards in the group hand, whether the group hand was reported or not, and the evaluator's guesses. After everyone is done observing the screen, the round is over and a new round begins.

Evaluator's payoff

The evaluator is paid for the **accuracy** of their guesses. The evaluator gets paid for either the accuracy of Guess A or for the accuracy of Guess B, *with equal chance*. **The evaluator earns more when the guess is closer to the value in the drawn card.** Specifically, Table 16a presents the evaluator's payoffs in all possible scenarios.

Group members' payoffs

Each group member is rewarded based on the evaluator's guess of their own respective value. The higher the evaluator's guess of a group member's value, the more that group member earns. **The group member earns more when the evaluator's guess of their value is higher, regardless of the value in the drawn card.** Specifically, Table 16b presents the group member's payoffs in all possible scenarios.

Additional information about payoffs

Regardless of your role, you will be paid according to your points in 1 round chosen at random, in addition to a show-up fee. Points will be exchanged to US dollars at a rate of 10 points to 1 dollar.

Practice rounds

Before the beginning of the experiment, you will play 2 practice rounds. These rounds are meant for you to familiarize yourselves with the screens. All the choices made in the practice rounds are unpaid and have no relation to the paid 30 rounds. These are for illustrative purposes only and they do not affect the actual experiment.

(a) Evaluator's payoff

	Guess 0	Guess 0.5	Guess 1	Guess 1.5	Guess 2	Guess 2.5	Guess 3	Guess 3.5	Guess 4	Guess 4.5	Guess 5	Guess 5.5	Guess 6	Guess 6.5	Guess 7	Guess 7.5	Guess 8	Guess 8.5	Guess 9	Guess 9.5	Guess 10
Value 0	110	109	106	103	99	95	90	85	80	74	68	62	56	50	43	36	29	22	15	7	0
Value 1	106	109	110	109	106	103	99	95	90	85	80	74	68	62	56	50	43	36	29	22	15
Value 2	99	103	106	109	110	109	106	103	99	95	90	85	80	74	68	62	56	50	43	36	29
Value 3	90	95	99	103	106	109	110	109	106	103	99	95	90	85	80	74	68	62	56	50	43
Value 4	80	85	90	95	99	103	106	109	110	109	106	103	99	95	90	85	80	74	68	62	56
Value 5	68	74	80	85	90	95	99	103	106	109	110	109	106	103	99	95	90	85	80	74	68
Value 6	56	62	68	74	80	85	90	95	99	103	106	109	110	109	106	103	99	95	90	85	80
Value 7	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110	109	106	103	99	95	90
Value 8	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110	109	106	103	99
Value 9	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110	109	106
Value 10	0	7	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110

(b) Group member's payoff

	Guess 0	Guess 0.5	Guess 1	Guess 1.5	Guess 2	Guess 2.5	Guess 3	Guess 3.5	Guess 4	Guess 4.5	Guess 5	Guess 5.5	Guess 6	Guess 6.5	Guess 7	Guess 7.5	Guess 8	Guess 8.5	Guess 9	Guess 9.5	Guess 10
Value 0	0	7	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110
Value 1	0	7	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110
Value 2	0	7	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110
Value 3	0	7	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110
Value 4	0	7	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110
Value 5	0	7	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110
Value 6	0	7	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110
Value 7	0	7	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110
Value 8	0	7	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110
Value 9	0	7	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110
Value 10	0	7	15	22	29	36	43	50	56	62	68	74	80	85	90	95	99	103	106	109	110

Table 16