Too Good to be True – Individual and Collective Decision-Making with Misleading Signals*

Sebastian Fehrler[£] Anna Hochleitner[†] Moritz Janas[‡]

[£] University of Bremen, SOCIUM
[†] Norwegian School of Economics, FAIR and SNF, Bergen
[‡] University of Gothenburg

November 11, 2025

Abstract

In many situations, an abundance of misleading evidence—such as fake customer reviews—can lead to false or deceptive conclusions. We experimentally investigate individual and collective decision-making in an information structure in which signals can be correlated, depending on the state of the world. In this setting, too much evidence pointing in one direction has the potential to mislead, necessitating a level of sophistication for rational decision-making. Overall, participants' performance is poor with only small differences in collective and individual decision-making accuracy. Interestingly, the more complex environment tends to encourage greater honesty within heterogeneous groups than a benchmark setting with independent signals, thus corroborating a rather subtle game-theoretic prediction.

Keywords: individual choice, collective choice, correlated signals, information aggregation, communication, voting, committees.

JEL: C92, D23, D71.

*We would like to thank Urs Fischbacher, Niall Hughes, Maik Schneider, the research groups at TWI, CBID and NYUAD, and participants at several seminars, workshops and conferences for valuable comments. The study obtained ethical approval by the GfeW (No. 6aJqoedH). A pre-analysis plan of the experiment was registered at the AEA registry as AEARCTR-0006323. We gratefully acknowledge financial support from Tamkeen under NYU Abu Dhabi Research Institute Award CG005, from the Juniorprofessorenprogramm Baden-Württemberg, and from the Federal Ministry of Labour and Social Affairs (Germany) through the funding network Interdisciplinary Social Policy Research (BMAS-FIS). Finally, we thank Aigerim Shintayeva for excellent research assistance.

1 Introduction

We explore how potentially misleading information affects individual and collective decision-making. As modern society is built on learning and exchange of information, we are exposed to a complex flux of information on a daily basis. This makes the formation of beliefs and opinions a challenging task by itself and even more so if the signals we receive can be correlated in some states of the world. Usually, we are more likely to hold a given belief, the more information we can collect in its favor. However, if signals are correlated this can generate situations where too much evidence is "too good to be true" and thus misleading. In other words, the absence of dissent or contradicting information can cast doubt on the accuracy of signals (see e.g., Gunn et al., 2016).

An example for this logic are online reviews. While in general a larger number of positive reviews increases the chances that the product is good, there may be a concern that reviews can be fabricated. If a product thus receives only the highest grades, there may be a danger that reviews are not independent but correlated. In line with this, studies have shown that an average rating of 5.0 out of 5 stars can be perceived as less trustworthy than a rating of 4.8, as it is unlikely that many independent reviews result in *only* 5 star reviews (Maslowska et al., 2017; Smirnov and Starkov, 2022; Collinger, 2016; Skillings, 2022).

While too many identical signals may thus be a warning sign, previous research has shown that in practice individuals tend to struggle to correctly account for correlated information. Instead, they often suffer from *correlation neglect*, treating correlated information as independent (see e.g., Keren and Lewis, 1994; Enke and Zimmermann, 2019; Moser and Wallmeier, 2021). This has far reaching consequences both in the political domain, e.g., for the formation of political preferences and voting outcomes (see e.g., Glaeser and Sunstein, 2009; Levy and Razin, 2015; Ortoleva and Snowberg, 2015; Denter et al., 2021), as well as for decisions in management and finance, e.g., for investment choices (Laudenbach et al., 2023), understanding financial contagion (Demange, 2018) or demand forecasts (Kremer et al., 2016).

One limitation is that the current empirical evidence is mainly based on individual decision-making (see e.g., Kallir and Sonsino, 2009; Eyster and Weizsäcker, 2016; Enke and Zimmermann, 2019; Cason et al., 2020). In this paper, we examine experimentally how individual and group decisions change in the presence of correlated information. Understanding how groups respond to potentially misleading information is highly relevant as many important decisions are made by committees or juries. This relevance is underscored by a substantial body of both theoretical and experimental work on information aggregation and collective decision-making (see e.g., Feddersen and Pesendorfer, 1998; Guarnaschelli et al., 2000; Coughlan, 2000; Goeree and Yariv, 2011; Bouton et al., 2016, 2018). Our paper connects these two strands of literature, combining correlated information structures in which too much evidence for one option can be too good to be true with a jury design.[1]

In the experiment, participants have to make a binary choice that leads to higher payoffs if it matches the state of the world. The two possible states of the world are ex-ante equally likely. This choice is either taken individually or within groups. Across treatments, we then vary whether the information participants receive can be correlated or not. To give a real world example of such a decision problem, let us stay with the example of fabricated online reviews and consider customers who want to take an informed purchasing decision.[2] The product is either good or bad. Before taking a decision, customers receive a number of imperfectly informative signals for the quality of the product in the form of online reviews. Without correlation, the decision is straightforward. The more positive reviews are observed, the more likely the product is good and the more likely it should be bought. In the treatment with correlated information, participants know that in the case of a bad product, reviews can be fabricated such that all signals indicate that the product is good, complicating the inference from signals. A similar challenge arises for an individual decision-maker who relies on information from experts or reports that may appear independent but are in fact shaped by a common influence, such as a special interest group or lobbyist.

The intuition for allowing correlation only in the bad state of the world is as follows. Both good and bad companies face potential reputational costs from manipulation. However, a company with a bad product has more to gain from generating fake positive reviews, while a company with a good product risks undermining an otherwise strong position. As a result, manipulation is more attractive in the bad than in the good state. A similar intuition applies to the lobbying example. Assume a policy can either serve the broad public interest (good state) or primarily advance narrow interests (bad state). In both states, manipulation is conceivable in principle. However, when a policy primarily benefits a small group of people, manipulation is easier and hence cheaper because collective action problems are easier to overcome. To keep the environment tractable for participants, we implement the extreme case in which correlation can occur only in the bad state. This simplification ensures that participants face the easiest possible detection problem. Arguably, if participants fail to detect correlation in this easier case, our results represent a lower bound to decision errors in environments where differences are more subtle and correlation is possible in both states, though not equally likely.

Using the data from the individual treatments as a benchmark for the prevalence of

[1] This set-up is related to the *hidden profile paradigm* (Stasser and Titus, 1985) in psychology, where groups have to piece together individually received information in order to find an optimal choice. In line with the notion of correlation neglect, groups typically do not achieve optimal outcomes as they focus too much on shared information (see Lightle et al., 2009; Sohrab et al., 2015, for an overview).

[2] See e.g. Mayzlin et al. (2014) or Luca and Zervas (2016) on fabricated online reviews.

correlation neglect in our set-up, we explore group decision-making in separate treatments. Participants are placed into groups of three, with each of them receiving one signal. Before deciding collectively, group members can communicate their signals to the rest of the group. Decisions are then made by simple majority rule. We further explore differences between homogeneous and heterogeneous groups, where some members are known to have a bias toward one of the two possible decisions. This allows us to test the effect of correlated information on how truthfully participants share their signals under both common and conflicting interests. While heterogeneous preferences make truthful sharing of information impossible with independent signals, we show theoretically that under a correlated information structure truth-telling can be achieved in equilibrium.

To provide an example of collective decision-making under correlated information, consider a tenure committee evaluating a candidate or a parliamentary committee assessing a policy proposal. Each member receives signals—for example, by reading different papers of the candidate or consulting different expert reports. While the signals may appear independent, in practice they could be influenced by a common source. For instance, experts may all have been exposed to the same special interest group or lobbyist. The concern that unanimity may reflect coordinated testimony rather than independent judgment has long been recognized; the Babylonian Talmud, for example, cautions that if a defendant is unanimously convicted by all judges, they should be acquitted (Sanhedrin 17a; see Epstein and Slotki, 1961). The challenge for modern committees is similar: members must aggregate signals that may be correlated, and detecting this correlation is essential for making sound decisions.[3]

Our results show that a significant share of individuals fails to understand the correlation structure in our set-up. This bias of not accounting for the fact that too much information for one state of the world can be too good to be true also extends to groups. In fact, we find that groups are not only failing to outperform individuals in the experiment, but the bias also persists when allowing for direct communication (open chat) between group members. In principle, this set-up allows individuals who understand the underlying correlation structure to explain it to their group members and thus to improve outcomes. However, we find that participants make very limited use of this option.

Finally, we show that the combination of biased preferences and correlated signals can lead to interesting interactions. For individuals, biased preferences can mitigate the consequences of correlation neglect if the bias for one state of the world counteracts the bias in information processing. For heterogeneous groups, correlated information increases

[3]This challenge becomes even more complex in the presence of biases. For example, a biased tenure committee member might misrepresent the observed quality of a reviewed paper, thereby impeding information aggregation (as discussed in Austen-Smith and Feddersen, 2009). Our experimental design captures such scenarios by including treatments with heterogeneous preferences.

incentives for truth-telling under correlated information, and we, indeed, observe more truthful information sharing than without correlation.

The remainder of this paper is structured as follows. In Section 2, we discuss the related literature. Section 3 describes the experimental setting in which we study correlation neglect, before Section 4 derives theoretical predictions for the behavior in different treatments. In Section 5, we present our results, while Section 6 concludes.

2 Related literature

Our paper contributes both to the literature on correlation neglect, as well as to work on information aggregation in groups. Previous research has analyzed correlation neglect in various domains such as social learning (DeMarzo et al., 2003; Golub and Jackson, 2010; Brandts et al., 2015; Gagnon-Bartsch and Rabin, 2022; Bohren, 2016; Chandrasekhar et al., 2020; Dasaratha and He, 2020), financial decision-making (Hedesström et al., 2006; Kallir and Sonsino, 2009; Ellis and Piccione, 2017; Laudenbach et al., 2023; Levy and Razin, 2022), as well as voting and the formation of political opinions (Glaeser and Sunstein, 2009; Ortoleva and Snowberg, 2015; Levy and Razin, 2015; Denter et al., 2021).[4] A lot of the work on correlation neglect is theoretical or focuses on correlation in a specific domain, as for instance correlation in the return between different assets when studying financial decision making (see e.g., Laudenbach et al., 2023).

Empirical support for correlation neglect in a more general set-up comes from Enke and Zimmermann (2019), who find that individuals hold this bias even in a relatively transparent environment. Our paper differs from previous studies by focusing on a specific case of correlated information structure. The correlation structure in Enke and Zimmermann (2019), for instance, studies a setting where signals can share a common information source. We show that also in a context where too much evidence in favor of one outcome is too good to be true, individuals struggle to correctly interpret the given information. Moreover, we find that correlation neglect is highly persistent. It does not only affect individual but also group decision-making, even when allowing for direct communication between group members.

While previous research has mainly focused on individual decisions under correlated information, one recent exception is a parallel paper by Monteiro et al. (2022). The authors adapt the setting of Enke and Zimmermann (2019) to groups of two. Contrasting to our results, they find that groups outperform individuals if there is at least one rational group member. There are two key differences to our paper. First, they use a different cor-

[4] Correlation neglect has also been linked to other biases such as overprecision (Moore and Healy, 2008; Hossain and Okui, 2024) or confirmation bias (Rabin and Schrag, 1999; Andreoni and Mylovanov, 2012; Ortoleva and Snowberg, 2015).

relation structure, which could lead to differences in how individuals process information. Second, the voting rule and deliberation processes differ, as there is no straw poll stage in Monteiro et al. (2022) and decisions need to be unanimous. More research is needed to understand how these details matter for decisions under correlated information.

When it comes to the consequences of correlation neglect, a key concern is that double counting correlated information as independent can increase the likelihood of developing confident, but wrong long-run beliefs (see e.g., Eyster and Rabin, 2010, 2014) and make individuals more prone to manipulation and persuasion (Currarini et al., 2020; Levy et al., 2022). In the financial market, this can lead to unexploited opportunities for risk hedging (Eyster and Weizsäcker, 2016), wrong predictions of stock prices (Spiwoks and Bizer, 2018) and suboptimal diversification decisions (Gubaydullina et al., 2015). In the political sphere, theoretical papers have linked correlation neglect to polarization and extremism (Glaeser and Sunstein, 2009; Ortoleva and Snowberg, 2015). Misinterpreting correlated signals as new information thereby strengthens pre-existing opinions and increases confidence in own beliefs.

While most studies thus conclude that correlation neglect can lead to suboptimal decisions, Levy and Razin (2015) develop a model where it can actually improve collective decisions in the presence of ideological preferences. The key idea is that as voters suffering from this bias place a higher value on information, their choices are more informed and less based on ideology. Similarly, Hughes et al. (2023) show that heterogeneous committees that suffer from correlation neglect can outperform their rational counterparts as they are more likely to share signals truthfully. Our study provides an example of when correlated information – rather than correlation neglect – can improve outcomes, by increasing information aggregation in groups with conflicting interests.

Finally, our paper contributes to the literature on decision-making in juries and committees. According to Condorcet's jury theorem groups are expected to make better decisions than individuals, as they aggregate individual signals, which in turn reduces the likelihood of errors (see Young, 1988). However, an established finding is that voting alone can be insufficient for reaching full information aggregation if there are incentives for strategic voting (Feddersen and Pesendorfer, 1998; Guarnaschelli et al., 2000). In homogeneous juries, full information aggregation can be achieved by allowing for communication (Coughlan, 2000), which also reduces the importance of the voting rule (Gerardi and Yariv, 2007; Goeree and Yariv, 2011). If preferences are heterogeneous, the requirements for honest communication are tighter and tend to require uncertainty about preferences (Austen-Smith and Feddersen, 2006). However, Le Quement and Yokeeswaran (2015) and Le Quement and Marcin (2020) show that even if preference types are publicly known, group deliberation before the voting stage can lead to equilibria with partial truth-telling. In Le Quement and Marcin's (2020) set-up, the majority group has a bias for one state of the world, while the minority is unbiased. We adopt this case in our heterogeneous group treatment, allowing us to test their equilibrium prediction experimentally. We extend their set-up by introducing correlation and showing that under correlated information, incentives for truth-telling arise. These incentives can arise due to a second situation in which a jury member can be pivotal and in which there is no conflict of interest. In this sense our paper is related to a paper by Feddersen and Gradwohl (2020), in which multiple pivot events, in a quite different set-up to ours, also restore incentives for truthtelling.[5]

3 Experimental design and procedures

3.1 Basic set-up

In the environment in which we study correlation neglect, decision-makers have to take a binary decision under uncertainty, as in a jury setting (e.g. Feddersen and Pesendorfer, 1998; Coughlan, 2000). In the game, there are two states of the world $S \in \{\rho, \beta\}$ that are ex-ante equally likely and individuals have to guess in which state they are in and cast a vote $v \in \{\rho, \beta\}$. If their decision matches the true state of the world they receive a positive payoff π, otherwise their payoff equals zero. While the true state of the world is unknown, they receive n signals prior to their decision, which can take the value $r(ed)$ or $b(lue)$. The probability of receiving r or b depends on the underlying state of the world. Signals are imperfectly informative, such that both $Pr(r|\rho)$ and $Pr(b|\beta)$ lie between 0.5 and 1. Moreover, in the basic set-up, signals are equally accurate in both states of the world, i.e., $Pr(r|\rho) = Pr(b|\beta)$ and are drawn independently with replacement. The posterior probability for each state after receiving n signals can then be calculated using Bayes' rule.

In our experiment, the two states of the world ρ and β are represented by a blue and a red urn and signals by colored balls that are drawn from the randomly selected urn (see e.g., Guarnaschelli et al., 2000). The red urn contains three red and two blue balls, while the blue urn contains three blue and two red balls, which is equivalent to $Pr(r|\rho) = Pr(b|\beta) = 3/5$. As we explain in the following, this set-up is varied between treatments.

3.2 Treatment variations

To explore how potentially misleading information affects individual decisions as well as information aggregation and decision-making in homogeneous and heterogeneous groups, we implement a $2 \times 2 \times 2$ between-subject design. We vary the information structure

[5]More generally, our paper also adds to the empirical literature on strategic communication in committees (e.g., Hansen et al., 2018; Fehrler and Hughes, 2018; Fehrler and Janas, 2021; Fehrler and Hahn, 2023; Mattozzi and Nakaguma, 2023; Breitmoser and Valasek, 2024a; Renes and Visser, 2024).

(correlated vs. uncorrelated), the decision unit (individual vs. group), and whether individuals have a preference for a certain state of the world (bias vs. no bias). For groups the no bias treatment implies that preferences are homogeneous, while they are heterogeneous in the bias treatment. Table 1 gives an overview of the different treatments.

Note that for the cell *Gr-NoBias-Corr*, we run an additional variation where we provide participants with further communication possibilities to test whether this can mitigate correlation neglect. More details on this variation are provided in Section 3.2.2.

Table 1: Treatment overview

Individual	No bias	Bias
No correlation	*Ind-NoBias-NoCorr*	*Ind-Bias-NoCorr*
Correlation	*Ind-NoBias-Corr*	*Ind-Bias-Corr*

Group	No bias	Bias
No correlation	*Gr-NoBias-NoCorr*	*Gr-Bias-NoCorr*
Correlation	*Gr-NoBias-Corr (2)*	*Gr-Bias-Corr*

Notes: We implemented two versions of Gr-NoBias-Corr. One with and one without an additional communication option via a free-form chat.

3.2.1 Information structure

We compare decisions across two scenarios:

1. No correlation: Signals in both states of the world are drawn independently, as described in Section 3.1.

2. Correlated signals: Signals in β are drawn independently, while in ρ signals are only independent with a probability of 0.5. With probability 0.5 voters receive the b signal n times. This correlation structure is common knowledge.[6]

While in Scenario 1 a state of the world becomes more likely the more signals are received in its favor, this simple logic does no longer hold in Scenario 2, where too much evidence for one outcome can be too good to be true. More concretely, if all signals in Scenario 1 are blue, it is more likely that the true state of the world is β. If correlation is possible,

[6]In the experiment we randomize whether correlation is possible in ρ or β. Results do not depend on the color of the state in which we allow for correlation (see Appendix E). For ease of presentation, we thus recode our data such that ρ is always the state in which signal correlation may occur.

in contrast, three blue signals should raise suspicion. In fact, for $n = 3$ blue signals, it is now more likely that the true state of the world is ρ.[7] Section 4 provides a more detailed analysis of posterior probabilities and optimal voting decisions across treatments.

3.2.2 Individual versus group decision-making

While our main interest is to evaluate the effect of misleading signals on group decisions, we use the individual treatments as a benchmark for our analysis. The absence of interactions with other participants provides a simpler environment that we can use as a starting point to test for correlation neglect. In the individual treatments, each participants receives $n = 3$ signals and has to decide for either ρ or β based only on their private information.

In the case of a collective decision, participants interact in groups of three that remain fixed throughout the experiment. In contrast to the individual treatments, each group member now receives only one signal, leaving the total number of signals for a decision unchanged. Prior to the voting stage, participants can share their private signal with the other group members through a modified straw poll (Coughlan, 2000), in which each of them can send a message m from the set $\{r, b, w\}$. In addition to *red* or *blue* messages, we allow participants to not reveal any information by sending a *white* message (communicated in the experiment as "I drew a red/blue ball" and "I prefer not to reveal"). We chose to allow for this option as even in strategic settings people tend to show a strong preference for honesty (e.g., Le Quement and Marcin, 2020; Breitmoser and Valasek, 2024b). The white message thus offers participants the opportunity to hide information from group members without lying directly. After messages are exchanged, each group member votes for either ρ or β and the group decision D is determined by simple majority rule. For group treatments we thus get two outcomes: individual votes and group decisions. For individual treatments, there is no voting stage and we only measure the individual vote for ρ or β, which by definition is identical to the individual's decision ($v = D$).

While participants in the main group treatments cannot engage in any exchange beyond the predefined messages, we allow for additional communication in the *Gr-NoBias-Corr-Communication* treatment. After observing the messages of their group members and before casting their final vote, individuals are asked to send a recommendation regarding the group decision and explain their recommendation via a free-form chat. The motivation for this treatment variation is to test whether open communication can reduce the effect of correlation neglect. If only a share of participants suffers from correlation

[7] Our experiment focuses on the case of perfect correlation in order to simplify calculations for participants.

neglect, rational participants might explain the correlation structure to their fellow group members and thus improve voting outcomes.

3.2.3 Biased preferences

To account for the fact that people rarely share the exact same (common value) objective and instead often have biases towards different states, we exogenously introduce a bias in some treatments. While in treatments without a bias all individuals earn the same payoff for any correct voting decision (£0.20), in treatments with a bias, payoffs for correctly choosing ρ or β differ (as in Le Quement and Marcin, 2020). In particular, individuals receive £0.25 for a correct choice of their preferred state and £0.15 for a correct choice of the other state. Table 2 provides an overview of the payoff structure. If there is no correlation, the preferred state is chosen randomly. In cases with correlation, the bias is always in favor of β, the state where no correlation can occur.

Table 2: Payoff structure (in £)

	Bias for red		Bias for blue		No bias	
	$S=\rho$	$S=\beta$	$S=\rho$	$S=\beta$	$S=\rho$	$S=\beta$
$D=\rho$	0.25	0	0.15	0	0.2	0
$D=\beta$	0	0.15	0	0.25	0	0.2

The introduction of biased preferences is particularly interesting in the group treatments. While interests are aligned among group members without a bias, the bias allows us to look at heterogeneous groups. More specifically, groups always consist of one unbiased and two biased members who favor the same state. Preference types are publicly known to all individuals (as in Le Quement and Yokeeswaran, 2015; Le Quement and Marcin, 2020), introducing incentives for strategic communication. The group treatments thus allow us to compare the extent of truth-telling across both homogeneous and heterogeneous groups, as well as under different information structures.

3.3 Experimental procedures

We programmed the experiment using LIONESS (Giamattei et al., 2020) and collected data for all treatments online via Prolific between August and December 2020.[8] To minimize comprehension problems, we excluded non-native English speakers from the study. In each treatment, instructions (see Appendix G) were followed by control questions, and

[8]See Palan and Schitter (2018) for a discussion on the use of Prolific for scientific studies.

six incentivized rounds of the game without feedback in-between. For the group treatments we used a partner matching, where individuals remained in the same group for all six rounds.

In an ex-post questionnaire, we elicited socio-economic characteristics, risk preferences (Falk et al., 2023), flexible thinking skills (Stanovich and West, 1997), the participants' understanding of probabilities, and cognitive reflection (Frederick, 2005) to elicit possible covariates of correlation neglect. Cognitive reflection has already been shown to correlate with a wide range of other heuristics and biases (Toplak et al., 2011). Both the cognitive reflection test and the understanding of probabilities were incentivized. Finally, we asked participants an open question about the strategy they used in the main experiment and provide feedback on their payoffs.

In total, we recruited about 100 participants for each individual treatment and 300 participants (100 groups) for each group treatment, thereby keeping the number of independent observations constant across treatments. These sample sizes were pre-registered and based on power calculations. On average, participants needed 18 minutes to complete the experiment and earned £8.43 per hour. Group treatments lasted slightly longer than individual treatments (20 vs. 15 minutes), which was accounted for by a higher show-up fee.

4 Equilibrium predictions and behavioral conjectures

In this section, we describe the payoff maximizing strategies for the individual choice and the equilibrium strategies for the collective choice treatments, which we formally derive in Appendix B. We conjecture deviations from these benchmarks due to prominent behavioral biases in the literature, such as correlation neglect or the tendency to communicate truthfully even if there is an incentive to lie strategically or to stay silent. We first focus on the simpler case without biased preferences before analyzing the treatments with bias.[9]

4.1 Treatments without biased preferences

Table 3 shows the posterior probability for each state conditional on the number of red and blue signals as well as the respective optimal vote. The left part shows predictions for the case without correlation, the right part for the case with correlation in ρ. In the absence of a bias, individuals maximize their expected payoff by voting for the blue state if the updated probability for $S = \beta$ exceeds 50%, and for the red state otherwise.

[9] Note that hypotheses are restructured compared to the pre-analysis plan to improve readability. To make this transparent and to highlight that they are not all derived from the benchmark model of perfectly rational behavior, we call the hypotheses "conjectures" in the main text. Appendix A presents the hypotheses in the original wording and order.

Table 3: Posterior probability for β and ρ with and without correlation

Blue signals	No correlation			Correlation in ρ		
	$P(S=\beta)$	$P(S=\rho)$	v^*	$P(S=\beta)$	$P(S=\rho)$	v^*
0	0.23	0.77	ρ	0.37	0.63	ρ
1	0.4	0.6	ρ	0.57	0.43	β
2	0.6	0.4	β	0.75	0.25	β
3	0.77	0.23	β	0.29	0.71	ρ

Notes: v^* indicates the payoff maximising vote for an individual given the posterior probabilities for β and ρ. If the posterior probability for $S=\beta$ exceeds 50% an individual should vote β, and ρ otherwise. With correlation, signals are drawn independently in β, while in ρ individuals see 3 blue signals with 50% probability.

We can see that the optimal vote after three blue signals differs between the two information structures. While without correlation three blue signals prescribe a vote for β, they are too good to be true under correlation and should lead to a vote for ρ. In addition, there is a second — arguably more subtle — case where the optimal vote differs between the correlated and uncorrelated information structure. As the correlation in state ρ makes it highly likely to see only blue signals, already one blue signal is enough to guess that the true state of the world is β. Without correlation, in contrast, only one blue signal makes it more likely that the true state is ρ.

In line with previous research, we expect some individuals to suffer from correlation neglect and to wrongly treat signals as if they were independent (as in the state with no correlation). We thus expect more suboptimal decisions in treatments with correlation after one and three blue signals than after zero or two. This also implies that the total share of suboptimal decisions will be higher than in treatments without correlation. Comparing scenarios with and without correlation thus allows us to assess the importance of correlation neglect for accurate decision-making.

Conjecture 1: ***(Correlation neglect)*** *In treatments without a bias, correlation neglect results i) in more suboptimal decisions and votes after one and three blue signals than after zero and two blue signals (within the treatments NoBias-Corr) and ii) in more suboptimal decisions and votes in treatments with (NoBias-Corr) than in treatments without correlation (NoBias-NoCorr).*[10]

Conjecture 1 applies to both individual and group treatments. As there is no conflict of interest within groups without biased preferences, no group member can be better off by

[10] Obviously, no such errors occur with perfect rationality and selfish payoff maximization.

misreporting their private signal.[11] Group members are thus predicted to first share their private signal truthfully, and then update the beliefs about the state of the world given their own signal and the other two messages. This is true independent of the information structure. Whether the amount of suboptimal decisions is larger in individual or group treatments is an empirical question and ultimately depends on the share of correlation neglecters in the population. The majority rule implies that the group composition is crucial. If correlation neglecters are in the minority in most groups, group votes will be superior to individual decisions and vice versa if they are in the majority. If there is noise in the communication stage and not every group member reports the private signal honestly, correlated signals are harder to detect for groups even in the absence of correlation neglect.

Finally, the additional group treatment with free-form communication allows us to explore whether a group setting with open deliberation can reduce the extent of correlation neglect. If individuals do not suffer from correlation neglect, but suspect that other group members might do so, it is possible to send a message explaining the correlation structure. We thus expect to see less correlation neglect in the free-form communication treatment.

Conjecture 2: ***(Communication)*** *The share of suboptimal decisions under correlated information is lower when allowing for additional communication among group members (Gr-NoBias-Corr-Communication) compared to the same treatment without communication (Gr-NoBias-Corr).*[12]

4.2 Treatments with biased preferences

4.2.1 Individual decisions

For the individual treatments we can abstract from any strategic considerations and focus on how biased preferences change the threshold of voting for β or ρ. While the posterior probabilities remain identical to Table 3, the threshold to vote in favor of the preferred state shifts from 50% to 37.5%. Recall that in the correlated information treatments the bias is always in favor of β, the state where no correlation can occur, implying that individuals should vote B if the posterior probability for $S = \beta$ surpasses 37.5%.

[11] When group members have aligned preferences, rational members cannot benefit from misreporting their signal, assuming that all rational agents believe others to be rational as well. Appendix B.3 discusses the relaxation of this assumption.

[12] No such effect would occur with perfect rationality, selfish payoff maximization and common knowledge thereof, as no suboptimal decision would occur in the first place.

Table 4: Optimal voting decisions with and without biased preferences

Blue signals	No correlation			Correlation in ρ		
	$P(S=\beta)$	unbiased	biased	$P(S=\beta)$	unbiased	biased
0	0.23	ρ	ρ	0.37	ρ	ρ
1	0.4	ρ	β	0.57	β	β
2	0.6	β	β	0.75	β	β
3	0.77	β	β	0.29	ρ	ρ

Notes: The table shows the payoff maximising vote for unbiased and biased individuals (with a bias for β) given the posterior probabilities for β and ρ. An unbiased individual should vote β if the posterior probability for $S=\beta$ exceeds 50%, an individual with a bias for β if it exceeds 37.5%. Signals are drawn independently in β, while in ρ individuals see 3 blue signals with 50% probability.

Table 4 compares optimal decisions in the case without and with biased preferences for β. We can see that in treatments without correlation, it already becomes optimal to vote for β for biased individuals after only one blue signal. For the treatments with correlation, the optimal voting decision is unaffected by the bias.

What does this mean for correlation neglecters? Correlation neglecters are expected to behave as in the treatment without correlation. As can be seen from Table 4, this means that with biased preferences correlation neglecters would only take a suboptimal decision in one scenario (after three blue signals), as opposed to two scenarios in the case without a bias (after one and three blue signals). Biased preferences might thus dampen the negative effect of correlation on decision-making.

Conjecture 3: ***(Biased preferences)*** *i) A bias for β increases the share of β decisions in treatments without correlation (Ind-Bias-NoCorr vs Ind-NoBias-NoCorr). ii) In the bias treatments, the increase in the error rate from introducing correlation will be smaller than in the treatments without bias (difference-in-differences: Ind-Bias Corr vs NoCorr & Ind-NoBias Corr vs NoCorr).*[13]

4.2.2 Collective decisions and strategic incentives

When preferences are biased, individuals have incentives for strategic communication. Recall that in heterogeneous groups, there is always a minority of one player who is

[13]Regarding ii): No such effect would occur with perfect rationality, selfish payoff maximization, as no suboptimal decision would occur in the first place.

unbiased, while the other two players have a bias for β. As the decision is made by simple majority rule, and the biased subjects always represent the majority, biased subjects have no incentive to misreport their signal to the other group members. The unbiased player, in contrast, is the minority and can be overruled by the biased majority. For unbiased players, there can thus be an incentive to strategically misreport their signal.

We first analyze how the misaligned incentives affect equilibrium communication in the case without correlation. From Table 4 we see that an unbiased group member would vote β after two or more blue signals, while biased members already want to vote β after one blue signal. If anyone reports a blue signal, and assuming truthful communication, the majority members will thus vote β, implying a group decision for β independent of the vote of the minority member. A minority member is only pivotal when they are the only one in the group with a blue signal. If a minority member hides their blue signal, the majority members do not know whether there are one or zero blue signals. In that case the expected payoff is higher for majority members if they vote ρ. The optimal strategy for minority members is thus to always hide their signal, which can be achieved by babbling or by always sending the white message. This strategic incentive for the minority member leads to a lower level of truth-telling (truthful revelation) in heterogeneous groups than in the case without biases.

Let us now consider the treatment with correlation. From Table 4, we see that the possibility of a correlated information structure removes the conflict of interest between biased and unbiased group members. Minority members thus have no incentive to hide their signal anymore and we expect more truth-telling than in the case without correlation. While this aligned interest depends on the chosen parameters of the game, there is an additional motive for more truth-telling under correlated information. Under correlated information, minority members are not only pivotal in the case of them receiving the only blue signal, but also when everyone receives a blue signal. Not revealing their signals honestly thus poses the threat of not being able to detect correlation. In expectation, the loss from not recognizing the correlated state of the world outweighs the gains from hiding one blue signal, making babbling or hiding a suboptimal strategy for minority members.

A further important point to consider for strategic incentives is the presence of correlation neglecters. The equilibrium predictions above assume that all individuals are perfectly rational. If minority members themselves suffer from correlation neglect, they would act as if there was no correlation and decide to always hide their signal. Similarly, if minority members understand the correlation structure, but expect that some majority members suffer from correlation neglect, the second pivotal element (detecting correlation after three blue signals) becomes less important. The presence of correlation neglect might thus reduce the gap in truth-telling between the cases with and without correlation. Nevertheless, as long as not all individuals suffer from correlation neglect, the amount of truth-telling will be higher than in the treatment without correlation.

Conjecture 4: ***(Truth-telling)*** *i) Without correlation, heterogeneous groups have lower levels of truth-telling compared to homogeneous groups (Gr-NoBias-NoCorr vs Gr-Bias-NoCorr). ii) Due to changed incentives for the minority member, there is more truth-telling in heterogeneous groups with than without correlation (Gr-Bias-Corr vs Gr-Bias-NoCorr).*[14]

5 Results

In the following, we report the results for testing our conjectures. The main text contains graphs plotting the data, while the corresponding regression analyses are relegated to Appendix C.[15]

5.1 Treatments without biased preferences

5.1.1 Existence of correlation neglect

The payoff-maximizing choice for individuals in the treatments without bias and without correlation is straightforward. Individuals maximize their expected payoff by choosing the state in line with the majority of signals. In the treatments with correlation, however, the updated probabilities change such that it is optimal to choose red only in case of zero or three blue signals. Figure 1 plots the share of individuals' decisions not in line with the expected payoff maximizing option given the signals (*ex-ante* suboptimal). The figure shows that without correlation there are very few suboptimal decisions (8.8% on average). We see slightly higher error rates after one or two blue signals than after zero or three, which is consistent with a random utility model since the probabilities are closer to 50% in these cases and the differences in expected utilities are therefore smaller.

With correlation, however, the share of suboptimal decisions increases significantly (53.8% on average, $p < 0.01$), which is in line with the predictions summarized in Conjecture 1 ii). Correlation neglect predicts suboptimal decisions in situations with one or three blue signals. In line with Conjecture 1 i), we find that these situations have a higher share of suboptimal decisions than decisions after zero or two blue signals ($p < 0.01$). Both results suggest that there is a substantial share of individuals that suffer from correlation neglect and are confirmed in a regression analysis (see Table C.1 in Appendix

[14] Unlike Conjectures 1-3, Conjecture 4 is derived from an equilibrium analysis assuming perfect rational and selfish behavior and common knowledge. See Appendix B for the formal derivation and an additional analysis where these assumptions are relaxed.

[15] For the individual treatments, we report standard errors that are clustered at the individual level, for the group treatments we cluster at the group level.

Figure 1: Share suboptimal decisions in individual treatments without bias

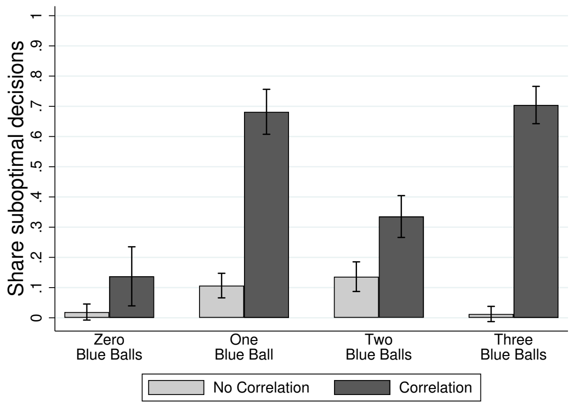

Notes: The figure shows the share of votes/decisions that are not in line with the expected payoff maximizing option given the signals (and are thus suboptimal). Comparing cases with and without correlation we see that correlation leads to a substantially higher share of suboptimal decisions, in particular after one and three signals. Whiskers represent 95% confidence intervals.

C).[16]

To assess whether participants' decisions are driven by correlation neglect rather than other forms of error such as imperfect Bayesian updating, we econometrically disentangle the relative importance of different factors in individual choices.

Following the approach of Breitmoser and Valasek (2024b), our specification allows subjects to place weight λ on the expected payoffs associated with each decision and weight κ on a simple heuristic that favors a majority of blue signals.[17] This heuristic is very intuitive and consistent with the behavior of correlation neglecters. It would also capture over-weighing of signals in the updating process as observed in Breitmoser and Valasek (2024b).

Allowing for logistic errors, this yields the probability of choosing blue conditional on the number of blue and red signals s, and Indicators for a blue majority of signals $\mathbb{I}_{s=BM}$, and a red majority of signals $\mathbb{I}_{s=RM}$:

[16] We also observe slightly higher error rates after zero and two blue signals as compared to the treatment without correlation. While this cannot be explained by random errors alone (for zero blue signals the probability for the state being blue is closer to 50% in the correlation treatment but it is further away from 50% after two blue signals), it seems that the overall higher complexity of the treatment leads to higher error rates after any signal combination. However, the error rates are much higher for the situations where correlation might play a role, which is after one or three blue signals. Note that for a correlation neglecter the (wrongly computed) expected probability of the state being blue is again closer to 50% after two blue signals than after zero, which might explain the higher error rate in this situation.

[17] Throughout the paper, we use the term correlation neglect to refer to the cognitive bias of treating correlated signals as independent. One common behavioral manifestation of this bias, observed in our data, is that subjects rely on a simple *majority heuristic*, choosing the option supported by most signals even when these may be correlated.

$$\Pr(D = \beta \mid s) = \frac{\exp\{\lambda \cdot E\Pi(D = \beta \mid s) + \kappa \cdot (\mathbb{I}_{s=BM})\}}{\exp\{\lambda \cdot E\Pi(D = \beta \mid s) + \kappa \cdot (\mathbb{I}_{s=BM})\} + \exp\{\lambda \cdot E\Pi(D = \rho \mid s) + \kappa \cdot (\mathbb{I}_{s=RM})\}} \tag{1}$$

Letting $dE\Pi(s) = E\Pi(D = \rho \mid s) - E\Pi(D = \beta \mid s)$ denote the difference in expected payoffs between choosing red and blue given the signals, this can be rearranged into the following logit regression model:

$$\Pr(D = \beta \mid s) = \frac{1}{1 + \exp\{\lambda \cdot dE\Pi(s) - 2\kappa \cdot [(\mathbb{I}_{s=BM} - 0.5)]\}} \tag{2}$$

Table 5 reports results from regressions of *'Decision $D = \beta$'* on the expected payoff difference and an indicator for whether blue is the signal majority, focusing first on treatments without bias and without correlation (columns 1 and 2, respectively). Column (3) presents results from the joint regression. Comparing the log-likelihoods and model fit values suggests that the expected payoff difference alone (λ) accounts well for the observed choice patterns. Column (3) also shows that while the coefficient on the expected payoff difference is highly significant, the coefficient on the blue majority indicator is also significant. This form of *overshooting* stems from the cases with 1 or 2 blue signals: while the expected payoff difference in these cases is close to zero, many subjects change their choices and the binary majority indicator captures this well.

In treatments with potentially correlated signals, individuals affected by correlation neglect are expected to follow the majority of signals due to wrongly computed expected payoff differences instead of the actual expected payoff differences. In contrast, rational Bayesian subjects should base their decisions solely on the actual expected payoff differences given the signals. Repeating the logistic regressions for treatments with correlated signals (columns 4–6) reveals that the majority heuristic ("Majority Blue") explains the choice patterns substantially better than the actual expected payoff difference. Not only is the coefficient on the expected payoff difference insignificant, but the AIC and BIC values also indicate better model fit when the majority heuristic is included, with no additional explanatory power from expected payoffs.

To confirm that this result is driven by the relevant *'too good to be true'* case, column (7) of Table 5 repeats the regression excluding the situation with three blue signals. In this specification, the expected payoff difference regains explanatory power, while the coefficient on the Majority Blue dummy becomes insignificant. Another way to test

Table 5: Logistic Regressions of Individual Decisions

	Ind-NoBias-NoCorr			Ind-NoBias-Corr					
							#Blue signals $\neq 3$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Exp. payoff diff.	9.383***		6.187***	-0.127		-0.041	2.690**		
	(1.113)		(1.545)	(0.256)		(0.272)	(1.092)		
Exp. payoff diff. (corr. neglect)								2.377***	1.118*
								(0.412)	(0.673)
Majority Blue		4.656***	1.498**		1.760***	1.759***	0.479		1.027**
		(0.435)	(0.728)		(0.286)	(0.285)	(0.552)		(0.490)
constant	-0.116	-2.459***	-0.867**	0.183**	-0.979***	-0.977***	-1.139***	-0.176**	-0.667**
	(0.114)	(0.256)	(0.348)	(0.072)	(0.196)	(0.195)	(0.210)	(0.080)	(0.274)
Log Likelihood	-169.122	-178.786	-167.800	-413.275	-365.980	-365.964	-232.798	-232.798	-367.185
Pseudo R^2	0.590	0.567	0.593	0.0004	0.115	0.115	0.128	0.128	0.112
AIC	342.243	361.572	341.600	830.550	735.960	737.928	471.597	738.370	733.875
BIC	351.037	370.366	354.790	839.344	744.754	751.119	483.472	747.164	747.066
N	600	600	600	600	600	600	387	600	600

Notes: Logistic regressions with *'Decision $D = \beta$'* as the dependent variable in the individual treatments without bias and without correlated signal structure (columns 1 to 3), and without bias and with correlated signal structure (columns 4 to 9). Column (7) only contains observations with 0, 1, or 2 Blue signals. Exp. payoff diff. is the (correct) expected payoff when deciding for β minus the expected payoff when deciding for ρ for a given treatment and a given number of blue signals. Exp. payoff diff. (corr. neglect) is the (potentially false) expected payoff assuming no correlation when deciding for β minus the expected payoff when deciding for ρ and a given number of blue signals. Majority Blue is a dummy variable which is 1 if the majority of signals are blue. Standard errors clustered at the subject level and depicted in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

whether individuals in our data suffer from correlation neglect is to examine whether a miscalculated expected payoff, consistent with correlation neglect, can explain variation in choices. In columns (8) and (9) of Table 5, we therefore regress the decision to choose blue in the correlated-treatment setting on the expected payoff difference from the no-correlation treatments. Comparing the model fit values of columns (4) and (8) reveals that these miscomputed expected payoffs following the neglect of correlation explain the data better than the correctly calculated probabilities in the correlation treatment. Adding the "Majority Blue" dummy in column (9) does not further improve the model fit.

The results of the logistic regression in Table 5 thus support the conclusion that, in treatments with a correlated signal structure, behavioral manifestations of correlation neglect, such as reliance on a majority heuristic or the use of wrongly computed expected payoffs, predict observed choices better than the actual expected payoff difference.[18]

To understand if certain individuals are more likely to suffer from correlation neglect, Table C.2 reports correlations between individual characteristics and suboptimal decision-making after three blue signals. Firstly, the regressions show that a misunderstanding of the game is not a likely reason for errors, as the coefficients for correctly answering the

[18] Accounting for the panel structure by including random effects at the subject level does not affect the significance of any coefficients reported in Table 5.

comprehension questions are low and far from significant. Secondly, the table shows that the level of cognitive reflection as measured by the three-item cognitive reflection test (CRT, Frederick, 2005) is a strong predictor of correlation neglecting behavior. Individuals who perform well in the CRT are highly significantly less likely to exhibit correlation neglect. Third, while a self-reported measure of critical thinking remains insignificant, performance in knowledge questions about probabilities and Bayesian updating is also a strong predictor of correlation neglect in the task.

When examining group treatments, Figure 2 shows a very similar pattern as for the individual treatments. There are far more suboptimal votes (and group decisions) in the treatments with correlation than without ($p < 0.01$). This is again driven by the situations where correlation neglect matters: in rounds with one or three blue signals. These results are confirmed by the regression analysis in Table C.1. Conjecture 1 therefore holds for both individual and group decisions. Figure 2 further shows, that in the situations where correlation neglect matters, the share of suboptimal group decisions is even higher than the share of suboptimal votes. This suggests, that the majority rule even aggravates the situation, as the high share of suboptimal votes overrule the low share of correct votes. In the group treatment, we observe a similar pattern as in the individual treatments. Groups with at least one member who achieved a perfect CRT score (27% of participants answered all three questions correctly) are weakly significantly less likely to choose blue after three blue signals ($p = 0.098$) in the treatment with potential correlation

Figure 2: Share suboptimal decisions in group treatments without bias

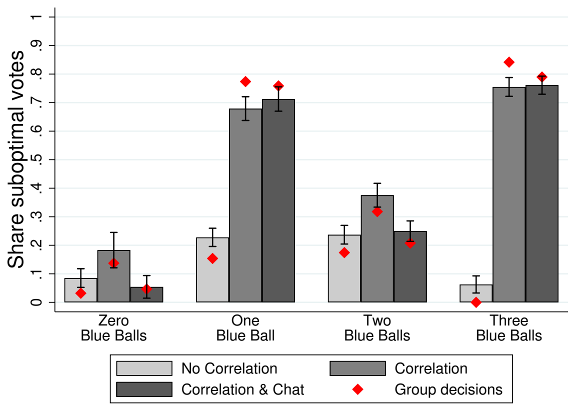

Notes: The bars show the share of group votes that are not in line with the expected payoff maximizing option given the signals (and are thus suboptimal). The red markers depict the share of suboptimal group decisions that are in the end implemented. Votes are aggregated into group decisions by simple majority rule. Comparing cases with and without correlation we see again that correlation leads to a substantially higher share of suboptimal decisions, in particular after one and three signals. This is true for both the case with and without a chat. Whiskers represent 95% confidence intervals.

and aligned preferences.[19] Section 5.1.3 analyzes the differences between individual- and group-decisions in more detail.

Finding 1: *Our results are in line with Conjecture 1. There are more suboptimal votes and decisions in treatments with correlation than in treatments without correlation. This is true for both individual and group treatments.*

5.1.2 The role of communication

To see whether communication decreases the number of suboptimal decisions, we implement a treatment with the option to send a voting recommendation together with a free-form message after the straw poll. The results for this treatment are shown by the dark bars in Figure 2. Even though, on average, the chat option leads to slightly fewer suboptimal votes compared to votes without communication ($p = 0.069$), this improvement is driven by zero and two blue signals, i.e, cases where correlation neglect plays no role. This pattern is confirmed by a regression analysis (see Table C.1). Our prediction that we state in Conjecture 2 can therefore not be confirmed.

Finding 2: *Our results are not in line with Conjecture 2. Additional communication among group members does not lead to fewer suboptimal votes or decisions in the cases relevant for correlation neglect.*

We now explore the content of the messages sent in this treatment to shed light on potential reasons why communication did not succeed in reducing the extent of correlation neglect in groups with aligned preferences.[20]

The general idea was that free-form communication, in addition to the voting recommendation, allows for an improvement in decision quality. Individuals who are no correlation neglecters can explain their reasoning to correlation neglecters, thereby reducing suboptimal group decisions. Overall, in the *Gr-NoBias-Corr-Communication* treatment, the correlation mechanism (where, with a 50% probability, all signals are blue in the red state) is only referred to in 110 out of the 1,134 cases where it should matter for the optimal decision (9.70%).[21] To see examples of how individuals explain the correlation structure to their group members see Appendix F. As subjects play in fixed group matchings for six periods, we can explore whether the mentioning of the correlation mechanism

[19] In contrast, the proportion of men in a group, the average understanding of probabilities and the average number of errors in the comprehension questions do not significantly affect group decisions in this case.

[20] The coding of all the messages was conducted by an English speaking research assistant who was not informed about any research question addressed in this paper.

[21] In general, among the 1836 messages sent, 54.19% of the messages contain some form of explanation about the probabilities or the recommendation given to the group.

improved group decision accuracy over the following periods. To do so, we construct a binary variable that takes the value one if the correlation structure was mentioned in the current or at least one of the prior periods and zero otherwise. When regressing suboptimal voting on this variable, we find that receiving an explanation about the correlation structure leads to a 4.3 percentage point decrease in suboptimal voting ($p = 0.255$). When restricting the analysis to instances where the correlation matters for the correct vote, i.e., for one or three blue signals, this number increases to 8.7 percentage points ($p = 0.064$). This finding indicates that communication is sometimes used as anticipated but that the overall number of individuals engaging in it is too low to make a detectable difference on the aggregate level. Interestingly, the open chat communication does improve decision accuracy overall ($p = 0.0275$). However, this effect mainly stems from better decisions in situations, in which the correlation structure does not change the optimal behavior (that is, after zero, or two blue signals).[22]

5.1.3 Individual vs group decisions

While we find that both individuals and groups suffer from correlation neglect, an interesting question is whether the latter do better or worse than individuals. This is particularly relevant as in practice many important decisions are taken by groups like juries or committees. Figure 3 plots the share of suboptimal decisions and votes in treatments without a bias. This allows for an overall comparison of decisions in the individual treatment, with votes and group decisions in the group treatment.

Without correlation, votes in groups are significantly worse than decisions in the individual treatment ($p < 0.01$). This can be plausibly attributed to the additional signal sharing stage in the group setting. Even when preferences are aligned, errors may arise as information about private signals must first be shared among group members before individual voting takes place. When looking at group decisions, the gap to individual decisions closes, but even group decisions keep being weakly significantly worse than individual decisions ($p < 0.1$). The improvement relative to votes is likely driven by the majority rule — mistakes by a few individuals are outweighed by the optimal choices of the majority, mitigating suboptimal voting outcomes.

For treatments with correlation, Figure 3 reveals a different pattern. While votes in group treatments without free-form communication are still slightly worse than decisions in individual treatments ($p < 0.1$), group decisions now even increase instead of mitigate this suboptimality, leading to a significant difference between individual and group decisions ($p < 0.05$). This result can be attributed to the notably high baseline level of

[22] The group-decision accuracy between the Chat and the No-Chat treatment is significantly higher for zero and two blue signals ($p < 0.1$), and not significantly different for one or three blue signals ($p > 0.1$).

Figure 3: suboptimal decisions/votes across individual versus group treatments without bias

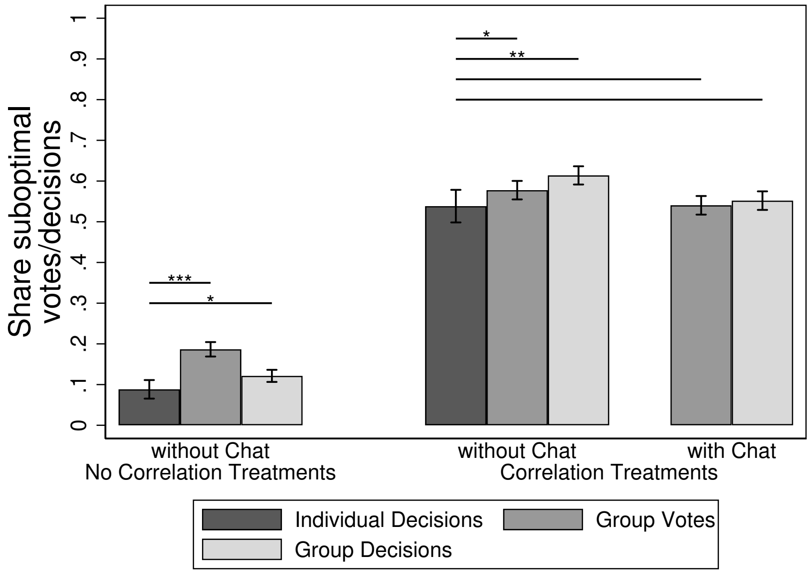

Notes: The figure compares the share of votes/decisions that are not in line with the expected payoff maximizing option given the signals (and are thus suboptimal) between individual and group treatments (both with and without a chat option). Stars indicate significance of difference between means, calculated via one-sided t-tests with clustering the standard errors on the individual/group level. *** (**/*) indicate p-values below 0.01 (0.05/ 0.1).

suboptimal decisions. With over 50% incorrect votes, the majority rule does not rectify the issue; instead, it exacerbates the magnitude of suboptimal decisions.

Finally, Figure 3 shows that if group members are allowed to communicate via chats the gap between individual decisions and group outcomes becomes insignificant. The ability to explain the correlation structure thus alleviates the additional problem associated with group decision-making described earlier. However, even with this communication option, groups still fail to outperform individuals. Table C.3 reports regression results for the overall vote and decision quality differences between individual- and group-decisions without biased preferences.

Finding 3: *Without communication, groups perform worse than individuals, both in cases with and without correlation. If groups can communicate, there is no difference between individual and group outcomes.*

To further investigate the differences in suboptimal decisions between groups and individuals, we ran an additional treatment as a robustness check.[23] We implemented a group treatment without bias, with correlation, and with the opportunity for subjects to en-

[23]We thank our referees for suggesting this treatment.

gage in free-form communication visible to all group members before making their voting decisions. The key novelty of this treatment is that all three signals within a group are publicly available and visible to each member, rather than each individual receiving only one privately observed signal. This design eliminates potential frictions in suboptimal voting that may arise due to the signal-sharing stage.

The results indicate, first, that free-form communication combined with public signals does not result in significantly better voting outcomes compared to free-form communication with private signals. This suggests that frictions in the signal-sharing stage do not have a substantial impact on decision accuracy. Second, we find no significant differences in suboptimal votes between this new treatment and the individual treatment without free-form communication, indicating that the additional free-form messaging stage does not substantially improve decision accuracy. Appendix D provides a detailed description of this treatment and its results.

5.2 Biased preferences in individual treatments

For individual decision-making, the implemented bias changes the expected payoff maximizer's decision rule for the number of blue signals needed to decide in favor of β.

In individual treatments without correlation, this change affects the case of one blue signal. Without a bias an individual should choose ρ after receiving only one blue signal, while with a bias one blue signal should already be enough to choose β. As Figure 4(a) shows, biased individuals indeed choose β significantly more often than unbiased individuals after receiving one blue signal ($p < 0.01$). As expected, there is no significant difference for any other signal combination. Table C.4 confirms that with bias, there is a significantly higher share of β decisions in no correlation treatments than without a bias. Our data thus supports the prediction we made in Conjecture 3 i). However, even though biased preferences lead to a significantly larger share of decisions for β after one blue signal, the absolute share is still very low, with a vast majority of biased individuals still deciding for ρ.

Figure 4: Individual decisions in favor of β in treatments without correlation (top) and suboptimal individual decisions in treatments without (bottom left) and with correlation (bottom right)

(a) No Correlation

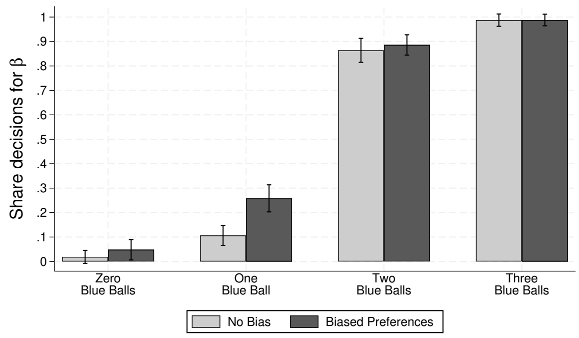

(b) No Correlation

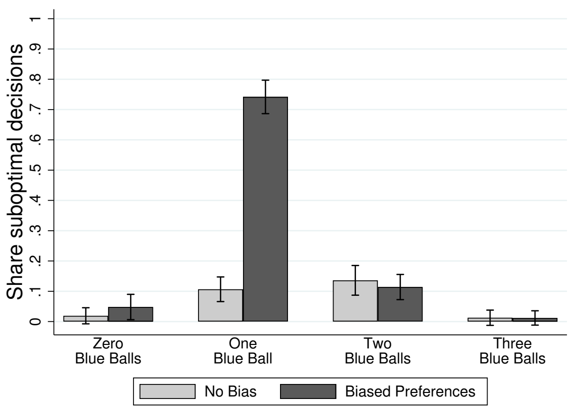

(c) Correlation in ρ

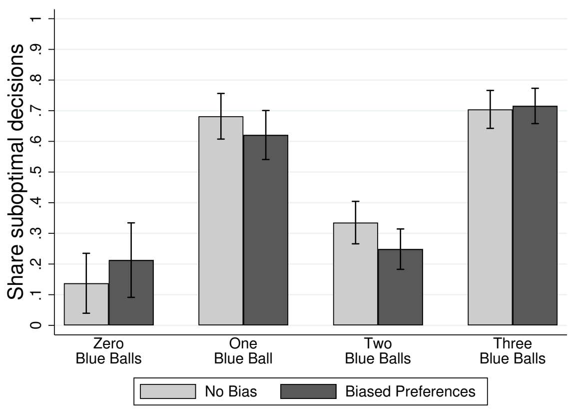

Notes: Panel (a) compares the share of decisions in favor of β between treatments with and without bias. Individuals choose β significantly more often in the case of one blue signal. Panels (b) and (c) show the share of decisions that are not in line with the expected payoff maximizing option given the signals (and are thus suboptimal). Panel (b) focuses on treatments without correlation, panel (c) on treatments with correlation in ρ. Whiskers represent 95% confidence intervals.

In individual treatments with correlation, biased preferences should not alter the decisions for rational expected payoff-maximizers. They would choose β after one blue signal independent of the bias. Individuals who suffer from correlation neglect, in contrast, are expected to choose ρ after one blue signal without biased preferences but switch to choosing β if they have biased preferences. In other words, the biased preferences make a correlation neglecter behave more similarly to an expected payoff-maximizer. We thus expect a lower share of suboptimal decisions with biased preferences. This is expected to be driven by decisions after one blue signal. The behavior after three blue signals is expected to remain unchanged. Figure 4(c) shows that the experimental data points in the direction of this prediction. There are fewer suboptimal decisions after one blue signal for biased individuals, but this difference is not statistical significant ($p > 0.1$).

The second part of Conjecture 3 predicts a difference-in-differences effect between biased and unbiased preferences when comparing treatments with and without correlation (see Figure 4(b) vs (c) and Table C.4). While both correlated signals and being biased increases the total share of suboptimal decisions, the negative effect of correlation is substantially weaker in the presence of biased preferences.

Finding 4: *Our results are in line with Conjecture 3. i) Biased preferences increase the share of decisions for β in individual treatments without correlation. ii) In the bias treatments the increase in the error rate from introducing correlation is smaller than in the treatments without bias.*

5.3 Biased preferences in groups

5.3.1 Truth-telling

In groups without a bias, the preferences of all members are aligned and in theory, when assuming rationality and common knowledge thereof, no-one can benefit by deviating from truthfully reporting the privately received signal. When introducing biased preferences in the case of no correlation, Conjecture 4 i) predicts that truth-telling will decrease. This should be driven by the unbiased minority member who can increase their expected payoff by not revealing their private signal.

Figure 5 plots the shares of messages that are not equal to the signal received (either an *w(hite)* or a *wrong* message) across treatments with and without correlation as well as in homogeneous and heterogeneous groups. For heterogeneous groups, we moreover distinguish between majority (biased) and minority (unbiased) group members. In line with Conjecture 4 i), the figure and regression Table C.5 show that in treatments without correlation there are more dishonest messages in heterogeneous than in homogeneous groups ($p < 0.05$). This is true even though the proportion of dishonest reports is quite low even among the unbiased minority, who we predicted would not report their signals truthfully in equilibrium.

From Conjecture 4 ii) we moreover expect that for heterogeneous groups, correlation decreases dishonest reporting from unbiased minority members relative to the case without correlation. Intuitively, introducing correlation re-aligns incentives between the unbiased and biased group members.[24] Figure 5 confirms a significant difference in the share of dishonest messages between the correlation and the no-correlation treatments for

[24] An alternative explanation for why the combination of correlated signals and biased preferences could increase truth-telling is the larger level of complexity in this treatment. An intuitive reaction to this complexity could be to avoid strategic considerations and instead use the heuristic of always telling the truth.

Figure 5: Share of dishonest messages in the group treatments

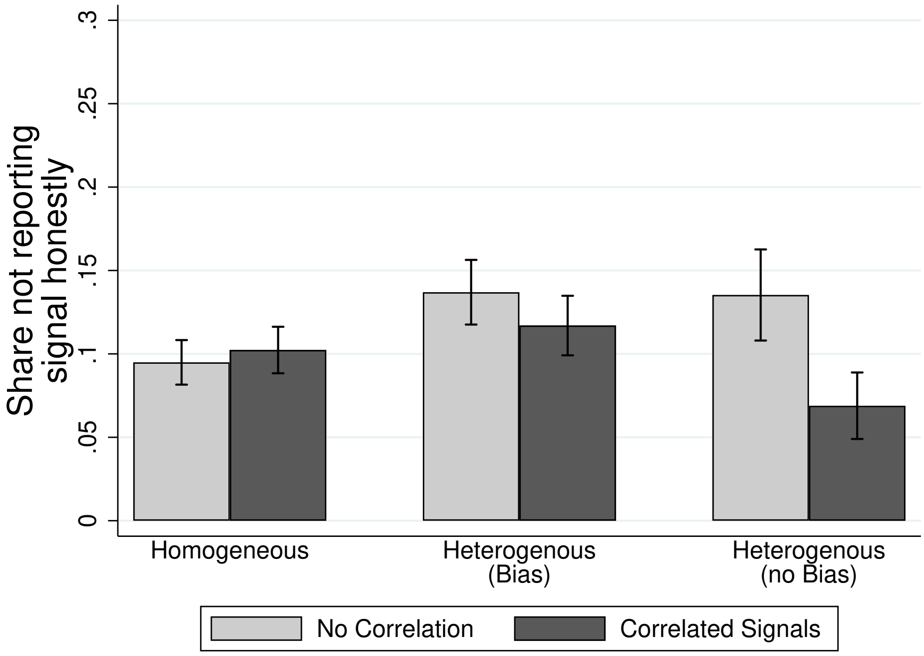

Notes: The figure shows the share of individuals in group treatments that do not report their signal truthfully, comparing treatments with and without correlation. Groups are either homogeneous or heterogeneous. In the homogeneous group treatments, preferences are always unbiased. In the heterogeneous group treatments, majority members have a bias in favor of β, minority members are unbiased. We thus report results separately for individuals with and without a bias. Whiskers represent 95% confidence intervals.

unbiased members ($p < 0.01$). In contrast, there is no statistically significant difference for the biased majority members and members in homogeneous groups (without biases) ($p > 0.1$, see regressions in Table C.5). We thus confirm that, while correlated information structures themselves can be harmful for decision-quality, these structures can also have positive side-effects on information aggregation in heterogeneous groups. The increase in honest reporting is, however, much smaller than expected from our equilibrium analysis.

Finding 5: *Our results are in line with Conjecture 4. i) Biased preferences lead to more dishonesty in groups. ii) Correlated signals increase honesty in heterogeneous groups. This increase in honest messages is driven by the (unbiased) minority individuals.*

To examine how the increase in honest reporting can influence decision quality, we compare the share of suboptimal votes of individuals who received two honest messages with the share of suboptimal votes of individuals who received at least one dishonest or uninformative message. Combining all group treatments, individuals who received two honest messages vote on average suboptimally in 42.42% of cases, while individuals who did not receive two honest messages voted suboptimally in 48.08% of cases. This difference is statistically significant ($p = 0.029$) and underscores how improved information aggregation can translate into better decision-making.

5.3.2 Comparing voting accuracy in homogeneous and heterogeneous groups

Finally, we explore voting accuracy within homogeneous and heterogeneous groups. From Figure 2 we already know that for homogeneous groups, voting accuracy is significantly lower in correlated compared to uncorrelated treatments, and that this difference is mainly driven by the scenarios with one and three blue signals. Here, we explore whether there are systematic differences between homogeneous and heterogeneous groups, as well as between minority (unbiased) and majority (biased) group members within the latter.

Figure 6 depicts the share of suboptimal votes across all group treatments by number of blue signals. As we can see, after zero and two blue signals, where neither correlation nor group heterogeneity should matter, differences between treatments are not very pronounced, even though correlation already slightly decreases voting accuracy in case of two blue signals ($p < .1$ for each of the three pairwise comparisons).

Figure 6: Share of suboptimal votes within groups

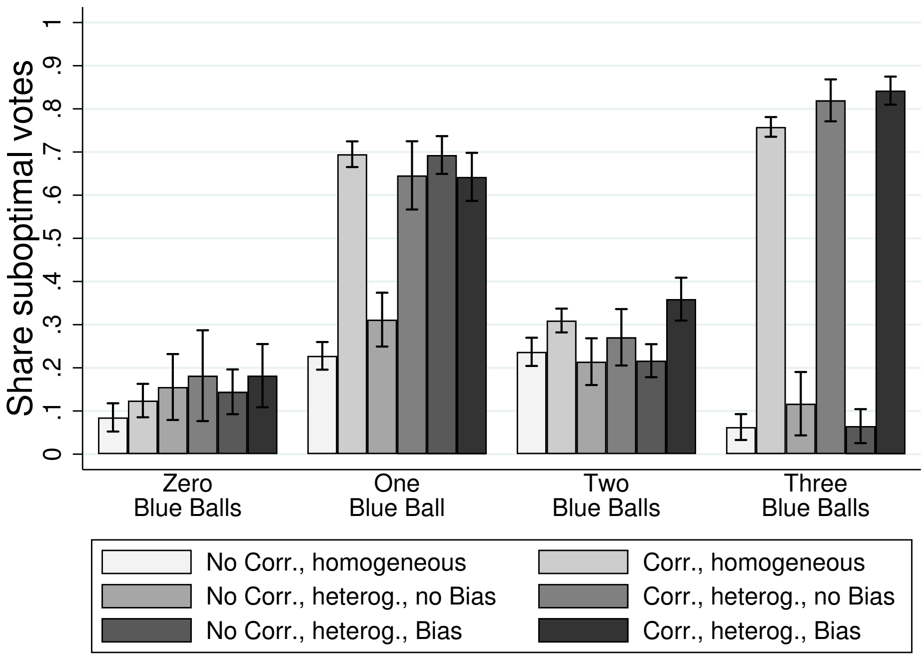

Notes: The figure shows the share of votes within groups that are not in line with the expected payoff maximizing option given the signals (and are thus suboptimal), comparing treatments with and without correlation. Groups are either homogeneous or heterogeneous. In the homogeneous group treatments, preferences are always unbiased. In the heterogeneous group treatments, majority members have a bias in favor of β, minority members are unbiased. We thus report results separately for individuals with and without a bias. Whiskers represent 95% confidence intervals.

After one and three blue signals, we see the strong effect of correlation on voting accuracy in homogeneous groups that we already saw in Figure 2. After three blue signals, the same pattern holds for all members of heterogeneous groups and is even slightly more pronounced. After one blue signal, in contrast, minority and majority members show different reactions to correlation. Note that this is also the scenario for which incentives between minority and majority members are misaligned. We find that the biased preferences of majority members already lead to a high share of suboptimal votes without correlation that is similar to the one with correlation. For the unbiased minority members on the other hand, correlation increases suboptimal votes just as in homogeneous groups.[25] Considering group decisions instead of group votes leads to the same conclusions (see Figure E.2 in Appendix E).

Overall, we thus find that heterogeneous groups perform worse than homogeneous groups. Within heterogeneous groups, we again find that voting accuracy is lower in the treatment with correlation compared to the treatment without correlation. This is true even though we showed above that correlation increases truth-telling within heterogeneous groups. The higher incentives for truth-telling do not offset the negative effect of the increased decision complexity in correlated information structures.

6 Discussion and conclusion

For centuries, scholars have pondered scenarios where overwhelming evidence is suspect. Today many of us are confronted with fabricated customer reviews, fake news and information bubbles on a daily basis – all situations in which the signals about some state of the world might be misleading. The difficulty in making accurate decisions in such information structures affects both the realms of individual decisions and that of collective choices.

We study a setting in which correlated signals occur in one state of the world but not in the other. Consider fabricated reviews for products from dishonest sellers, contrasted with independent reviews from honest sellers, or manipulated evidence from criminal and guilty defendants, contrasted with independent evidence from honest and innocent defendants, and so forth. In other scenarios, correlated signals might occur in both states of the world. For example, this could be due to faulty signal production. Instead of indicating "guilt" or "innocence" or "positive" or "negative," as it might appear at face value, these signals may actually indicate a measurement problem. Think of a large number of Covid tests for a test center that all indicate "positive" when used. Rather than concluding that everybody in town is ill, a much more likely correct conclusion would be that the tests are faulty. An early awareness of this problem can be found in the Babylonian Talmud which

[25] t-tests reveal $p < 0.01$ for each of the three pairwise comparisons with three blue signals, and $p < .01$ for two of the pairwise comparisons with one blue signal. The difference in suboptimal voting between treatments with correlation and treatments without correlation for biased group-members of heterogeneous groups is not significant.

includes a rule stating that a defendant should not be convicted if all judges unanimously believe in their guilt. It seems reasonable to interpret this rule as based on suspicion of faulty evidence collection. What initially appears to be overwhelming evidence might, if interpreted correctly, actually reveal flaws in the evidence itself.

In addition to the increased difficulty in making rational choices, such scenarios (with possible correlation in one or more states) can give rise to a rather subtle game-theoretic prediction: more truthful communication in heterogeneous groups. This effect arises partly from the common interest in detecting correlation and shows a new mechanism through which incentives for truth-telling can be restored among jury members (see also Feddersen and Gradwohl, 2020). While we do discover evidence of this subtle effect, its magnitude is small. However, what looms large are the error rates in both individual and collective decision-making when detecting correlation becomes crucial for rational decision-making. These results align with earlier findings in the literature on correlation neglect and underscore how easily people can be misled in relatively straightforward situations (see e.g., Keren and Lewis, 1994; Enke and Zimmermann, 2019; Moser and Wallmeier, 2021). Building on this research, we show that this bias is not only present in individual, but also in group decisions. Our results thus highlight the importance of correlation neglect across different settings.

Without the opportunity to communicate freely, groups perform even worse than individuals in our experiment. Once participants are allowed to use a chat interface, performance improves but does not exceed that of individuals. These results are comparable to findings from research on the hidden-profile paradigm in social psychology (see Lightle et al., 2009; Sohrab et al., 2015, for an overview). What is consistently found in these studies is that groups have a very hard time to find the hidden correct answer (state of the world), as this requires discounting redundant pieces of information (signals). Communication often does not help much in these settings either.

More education on how to rationally update beliefs in the presence of potentially misleading signals, or at least efforts to increase awareness of such situations, appear to be the most obvious policy conclusions from our findings. A natural next step for research would be to endogenize information acquisition and study how many signals people want to buy to detect possible correlation. Further, it could be interesting to study how people respond to strategic rather than exogenous manipulation of signals. We leave these important questions for future endeavors.

References

Andreoni, J. and T. Mylovanov (2012). Diverging opinions. *American Economic Journal: Microeconomics 4*(1), 209–232.

Austen-Smith, D. and T. J. Feddersen (2006). Deliberation, preference uncertainty, and voting rules. *American Political Science Review 100*(2), 209–217.

Austen-Smith, D. and T. J. Feddersen (2009). Information aggregation and communication in committees. *Philosophical Transactions of the Royal Society of London B: Biological Sciences 364*(1518), 763–769.

Bohren, J. A. (2016). Informational herding with model misspecification. *Journal of Economic Theory 163*, 222–247.

Bouton, L., M. Castanheira, and A. Llorente-Saguer (2016). Divided majority and information aggregation: Theory and experiment. *Journal of Public Economics 134*, 114–128.

Bouton, L., A. Llorente-Saguer, and F. Malherbe (2018). Get rid of unanimity rule: The superiority of majority rules with veto power. *Journal of Political Economy 126*(1), 107–149.

Brandts, J., A. E. Giritligil, and R. A. Weber (2015). An experimental study of persuasion bias and social influence in networks. *European Economic Review 80*, 214–229.

Breitmoser, Y. and J. Valasek (2024a). Strategic communication in committees with expressive payoffs. *The RAND Journal of Economics*.

Breitmoser, Y. and J. Valasek (2024b). Why do committees work? *NHH Dept. of Economics Discussion Paper*.

Cason, T. N., T. Sharma, and R. Vadovič (2020). Correlated beliefs: Predicting outcomes in 2× 2 games. *Games and Economic Behavior 122*, 256–276.

Chandrasekhar, A. G., H. Larreguy, and J. P. Xandri (2020). Testing models of social learning on networks: Evidence from two experiments. *Econometrica 88*(1), 1–32.

Collinger, T. (2016). You don't want a 5-star review. *Northwestern*, available at: https://news.northwestern.edu/stories/2016/01/opinion–techcrunch–reviews.

Coughlan, P. J. (2000). In defense of unanimous jury verdicts: Mistrials, communication, and strategic voting. *American Political Science Review 94*(02), 375–393.

Currarini, S., G. Ursino, and A. K. Chand (2020). Strategic transmission of correlated information. *Economic Journal 130*(631), 2175–2206.

Dasaratha, K. and K. He (2020). Network structure and naive sequential learning. *Theoretical Economics 15*(2), 415–444.

Demange, G. (2018). Contagion in financial networks: a threat index. *Management Science 64*(2), 955–970.

DeMarzo, P. M., D. Vayanos, and J. Zwiebel (2003). Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics 118*(3), 909–968.

Denter, P., M. Dumav, and B. Ginzburg (2021). Social connectivity, media bias, and correlation neglect. *The Economic Journal 131*(637), 2033–2057.

Ellis, A. and M. Piccione (2017). Correlation misperception in choice. *American Economic Review 107*(4), 1264–92.

Enke, B. and F. Zimmermann (2019). Correlation neglect in belief formation. *The Review of Economic Studies 86*(1), 313–332.

Epstein, I. and J. J. Slotki (1961). *The Babylonian Talmud: Seder Nezikin*, Volume 13. Soncino Press.

Eyster, E. and M. Rabin (2010). Naive herding in rich-information settings. *American Economic Journal: Microeconomics 2*(4), 221–243.

Eyster, E. and M. Rabin (2014). Extensive imitation is irrational and harmful. *The Quarterly Journal of Economics 129*(4), 1861–1898.

Eyster, E. and G. Weizsäcker (2016). Correlation neglect in portfolio choice: Lab evidence. *Available at SSRN 2914526*.

Falk, A., A. Becker, T. Dohmen, D. Huffman, and U. Sunde (2023). The preference survey module: A validated instrument for measuring risk, time, and social preferences. *Management Science 69*(4), 1935–1950.

Feddersen, T. and R. Gradwohl (2020). Decentralized advice. *European Journal of Political Economy 63*, 101871.

Feddersen, T. and W. Pesendorfer (1998). Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting. *American Political Science Review 92*(01), 23–35.

Fehrler, S. and V. Hahn (2023). Committee decision-making under the threat of leaks. *The Journal of Politics 85*, 1107–1122.

Fehrler, S. and N. Hughes (2018). How transparency kills information aggregation: theory and experiment. *American Economic Journal: Microeconomics 10*(1), 181–209.

Fehrler, S. and M. Janas (2021). Delegation to a group. *Management Science 67*(6), 3714–3743.

Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives 19*(4), 25–42.

Gagnon-Bartsch, T. and M. Rabin (2022). Naive social learning, mislearning, and unlearning. *Mimeo*.

Gerardi, D. and L. Yariv (2007). Deliberative voting. *Journal of Economic Theory 134*(1), 317–338.

Giamattei, M., K. S. Yahosseini, S. Gächter, and L. Molleman (2020). Lioness lab: a free web-based platform for conducting interactive experiments online. *Journal of the Economic Science Association 6*(1), 95–111.

Glaeser, E. L. and C. R. Sunstein (2009). Extremism and social learning. *Journal of Legal Analysis 1*(1), 263–324.

Goeree, J. K. and L. Yariv (2011). An experimental study of collective deliberation. *Econometrica 79*(3), 893–921.

Golub, B. and M. O. Jackson (2010). Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics 2*(1), 112–149.

Guarnaschelli, S., R. D. McKelvey, and T. R. Palfrey (2000). An experimental study of jury decision rules. *American Political Science Review 94*(2), 407–423.

Gubaydullina, Z., M. Spiwoks, et al. (2015). Correlation neglect, naïve diversification, and irrelevant information as stumbling blocks for optimal diversification. *Journal of Finance and Investment Analysis 4*(2), 1–19.

Gunn, L. J., F. Chapeau-Blondeau, M. D. McDonnell, B. R. Davis, A. Allison, and D. Abbott (2016). Too good to be true: when overwhelming evidence fails to convince. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 472*(2187), 20150748.

Hansen, S., M. McMahon, and A. Prat (2018). Transparency and deliberation within the fomc: A computational linguistics approach. *The Quarterly Journal of Economics 133*, 801–870.

Hedesström, T. M., H. Svedsäter, and T. Gärling (2006). Covariation neglect among novice investors. *Journal of Experimental Psychology: Applied 12*(3), 155.

Hossain, T. and R. Okui (2024). Belief formation under signal correlation. *Games and Economic Behavior 146*, 160–183.

Hughes, N., Z. U. Khan, and F. Mengel (2023). Diversity in committees. *SSRN Working Paper 4495035*.

Kallir, I. and D. Sonsino (2009). The neglect of correlation in allocation decisions. *Southern Economic Journal 75*(4), 1045–1066.

Keren, G. and C. Lewis (1994). The two fallacies of gamblers: Type i and type ii. *Organizational Behavior and Human Decision Processes 60*(1), 75–89.

Kremer, M., E. Siemsen, and D. J. Thomas (2016). The sum and its parts: Judgmental hierarchical forecasting. *Management Science 62*(9), 2745–2764.

Laudenbach, C., M. Ungeheuer, and M. Weber (2023). How to alleviate correlation neglect in investment decisions. *Management Science 69*(6), 3400–3414.

Le Quement, M. T. and I. Marcin (2020). Communication and voting in heterogeneous committees: An experimental study. *Journal of Economic Behavior and Organization 174*, 449–468.

Le Quement, M. T. and V. Yokeeswaran (2015). Subgroup deliberation and voting. *Social Choice and Welfare 45*(1), 155–186.

Levy, G., I. M. d. Barreda, and R. Razin (2022). Persuasion with correlation neglect: a full manipulation result. *American Economic Review: Insights 4*(1), 123–138.

Levy, G. and R. Razin (2015). Correlation neglect, voting behavior, and information aggregation. *American Economic Review 105*(4), 1634–1645.

Levy, G. and R. Razin (2022). Combining forecasts in the presence of ambiguity over correlation structures. *Journal of Economic Theory 199*, 105075.

Lightle, J. P., J. H. Kagel, and H. R. Arkes (2009). Information exchange in group decision making: The hidden profile problem reconsidered. *Management Science 55*(4), 568–581.

Luca, M. and G. Zervas (2016). Fake it till you make it: Reputation, competition, and yelp review fraud. *Management Science 62*(12), 3412–3427.

Maslowska, E., E. C. Malthouse, and S. F. Bernritter (2017). Too good to be true: the role of online reviews' features in probability to buy. *International Journal of Advertising 36*(1), 142–163.

Mattozzi, A. and M. Y. Nakaguma (2023). Public versus secret voting in committees. *Journal of the European Economic Association 21*, 907–940.

Mayzlin, D., Y. Dover, and J. Chevalier (2014). Promotional reviews: An empirical investigation of online review manipulation. *American Economic Review 104*(8), 2421–2455.

Monteiro, S., M. Sutter, and M. Praxmarer (2022). Correlation neglect in belief formation: Team versus individual decision-making. *PhD dissertation chapter, University of Cologne*.

Moore, D. A. and P. J. Healy (2008). The trouble with overconfidence. *Psychological Review 115*(2), 502–517.

Moser, J. and N. Wallmeier (2021). Correlation neglect in voting decisions: An experiment. *Economics Letters 198*, 109656.

Ortoleva, P. and E. Snowberg (2015). Overconfidence in political behavior. *The American Economic Review 105*(2), 504–535.

Palan, S. and C. Schitter (2018). Prolific.ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance 17*, 22–27.

Rabin, M. and J. L. Schrag (1999). First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics 114*(1), 37–82.

Renes, S. and B. Visser (2024). Strategic expert committees and the markets that assess them: A laboratory experiment. *Management Science*.

Skillings, J. (2022). When is 4.8 greater than 5.0? the world of fake reviews. *YouCanBeFound*, available at: https://www.youcanbefound.com/when–is–4–8–greater–than–5–0–the–world–of–fake–reviews/.

Smirnov, A. and E. Starkov (2022). Bad news turned good: reversal under censorship. *American Economic Journal: Microeconomics 14*(2), 506–560.

Sohrab, S. G., M. J. Waller, and S. Kaplan (2015). Exploring the hidden-profile paradigm: A literature review and analysis. *Small Group Research 46*(5), 489–535.

Spiwoks, M. and K. Bizer (2018). Correlation neglect and overconfidence an experimental study. *Journal of Applied Finance and Banking 8*(3), 75–86.

Stanovich, K. E. and R. F. West (1997). Reasoning independently of prior belief and individual differences in actively open-minded thinking. *Journal of Educational Psychology 89*(2), 342.

Stasser, G. and W. Titus (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology 48*(6), 1467–1478.

Toplak, M. E., R. F. West, and K. E. Stanovich (2011). The Cognitive Reflection Test as a predictor of performance on heuristics-and-biases tasks. *Memory & Cognition 39*(7), 1275–1289.

Young, H. P. (1988). Condorcet's theory of voting. *American Political Science Review 82*(4), 1231–1244.

Appendix

A Pre-registered hypotheses

Below are the pre-registered hypotheses for the experiment. In the main text, we summarize H1, H2, H6 and H8 in Conjecture 1 (Correlation Neglect), focusing on comparing the suboptimality of decisions across treatments with and without correlation. Conjecture 2 in the main text on the effect of communication in groups is equivalent to H9. When it comes to biased preferences, Conjecture 3 in the main text is derived from H3, H6, and H7. Finally, Conjecture 4 (truth-telling in groups) in the main text combines H4 and H5.

H1 (*Individual Treatments*) The share of subjects deciding for blue after observing three blue signals will be higher in the *NoCorrelation* treatments than in the *Correlation* treatments.

H2 (*Individual Treatments*) The share of unbiased subjects deciding for blue after observing one blue signal will be higher in the *Correlation* treatments than in the *NoCorrelation* treatments.

H3 (*Individual Treatments*) In the *NoCorrelation* treatments the share of unbiased subjects deciding for blue after observing one blue signal will be lower than the share of biased subjects deciding for blue after observing one blue signal.

H4 (*Group Treatments*) In the *NoCorrelation* treatments the share of unbiased subjects who honestly report the signal they observed will be higher in the *NoBias* treatments than in the *Bias* treatments.

H5 (*Group Treatments*) In the *Bias* treatments the share of unbiased subjects who honestly report the signal they observed will be higher in the *Correlation* treatments than in the *NoCorrelation* treatments.

H6 (*Individual no Bias Treatment with Correlation*) The share of wrong decisions will be higher after either one or three blue signals than with zero or two blue signals when correlation can only occur in the red state.

H7 (*Individual Bias Treatment with Correlation*) The share of wrong decisions will be higher after three blue signals than with fewer blue signals when correlation can only occur in the red state.

H8 (*Correlation Neglect – Individual Treatments*) The share of wrong decisions is higher in Individual with Correlation than in the Individual Treatments without Correlation in the situations with three red signals.

H9 (*Free form communication*) The share of correct decisions is higher in the *Gr-NoBias-Corr-Message* treatment than in the *Gr-NoBias-Corr* treatment in the situations with two or three red signals.

B Equilibrium Analysis

In parts B.1 and B.2, we present the equilibrium predictions for the group treatments under the (standard) assumptions of full rationality, selfish expected payoff maximization and common knowledge thereof. We relax this assumption for a subset of (correlation-neglecting) players in part B.3.

Recall that after receiving their individual signal $s \in \{r, b\}$ about the state of the world $S \in \{\rho, \beta\}$, group members first send a message $m \in \{r, b, w\}$, then observe the other group members' messages and finally cast a vote $v \in \{\rho, \beta\}$. The group decision $D \in \{\rho, \beta\}$ is determined by simple majority rule.

The solution concept is Perfect Bayesian Equilibrium with the following refinements: i) robustness to trembles at the voting stage (to rule out implausible equilibria where nobody is pivotal and votes for the option with the lower expected payoff), ii) highest possible number of group members communicate truthfully (to rule out babbling equilibria, when information can be shared), iii) no inverted language (to rule out equilibria where "b" means "r").

B.1 No-Bias Treatments

Proposition 1: ***(No-Bias Treatments)*** *Group members will i) truthfully reveal their signal in the straw poll, and ii) vote for the color that matches the true state of the world with the highest probability given the information shared in the straw poll.*

Proof. We begin with the voting stage. In case of pivotality, it is optimal to vote for the color that matches the true state of the world with the highest probability, as the reward for a correct decision is the same in either state, but with the described voting behavior nobody is ever pivotal. However, all other voting behaviors that never lead to anyone being pivotal (such as, always voting for β) are ruled out by refinement i), which re-introduces a small probability of pivotality. Now we turn to the straw poll. Given that every other player believes the messages and votes as described, it is optimal to tell the truth as all players share the same preferences. Given that everybody tells the truth it is optimal to believe the messages. $\square$

B.2 Bias Treatments (and Conjecture 4)

Proposition 2: ***(Bias Treatment without Correlation)*** *In equilibrium, i) the majority group members will truthfully reveal their signal in the straw poll, while the minority member will not share any information, and ii) the majority members will vote for β if they reported at least one b signal in the straw poll (ignoring the minority member's message) and else for ρ, while the minority subject will vote for the color that is indicated by the majority of the two reported signals of the majority subjects and its own signal.*

Proof. We begin by showing that there is no equilibrium in which all group members tell the truth. If every member always told the truth and then voted according to the information contained in all three signals, the minority member's message would be pivotal when their signal was b and the two majority members' signals r. However, in this situation the minority member would benefit from sending the wrong message r. Hence, no equilibrium exists with all three members telling the truth.

Next, we show that the behaviors described in the proposition are part of an equilibrium with two members telling the truth. We begin with the voting stage. The probability that the state of the world is β is $\approx 31\%$ after two (truthfully reported) r signals, and hence below the threshold of 37.5% above which a biased subject would want to vote for β. After one (truthfully reported) b and one (truthfully reported) r signal or two (truthfully reported) r signals the probability that the state of the world is β is 50% and hence above the threshold. The minority member is never pivotal if the majority members vote in the same way. However, refinement i) makes voting as described the only option that survives. The same holds for the described majority subjects' voting behavior in the situations where they are not pivotal without the refinement. We now turn to the straw poll. As the minority subject's message is ignored, babbling is a best response. As a majority member, you are pivotal in the straw poll when the other majority member has an r signal and you have a b signal. In this case, falsely reporting r will lead to a group decision for ρ with positive probability (with certainty if you also vote for ρ), which is clearly worse than telling the truth and ensuring a group decision for β. $\square$

Proposition 3: ***(Bias Treatment with Correlation)*** *Group members will i) truthfully reveal their signal in the straw poll, and ii) vote for the color that gives the highest expected payoff given the information shared in the straw poll.*

Proof. As can be easily verified, there is no conflict of interest regarding the correct group decision after any number of b signals. Hence, there is also no incentive to lie or withhold information in the straw poll or vote differently than described in the proposition. $\square$

Conjecture 4: *i) Without correlation, heterogeneous groups have lower levels of truth-telling compared to homogeneous groups. ii) Due to changed incentives for the minority member, there is more truth-telling in heterogeneous groups with than without correlation.*

Proof. This follows immediately from Propositions 1-3. ☐

B.3 The role of beliefs about correlation neglecters

Sections B.1 and B.2 derive predictions under the assumption that rational agents believe all other group members are rational. We next discuss how relaxing this assumption affects the existence of a truth-telling equilibrium. The following schematic diagram shows the different stages of the game.

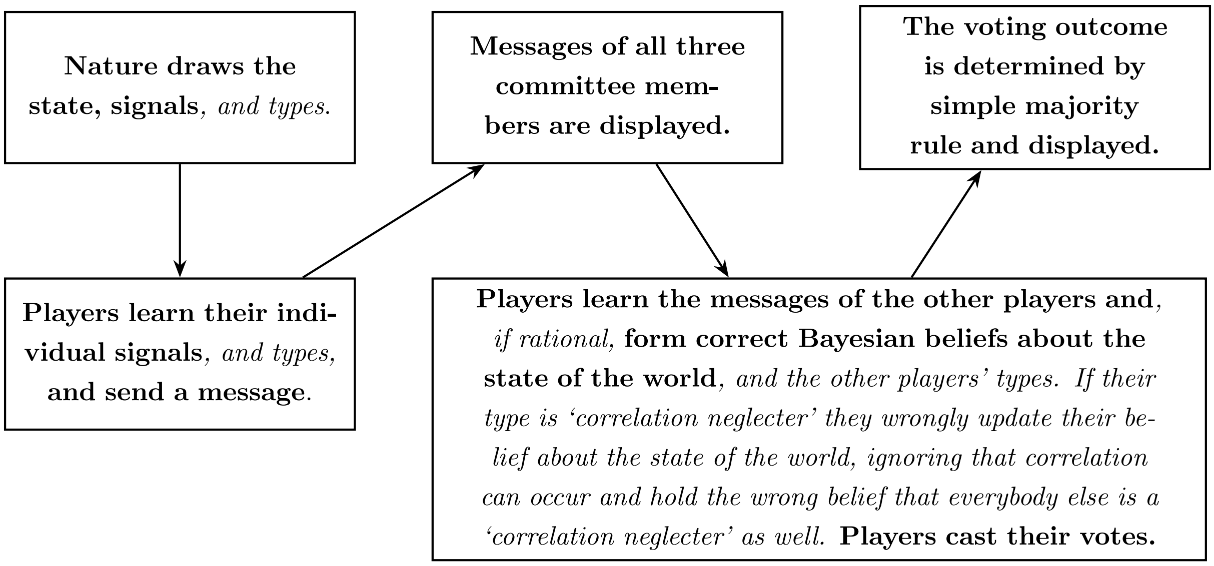

Notes: The text in *italics* is what is new in the setup when considering rational and correlation-neglecting types of players. The setup with standard rationality assumptions is in **bold** text.

No-Bias Treatments When group members have equal payoffs and there is no possibility of a correlated signal structure, both rational members and those exhibiting correlation neglect prefer the group decision to be β whenever two or three signals are b. The absence of misaligned preferences makes truthful reporting a dominant strategy for each member. However, if signals can be correlated, the preferred group decision for a given signal distribution may differ between rational members and those with correlation neglect. We assume that individuals with correlation neglect act as if a correlated signal structure cannot occur, and vote for the state that seems more likely with the wrongly computed probabilities. This is the state indicated by the majority of signals. In contrast, rational group members prefer the group decision to be β with one or two b signals and ρ otherwise. If rational members assume all others are also rational, truthfully reporting private signals is the dominant strategy. However, if they believe that some group members may exhibit correlation neglect with positive probability, truthful reporting is not necessarily an equilibrium strategy anymore.

Let us assume each rational group member believes another group member exhibits correlation neglect with a prior probability $0 \leq \alpha \leq 1$, and that individuals exhibiting correlation neglect truthfully report their signals and assume all received messages about signals are true.[26] When privately observing a b signal, a rational group member considers that the other group members may have zero, one, or two b signals. If both other members exhibit correlation neglect, misreporting one's signal as r changes the group outcome only if the other members hold one r and one b signal. In this case, however, both the rational member and the members exhibiting correlation neglect prefer the group decision to be β, making misreporting unprofitable. Thus, for any $0 \leq \alpha \leq 1$, it holds that $EU_i(m_i = r \mid s_i = b) < EU_i(m_i = b \mid s_i = b)$.

However, if a rational group member privately observes a r signal, a higher α may affect the existence of an equilibrium with truthful reporting. When one of the other members has a b signal, the rational member prefers the group decision to be β, whereas those exhibiting correlation neglect would vote for ρ if only one b message is observed, creating an incentive to misreport. Conversely, if both other members have a b signal, misreporting as b risks shifting the group decision to ρ if the other members are rational—despite β being the preferred outcome given the true signal structure. To find the threshold value of α that leads rational group members to prefer misreporting their r signal, we calculate $E\Pi(m_i = r \mid s_i = r, \alpha)$ and $E\Pi(m_i = b \mid s_i = r, \alpha)$. As we focus on the truth-telling aspect of the equilibria, the following equations assume optimal voting given the true information about one's own signal and the belief that other messages are truthful.[27] For individual i in a group with members j and k, the expected utilities are as follows.

$$\begin{aligned}
E\Pi(m_i = r \mid s_i = r, \alpha) = {} & \Pr(s_j = r, s_k = r \mid s_i = r) \cdot E\Pi(D = \rho \mid s_i = r, s_j = r, s_k = r) \\
& + \Pr(s_j = r, s_k = b \mid s_i = r) \cdot \big[\alpha^2 E\Pi(D = \rho \mid s_i = r, s_j = r, s_k = b) \\
& \quad + 2\alpha(1-\alpha) E\Pi(D = \beta \mid s_i = r, s_j = r, s_k = b) \\
& \quad + (1-\alpha)^2 E\Pi(D = \beta \mid s_i = r, s_j = r, s_k = b)\big] \\
& + \Pr(s_j = b, s_k = r \mid s_i = r) \cdot \big[\alpha^2 E\Pi(D = \rho \mid s_i = r, s_j = b, s_k = r) \\
& \quad + 2\alpha(1-\alpha) E\Pi(D = \beta \mid s_i = r, s_j = b, s_k = r) \\
& \quad + (1-\alpha)^2 E\Pi(D = \beta \mid s_i = r, s_j = b, s_k = r)\big] \\
& + \Pr(s_j = b, s_k = b \mid s_i = r) \cdot E\Pi(D = \beta \mid s_i = r, s_j = b, s_k = b) \qquad (3)
\end{aligned}$$

[26]By assumption, individuals suffering from correlation neglect are unaware of their bias and therefore do not consider the possibility that others might behave more rationally than themselves.

[27]As all other group members, rational or not, tell the truth, nothing can be learned about their types in the communication stage, and the rational (Bayesian) posterior belief about any other group member's type is equal to the prior $\hat{\alpha} = \alpha$.

$$
\begin{aligned}
E\Pi(m_i = b \mid s_i = r, \alpha) = \Pr(s_j = r, s_k = r \mid s_i = r) \cdot \big[&\alpha^2 E\Pi(D = \rho \mid s_i = r, s_j = r, s_k = r) \\
&+ 2\alpha(1-\alpha) E\Pi(D = \rho \mid s_i = r, s_j = r, s_k = r) \\
&+ (1-\alpha)^2 E\Pi(D = \beta \mid s_i = r, s_j = r, s_k = r)\big] \\
+ \Pr(s_j = r, s_k = b \mid s_i = r) &\cdot E\Pi(D = \beta \mid s_i = r, s_j = r, s_k = b) \\
+ \Pr(s_j = b, s_k = r \mid s_i = r) &\cdot E\Pi(D = \beta \mid s_i = r, s_j = b, s_k = r) \\
+ \Pr(s_j = b, s_k = b \mid s_i = r) \cdot \big[&\alpha^2 E\Pi(D = \beta \mid s_i = r, s_j = b, s_k = b) \\
&+ 2\alpha(1-\alpha) E\Pi(D = \beta \mid s_i = r, s_j = b, s_k = b) \\
&+ (1-\alpha)^2 E\Pi(D = \rho \mid s_i = r, s_j = b, s_k = b)\big] \qquad (4)
\end{aligned}
$$

Using equations (3) and (4), it is straightforward to calculate the value of α that allows to sustain the truth-telling equilibrium as $E\Pi(m_i = b \mid s_i = r) < E\Pi(m_i = r \mid s_i = r)$. For the parameters chosen in the experiment, this threshold value is 0.6307. For $\alpha > 0.6307$, rational group-members would profit from misreporting their red signal, and a truth-telling equilibrium cannot be sustained.

Bias treatments In treatments with biased preferences, where rational individuals prefer $D = \beta$ in case of one, two, or three b signals (without correlated signals) or after one or two b signals (with correlated signals), relaxing the assumption that rational group members assume all other group members are rational does not affect the existence of a truth-telling equilibrium. Without correlated signals, beliefs about others' correlation neglect do not influence expected payoffs. With correlated signals, for any signal combination, the group decision aligns with the ex-ante preferred decision for both minority and majority group members if all group members are rational and truthfully report their private signals. However, when a group member truthfully reports their signal while others exhibit correlation neglect, the group decision differs if the member reports a b signal and both other group members also observe a b signal. In this case, misreporting one's signal does not make the individual pivotal, as the decision remains unchanged when both other group members observe a b signal. Moreover, misreporting a b signal when both other group members suffer from correlation neglect alters the decision (to an inferior outcome) if both others have received a r signal. Thus, in treatments with biased preferences and the potential for correlated signals, it holds for all $0 \leq \alpha \leq 1$ that $E\Pi(m_i = b \mid s_i = b) > E\Pi(m_i = r \mid s_i = b)$ and $E\Pi(m_i = r \mid s_i = r) > E\Pi(m_i = b \mid s_i = r)$. Consequently, the truth-telling equilibrium outlined in Section B.2 remains unaffected for different values of α.

C Results - Regressions

C.1 Treatments without biased preferences

Table C.1: Suboptimal decisions and votes

	Individual decisions		Votes in group treatments			
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment: Correlated signals	0.450***	0.197***	0.372***	0.080***	0.391***	0.141***
	(0.030)	(0.038)	(0.018)	(0.023)	(0.021)	(0.027)
1 or 3 blue signals		-0.012		-0.009		-0.009
		(0.018)		(0.021)		(0.021)
Treatment: Correlated signals × 1 or 3 blue signals		0.415***		0.470***		0.400***
		(0.057)		(0.032)		(0.038)
Treatment: Communication					-0.037	-0.118***
					(0.024)	(0.031)
1 or 3 blue signals × Treatment: Communication						0.138***
						(0.048)
Constant	0.088***	0.095***	0.187***	0.191***	0.187***	0.191***
	(0.017)	(0.019)	(0.013)	(0.016)	(0.013)	(0.016)
Observations	1200	1200	5508	5508	5508	5508
Clusters	200 Ind.	200 Ind.	306 Grp	306 Grp	306 Grp	306 Grp
R^2	0.235	0.325	0.126	0.260	0.127	0.264

Notes: OLS regressions for suboptimal individual decisions (columns 1 and 2), and suboptimal votes in the group (columns 3 to 6). Standard errors are clustered at the individual level (columns 1 and 2) or group level (columns 3 to 6) and depicted in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.2: Predicting suboptimal decisions after three blue signals in treatments with correlation

	(1)	(2)	(3)	(4)	(5)	(6)
Mistakes in understanding questions	0.035				-0.003	-0.004
	(0.036)				(0.029)	(0.029)
CRT score		-0.104***			-0.068**	-0.068**
		(0.027)			(0.028)	(0.028)
Probability quiz score			-0.162***		-0.129***	-0.127***
			(0.033)		(0.036)	(0.037)
Critical Thinking scale (self-reported)				-0.018	0.054	0.056
				(0.056)	(0.052)	(0.053)
Male						-0.034
						(0.059)
Age						0.001
						(0.002)
Treatment: bias						0.042
						(0.060)
constant	0.691***	0.859***	0.893***	0.778***	0.755***	0.701***
	(0.038)	(0.039)	(0.038)	(0.208)	(0.194)	(0.226)
Observations	452	452	452	452	452	452
Clusters	187 Ind.	187 Ind.	187 Ind.	187 Ind.	187 Ind.	187 Ind.
R^2	0.004	0.070	0.100	0.001	0.124	0.128

Notes: OLS regressions for choosing blue ($\in \{0,1\}$) after three blue signals in treatments with potential correlation. *Mistakes in understanding questions* depicts the number of mistakes in the question about the game (between 0 and 6). *CRT score* depicts the number of correctly answered CRT questions (between 0 and 3). *Probability quiz score* depicts the number of correctly answered questions about probabilities and Bayesian updating (between 0 and 3). *Critical Thinking scale* (from 0 to 5) refers to the level of critical thinking, measured via self-reported strength of agreement to five statements. Std. errors are clustered at the subject level and depicted in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.3: Differences in suboptimal decisions/votes between groups and individuals without biased preferences

	No Correlation		Correlation Treatments		NoCorr + Corr		NoCorr + Corr +chat	
	Vote	Decision	Vote	Decision	Vote	Decision	Vote	Decision
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Group Treatment	0.098***	0.033	0.039	0.076**	0.098***	0.033	0.098***	0.033
	(0.022)	(0.023)	(0.029)	(0.033)	(0.022)	(0.023)	(0.022)	(0.023)
Treatment: correlated signals					0.450***	0.450***	0.450***	0.450***
					(0.030)	(0.030)	(0.030)	(0.030)
Group $\times$ correlated signals					-0.059	0.043	-0.059	0.043
					(0.036)	(0.040)	(0.036)	(0.040)
Group $\times$ Chat $\times$ correlated signals							-0.037	-0.065**
							(0.024)	(0.032)
cons.	0.088***	0.088***	0.538***	0.538***	0.088***	0.088***	0.088***	0.088***
	(0.017)	(0.017)	(0.024)	(0.024)	(0.017)	(0.017)	(0.017)	(0.017)
Observations	2454	2454	2418	2418	4872	4872	6708	6708
Clusters	203	203	201	201	404	404	506	506
R^2	0.013	0.002	0.001	0.004	0.182	0.257	0.152	0.207

Notes: OLS regressions for suboptimal votes and decisions in the different treatments without bias. Std. errors depicted in parentheses and clustered at the subject level for individual treatments and group level in the group treatments. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.2 Biased preferences in individual treatments

Table C.4: Decisions with and without biased preferences

	Decision for β treatments without correlation (1)	(2)	suboptimal decisions (3)	(4)
Treatment: bias	0.087***	0.069***	0.234***	0.236***
	(0.025)	(0.017)	(0.024)	(0.024)
Treatment: correlated signals			0.450***	0.429***
			(0.030)	(0.030)
Treatment: bias × Treatment: correlated signals			-0.251***	-0.259***
			(0.042)	(0.042)
1 blue signal		0.150***		0.418***
		(0.023)		(0.028)
2 blue signals		0.839***		0.076***
		(0.024)		(0.021)
3 blue signals		0.953***		0.322***
		(0.015)		(0.032)
constant	0.448***	-0.001	0.088***	-0.135***
	(0.019)	(0.015)	(0.017)	(0.022)
Observations	1260	1260	2460	2460
Clusters	210 Ind.	210 Ind.	410 Ind.	410 Ind.
R^2	0.007	0.607	0.140	0.256

Notes: OLS regressions for deciding β (columns 1 and 2) and deciding suboptimally given the signals (columns 3 and 4). All variables are dummy variables indicating the treatment subjects are in or the number of blue signals in the round. Standard errors are clustered at the individual level and depicted in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.3 Biased preferences in groups

Table C.5: Dishonest reporting

	All groups		Heterogeneous groups		
	No correlation	Correlation	All	Minority	Majority
	(1)	(2)	(3)	(4)	(5)
Treatment:	0.041^{**}	-0.001			
heterogeneous groups	(0.019)	(0.018)			
Treatment:			-0.035^{*}	-0.066^{**}	-0.020
correlated signals			(0.019)	(0.026)	(0.022)
Constant	0.095^{***}	0.102^{***}	0.136^{***}	0.135^{***}	0.137^{***}
	(0.012)	(0.014)	(0.014)	(0.021)	(0.016)
Observations	3672	3690	3690	1230	2460
Clusters	204 Groups	205 Groups	205 Groups	205 Groups	205 Groups
R^2	0.004	0.000	0.003	0.012	0.001

Notes: OLS regressions for not reporting the true signal (and either reporting an empty or a wrong signal). Data from the treatment with chat opportunity is not included in this table. Column (1) contains data of the group treatments without correlation possibility. Column (2) contains data of the group treatments with correlation possibility. Columns (3)-(5) contain data of the group treatments with heterogeneous groups. Column (4) only includes data from the unbiased minority individuals, and column (5) only includes data from the biased majority individuals within a group. Standard errors are clustered at the group level and depicted in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D Communication treatment with public signals

In December 2024, we collected data from 300 subjects in an additional treatment. This treatment closely mirrors the Gr-NoBias-Corr treatment with open communication, except that each group member observed all three signals on their screen. By making the signals public, this design eliminates the possibility of adding noise to the decision-making process during the signal-sharing stage. We pre-registered the design, sample size, and hypotheses as an add-on to the original pre-registration (AEARCTR-0006323).

To assess whether the public availability of signals improves group decision-making, we compare the share of suboptimal decisions (given the group's signals) between the public and private signal treatments. Column (1) in Table D.1 shows no significant improvement in decision-making in the public signals treatment. This absence of a significant difference suggests that noise introduced during the signal-sharing stage is unlikely to be a major source of inefficiency in this setting with aligned preferences. Additionally, to evaluate whether free-form chat facilitates a better understanding of the correlation structure and enhances decision-making, Column (2) in Table D.1 compares suboptimal decision-making in the individual treatment (where all three signals are observed) with the group treatment (where all three signals are observed, and decisions are made through voting). These results also reveal no significant differences in suboptimal voting or decision-making behavior. Thus, the findings from this additional treatment suggest that open communication in groups with aligned preferences is not sufficient to enhance decision-making quality.

Table D.1: Suboptimal voting/decision-making

	Group treatments with communication (1)	Three signals observed (2)
Public Signals	0.028	
	(0.026)	
Group Treatment		0.030
		(0.030)
Constant	0.540***	0.538***
	(0.018)	(0.024)
Observations	3690	2454
Clusters	205 Groups	203 Groups/ Ind.
R^2	0.001	0.000

Notes: OLS regressions for voting against the payoff-maximizing option given the signals (column 1) and for voting or deciding against the payoff-maximizing option given the signals (column 2). Column (1) contains data of the group treatments without bias, with correlation possibility, and with free-form communication. Column (2) contains data of the group treatment with public signals, without bias, with correlation possibility, and with free-form communication, as well as data of the individual treatment without bias and with correlation possibility. Standard errors are clustered at the group level (at the group and individual level in column 2) and depicted in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E Robustness checks

In the experiment we randomize whether correlation is possible in ρ or β. In this section, we show that the participants' behavior does not depend on the color of the state in which we allow for correlation.

Figure E.1: Suboptimal individual decisions and group votes in treatments with and without correlation.

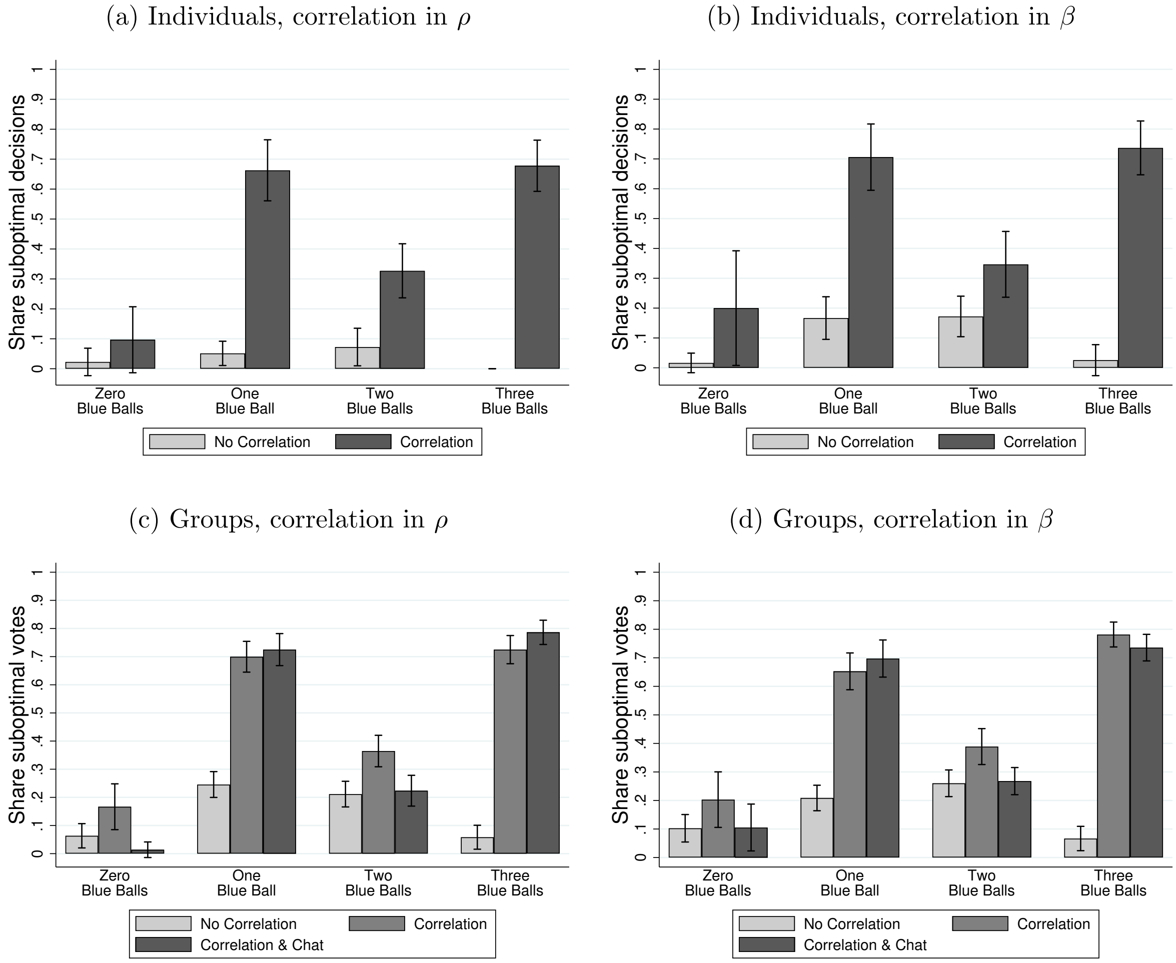

Notes: This figure replicates Figures 1 and 2 separately for the sessions that had correlation in state ρ (left) or in state β (right). For better comparison, the coding of the underlying variables follows the coding of the rest of the paper. Whiskers represent 95% confidence intervals.

For individual decision-making, the top panels of Figure E.1 provide an overview over the individual treatments with and without correlation, and all signal distributions. The bottom panels plot the same for groups and further differentiate between the Communication and the No-Communication treatments. The Figure shows qualitatively very similar results for both versions of the game.

Figure E.2: Share of suboptimal decisions within groups

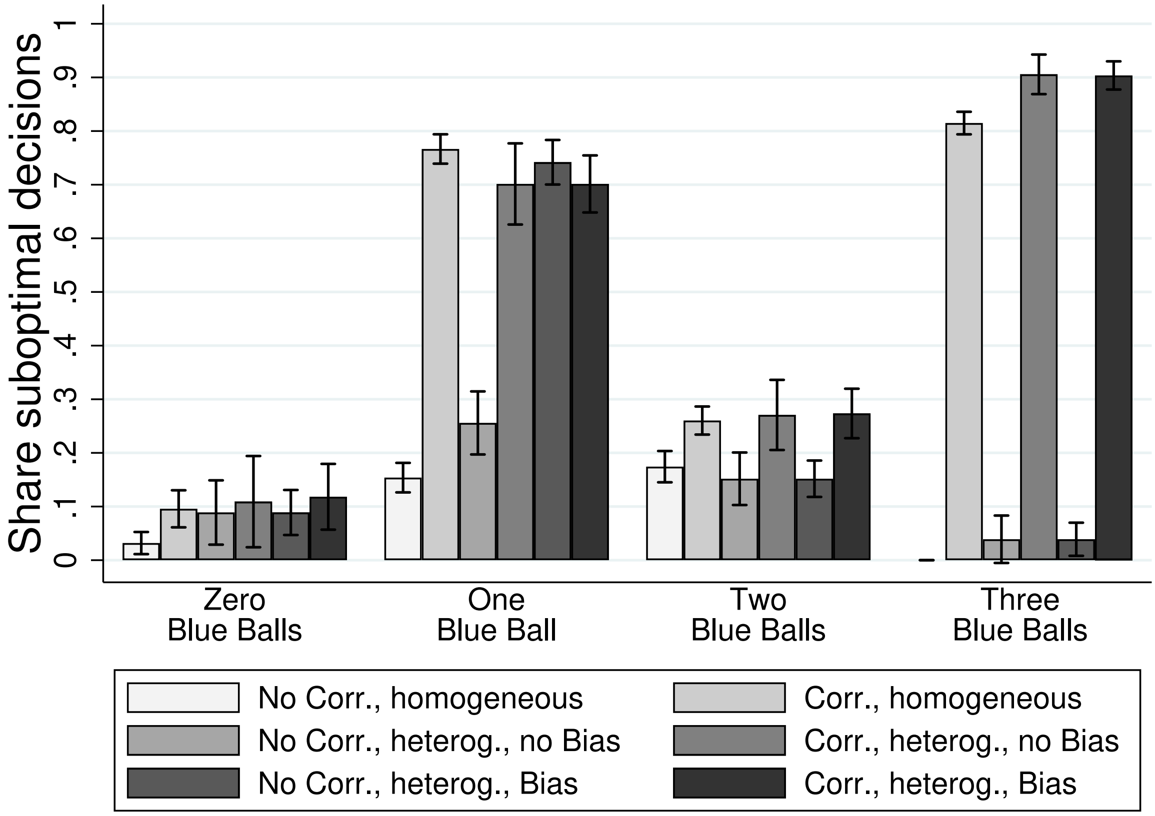

Notes: In the homogeneous group treatments, no group member is biased. In the heterogeneous group treatments, a majority is biased towards blue, a minority is unbiased. Whiskers represent 95% confidence intervals.

Figure E.2 replicates Figure 6 with group decisions instead of group votes. The figure shows that all conclusions in Section 5.3.2 also hold for group decisions.

F Chat data - Examples for how participants explain correlation

Below we provide examples for how individuals discussed the correlation structure after seeing three identical signals. In the experiment we randomized whether correlation was possible in the blue or red urn, we thus provide separate examples for each case.

Table F.1: Correlation in the blue urn: Examples where participants recommended to vote for the blue urn after seeing three red signals

ID	Argument
1653	the blue urn has a 50% chance to only select red balls, so the probability is higher that the blue urn was picked
1674	we all got red, in the blue urn there is 50% only red
1686	Blue urn 50% chance of only red balls
1692	I'll go blue now because of the 50% only red draw
1736	Probability of getting 3 red from the red urn is $(3/5)^3$. 21.6%. Probability of getting 3 red from the blue urn is 50%. Blue urn is more likely.
1753	Blue urn has 50% red balls
1782	50% probability that only red balls are drawn
1783	Blue urn is more likely because 50% chance to be always red, plus random chance

Table F.2: Correlation in the red urn: Examples where participants recommended to vote for the red urn after seeing three blue signals

ID	Argument
1800	Since there's a 50% chance the red urn only produces blue balls and we all drew blue, that seems like a good bet.
1810	50% chance
1812	It's likely that we got the 50% chance on the red urn to only have blue drawn.
1815	50% chance
1849	Red has 50% chance of blue and we all got blue.
1851	All blue from red 50%, all blue from blue is less chance.
1852	Red urn has a 50% probability of a blue ball drawn and we got all red, which makes blue more likely.
1919	I think statistically it's probably more likely it's the red urn with the probability of 50% of draws always being a blue ball.
1928	Greater chance of blue balls from red because of the 50% chance rule.

G Instructions and decision screens

Figure G.1: Welcome Screen for participants

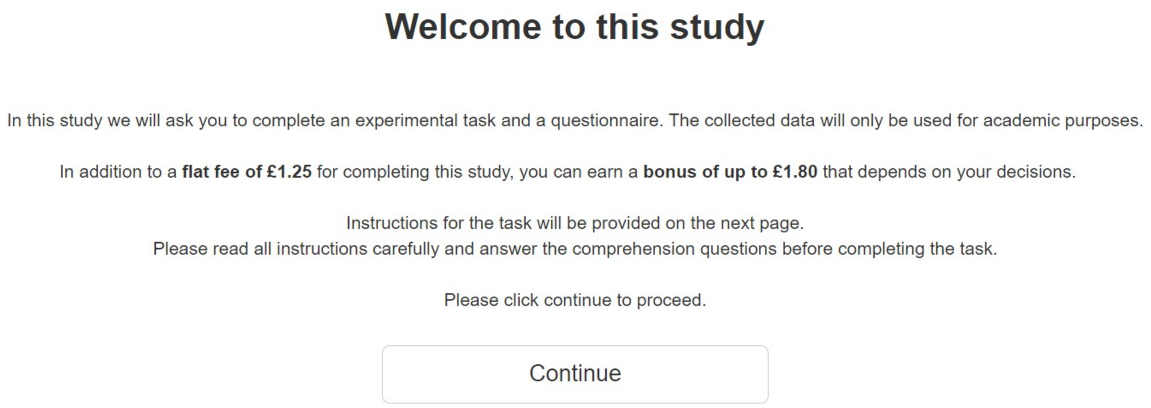

Welcome to this study

In this study we will ask you to complete an experimental task and a questionnaire. The collected data will only be used for academic purposes.

In addition to a **flat fee of £1.25** for completing this study, you can earn a **bonus of up to £1.80** that depends on your decisions.

Instructions for the task will be provided on the next page.
Please read all instructions carefully and answer the comprehension questions before completing the task.

Please click continue to proceed.

Continue

G.1 Instructions

Figure G.2: Instructions for the *Ind-NoBias-NoCorr* treatment.

Instructions

Please read the instructions carefully.

You are asked to play this task for 6 rounds.

In each round 3 balls are either drawn from a red or a blue urn (see pictures left and right). Each urn is picked with a probability of 50% and your challenge is to find out which of the two urns has been selected.

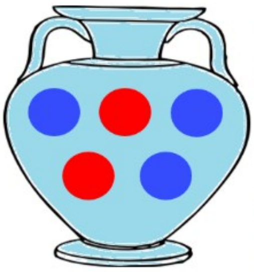

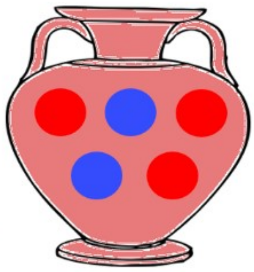

To help you make an informed decision we provide some additional information about the urns:

1. The red urn contains **3 red balls and 2 blue balls**.
2. The blue urn contains **2 red balls and 3 blue balls**.

After one of the urns is chosen (with 50%), we draw 3 balls from the chosen urn. Each time a ball is drawn it is put back into the urn. It is thus possible to see the same ball several times. The balls are always drawn randomly in both urns. After the balls have been drawn, two buttons will appear through which you can indicate whether you think that the balls come from the red or the blue urn. Below you see a picture of how the decision screen will look like.

If you chose the correct urn you will receive a bonus of **£0.20** for that round. In the end, all the money you earned from the different rounds will be summed up. After completing the questionnaire you will receive feedback on how many of your guesses were correct.

Figure G.3: Instructions for the *Ind-NoBias-Corr* treatment.

Instructions

Please read the instructions carefully.

You are asked to play this task for 6 rounds.

In each round 3 balls are either drawn from a red or a blue urn (see pictures left and right). Each urn is picked with a probability of 50% and your challenge is to find out which of the two urns has been selected.

To help you make an informed decision we provide some additional information about the urns:

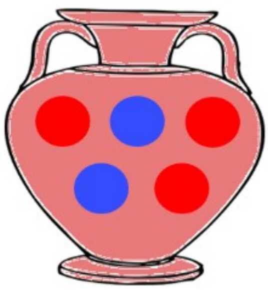

1. The red urn contains **3 red balls and 2 blue balls**.
2. The blue urn contains **2 red balls and 3 blue balls**.

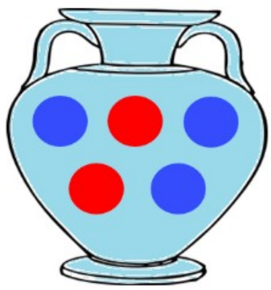

After one of the urns is chosen (with 50%), we draw 3 balls from the chosen urn. Each time a ball is drawn it is put back into the urn. It is thus possible to see the same ball several times. The balls are drawn in the following way:

1. Red urn: Each of the 3 balls are chosen **randomly**.
2. Blue urn:
   - With a probability of **50%** balls are chosen **randomly**.
   - With a probability of **50% only red balls** can be drawn.

After the balls have been drawn, two buttons will appear through which you can indicate whether you think that the balls come from the red or the blue urn. Below you see a picture of how the decision screen will look like.

If you chose the correct urn you will receive a bonus of **£0.20** for that round. In the end, all the money you earned from the different rounds will be summed up. After completing the questionnaire you will receive feedback on how many of your guesses were correct.

Figure G.4: Instructions for the *Ind-Bias-NoCorr* treatment.

Instructions

Please read the instructions carefully.

You are asked to play this task for 6 rounds.

In each round 3 balls are either drawn from a red or a blue urn (see pictures left and right). Each urn is picked with a probability of 50% and your challenge is to find out which of the two urns has been selected.

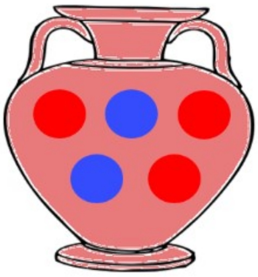

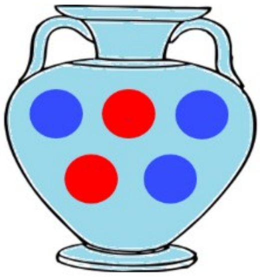

To help you make an informed decision we provide some additional information about the urns:

1. The red urn contains **3 red balls and 2 blue balls**.
2. The blue urn contains **2 red balls and 3 blue balls**.

After one of the urns is chosen (with 50%), we draw 3 balls from the chosen urn. Each time a ball is drawn it is put back into the urn. It is thus possible to see the same ball several times. The balls are always drawn randomly in both urns. After the balls have been drawn, two buttons will appear through which you can indicate whether you think that the balls come from the red or the blue urn. Below you see a picture of how the decision screen will look like.

Bonus structure
If you chose the correct urn you will receive a bonus for that round. However, both urns do not pay the same:

1. Correct decision for red: **£0.25**.
2. Correct decision for blue: **£0.15**.

In the end, all the money you earned from the different rounds will be summed up. After completing the questionnaire you will receive feedback on how many of your guesses were correct.

Figure G.5: Instructions for the *Ind-Bias-Corr* treatment.

Instructions

Please read the instructions carefully.

You are asked to play this task for 6 rounds.

In each round 3 balls are either drawn from a red or a blue urn (see pictures left and right). Each urn is picked with a probability of 50% and your challenge is to find out which of the two urns has been selected.

To help you make an informed decision we provide some additional information about the urns:

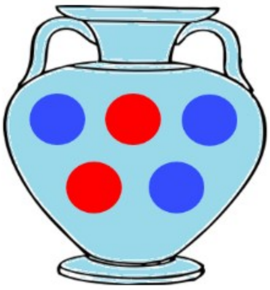

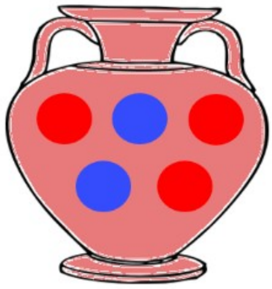

1. The blue urn contains **2 red balls and 3 blue balls**.
2. The red urn contains **3 red balls and 2 blue balls**.

After one of the urns is chosen (with 50%), we draw 3 balls from the chosen urn. Each time a ball is drawn it is put back into the urn. It is thus possible to see the same ball several times. The balls are drawn in the following way:

1. Blue urn: Each of the 3 balls are chosen **randomly**.
2. Red urn:
   - With a probability of **50%** balls are chosen **randomly**.
   - With a probability of **50% only blue balls** can be drawn.

After the balls have been drawn, two buttons will appear through which you can indicate whether you think that the balls come from the red or the blue urn. Below you see a picture of how the decision screen will look like.

Bonus structure
If you chose the correct urn you will receive a bonus for that round. However, both urns do not pay the same:

1. Correct decision for blue: **£0.25**.
2. Correct decision for red: **£0.15**.

In the end, all the money you earned from the different rounds will be summed up. After completing the questionnaire you will receive feedback on how many of your guesses were correct.

Figure G.6: Instructions for the *Gr-NoBias-NoCorr* treatment

<u>Instructions</u>

Please read the instructions carefully.

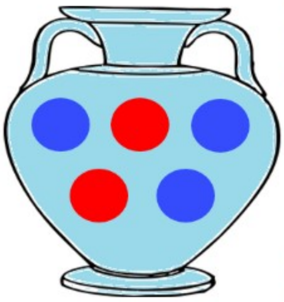

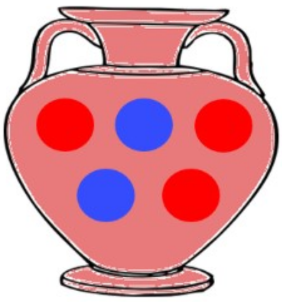

You are asked to play this task for 6 rounds with two other participants.

In each round 3 balls are either drawn from a red or a blue urn (see pictures left and right). Each urn is picked with a probability of 50% and your challenge is to find out which of the two urns has been selected. On a first screen, each participant will be shown a single ball and has the chance to share this information with the other participants in the group. On a second screen, you will then see the colour of your own draw, plus the colours reported by the other participants.

To help you make an informed decision we provide some additional information about the urns:

1. The blue urn contains **2 red balls and 3 blue balls**.
2. The red urn contains **3 red balls and 2 blue balls**.

After one of the urns is chosen (with 50%), we draw 3 balls from the chosen urn. Each time a ball is drawn it is put back into the urn. It is thus possible that you and the other participants in your group will see the same ball. The balls are always drawn randomly in both urns. After you draw your ball, you can send a message to the participants in your group. The message can be 'I drew a red ball', 'I drew a blue ball' or 'I prefer not to reveal'. Below you see a picture of how this screen will look like.

.

On the next screen you will see both the colour of your ball and the balls reported by the participants in your group. A white ball thereby means that a participant chose not to reveal their colour. Your task is then to make a guess whether the balls come from the red or blue urn. You can see a picture of the decision screen below.

The decision for an urn will be determined by **majority vote**. This means that if at least two people decide for an urn, this is taken as the group's decision. If the group decision is correct, everyone in the group will receive a bonus of **£0.20** for that round. In the end, all the money you earned from the different rounds will be summed up. After completing the questionnaire you will receive feedback on how many of your guesses were correct.

Figure G.7: Instructions for the *Gr-Bias-NoCorr* treatment

Instructions

Please read the instructions carefully.

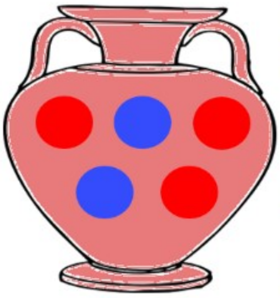

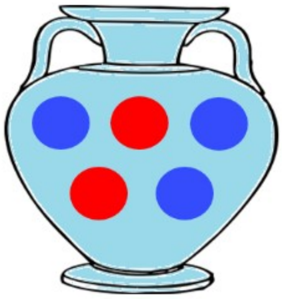

You are asked to play this task for 6 rounds with two other participants.

In each round 3 balls are either drawn from a red or a blue urn (see pictures left and right). Each urn is picked with a probability of 50% and your challenge is to find out which of the two urns has been selected. On a first screen, each participant will be shown a single ball and has the chance to share this information with the other participants in the group. On a second screen, you will then see the colour of your own draw, plus the colours reported by the other participants.

To help you make an informed decision we provide some additional information about the urns:

1. The red urn contains **3 red balls and 2 blue balls**.
2. The blue urn contains **2 red balls and 3 blue balls**.

After one of the urns is chosen (with 50%), we draw 3 balls from the chosen urn. There is one urn selected per group for each round. Each time a ball is drawn it is put back into the urn. It is thus possible that you and the other participants in your group will see the same ball. The balls are always drawn randomly in both urns. After you draw your ball, you can send a message to the participants in your group. The message can be 'I drew a red ball', 'I drew a blue ball' or 'I prefer not to reveal'. Below you see a picture of how this screen will look like.

On the next screen you will see both the colour of your ball and the balls reported by the participants in your group. A white ball thereby means that a participant chose not to reveal their colour. Your task is then to make a guess whether the balls come from the red or blue urn. You can see a picture of the decision screen below.

Bonus structure
The decision for an urn will be determined by **majority vote**. This means that if at least two people decide for an urn, this is taken as the group's decision. If your group chose the correct urn everyone will receive a bonus for that round. However, both urns do not pay the same for all participants. One participant will receive **£0.20** for a correct group decision independent of the urn colour, while for the other two:

1. Correct decision for red: **£0.25**.
2. Correct decision for blue: **£0.15**.

You will learn which of the payoffs is relevant for you after you have been matched with two other participants. In the end, all the money you earned from the different rounds will be summed up. After completing the questionnaire you will receive feedback on how many of your guesses were correct.

Figure G.8: Instructions for the *Gr-NoBias-Corr* treatment

Instructions

Please read the instructions carefully.

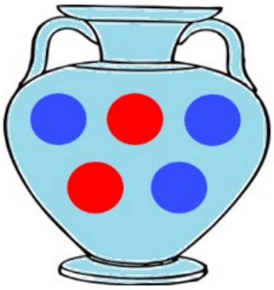

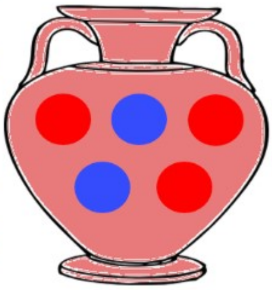

You are asked to play this task for 6 rounds with two other participants.

In each round 3 balls are either drawn from a red or a blue urn (see pictures left and right). Each urn is picked with a probability of 50% and your challenge is to find out which of the two urns has been selected. On a first screen, each participant will be shown a single ball and has the chance to share this information with the other participants in the group. On a second screen, you will then see the colour of your own draw, plus the colours reported by the other participants.

To help you make an informed decision we provide some additional information about the urns:

1. The red urn contains **3 red balls and 2 blue balls**.
2. The blue urn contains **2 red balls and 3 blue balls**.

After one of the urns is chosen (with 50%), we draw 3 balls from the chosen urn. There is one urn selected per group for each round. Each time a ball is drawn it is put back into the urn. It is thus possible that you and the other participants in your group will see the same ball. The balls are drawn in the following way:

1. Red urn: Each of the 3 balls are chosen **randomly**.
2. Blue urn:
   - With a probability of **50%** balls are chosen **randomly**.
   - With a probability of **50% only red balls** can be drawn.

After you draw your ball, you can send a message to the participants in your group. The message can be 'I drew a red ball', 'I drew a blue ball' or 'I prefer not to reveal'. Below you see a picture of how this screen will look like.

On the next screen you will see both the colour of your ball and the balls reported by the participants in your group. A white ball thereby means that a participant chose not to reveal their colour. Your task is then to make a guess whether the balls come from the red or blue urn. You can see a picture of the decision screen below.

The decision for an urn will be determined by **majority vote**. This means that if at least two people decide for an urn, this is taken as the group's decision. If the group decision is correct, everyone in the group will receive a bonus of **£0.20** for that round. In the end, all the money you earned from the different rounds will be summed up. After completing the questionnaire you will receive feedback on how many of your guesses were correct.

Figure G.9: Instructions for the *Gr-Bias-Corr* treatment.

Instructions

Please read the instructions carefully.

You are asked to play this task for 6 rounds with two other participants.

In each round 3 balls are either drawn from a red or a blue urn (see pictures left and right). Each urn is picked with a probability of 50% and your challenge is to find out which of the two urns has been selected. On a first screen, each participant will be shown a single ball and has the chance to share this information with the other participants in the group. On a second screen, you will then see the colour of your own draw, plus the colours reported by the other participants.

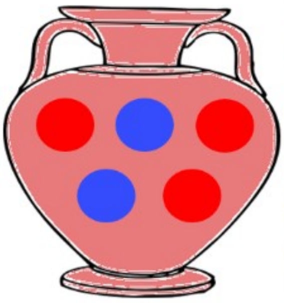

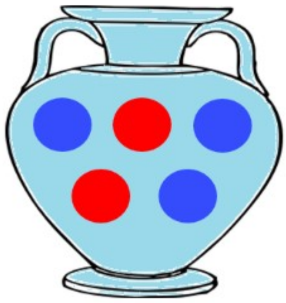

To help you make an informed decision we provide some additional information about the urns:

1. The red urn contains **3 red balls and 2 blue balls**.
2. The blue urn contains **2 red balls and 3 blue balls**.

After one of the urns is chosen (with 50%), we draw 3 balls from the chosen urn. There is one urn selected per group for each round. Each time a ball is drawn it is put back into the urn. It is thus possible that you and the other participants in your group will see the same ball. The balls are drawn in the following way:

1. Red urn: Each of the 3 balls are chosen **randomly**.
2. Blue urn:
    - With a probability of **50%** balls are chosen **randomly**.
    - With a probability of **50% only red balls** can be drawn.

After you draw your ball, you can send a message to the participants in your group. The message can be 'I drew a red ball', 'I drew a blue ball' or 'I prefer not to reveal'. Below you see a picture of how this screen will look like.

On the next screen you will see both the colour of your ball and the balls reported by the participants in your group. A white ball thereby means that a participant chose not to reveal their colour. Your task is then to make a guess whether the balls come from the red or blue urn. You can see a picture of the decision screen below.

Bonus structure
The decision for an urn will be determined by **majority vote**. This means that if at least two people decide for an urn, this is taken as the group's decision. If your group chose the correct urn everyone will receive a bonus for that round. However, both urns do not pay the same for all participants. One participant will receive **£0.20** for a correct group decision independent of the urn colour, while for the other two:

1. Correct decision for red: **£0.25**.
2. Correct decision for blue: **£0.15**.

You will learn which of the payoffs is relevant for you after you have been matched with two other participants. In the end, all the money you earned from the different rounds will be summed up. After completing the questionnaire you will receive feedback on how many of your guesses were correct.

Figure G.10: Instructions for the *Gr-NoBias-Corr-Chat* treatment.

Instructions

Please read the instructions carefully.

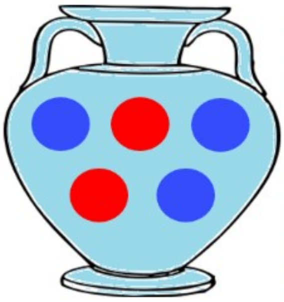

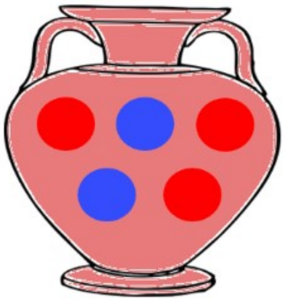

You are asked to play this task for 6 rounds with two other participants.

In each round 3 balls are either drawn from a red or a blue urn (see pictures left and right). Each urn is picked with a probability of 50% and your challenge is to find out which of the two urns has been selected. On a first screen, each participant will be shown a single ball and has the chance to share this information with the other participants in the group. On a second screen, you will then see the colour of your own draw, plus the colours reported by the other participants.

To help you make an informed decision we provide some additional information about the urns:

1. The blue urn contains **2 red balls and 3 blue balls**.
2. The red urn contains **3 red balls and 2 blue balls**.

After one of the urns is chosen (with 50%), we draw 3 balls from the chosen urn. There is one urn selected per group for each round. Each time a ball is drawn it is put back into the urn. It is thus possible that you and the other participants in your group will see the same ball. The balls are drawn in the following way:

1. Blue urn: Each of the 3 balls are chosen **randomly**.
2. Red urn:
    - With a probability of **50%** balls are chosen **randomly**.
    - With a probability of **50% only blue balls** can be drawn.

After you draw your ball, you can send a message to the participants in your group. The message can be 'I drew a red ball', 'I drew a blue ball' or 'I prefer not to reveal'. Below you see a picture of how this screen will look like.

On the next screen you will see both the colour of your ball and the balls reported by the participants in your group. A white ball thereby means that a participant chose not to reveal their colour. Your task is then to make a guess whether the balls come from the red or blue urn. Before making your own choice you can recommend a choice to the other players in your team and will also see their recommendations. You can see a picture of the final decision screen below.

The decision for an urn will be determined by **majority vote**. This means that if at least two people decide for an urn, this is taken as the group's decision. If the group decision is correct, everyone in the group will receive a bonus of **£0.20** for that round. In the end, all the money you earned from the different rounds will be summed up. After completing the questionnaire you will receive feedback on how many of your guesses were correct.

Figure G.11: Control questions.

Comprehension Questions

In order to move on please answer the following questions:

1) How many red balls are in the red urn/ how many blue balls are in the blue urn?

- Five
- Three
- Two

2) Is it more likely that one of the urns is selected?

- The red urn is more likely to be selected.
- Each urn is selected with 50%.
- The blue urn is more likely to be selected.

3) If the blue urn is selected, how are balls drawn from this urn?

- Always at random.
- 50% at random, 50% only red balls are drawn.
- 50% at random, 50% only blue balls are drawn.

4) If the red urn is selected, how are balls drawn from this urn?

- Always at random.
- 50% at random, 50% only red balls are drawn.
- 50% at random, 50% only blue balls are drawn.

5) Given that you chose the correct urn, how does your payoff differ between red and blue urn?

- Correct choice for blue: £0.25; Correct choice for red: £0.15
- Correct choice for blue: £0.15; Correct choice for red: £0.25
- No difference.

G.2 Decision screens

Example of a decision screen in an individual treatment

Figure G.12: Decision-making Screen in the *Ind-Bias-Corr* treatment.

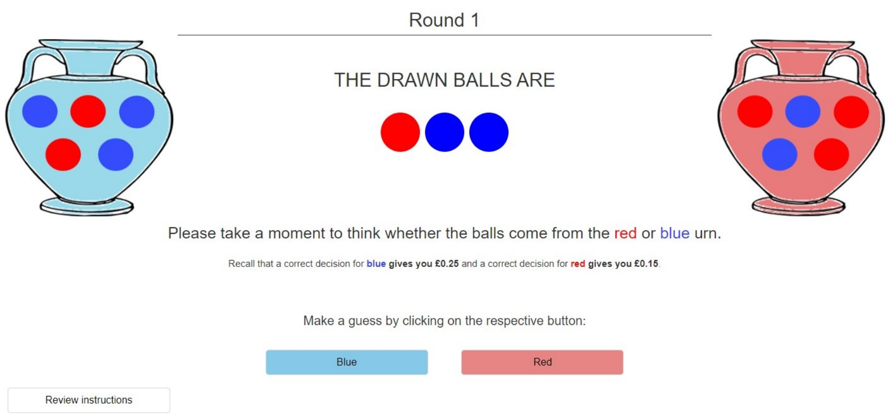

Example of decision screens in a group treatment

Figure G.13: Signal sharing screen in the *Gr-Bias-Corr* treatment (unbiased player).

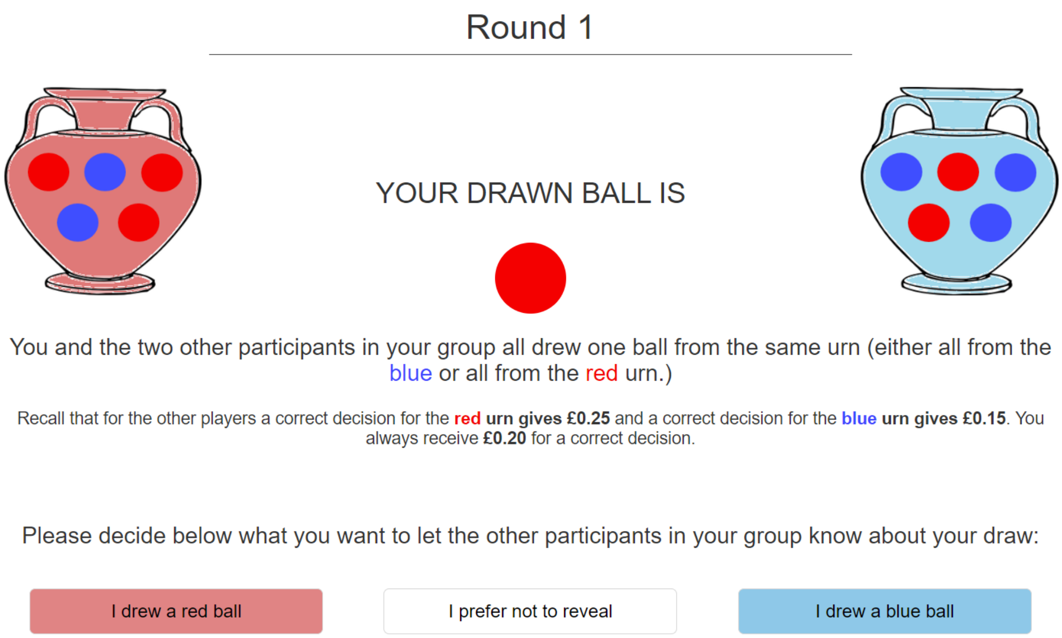

Figure G.14: Decision-making screen in the *Gr-Bias-Corr* treatment (biased player).

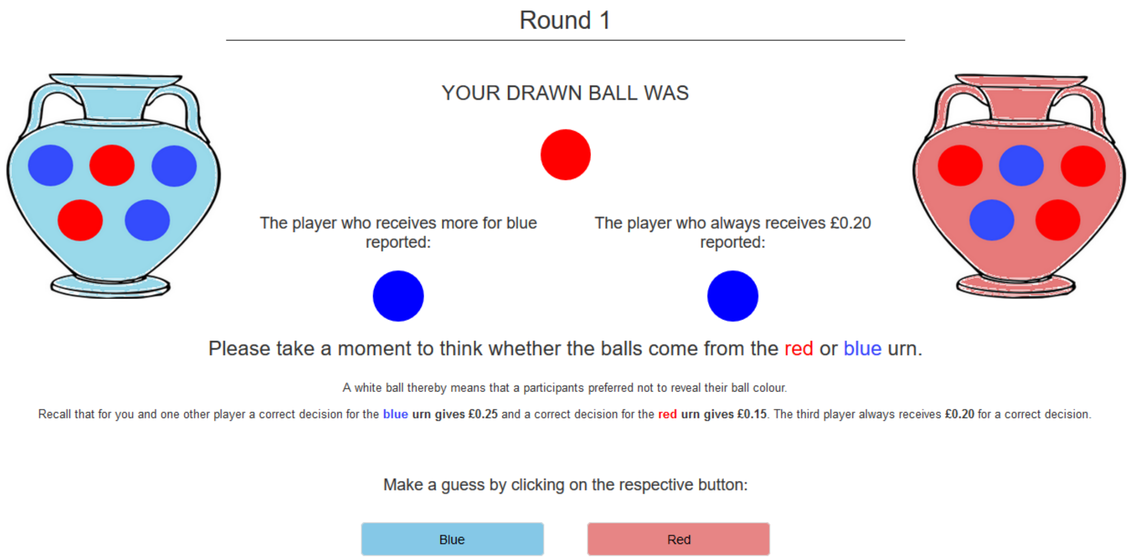

Figure G.15: Communication screen in the *Gr-NoBias-Corr-Communication* treatment

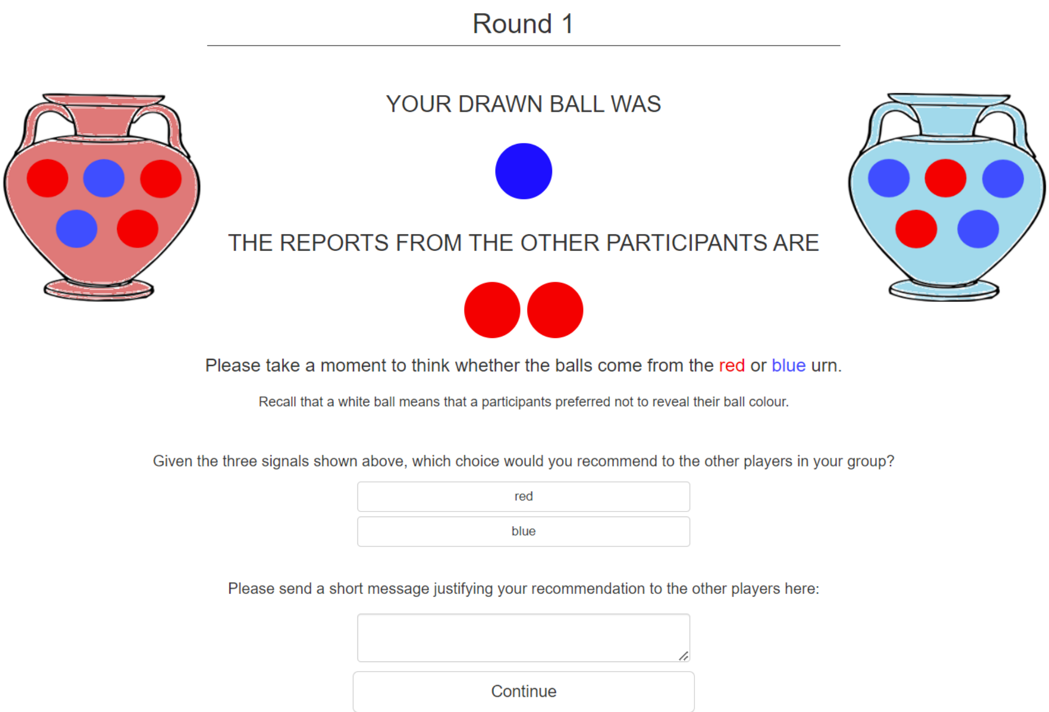

G.3 Questionnaire

Figure G.16: CRT questionnaire.

Questionnaire - Part 1/2

Please answer the six questions below. For each correct answer you receive an additional bonus of £0.10.

1. A pen and a case cost £1.10 in total. The case costs £1.00 more than the pen. How much does the pen cost? ________ cents

2. If it takes 5 machines 5 minutes to make 5 screws, how long would it take 100 machines to make 100 screws? ________ minutes

3. In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake? ________ days

Figure G.17: Probability quiz.

4. A normal, six sided dice has four green and two red sides. The dice is thrown 20 times and each time somebody takes note whether it shows a green (G) or red (R) side. If you had to bet on one of the following sequences, which one would you choose?

   - RGRRR
   - GRGRRR
   - GRRRRR

5. A study is conducted to test the effectiveness of a new medication. The results show: if the medication is taken, 200 test persons show an improved condition, 75 don't. If the medication is not taken 50 test persons show an improved condition, 15 don't. Was the medication successful?

   - Yes
   - No

6. A bag contains 20 blue and 10 green marbles. You have to draw out a marble six times. Each time the drawn marble has to be replaced in the bag (you always draw a marble from 30 marbles, 20 blue and 10 green). Imagine that you receive 10 Euros for each time you correctly predict the colour of the drawn marble. Which of the following is the best choice to win as much as possible?

   - Predict randomly blue or green
   - Any combination of 4 blue and 2 green
   - Always predict blue

Questionnaire - Part 2/2

1) Please indicate below how difficult you found the urn task.

not at all difficult ○ ○ ○ ○ ○ very difficult

2) Can you briefly describe how you made your decision for one urn or the other/ which factors you considered for that decision?

3) Please use the scale below to indicate how much you agree with the following statements, where choosing the far left means '*I agree a lot*', the mid point means '*I neither agree nor disagree*' and choosing the far right means '*I disagree a lot*'. Please be as honest and accurate as you can throughout.

If I think longer about a problem I will be more likely to solve it.

I agree a lot ○ ○ ○ ○ ○ I disagree a lot

Intuition is the best guide in making decisions.

I agree a lot ○ ○ ○ ○ ○ I disagree a lot

People should always take into consideration evidence that goes against their beliefs.

I agree a lot ○ ○ ○ ○ ○ I disagree a lot

Considering too many different opinions often leads to bad decisions.

I agree a lot ○ ○ ○ ○ ○ I disagree a lot

Changing your mind is a sign of weakness.

I agree a lot ○ ○ ○ ○ ○ I disagree a lot

4) Please tell me, in general, how willing or unwilling you are to take risks?
Choosing the response to the far left means '*Completely unwilling* to take risks' and choosing the response to the far right means you are '*Very willing* to take risks'.
Pick a response below to indicate where you fall on this scale.

Completely unwilling ○ ○ ○ ○ ○ ○ ○ ○ ○ ○ ○ Very willing

5) What is your age?

6) What is your gender?

7) What is your first language?

8) What is the highest degree or level of school you have completed?

- No formal qualifications
- Secondary school/GCSE
- College/A levels
- Undergraduate degree (BA/BSc/other)
- Graduate degree (MA/MSc/MPhil/other)
- Doctorate degree (PhD/MD/other)
- Prefer not to say

G.4 Feedback Screens

Figure G.18: Feedback about decisions and urns.

Below you can find an overview of your performances in the different parts. Continue to connect to Prolific and claim your payment.

Urn task

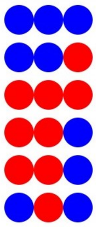

Round	The drawn balls were	Your choice was	The correct urn was	Your bonus
1		red	red	£0.15
2		blue	blue	£0.25
3		red	red	£0.15
4		blue	blue	£0.25
5		blue	blue	£0.25
6		blue	blue	£0.25

Figure G.19: Feedback about CRT and probability quiz questions.

Questionnaire

Question	Your answer	Correct answer	Your bonus
A pen and a case cost £1.10 in total. The case costs £1.00 more than the pen. How much does the pen cost?	5	5 cents	£0.10
If it takes 5 machines 5 minutes to make 5 screws, how long would it take 100 machines to make 100 screws?	5 minutes	5 minutes	£0.10
In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake?	5 days	47 days	£0.00
A normal, six sided dice has four green and two red sides. The dice is thrown 20 times and each time somebody takes note whether it shows a green (G) or red (R) side. If you had to bet on one of the following sequences, which one would you choose?	RGRRR	RGRRR	£0.10
A study is conducted to test the effectiveness of a new medication. The results show: if the medication is taken, 200 test persons show an improved condition, 75 don't. If the medication is not taken 50 test persons show an improved condition, 15 don't. Was the medication successful?	No	No	£0.10
A bag contains 20 blue and 10 green marbles. You have to draw out a marble six times. Each time the drawn marble has to be replaced in the bag (you always draw a marble from 30 marbles, 20 blue and 10 green). Imagine that you receive 10 Euros for each time you correctly predict the colour of the drawn marble. Which of the following is the best choice to win as much as possible?	Any combination of 4 blue and 2 green	Always predict blue	£0.00

In total you thus receive a bonus of **£1.70**.

Before you go: Any feedback from your side regarding this study is highly appreciated. Please leave comments below