

Learning Through Imitation: An Experiment^{*}

Marina Agranov[†] Gabriel Lopez-Moctezuma[‡]

Philipp Strack[§] Omer Tamuz[¶]

May 17, 2026

Abstract

We compare how well agents aggregate information in two repeated social learning environments. In the first setting agents have access to a public data set. In the second they have access to the same data, and also to the past actions of others. Despite the fact that actions contain no additional payoff-relevant information, and despite potential herd behavior, free riding and information overload issues, observing and imitating the actions of others leads agents to take the optimal action more often in the second setting. We also investigate the effect of group size, as well as a setting in which agents observe private data and others' actions.

^{*}This work was presented at Princeton University, MidExLab virtual seminar, Chicago Harris, UCSB, NYU, University of Michigan, Osaka University, Texas A&M, Stanford University, CREST, University of Hamburg, University of Maryland, Caltech, Essex University, NYU Abu Dhabi, University of Southampton, Kansas University, Shanghai Jiao Tong University, and GAMES 2020. We thank the participants of these seminars and conferences and Carlo Cusumano for their helpful comments.

[†]Caltech and NBER. Email: magranov@hss.caltech.edu. Experiments were funded by the grant received from the Ronald and Maxine Linde Institute of Economic and Management Sciences at Caltech.

[‡]Caltech. Email: glmoctezuma@caltech.edu.

[§]Yale University. Email: philipp.strack@yale.edu.

[¶]Caltech. Email: tamuz@caltech.edu. Omer Tamuz was supported by a grant from the Simons Foundation (#419427), a Sloan fellowship, a BSF award (#2018397), and a National Science Foundation CAREER award (DMS-1944153).

1 Introduction

A large literature in economics and finance has shown that imitation in social learning leads to inefficiencies: agents who imitate do not reveal their private information, giving rise to information cascades and herd behavior. This point has been made both theoretically and experimentally in a variety of settings, including those with rational agents (Banerjee, 1992; Bikhchandani et al., 1992; Anderson and Holt, 1997; Çelen and Kariv, 2004; Harel et al., 2021) and those that consider behavioral biases and heuristics (Enke and Zimmermann, 2019; Eyster et al., 2018).

We study a social learning setting in which all information is *public*, and ask whether agents perform better or worse when given the opportunity to imitate the actions of others. Although these actions do not contain additional payoff-relevant information beyond the public raw data, and despite their potential to create herd behavior, free riding, and information overload, we find that agents perform better when allowed to observe others' actions.

In our experiment, participants play the following game: An urn is chosen to be either majority green or majority red, each with probability one half. In the former case the urn contains six green balls and four red balls, whereas in the latter case the numbers are reversed. A signal is a random draw (with replacement) of a ball from the urn. A game consists of twenty periods. In each period, each member of a group of eight participants observes some information and then has to guess the majority color of the urn. Participants are rewarded for guessing correctly.

In each period, each participant observes a signal. Furthermore, depending on the treatment, participants observe some additional information. In the SIGNALS treatment, participants observe others' signals. In the ALL treatment, they likewise observe all the others' signals, and also all the others' past actions. Thus, in both treatments all signals are public, but actions are public only in the latter.

The comparison between the SIGNALS and ALL treatments sheds light on the usefulness of observing other people's decisions, in addition to the raw public data. Rational agents are expected to perform identically in these two treatments, since both offer agents exactly the same information about the state; in both cases all agents can see all the signals available to society. Some behavioral heuristics suggest that agents should do better when observing only signals, as observing the actions of others can—through information overload and correlation neglect—lead to wrong conclusions, potentially through herd behavior and groupthink. Information overload has been known to interfere with information processing and may lead to sub optimal decisions (Caplin et al., 2011; Scheibehenne and Greifeneder, 2010). Since the ALL treatment features twice as many pieces of information as the SIGNALS

treatment, these phenomena suggest that performance in the ALL treatment should be worse. Furthermore, the literature has documented a tendency of people to neglect correlation among dependent pieces of information (Enke and Zimmermann, 2019). In the SIGNALS treatment, correlation neglect plays no role because the signals of other group members are independent conditional on the state of nature. However, in the ALL treatment, correlation neglect might be detrimental since others' actions are not independent and regarding them as independent will hinder learning. Finally, pure imitation by all agents is clearly inefficient, as it leaves no room for the signals to inform the actions.

Nevertheless, we find that agents perform better in the ALL treatment than in the SIGNALS treatment: they are better off when they are privy to others' actions, in addition to the information that those actions are based on. A possible explanation for this finding is that observing the actions of others allows subjects to benefit from the aggregation of information done by other subjects. We observe that subjects rely on the actions of others more heavily when signals are "weak", by which we mean relatively uninformative (i.e., roughly the same number of red and green balls have been drawn) and it is thus harder to determine which color was drawn more. Individual level analysis shows that both high-IQ and low-IQ subjects condition on the actions of others to a similar extent, and do so almost exclusively when signals are weak. Both groups benefit from imitating others.

Having documented the benefits of imitation in a repeated social learning setting with *public* information, we complement our findings by studying a similar setting with *private* information. We focus on the effects of information structures and group size and conduct two additional treatments. In the NO INFO treatment each participant observes her own signals only, and does not see others' signals or actions. In the ACTIONS treatment, subjects again observe their own signals and not the signals of others, but do observe the past actions of others.

The comparison between the NO INFO and ACTIONS treatments reveals to what extent people are capable of extracting the information contained in their peers' private signals by observing their actions, in a repeated setting. Theory predicts that such inference is difficult and has limited benefits (Harel et al., 2021), which leads to the natural question of how well people perform in practice in such settings. We find that participants perform significantly better in the ACTIONS treatment as compared with the NO INFO treatment, despite the complexity of the information extraction problem.

Finally, we compare how changing the group size affects results, by repeating these experiments with groups of only four agents. While in the ALL treatment larger groups perform better, in the ACTIONS treatment performance is identical across group sizes. This result qualitatively matches a theoretical result due to Harel et al. (2021), who show that myopic Bayesian agents, in the long-run, do not perform significantly better in larger groups.

We conclude the analysis by presenting a simple behavioral model that captures the main features of our experiment. In this model, agents noisily respond to a weighted sum of two stimuli: signals and others' actions. Estimating this model on our experimental data reproduces all of our main findings. A natural direction for future research is a formal analysis of the model, including proofs of the conjectures we put forward.

We proceed as follows. Related literature is surveyed in Section 2. Section 3 contains the detailed description of our experiment and Section 4 outlines theoretical predictions for each of the treatments. Section 5 reports the experimental results: Section 5.1 focuses on the setting with public information (ALL versus SIGNALS treatments) and Section 5.2 considers settings with private information (the ACTIONS treatment) and documents group size effects. Section 6 presents the behavioral model and simulation results. Finally, we offer some discussion of implications of our experimental findings in Section 7.

2 Related Literature

Our paper contributes to a large literature on social learning. Most of the theoretical literature focuses on sequential settings, in which exogenously ordered agents act once after observing some statistics about their predecessors' actions and their own private signal (see for example Banerjee, 1992; Bikhchandani et al., 1992; Smith and Sørensen, 2000, or a recent survey by Bikhchandani et al., 2021).

The experimental literature on sequential social learning is vast, and a complete survey is beyond the scope of this paper. The first experimental paper that documents informational cascades is Anderson and Holt (1997). Since then, a fruitful literature in experimental economics has studied the determinants and the limitations of herding behavior in sequential settings. Hung and Plott (2001) report the robustness of informational cascades under different rewards systems. Çelen and Kariv (2004) elicit subjects' beliefs in order to identify the source of imitation behavior. Kubler and Weiszacker (2004) allow subjects to purchase a private signal for a small fee. Angrisani et al. (2005) embed this game in a financial market setting. Goeree, Palfrey, Rogers, and McKelvey (2007) look at longer-horizon games. Ziegelmeyer et al. (2010) study the fragility of information cascades patterns using groups of differently informed agents. Weiszacker (2010) conducts a meta-study of sequential social learning experiments and finds that subjects follow their own private information more frequently than is empirically optimal.¹ Duffy, Hopkins, Kornienko, and Ma (2019) allow

¹Ziegelmeyer, March, and Krugel (2013) enlarge the data and use a modified methodology compared with Weiszacker (2010). They find that subjects tend to over-weigh their private signals but do so to a lesser degree than previously found. De Filippis et al. (2022) identify that such over-weighting occurs only when private signals are in conflict with social information. See also Angrisani et al. (2021) who study social learning in a continuous action space experiment and disentangle different theories that deliver over-weighting of private information. Eyster and Rabin (2014) investigate agents who naively believe that the actions of others solely

subjects to choose between private and public information prior to guessing the state, while [Duffy, Hopkins, and Kornienko \(2020\)](#) in addition vary the persistence of the state across periods.

In comparison to this work, the repeated social learning game we study is quite different from the sequential move games surveyed above as, in our case, the same group of agents interacts repeatedly. This setting allows us to study how the information and the actions of others impact the beliefs and actions in an arguably more realistic setting. Accordingly, our paper is related to a new and exciting experimental literature that compares social learning outcomes under different network structures ([Mueller-Frank and Neri, 2015](#); [Grimm and Mengel, 2020](#); [Chandrasekhar, Larreguy, and Xandri, 2020](#); [Dasaratha and He, 2021](#); [Choi, Gale, and Kariv, 2012](#); [Agranov, Gillen, and Persitz, 2025](#)). These papers vary the connections between agents in a group, i.e., who can observe whose actions, and ask how well subjects aggregate dispersed private information through repeated observations of neighbors' actions. Contrary to this literature, we fix the network structure to be complete, and vary the type of information our agents observe.²

Two recent laboratory games deserve special attention. The first one is [Eyster, Rabin, and Weizsacker \(2018\)](#) which compares social learning outcomes in a standard sequential move game and a cleverly constructed four-at-a-time move game. The sequential move game features participants who move one at a time after observing their predecessors' actions. The four-at-a-time game has four participants move in each period after observing predecessors from previous periods. Theoretically, Bayesian inference prescribes subjects to anti-imitate their predecessors in the latter but not the former game. The experimental data shows that subjects rarely anti-imitate in both games, which is at odds with the rational model. The four-at-a-time move game is interesting as it is one of the first environments in which, due to the incorrect processing of social information by participants, social learning is harmful on average, i.e., participants would do much better if they just ignored their predecessors actions and focused on their own private information instead.

The second paper is [Evdokimov and Garfagnini \(2020\)](#) who focus on groups with two players and vary whether both players, only one, or no one observes others' actions in a repeated social learning game. Similarly to our game, each agent gets a conditionally independent private signal in each period and reports her best guess about the state. However, contrary to our main results, the authors find no significant differences in the average quality of guesses across informational treatments. This difference between our papers indicates that group size plays an important role when aggregating social information.

From a theory perspective, our NO INFO, SIGNALS and ALL treatments are trivial to represent their private signals.

²Moreover, in our game, participants get private signals in each game round, while in the papers above private signals are distributed only once at the very beginning of the game.

analyze, as they are equivalent to a single agent problem for a Bayesian decision maker. In contrast, the ACTIONS treatment is difficult to analyze within a rational framework, because higher order beliefs play a significant role. We elaborate on this in Section 4. Vives (1993) studies the speed of learning in a similar setting, but with a continuum of agents and continuous signals. Harel et al. (2021) study a setting that is identical to ours, under the assumption that agents are Bayesian and myopic. They focus on the long-term rate of learning, and show that even if many more agents are added to the group, the speed of learning remains bounded. Huang et al. (2021) extend this result to a network setting and to forward-looking agents.

3 Experimental Design

Subjects are presented with the following scenario. There are two possible urns: a “red” urn and a “green” urn. The red urn contains six red and four green balls, and the green urn contains six green and four red balls. The color of the urn represents the state of nature (the superior policy, the best candidate for a job, the tastier croissant, etc.), which subjects are trying to learn by interacting repeatedly throughout the game. The structure and a sample of instructions are presented in sections D and E of the Online Appendix, respectively.

In all experiments, subjects play 10 games. At the start of each game, subjects are randomized into a group of eight subjects and one of the urns is chosen by a toss of a fair coin. This urn then remains unchanged for the duration of the game, and is used by all the participants of the game. Each game consists of 20 rounds. In each round, a subject guesses the color of the urn. We refer to these guesses as her actions. She then receives an independent draw (with replacement) from the chosen urn. The color of the drawn ball matches the color of the urn with probability 60%, which is therefore the precision of one’s private signal; we purposefully chose an environment with low precision signals because in this environment the signals and actions of others are particularly valuable.

Depending on the treatment, subjects may observe additional information at the end of each round, which may help them choose their actions for the follow-up rounds. At the end of a session, a random round of a random game is selected uniformly. Subjects are rewarded for guessing correctly in that round: the correct action earns \$20, while the incorrect one earns \$5.³

³Paying for a single round chosen at random (as opposed to paying for all rounds) eliminates hedging incentives, which might cause subjects to switch between guesses. Moreover, the standard “isolation” assumption—that participants treat each game as independent—coupled with myopic optimization assumption—that participants report their best guess in each round of a game—imply that incentive compatibility is preserved (Azrieli et al., 2018). In an alternative design, payment is based on the accuracy of last-round guess alone. In principle, such a payment design could induce a wide variety of behavior. For example, in the actions treatment (where agents observe others’ actions but not their signals) agents could in principle spend all periods but the last playing actions equal to their signals, thus revealing their signals to each other.

Our main variation between treatments is the information available to subjects at the time they choose their actions (in addition to their own signals). We consider four information structures:

- (1) The NO INFO treatment, in which subjects observe only their private signals in every round of a game.
- (2) The ACTIONS treatment, in which in addition to observing their own private signals, subjects also get to see the actions chosen by their group members in all previous rounds of the game.
- (3) The SIGNALS treatment, in which subjects observe both their own private signals and the private signals of their group members in all previous rounds of the game.
- (4) The ALL treatment, in which subjects observe their own private signals, their group members' private signals and the actions chosen by their group members in all previous rounds of the game.

We conduct two additional treatments with the same information structure as the ACTIONS and the ALL treatments, but with smaller groups of four members each. We call these treatments ACTIONS4 and ALL4. These treatments allow us to explore how fast small and large groups learn depending on the information available to their members. Note that when subjects take their first action in round 1, they have no additional information about the state except for the prior. This sequencing of events within a round allows for a clean comparison between information treatments as we describe below, where all the information is delivered at the end of the round before the next-round actions are taken.

Throughout the game, subjects have access to a table that keeps track of all signals and actions they observed in the past rounds of the game. This table presents information in an intuitive and visual way. This feature of the design ensures that our results are not affected by the memory subjects may have about the events that transpired during the game (see section E in the Online Appendix for screenshots).

At the end of the experiment, we elicit subjects' strategies with a series of open-ended questions as well as their beliefs about the fraction of correct actions of the other participants in the last round of the game in various treatments. Subjects are paid for the accuracy of their prediction in one randomly selected belief question. Subjects also report their gender and major, and complete a series of control tasks including risk attitudes elicitation using two investment tasks (Gneezy and Potters, 1997), IQ questions (ICAR, Condon and Revelle,

Alternatively, in any treatment, given the cognitive cost of guessing well, agents could just (say) always choose red until the last round, where they could make a good guess.

2014), and overconfidence.⁴ In the analysis of experimental data, we classify subjects into *low-IQ* and *high-IQ* based on their answers to six IQ questions. We use their self-described strategies and beliefs to measure how useful low-IQ and high-IQ subjects think observing different types of information is, and whether it relates to their behavior in the experiment.

Table 1: Experimental Design

Treatment	Group size	Information				# of sessions	# of subjects
		signals		actions			
		own	others	own	others		
NO INFO	1	yes	no	yes	no	4 sessions	80 subjects
ACTIONS	8	yes	no	yes	yes	8 sessions	136 subjects
SIGNALS	8	yes	yes	yes	no	3 sessions	82 subjects
ALL	8	yes	yes	yes	yes	8 sessions	152 subjects
ACTIONS4	4	yes	no	yes	yes	4 sessions	76 subjects
ALL4	4	yes	yes	yes	yes	4 sessions	80 subjects
Total						31 sessions	606 subjects

Table 1 summarizes the experimental design and number of participants. Experimental sessions were conducted at two locations: University of California in San Diego (UCSD) and the Ohio State University (OSU).⁵ The experiment was programmed and conducted with the oTree software (Chen et al., 2016) and was pre-registered in the AEA RCT registry (AEARCTR-0003315). Overall, 606 subjects participated in 31 sessions, and no subject participated in more than one session. The experiment lasted about 90 minutes. Subjects earned on average \$25.7, including a \$7 participation fee.

4 Theoretical Preliminaries

We model each game played by subjects in the experiment as follows. Denote by $\omega \in \{R, G\}$ the state of nature, which is chosen uniformly at random and stays fixed throughout each game. Denote by N the set of players. At the beginning of each round $t \in \{1, \dots, 20\}$, each player i chooses an action $a_{i,t} \in \{R, G\}$. She then observes a random signal $s_{i,t} \in \{R, G\}$ such that conditioned on the state ω , the probability that the signal equals the state (i.e., $s_{i,t} = \omega$) is 60%. Conditional on the state, the signals are independent, both between time periods and players. Depending on the treatment, the players observe additional information;

⁴See sections E.2, E.3 and E.4 in the Online Appendix for details on the strategies, beliefs and other controls, respectively.

⁵Because of the closure of physical labs due to COVID-19, our data collection process at UCSD was interrupted and we had to finish collecting the data online at OSU using the same subject pool of students who normally participate in laboratory experiments. NO INFO, ACTIONS, and 4 sessions of ALL treatments were conducted in the experimental laboratory at UCSD, while the remaining 4 sessions of ALL treatment and SIGNALS treatment were conducted online at OSU. In section F of the Online Appendix, we compare data collected at the physical lab at UCSD and in the online lab at OSU for the ALL treatment. We find that aggregate outcomes and individual level behavior of subjects in these two types of sessions are very similar to each other and statistically indistinguishable.

we denote by $I_{i,t}$ all that player i has observed before choosing her action $a_{i,t}$ at the beginning of round t . We use the terms action and guess interchangeably throughout the text.

Each player plays some number of games. At the conclusion of player i 's session one of these games is chosen uniformly at random, and in that game a round t is chosen uniformly at random. Player i receives a payoff equal to \$20 if she chose the correct action $a_{i,t} = \omega$, and to \$5 otherwise. We assume that players are Bayesian and that they have a strictly increasing utility for money. Under this assumption, behavior depends on the information players observe, but *not* on their preferences over risks.⁶

The NO INFO and SIGNALS treatments. In these treatments the information $I_{i,t}$ available to player i in round t consists of a collection of signals that she observed up until this round. Specifically, in the NO INFO treatment this collection consists of her own signals only, $I_{i,t} = (s_{i,\tau})_{\tau < t}$, while in the SIGNALS treatment this collection contains signals of all group members, $I_{i,t} = (s_{j,\tau})_{j \in N, \tau < t}$. A standard calculation shows that the action $a_{i,t}$ is optimal if it is equal to any color that appears in $I_{i,t}$ at least as often as the other. In other words, the optimal action $a_{i,t}$ is equal to the majority color observed in the past signals, and can be both in case of a tie. Note the distinction between the correct action, i.e., the one that matches the state, and the optimal action, which is the one that matches the majority of the observed signals.

Given that subjects observe eight times more signals in the SIGNALS treatment than in the NO INFO treatment, theory predicts that subjects should be better at guessing the state in the former case in every round of the game (except the first, where they have no information). In particular, given our parameters, in the second round of the game, subjects should guess the state correctly 60% of the time in the NO INFO treatment and 71% of the time in the SIGNALS treatment. In the last round of the game, subjects are expected to guess the state 81% of the time in the NO INFO treatment and 99% of the time in the SIGNALS treatment. In Table 2 we provide the predicted probability of choosing correctly in each round.

The ALL Treatment. In the ALL treatment, subject i in round t observes all the signals and actions of her group members up until round t , so that $I_{i,t} = (s_{j,\tau}, a_{j,\tau})_{j \in N, \tau < t}$. Since the past actions $(a_{j,\tau})$ can contain no further information than the signals $(s_{j,\tau})$, in this case too optimal behavior implies that $a_{i,t}$ is equal to any color that appears in $(s_{j,\tau})_{j \in N, \tau < t}$ at least as often as the other. In other words, for Bayesian subjects we expect identical results

⁶As there are only two outcomes, the possible payoff distributions are ranked in first order stochastic dominance. Therefore, the player will choose the same strategy under any preference that is monotone in first order stochastic dominance (e.g., cumulative prospect theory, etc.), as long as she forms beliefs using Bayes' rule.

Table 2: Probabilities of Correct Actions

t	NO INFO	ACTIONS	SIGNALS/ALL	ACTIONS4	ALL4
1	0.50	0.50	0.50	0.50	0.50
2	0.60	0.60	0.71	0.60	0.65
3	0.60	0.73	0.79	0.68	0.71
4	0.65	0.77	0.84	0.71	0.75
5	0.65	0.78	0.87	0.73	0.79
6	0.68	0.80	0.90	0.75	0.81
8	0.71	0.83	0.93	0.77	0.86
10	0.73	0.84	0.96	0.79	0.89
12	0.75	0.86	0.97	0.81	0.91
14	0.77	0.87	0.98	0.83	0.93
16	0.79	0.88	0.99	0.84	0.94
18	0.80	0.89	>0.99	0.85	0.95
20	0.81	0.90	>0.99	0.86	0.96

Notes: Predicted probability of correct actions for parameters used in the experiment. In the ACTIONS treatments we assume common knowledge of rationality and that players act myopically.

in the ALL and SIGNALS treatments, since subjects should ignore others' actions and base their decisions on the observed signals.

The ACTIONS Treatment. The situation is more complicated in the ACTIONS treatment. Here, the information available to subject i in round t consists of her private signals and of the actions of the other subjects in all previous rounds, i.e., $I_{i,t} = (s_{i,\tau}, a_{j,\tau})_{j \in N, \tau < t}$. To the best of our knowledge, this case is analytically intractable. Nevertheless, a simulation of equilibrium behavior is possible under the assumption that players act myopically, i.e. do not change their action to manipulate the future behavior of others.⁷ Formally, in this simulation we assume that in each period each player maximizes the probability of making a correct choice, i.e.

$$a_{i,t} = \begin{cases} R & \text{if } p_{i,t} > 0.5 \\ G & \text{if } p_{i,t} < 0.5 \end{cases} \quad (1)$$

where $p_{i,t}$ denotes the probability player i assigns to state R at the beginning of period t . When $p_{i,t} = 0.5$, the player is indifferent and we assume that she randomizes uniformly over the two actions. Each player computes this probability using Bayes' rule taking into account the actions of others. We assume common knowledge of rationality, so each player knows that the others are making their decisions according to (1). We thus assume that each player does not only know the signal generating process, but also the strategies of all other players, which is arguably a strong rationality requirement. In Table 2 we report the probabilities of correct actions under these assumptions. We provide these simulation results

⁷This assumption is commonly made in the literature, see e.g., Parikh and Krasucki (1990); Gale and Kariv (2003); Mossel et al. (2014); Harel et al. (2021) and seems plausible given the complexity of the environment.

as a benchmark for what Bayesian agents can hope to achieve in this setting; our results do not hinge on a comparison of the subjects’ behavior to this benchmark.

5 Results

We present results from our experiments in the following order. First, we focus on the public data setting, and compare the SIGNALS and ALL treatments (Section 5.1). Second, we study the private data setting (the ACTIONS treatment) and investigate group size effects (Section 5.2). As we go through the analysis, we summarize the main empirical findings as observations.

Data Analysis Approach. Our main analysis considers all ten games in all sessions.⁸ We investigate the effect of the treatment on the fraction of correct actions and on the consensus rates using regression analysis. To compare performance across treatments, we estimate a linear regression of the indicator of a correct action by a given participant i in a round t of game g on a round-specific treatment effect.⁹ To compare consensus rates across treatments, we regress the relative size of the majority in a given round t on a round-specific treatment effect.^{10,11}

Throughout our analysis, we classify the collection of agents’ signals $(s_{i,\tau})_{\tau < t}$ according to their strength, as given by the difference in the number of signals of one color relative to the other. We define five categories of signal strength: *Very Strong Green*, *Strong Green*, *Weak*, *Strong Red*, and *Very Strong Red*. These correspond, respectively, to the difference between the number of red and green signals being less than -26 , in $[-26, -15]$, in $[-14, 14]$,

⁸We see moderate learning across games within a session. Section G in the Online Appendix shows aggregate results subsetting our data for early (first 5) and late (last 5) games.

⁹Specifically, let v_{it}^g equal one if participant i guessed optimally in round t of game g , and zero otherwise. Let $\text{treat}_i \in \{\text{SIGNALS}, \text{ALL}\}$ denote the treatment assigned to participant i . We estimate the regression

$$v_{it}^g = \alpha + \beta \text{treat}_i + \gamma_t + \delta_t \text{treat}_i + \varepsilon_{it}^g,$$

where γ_t are round fixed effects and δ_t capture round-specific treatment effects, normalized so that $\delta_1 = 0$. Under this normalization, β measures the treatment effect in round 1, and δ_t captures the differential treatment effect in round t relative to round 1. Identification comes from within-participant variation across rounds and across games.

¹⁰Specifically, let c_t^g denote the share of participants in the majority subgroup in round t of game g , and let $\text{treat}^g \in \{\text{SIGNALS}, \text{ALL}\}$ denote the treatment assigned to participants in game g . We estimate the regression

$$c_t^g = \alpha + \beta \text{treat}^g + \gamma_t + \delta_t \text{treat}^g + \varepsilon_t^g,$$

where γ_t are round fixed effects and δ_t capture round-specific treatment effects, normalized so that $\delta_1 = 0$. Under this normalization, β measures the treatment effect in round 1, and δ_t captures the differential treatment effect in round t relative to round 1. Identification comes from within-game variation across rounds.

¹¹All standard errors are clustered at the session level to account for the inter-dependencies of observations that come from re-matching subjects within a session. Appendix A reports the average treatment effects across rounds as well as treatment effects in early (rounds 2 to 10) and late rounds (rounds 11 to 20).

in [15, 26], and above 26. The boundaries between these categories correspond to the 10th, 25th, 75th and 90th percentile of the distribution at the end of the game. For the small group treatments we adjust the intervals to match the same percentiles.¹²

In some discussions, we only distinguish between “weak” and “strong” signals by pooling together the *Strong* and *Very Strong* categories for each color. An alternative approach of defining the signal strength based on the *fraction of signals* instead of the difference in the number of signals of each color yields similar results. We discuss these two approaches in detail in Section 5.1.

5.1 The Public Data Setting and the Effect of Redundant Information

In this section we compare the ALL and SIGNALS treatments. Panel (a) in Figure 1 presents our main outcome of interest: how often subjects guess the state correctly in each round of the game, in comparison to the Bayesian benchmark.




Figure 1: Aggregate statistics in the ALL and SIGNALS treatments

Notes: Panel (a) presents the average frequency of correct actions in each round, averaged across games. Panel (b) depicts the evolution of consensus in each round, i.e., the relative size of the majority, averaged across games. For panel (b) we exclude cases with an equal number of green and red signals. Shaded regions represent 95% confidence intervals from 50% (darkest) to 95% (faintest) probability levels. Confidence intervals are constructed with a variance-covariance matrix clustered by session.

¹²In section H of the Online Appendix we show that the qualitative results reported in this section are robust to using different cutoffs of these categories.

In both treatments, the first action that subjects make is a coin toss. This is expected since this action is made before any signals are observed, so the probability of guessing correctly is one half. As the game progresses and more information arrives, we observe an upward trend in the likelihood of choosing the correct action.

The main finding of this paper is that subjects perform better in the ALL treatment, as compared to the SIGNALS treatment. This happens despite the aforementioned fact that the only difference between the two treatments is the observability of other players' actions, which are informationally redundant. In Online Appendix L, we show that the performance advantage of the ALL treatment over the SIGNALS treatment remains robust after conditioning on the distribution of prediction quality across participants. In particular, participants throughout the distribution perform better in the ALL treatment than in the SIGNALS treatment.

Across rounds, the percentage of correct actions is on average 5% larger in the ALL treatment, as compared to the SIGNALS treatment.¹³ In the next section, we explore in depth why this is the case. We note that both treatments significantly underperform relative to optimal behavior, which means that in both treatments at least some subjects deviate from the Bayesian benchmark.¹⁴

Panel (b) of Figure 1 measures the polarization of opinions, presenting the evolution of consensus rates. The consensus rate is defined as the relative size of the majority subgroup in a given round, based on members' actions. This rate varies between one half and one, where a consensus rate of one half indicates a maximally polarized group with half of the members choosing each of the two actions, while a consensus rate of one indicates the case where all members choose the same action in a particular round.

Overall, the consensus rates are increasing in both treatments as the game progresses, showing that group members' opinions tend to align the longer they interact with each other. This is to be expected, since actions become more correlated with the state, and hence with each other, as subjects gather more information. Nevertheless, consensus rates are significantly less than one, which would be the theoretical prediction for Bayesian agents.¹⁵

More interestingly, the consensus rates are significantly higher in the ALL treatment than in the SIGNALS treatment: subjects agree more when they see others' actions and raw data, than when they only see raw data. On average, consensus rates are 10% higher in the ALL treatment than in the SIGNALS treatment.¹⁶ This is consistent with subjects being influenced

¹³See panel (a) of Figure A.1 in Appendix A and columns (1) and (2) in Table A.1 for point estimates and standard errors.

¹⁴Online Appendix N analyzes the extent to which these deviations are driven by participants who behave as if random, ignoring all relevant information in their signals.

¹⁵Note that Panel (b) excludes cases in which the number of red and green signals is the same, for which theory does not provide a unique prediction on the size of the majority.

¹⁶Additional evidence is presented in panel (b) of Figure A.1 and average point estimates and standard

by the actions of others. In Online Appendix K, we complement these results by quantifying the improvement in performance from using the redundant information in others’ actions beyond the use of signals. Specifically, we compute the root mean squared error (RMSE) between the Bayesian benchmark probabilities of correct actions (as reported in Table 2) and model-predicted probabilities under two specifications: one in which participants’ behavior depends only on signals, and another in which it depends on both signals and others’ actions. We find that incorporating others’ actions significantly reduces RMSE, with a magnitude equivalent to approximately one-third of the performance improvement attributable to signals.

Finally, focusing on cases with strict majorities, we note that the probability that the majority is correct is extremely high and exceeds 90% on average in both the ALL and SIGNALS treatments. Furthermore, it is 4% higher in the ALL treatment, as compared to the SIGNALS treatment.¹⁷

***Observation 1:** People learn faster and develop more unified opinions when they observe each other’s signals and actions, compared to observing signals only.*

What Drives These Aggregate Results? To what extent do subjects’ actions follow the information contained in the signals they observe? We start with the simplest statistic, i.e., the second round behavior, which, when performed optimally, entails reporting the color of the majority of the eight signals observed in the first round of the game. In both treatments, a large fraction of around 84% of our subjects choose optimally in the second round.

As more and more signals are observed, a Bayesian subject would keep a tally of the signals and report the majority color. Panel (a) in Figure 2 shows the frequency with which subjects chose the color red as aggregated over the observed difference between the total number of red and green signals. Panel (b) shows the estimated probability of an optimal action, i.e., reporting the color of the majority of signals, as a function of signal strength.¹⁸

errors are in columns (3) and (4) in Table A.1 in Appendix A.

¹⁷The evidence is presented in Figure A.2 and columns (5) and (6) of Table A.1 in Appendix A.

¹⁸To recover the probabilities in panel (b) of Figure 2, we estimate a Bayesian logistic regression pooling the data from the SIGNALS and ALL treatments, where the probability of a red action, $Pr(a_{it} = R)$, is modeled as a function of treatment and signal strength, accounting for session-level random effects. Specifically,

$$\text{logit}(Pr(a_{it} = R)) = \beta_{\text{strength}_j} \text{treat}_i + \gamma_{\text{session}(i)}, \quad (2)$$

where $\text{treat}_i \in \{\text{SIGNALS}, \text{ALL}\}$ denotes the treatment assigned to participant i , $\beta_{\text{strength}_j} \sim \mathcal{N}(0, \sigma_{\text{strength}})$ is a random slope that varies across signal-strength categories j , and $\gamma_{\text{session}} \sim \mathcal{N}(0, \sigma_{\gamma})$ is a session-level random intercept. As a robustness check, Table I.2 in the Online Appendix reports similar estimates using linear probability models with standard errors clustered at the session level. We obtain similar results when additionally accounting for game, round, and participant effects.




Figure 2: Learning from signals in the SIGNALS and ALL treatments

Notes: Panel (a) depicts the fraction of red actions as a function of the difference between the number of red and green signals. The size of the dot corresponds to the number of observations in each bin. Panel (b) shows the estimated probability of an optimal action (i.e., reporting the color of the majority of signals) as a function of the treatment and signal strength. We estimate a Bayesian logistic regression of the probability of taking the red action on treatment that varies by signal strength and controlling for session random effects. We present the estimated posterior median and 95% confidence intervals. The categories in Panel (b) for signal strength are as defined in Section 5.

This figure suggests several insights. First, while panel (a) would look like a step function for a Bayesian (the green line), in the data we see a gradual increase in the frequency of choosing red. Thus, subjects respond to signals, and do so in the correct direction, but not perfectly. Second, the mistakes are much more frequent when signals are weak, i.e., when the majority signal color has occurred only slightly more than the other color. Because the expected utility is a function of the difference between the number of red and green signals, this is consistent with a Luce model of behavior, in which choice probabilities are related to the difference in expected payoffs (Luce, 1958). It is also consistent with standard models of perception where close-by states are harder to distinguish. Third, subjects systematically perform better in the ALL treatment across different signal strengths. In particular, when signals are weak subjects are 6% more likely to choose correctly in the ALL treatment compared to the SIGNALS treatment (73% for the SIGNALS versus 79% for the ALL treatment).




Figure 3: Learning from others' actions in the ALL treatment

Notes: This figure depicts the probability of choosing red as a function of the share of red actions of other group members. The estimates are obtained from a Bayesian logistic regression of subjects' actions on the share of others' actions in the previous round conditional on signal strength and session random effects. Shaded regions represent 95% confidence intervals from 50% (darkest) to 95% (faintest) probability levels.

Figure 3 provides a closer look at behavior in the ALL treatment. To quantify the informational value of observing others' actions, beyond the raw data, we estimate a Bayesian logistic regression in which a subject's action is regressed on the share of others' actions in the previous round. The effect of others' actions is allowed to vary across signal-strength categories, as previously defined.¹⁹ We find that when signals are strong or very strong, a subject's actions are barely influenced by those of others: regardless of what others do, subjects tend to follow the majority of signals. This can be seen by mostly flat lines in the bottom and top of Figure 3. In contrast, when the number of signals of each color is comparable, subjects rely heavily on the actions of others, as is indicated by the steep middle line, which shows a strong responsiveness of subjects to others' actions.

An alternative approach, in which one classifies signal strength by the *fraction* of signals

¹⁹Specifically, we estimate

$$\text{logit}(Pr(a_{it} = R)) = \beta_{\text{strength}_j} \text{share}_{a_{k \neq i, t-1}} + \gamma_{\text{session}(i)}, \quad (3)$$

where $\text{share}_{a_{k \neq i, t-1}}$ is the share of others' actions in round $t-1$. $\beta_{\text{strength}_j} \sim \mathcal{N}(0, \sigma_{\text{strength}})$ is a random slope for signal-strength category j , and $\gamma_{\text{session}} \sim \mathcal{N}(0, \sigma_{\gamma})$ is a session-level random intercept.

instead of the difference in the *number* of signals of each color yields similar results. This alternative approach is inspired by prior experiments, which found that some people rely on sample proportion over sample size in probability judgement tasks (see, e.g., Benjamin, 2019).²⁰

***Observation 2:** In both the SIGNALS and ALL treatments, subjects use signals to choose their actions, but do so less well than in the Bayesian optimum. In particular, subjects tend to make more mistakes when the tally of signals is close, i.e., when signals are not very informative. In the ALL treatment, subjects partially correct for these mistakes by relying on others' actions when signals are weak, which accelerates learning as compared to the SIGNALS treatment.*

Individual Level Analysis. Our results so far show that agents condition on the social signal, i.e., on the actions of others, which improves performance in the ALL treatment, as compared to the SIGNALS treatment. In this section we calculate a number of statistics describing a subject's behavior in our game and explore its joint distribution in the population. We analyze subjects' behavior by their level of IQ, given our auxiliary IQ measure.²¹

We classify subjects into low-IQ and high-IQ based on their answers to six IQ questions: three matrix reasoning questions, which are similar to Raven's Progressive Matrices, and three 3-D rotation questions (Chapman et al., 2019).²² The low-IQ subjects are those who answered at most half of the IQ questions correctly and the high-IQ subjects are the remaining group.²³ Despite its obvious coarseness, our IQ measure correlates with subjects' performance in the main game: high-IQ subjects are approximately 9% more likely to guess the state correctly than the low-IQ ones in both the ALL and the SIGNALS treatments

²⁰One may ask which of the two models provides a better fit to the data, i.e., whether participants behave as if they condition their actions on the fraction or on the number of signals of each color. Our data slightly favors the former specification, although the difference is minimal: the average fraction of correctly predicted guesses is 72% versus 71%. Figure H.2 in the Online Appendix replicates Figure 3 using the fraction of signals to measure signal strength and reports the in-sample fit, in terms of the fraction of correctly predicted guesses, for the two alternative approaches.

²¹Online Appendix M examines covariate balance across treatments. Using a paired matched sample to address any imbalances, we replicate the main result that participants perform better in the ALL treatment than in the SIGNALS treatment. Relative to the full sample, performance in the matched sample is higher in both treatments, with outcomes approaching the Bayesian benchmark, in particular in the ALL treatment.

²²The questions are drawn from the International Cognitive Ability Resource, a public domain intelligence measure (ICAR; Condon and Revelle (2014)). In each of the first three questions, participants are asked to determine which of the options completes a graphic pattern. In the remaining questions, participants are asked to identify which of the presented drawings of a cube were compatible with another drawing of a cube. In general, these questions have been shown to capture a variation of fluid intelligence in the general population and is becoming one of the popular measures of IQ on par with CRT tests and general knowledge tests.

²³The average number of correct answers is 3.7 in the ALL treatment and 3.4 in the SIGNALS treatment. This classification delivers roughly similar proportions of low-IQ subjects across treatments: 47% in the ALL treatment and 40% in the SIGNALS treatment.

($p < 0.01$).²⁴

Next, we investigate whether there are systematic differences in how low-IQ and high-IQ subjects process the information available to them. We start with the SIGNALS treatment and calculate two statistics for each subject: (i) *responsiveness to weak signals*: the probability that they choose optimally (i.e., with the majority of signals, as a Bayesian would) when signals are weak and (ii) *responsiveness to strong signals*: the probability that they choose optimally when signals are strong. This probability is computed via a Bayesian logistic regression of the subject’s action on our measure of signal strength. To compute responsiveness at the individual level, we exploit the variation across rounds and games for a given subject, which allows us to recover measures of responsiveness at the individual level via individual-specific random effects.²⁵ The top part of Figure 4 presents the kernel distributions of an individual responsiveness to signals in the SIGNALS and ALL treatments. The responsiveness is calculated as the posterior median of the probability of an optimal action by the participant. For the ALL treatment we interact it with the social signal measured by the actions of other group members in the previous round as in equation (3).²⁶ To isolate the responsiveness to raw signals, we set the number of others who take each action to be equally split between red and green.

Two patterns are apparent from Figure 4. First, in both treatments, the high-IQ participants are on average much more responsive to raw signals than the low-IQ ones, both when signals are weak and when they are strong. In fact, high-IQ subjects are very close to the Bayesian benchmark when it comes to strong signals, but some of them make a significant number of mistakes with weak signals. These mistakes are, however, less frequent than the mistakes of the low-IQ subjects, which explains why low-IQ subjects perform worse

²⁴Tables J.1 and J.2 in the Online Appendix show the estimates and clustered standard errors of a linear probability model of correct actions on the 6-question IQ measure and the low-IQ/high-IQ indicator, respectively. The relationship between accuracy and IQ is robust to controlling for game and round effects as well as for other subject characteristics.

²⁵Specifically, we estimate an individual-level specification of equation (2) for the SIGNALS treatment:

$$\text{logit}(Pr(a_{it} = R)) = \alpha_i + \beta_{i,\text{strength}_j} + \gamma_{\text{session}(i)}, \quad (4)$$

where α_i is an individual-specific random intercept, $\beta_{i,\text{strength}_j}$ is an individual-specific random slope that varies across signal-strength categories j , with $\begin{pmatrix} \alpha_i \\ \beta_{i,\text{strength}_j} \end{pmatrix} \sim \mathcal{N}(0, \Sigma)$ and $\gamma_{\text{session}} \sim \mathcal{N}(0, \sigma_\gamma)$ is a session-level random intercept. For a fully Bayesian participant i we expect to see 100% responsiveness for both weak and strong signals.

²⁶Specifically, the logit regression takes the form

$$\text{logit}(Pr(a_{it} = R)) = \alpha_i + \beta_{i,\text{strength}_j} \text{share}_{a_{k \neq i, t-1}} + \gamma_{\text{session}(i)}, \quad (5)$$

where $\text{share}_{a_{k \neq i, t-1}}$ is the share of others’ actions in round $t - 1$. $\beta_{i,\text{strength}_j}$ is an individual-specific random slope that varies across signal-strength categories j , with $\begin{pmatrix} \alpha_i \\ \beta_{i,\text{strength}_j} \end{pmatrix} \sim \mathcal{N}(0, \Sigma)$ and $\gamma_{\text{session}} \sim \mathcal{N}(0, \sigma_\gamma)$ is a session-level random intercept.

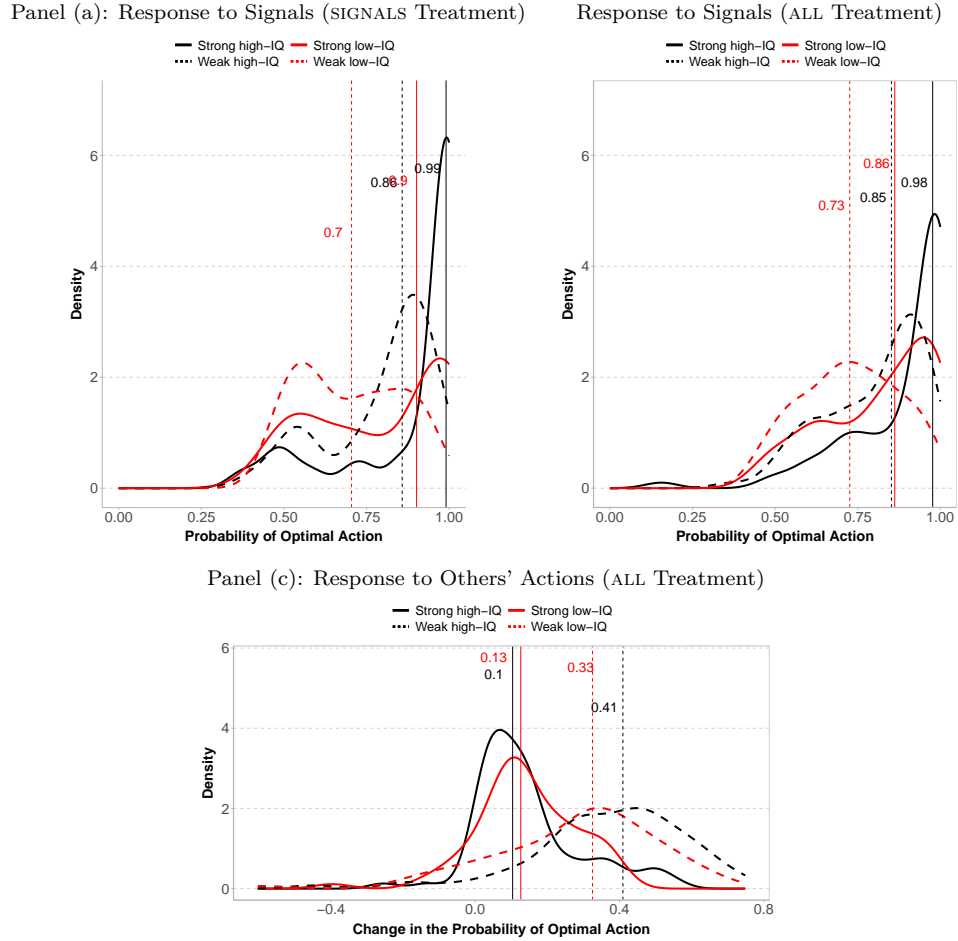


Figure 4: Responsiveness to signals and actions, individual level data

Notes: Kernel distributions of participants' responsiveness to weak and strong signals are presented in the top part of the graph (Panel (a) for SIGNALS and Panel (b) for ALL treatments). Responsiveness to signals is calculated based on regressions (4) and (5) for the SIGNALS and ALL treatments, respectively. Responsiveness is given by the probability that a participant's action is optimal. For the ALL treatment, we also set the actions of others group members in the previous round to be uninformative, i.e., split equally between green and red. Kernel distributions of participants' responsiveness to others' actions are presented in panel (c) for the ALL treatment. Responsiveness to others' actions is measured by the change in the probability of choosing the action of the majority of signals when all versus none of the other group members choose the majority of signals in the last round. The vertical lines and the numbers next to them are median responsiveness for each group.

than the high-IQ ones in the SIGNALS treatment.

Second, there is substantial heterogeneity in individual responsiveness to raw signals among participants in both treatments. For instance, about a quarter of high-IQ subjects in

the ALL treatment have estimated responsiveness of above 90% for both weak and strong signals, compared to only 7% of low-IQ subjects with similar estimates. This means that the social signal, measured by the average action of all group members in the ALL treatment, is informative since it contains a sizable fraction of correct actions.

Who uses this social signal, and does it affect one’s performance? Panel (c) of Figure 4 sheds light on this question. For both low-IQ and high-IQ participants, we compute individual responsiveness to social signals as the difference in the probability of choosing optimally when all others agree with the majority of signals minus this probability when all others disagree with the majority of signals.²⁷ As shown in the figure, both groups of participants rarely condition on others’ actions when signals are strong, but do so to a substantial degree when signals are weak. On average, high-IQ participants rely on others’ actions to a greater extent than low-IQ participants ($p = 0.004$).

Importantly, because the social signal is informative, conditioning on it improves performance for all participants, not only low-IQ individuals. Among low-IQ participants, those whose responsiveness to others’ actions under weak signals is above the group median guess the state correctly 86% of the time, compared to 70% for those below the median ($p < 0.001$). The same pattern holds for high-IQ participants, with success rates of 91% versus 81% for those above and below the median, respectively ($p = 0.001$).

In other words, a substantial fraction of participants effectively aggregate information, making the social signal informative and worth conditioning on.

***Observation 3:** The high-IQ subjects process the raw signals and perform well, regardless of whether social information is available. The low-IQ subjects do not process well the information contained in the signals, and hence perform badly when social information is not available. When social information is available, it is used by subjects in both groups and boosts the performance of those who use it.*

5.2 The Private Data Setting and the Effect of Group Size


We now turn our attention to the ACTIONS treatment, in which people observe each other’s actions but not others’ signals. We ask whether they can parse useful information from these choices. This task is complex since actions of group members are inherently correlated and making such inferences requires forming beliefs about how others act. Even under heroic assumptions about common knowledge of rationality among group members, theory is only able to characterize long run outcomes such as the speed of learning, but not how individuals

²⁷For example, if a participant chooses the correct action with probability 0.69 when all observed actions from the previous round align with the majority of signals, and with probability 0.42 when they disagree, then the implied responsiveness to social signals is 27 percentage points.

should act in the short term (Harel et al., 2021). Our experimental data is particularly useful in these circumstances as it provides the first step at documenting how people actually behave in this complex situation.

Panel (a) in Figure 5 compares the performance in the ACTIONS treatment with two benchmarks: the NO INFO and the SIGNALS treatments. The former benchmark is the lower bound of learning rates, which is what one expects to happen in the ACTIONS treatment if subjects ignore entirely each other’s actions and base their decisions only on the sequence of private signals they receive. The latter benchmark is the upper bound, in which the signals of all members are public; this is what would happen if subjects could perfectly infer the signals of others from their actions.

Figure 5: Frequency of correct actions, by information structure and group size



Notes: Both panels present the average frequency of correct actions in each treatment per each round, averaged across games. Shaded regions represent confidence intervals from 50% (darkest) to 95% (faintest) probability levels. Confidence intervals are constructed with a variance-covariance matrix clustered by session.

Panel (a) shows that, despite the complexity of the inference problem in the ACTIONS treatment, people are able to extract useful information from others’ actions and learn faster than they would without this information, as the comparison to the NO INFO treatment shows. For instance, notice that subjects chose the correct action with a higher probability in period 13 of the ACTIONS treatment than in period 20 of the NO INFO treatment. Thus, the benefit of observing others’ actions eventually exceeds the benefit of the additional 7 signals which are potentially revealed by the agents’ second period actions. The same holds more generally if we compare period t of the ACTIONS treatment with period $t + 7$ of the NO

INFO treatment. However, if people observe each other’s raw signals instead of actions then they learn even faster, as the comparison between ACTIONS and SIGNALS shows.²⁸ Both of these patterns are qualitatively consistent with theoretical predictions.

Panel (b) of Figure 5 explores the effect of group size on learning rates in the two information structures: one with private signals, the ACTIONS treatments, and another with public signals, the ALL treatments. Interestingly, larger groups learn faster only when both signals and the actions of others are public (ALL4 versus ALL treatments), but not when group members observe each other’s actions only (ACTIONS4 versus ACTIONS treatments).²⁹ The lack of a group size effect on learning in the ACTIONS treatment is reminiscent of the theoretical result described in Harel et al. (2021), according to which the speed of learning does not change with the group size. This is, however, quite a loose interpretation of the theoretical result, given that the theory refers to the long-run effects at the limit as time tends to infinity. With regards to the level of consensus across group sizes, we find smaller groups to be more aligned than larger groups, especially in early rounds of the game.³⁰

***Observation 4:** Larger groups learn faster than small groups when people have access to both others’ actions and others’ signals. However, when people can only rely on others’ actions, the group size does not affect the speed of learning.*

6 Behavioral model

In this section we propose a simple behavioral alternative to the Bayesian model, and show that it reproduces the main empirical patterns observed in the experiment.

The information potentially available to agent i at time t includes her own previous-round signals $(s_{i,\tau})_{\tau < t}$, the signals of others $(s_{j,\tau})_{j \neq i, \tau < t}$ and the actions of others $(a_{j,\tau})_{j \neq i, \tau < t}$. The first is available in all treatments, the second is available in the SIGNALS and ALL treatments, and the last is available in the ACTIONS and ALL treatments.

Our data shows that in the ALL treatment, subjects rely both on the signals, as well as the actions of others. Accordingly, we assume that agent i relies on two quantities when making a decision in round t : A_{it} , which depends on others’ actions, and S_{it} , that depends on the signals. Given S_{it} and A_{it} , we assume that agents choose the action R_{it} with probability

²⁸Table A.2 and panel (a) of Figure A.3 in Appendix A confirm that the difference in performance of ACTIONS vs NO INFO treatment is significantly higher in later rounds (9.6%) compared with early ones (5.3%) with $p < 0.1$. Table A.2 in Appendix A compares performance of NO INFO vs ACTIONS treatments as well as ACTIONS vs SIGNALS treatments separately for rounds 2 - 20 and rounds 11 - 20. In all pairwise comparisons we obtain significant difference between all treatments.

²⁹Table A.3 and panel (b) of Figure A.3 in Appendix A confirm these results. Specifically, the probability of a correct action is 10% higher in the ALL compared with ALL4 ($p < 0.01$), whereas this difference is not statistically different from zero when comparing ACTIONS4 vs ACTIONS.

³⁰See Figure I.1 and Table I.1 in the Online Appendix for details.

proportional to $\exp(\beta S_{it} + \gamma A_{it})$ for some $\beta, \gamma \geq 0$. That is, agents use a logistic function to noisily respond to these stimuli.

We now describe how A_{it} and S_{it} are calculated. When deciding on an action at time t , agent i chooses a random agent j , and records that agent’s action $A_{it} = a_{j,t-1}$ from the previous round; in this section, we identify the red actions and red signals with $+1$, and the green ones with -1 , so that $a \in \{-1, +1\}$. Agent i also calculates a normalized sum of signals

$$S_{it} = \frac{\sum_{k,\tau < t} s_{k,\tau}}{(\sum_{k,\tau < t} 1)^\psi},$$

where $\psi \geq 0$ is a third parameter of the model.

Since $s_{k,\tau} \in \{-1, +1\}$, S_{it} is equal to the sum of the signals (that is, the difference between the number of red and green signals) normalized by a function of the number of signals. We assume the general form $(\sum_{k,\tau \leq t} 1)^\psi$, with $\psi \geq 0$, which includes linear normalization ($\psi = 1$), square-root normalization ($\psi = \frac{1}{2}$), and no normalization ($\psi = 0$) as special cases.

From a Bayesian perspective, the sum of the signals is a sufficient statistic, regardless of the total number of signals; this would correspond to $\psi = 0$. However, empirical evidence shows that as the number of signals increases, a given net signal difference becomes less persuasive: two red signals alone are perceived as stronger evidence than one hundred and two red signals accompanied by one hundred green signals (Griffin and Tversky, 1992). We observe the same pattern in our data.

Allowing the exponent ψ to be estimated lets the data determine how strongly the signal sum should be discounted as the number of signals increases. A common approach in the literature (Benjamin, 2019) is to use the sample proportion, which in our framework corresponds to the restriction $\psi = 1$. Our experimental data allow us to relax this restriction and estimate the normalization parameter directly, thereby assessing which scaling of the signal history best fits observed behavior.

This model captures agents who (potentially) pay attention to the two sources of information, even though the actions are redundant, and use a noisy heuristic rather than optimizing subject to some cognitive or attention constraint. This is not an unreasonable heuristic, since imitating others is a potentially low-cognitive-cost alternative to counting signals. As we show, this three-parameter model fits our data remarkably well across all treatments.

6.1 Simulation of the Behavioral Model

We first simulate the ALL treatment for parameter values $\beta = 1/2$, $\gamma = 1$ and $\psi = \frac{1}{2}$. These values are chosen to illustrate a potential mechanism underlying the patterns observed in the data. We then use the experimental data to estimate the values of these parameters that best fit observed betting behavior across informational treatments.

In the SIGNALS treatment, actions are not observable, and thus we set $\gamma = 0$. The same applies to the NO INFO treatment, where the signal statistic S is constructed solely from an agent’s own signals, as others’ signals are not available. An analogous restriction applies in the ACTIONS treatment.

Panel (a) in Figure 6 shows simulated predictions for the probability of a correct guess in each round of the game under the ALL and SIGNALS treatments. These closely match the empirical patterns in panel (a) of Figure 1, showing that participants are more likely to make correct choices in the ALL treatment, where they observe others’ actions in addition to signals.

To build intuition for this result, note that the probability of a mistake depends on both the sign and the magnitude of $\beta S_{it} + \gamma A_{it}$. Assigning a strictly positive weight $\gamma > 0$ to A_{it} can reduce the overall probability of a mistake because actions are correlated with signals. As a result, including A_{it} often moves $\beta S_{it} + \gamma A_{it}$ in the correct direction (i.e., the sign of S_{it}), even though in some cases it may shift it in the wrong direction.

Panel (b) shows simulation results for the ACTIONS, NO INFO, and SIGNALS treatments, which closely match the patterns in panel (a) of Figure 5. These results indicate that participants are able to extract information from others’ actions to some extent, although less effectively than from observing signals directly.

Panel (c) reports simulated predictions for the effect of group size on the probability of a correct choice. In the ALL treatment, larger groups are more likely to choose correctly, whereas in the ACTIONS treatment, groups of size 4 and 8 exhibit very similar success probabilities. This pattern replicates the empirical findings in panel (b) of Figure 5.

Next, we show that the model not only matches the aggregate results, but also generates these predictions for the “correct” reasons. In Figure 7, we reproduce the empirical patterns in Figure 3, showing that participants are more responsive to others’ actions when signals are weak. This is intuitive, as the term $\exp(\beta S_{it} + \gamma A_{it})$ implies that when S_{it} is close to zero, variation in behavior is primarily driven by A_{it} , which captures the influence of others’ actions.

Figure 6: Simulation results: Evolution of the correct actions (frequency)





Figure 7: Simulation results: Learning from others' actions in the ALL treatment



6.2 Estimation of the Behavioral Model

Leveraging our experimental data, we estimate the parameters β , γ , and ψ that govern our behavioral model for each of the four information structures: NO INFO, ACTIONS, SIGNALS, and ALL. We describe the data likelihood for the ALL treatment, where participants observe both the history of signals and the past actions of others.

Let red signals and actions in the data be coded as +1 and green signals and actions as -1. For participant i in round t , define the net signal imbalance and the number of observed signals as $M_{it} \equiv \sum_{k \in N, \tau < t} s_{k\tau}$, $n_{it} \equiv \sum_{k \in N, \tau < t} 1$, and let

$$S_{it} = \frac{M_{it}}{(n_{it})^\psi}$$

denote the signal index. In the model, participant i also receives a social signal $A_{it} \in \{-1, +1\}$, obtained by sampling one of the other participants' actions from the previous round. Conditional on a realization of A_{it} , the probability that participant i chooses red is

$$P_{\beta, \gamma, \psi}(a_{it} = R \mid S_{it}, A_{it}) = \Lambda(\beta S_{it}^g + \gamma A_{it}),$$

where $\Lambda(x) = \text{logit}^{-1}(x)$.

The realized social signal A_{it} is latent from the econometrician's perspective: the data reveal the set of previous-round actions from which it is drawn, but not which action the participant samples in the model. We therefore integrate over the empirical distribution of possible realizations of A_{it} . Let

$$p_{it} = \frac{\sum_{j \neq i} \mathbf{1}\{a_{j,t-1} = R\}}{\sum_{j \neq i} 1}$$

denote the share of other participants who chose red in the previous round. Then

$$\mathbb{P}[A_{it} = +1 \mid (a_{j,t-1})_{j \neq i}] = p_{it}$$

and

$$\mathbb{P}[A_{it} = -1 \mid (a_{j,t-1})_{j \neq i}] = 1 - p_{it}.$$

Thus, the model-implied conditional choice probability can be written as

$$q_{\beta, \gamma, \psi}(M_{it}, n_{it}, p_{it}) = p_{it} \Lambda\left(\beta \frac{M_{it}}{n_{it}^\psi} + \gamma\right) + (1 - p_{it}) \Lambda\left(\beta \frac{M_{it}}{n_{it}^\psi} - \gamma\right), \quad (6)$$

which is a two-component finite mixture of logit choice probabilities, induced by uncertainty over the realization of the latent social signal A_{it} .

Identification of the model parameters follows from variation in both the signal history, summarized by (M_{it}, n_{it}) , and the previous-round actions of others, summarized by p_{it} . Let $q_0(M_{it}, n_{it}, p_{it})$ denote the population conditional choice probability. The behavioral model imposes the restriction $q_0(M_{it}, n_{it}, p_{it}) = q_{\beta, \gamma, \psi}(M_{it}, n_{it}, p_{it})$.

The parameter γ is identified from how choices vary with the share of others choosing red, p_{it} , holding signal information fixed. In particular, consider observations with balanced signals, $M_{it} = 0$. For any $n_{it} > 0$,

$$q_0(0, n_{it}, p_{it}) = p_{it} \Lambda(\gamma) + (1 - p_{it}) \Lambda(-\gamma).$$

Since $\Lambda(-\gamma) = 1 - \Lambda(\gamma)$, this can be written as

$$q_0(0, n_{it}, p_{it}) = \Lambda(-\gamma) + p_{it} [2\Lambda(\gamma) - 1].$$

Thus, for any two values $p_1 \neq p_2$ in the support of the data,

$$\frac{q_0(0, n_{it}, p_1) - q_0(0, n_{it}, p_2)}{p_1 - p_2} = 2\Lambda(\gamma) - 1.$$

Because $2\Lambda(\gamma) - 1$ is strictly increasing in γ , variation in p_{it} when the signal imbalance is zero identifies γ .

The parameter β is identified from the responsiveness of choices to the sign and magnitude of the signal imbalance M_{it} . Given γ , define

$$F_{\gamma,p}(z) \equiv p_{it} \Lambda(z + \gamma) + (1 - p_{it}) \Lambda(z - \gamma).$$

For every $p_{it} \in [0, 1]$, this function is strictly increasing in z . Therefore,

$$z(M_{it}, n_{it}) = \beta \frac{M_{it}}{n_{it}^\psi} = F_{\gamma,p}^{-1}(q_0(M_{it}, n_{it}, p_{it})).$$

Third, β and ψ are separately identified from variation in (M_{it}, n_{it}) . Since

$$z(M_{it}, n_{it}) = \beta \frac{M_{it}}{n_{it}^\psi},$$

for any $M_{it} \neq 0$ and two values $n_1 \neq n_2$ in the support of the data,

$$\frac{z(M_{it}, n_1)}{z(M_{it}, n_2)} = \left(\frac{n_2}{n_1}\right)^\psi.$$

Hence,

$$\psi = \frac{\log |z(M_{it}, n_1)| - \log |z(M_{it}, n_2)|}{\log n_2 - \log n_1}.$$

Once ψ is identified, β is identified from

$$\beta = \frac{z(M_{it}, n_{it}) n_{it}^\psi}{M_{it}}.$$

Equation (6) gives the model-implied conditional choice probability for treatments in which subjects observe others' actions, namely the ACTIONS and ALL treatments. In the NO INFO and SIGNALS treatments, subjects do not observe others' actions, and we therefore

impose $\gamma = 0$. The conditional choice probability then reduces to

$$q_{\beta,\psi}(M_{it}, n_{it}) = \Lambda\left(\beta \frac{M_{it}}{(n_{it})^\psi}\right).$$

Thus, the likelihood contribution of an observed action $y_{it} = \mathbf{1}\{a_{it} = R\}$ is Bernoulli with success probability $q_{\beta,\psi}(M_{it}, n_{it})$ in the NO INFO and SIGNALS treatments, and with success probability $q_{\beta,\gamma,\psi}(M_{it}, n_{it}, p_{it})$ in the ACTIONS and ALL treatments. We estimate the behavioral model separately for the NO INFO, ACTIONS, SIGNALS, and ALL treatments using a Bayesian approach implemented via Markov chain Monte Carlo (MCMC).³¹

Table 3 reports the estimated model parameters. The estimates of signal sensitivity, β , are positive in all treatments. They are also similar within pairs of treatments that differ only in whether subjects observe others' signals but share the same observability of actions: β is close to one in the NONE and SIGNALS treatments, and around 1.27 in the ACTIONS and ALL treatments. Consistent with our previous findings, participants respond to others' actions beyond the information contained in signals, as indicated by $\gamma > 0$. Sensitivity to actions is larger in the ACTIONS treatment, where others' actions convey information about unobserved private signals. Nevertheless, participants also place positive weight on others' actions in the ALL treatment, where this information is redundant because the underlying signals are already observed.

The estimates of ψ are closer to 1/2 than to 1 across treatments, indicating that participants appear to scale signal imbalances less aggressively than the sample proportion would imply. Thus, the data favor a normalization closer to a square-root scaling of the signal information rather than the sample proportion normalization commonly used in the literature (Benjamin, 2019).

Beyond the parameter values, we evaluate the effects of comparable changes in S_{it} and A_{it} on choice probabilities. In particular, in the ALL treatment, we find that a one-standard-deviation increase in S has an effect similar to that of a change in A from observing a green to a red action, corresponding to an increase of 0.3 in the probability of a red action.

Panel (a) in Figure 8 presents the predicted probability of a correct action in the SIGNALS and ALL treatments, computed from the behavioral model and averaged by round across games and participants. As in the simulation exercise in Figure 6, the behavioral model closely matches the experimental data in both treatments, with the performance gap driven by participants' use of others' actions in their decision-making.

Panels (b) and (c) replicate the patterns in Figure 5. The behavioral model captures the improvement in performance from incorporating others' actions relative to the NONE

³¹The parameter estimates are virtually identical when the model is estimated by maximum likelihood (MLE).

Table 3: Behavioral Model Parameters

Parameter	NONE	SIGNALS	ACTIONS	ALL
β	1.047 [0.960, 1.137]	1.018 [0.880, 1.174]	1.268 [1.187, 1.356]	1.273 [1.118, 1.434]
γ	– –	– –	0.968 [0.910, 1.027]	0.648 [0.582, 0.710]
ψ	0.524 [0.487, 0.558]	0.570 [0.537, 0.602]	0.516 [0.488, 0.544]	0.628 [0.598, 0.656]
$\Delta Pr(a_{it}^g = R)$ (1 SD in S)	0.253 [0.245, 0.261]	0.314 [0.307, 0.321]	0.265 [0.258, 0.272]	0.294 [0.283, 0.304]
$\Delta Pr(a_{it}^g = R)$ ($A = -1$ to $A = 1$)	– –	– –	0.449 [0.426, 0.473]	0.313 [0.283, 0.341]
N	15,200	15,580	25,840	28,880
Participants	80	82	136	152

Notes: Point estimates represent the median across posterior draws. 95% credible intervals are reported in square brackets.

treatment. In addition, at the estimated parameter values, the model successfully reproduces the empirical finding that larger groups perform better than smaller ones when both others' actions and signals are observable.

Figure 9 depicts the predicted probability of a red action from the behavioral model grouped by the share of observed red actions in the previous round and by signal strength categories, as defined in Section 5. The model replicates the observed pattern in Figure 3, whereby participants rely more heavily on others' actions when signals are *weak*, consistent with the increasing relative importance of γ as the informativeness of S_{it} declines.

The behavioral model replicates participants' behavior in the data without relying on heterogeneity in sensitivity to signals or actions. Nevertheless, the model is flexible enough to accommodate heterogeneity across participants' characteristics. Motivated by our analysis of IQ and the associated variation in behavior, we allow β and γ to vary across participants, capturing differences in their ability to extract information from signals and others' actions.³²

$$\beta_i = \exp(\beta_{\text{common}} + \beta_{\text{low-IQ}} \text{low-IQ}_i), \quad \gamma_i = \exp(\gamma_{\text{common}} + \gamma_{\text{low-IQ}} \text{low-IQ}_i),$$

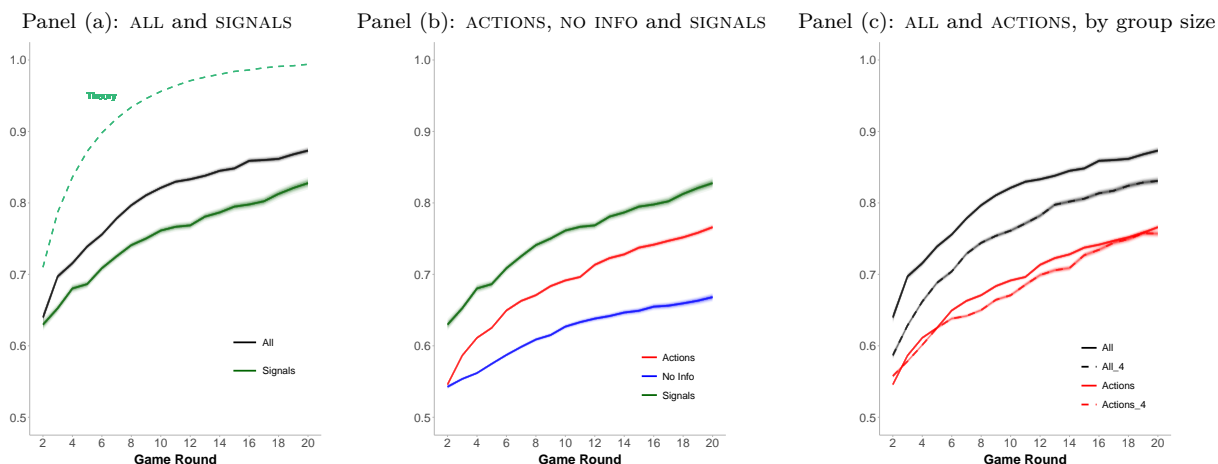
where $\exp(\cdot)$ ensures that $\beta_i, \gamma_i \geq 0$, as required by the behavioral model.³³

Figure 10 reports the results from estimating the heterogeneous behavioral model on data from the ALL treatment. Consistent with the analysis of individual responses by IQ, we

³²In the IQ heterogeneity specification, we hold ψ fixed across participants since we do not find ψ changes by IQ level.


³³The results do not rely on this specific functional form or prior assumptions; the posterior distributions of β and γ are concentrated on positive values.

Figure 8: Estimation results: Correct actions and learning from others



Notes: Panel (a) presents the predicted probability of a correct action in the SIGNALS and ALL treatments, computed from the behavioral model and averaged by round across games and participants. Panel (b) presents the predicted probability of a correct action in the SIGNALS, ACTIONS and NO INFO treatments, computed from the behavioral model and averaged by round across games and participants. Panel (c) presents the predicted probability of a correct action in the ALL, ALL4, ACTIONS and ACTIONS4 treatments, computed from the behavioral model and averaged by round across games and participants. Shaded regions represent 95% credible posterior intervals from 50% (darkest) to 95% (faintest) probability levels.

Figure 9: Simulation results: Learning from others' actions in the ALL treatment

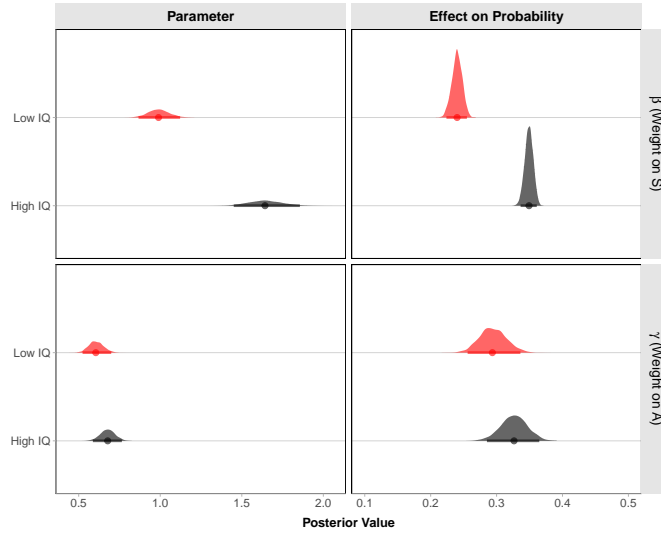


Notes: This figure depicts the probability of a red action computed from the behavioral model and grouped by the share of red actions in the previous rounds as observed in the data within each signal strength category.

find that *low-IQ* participants are substantially less sensitive to signal information in S than *high-IQ* participants ($\beta_{\text{low-IQ}} = 0.99$ vs. $\beta_{\text{high-IQ}} = 1.64$). This difference translates into a markedly smaller effect on betting probabilities from a one-standard-deviation increase in S for *low-IQ* participants (0.24) relative to *high-IQ* participants (0.35), with the 95% credible interval for the difference excluding zero.

By contrast, *low-IQ* and *high-IQ* participants place similar weight on others' actions ($\gamma_{\text{low-IQ}} = 0.61$ vs. $\gamma_{\text{high-IQ}} = 0.68$), with the 95% credible interval for the difference including zero. Consistent with the individual results summarized in Figure 4, this implies that the two groups respond similarly to a change in the observed action in the previous round from green to red (0.29 for *low-IQ* vs. 0.33 for *high-IQ*), again with the 95% credible interval for the difference including zero.

Figure 10: Behavioral Model Parameters (ALL treatment) as a function of IQ



Notes: The upper panels of this figure present posterior estimates of β and of the effect of S_{it} on choice probability by participants' IQ level. The lower panels present the posterior distribution of γ and of the effect of A_{it} on choice probability by participants' IQ level. Points represent posterior medians, and lines denote 95% credible intervals.

The focus on this paper is experimental, and so we leave for future work the in-depth theoretical analysis of this heuristic and its general properties. Both the simulations and estimation on our data suggest some natural conjectures.

First, in the presence of both signals and actions, imitation can be beneficial. For a given value of β , which governs responses to signals, some positive values of γ improve outcomes relative to the baseline case $\gamma = 0$, where others' actions are ignored. The intuition is that, because individual responses are noisy, past actions can contain useful additional

information. One might expect this benefit to diminish once too much weight is placed on past actions, since all information ultimately originates in the signals.

Second, the results from the ACTIONS treatment indicate that group size does not have a substantial effect on outcomes. This is consistent with the theoretical results of [Harel et al. \(2021\)](#), although that setting assumes Bayesian agents with a high degree of rationality. More generally, it is of interest to study the extent to which imitation remains useful under cognitive limitations and other departures from optimal behavior.

7 Discussion

The imitation of the actions, opinions and values of others is a natural human tendency. It is an integral part of how infants and adults learn, and of how information is disseminated through society. It can also lead to inefficient outcomes, as has been highlighted by the social learning literature over the past few decades. In our setting, which features repeated interaction between participants and gradual accumulation of public raw data as well as others' opinions, we show that participants imitate each other, and that imitation improves their performance.

This result was not a priori obvious. Optimal Bayesian behavior does not feature imitation, and behavioral economics does not offer a clear-cut prediction that imitation should improve efficiency, or indeed suggests the opposite. Clearly, it is highly inefficient for all agents to ignore the raw data and mimic each other, and so too much imitation is harmful. Hence the usefulness of imitation lies in moderation: both the signals and the actions have to play a role. We develop a simple behavioral model that reproduces these results in simulations. A fruitful avenue for future research is the further development and formal study of this model.

To study imitation we constructed a social learning environment with repeated actions. The experimental literature on such settings is small, and we believe that there is much room for further research. In particular, it would be interesting to understand the robustness of imitation to various interventions, such as the inclusion of a well-informed central information source, or of ideological agents whose actions are independent of the data, or a more general structure of heterogeneity in information quality of different agents. These questions are important in an age in which the dissemination of uncertain but critical information to the public is a major societal challenge.

A Additional Aggregate Results




Figure A.1: Differences between the ALL and SIGNALS treatments

Notes: Panel (a) presents the average difference in the frequency of correct actions in each treatment in each round, averaged across games. Panel (b) depicts the difference in consensus rates in each round. For panel (b) we exclude cases with equal number of green and red signals. Shaded regions represent confidence intervals from 50% (darkest) to 95% (faintest) probability levels. Confidence intervals are constructed with a variance-covariance matrix clustered by session.




Figure A.2: How often is the majority correct? ALL versus SIGNALS treatments

Notes: Panel (a) presents the frequency of correct actions by the majority in each treatment per round, averaged across games. Panel (b) presents the average difference in the frequency of correct actions by the majority between the two treatments. We exclude cases with an equal number of green and red signals. Shaded regions represent confidence intervals from 50% (darkest) to 95% (faintest) probability levels. Confidence intervals are constructed with a variance-covariance matrix clustered by session.

Table A.1: Treatment Effects for SIGNALS and ALL treatments

	<i>Dependent variable:</i>					
	Correct Actions		Consensus Rate		Majority is Correct	
	(1)	(2)	(3)	(4)	(5)	(6)
ALL (Baseline)	0.702*** (0.012)	0.701*** (0.013)	0.812*** (0.009)	0.809*** (0.011)	0.803*** (0.013)	0.808*** (0.014)
SIGNALS (Effect)	-0.051* (0.028)	-0.047* (0.028)	-0.095*** (0.028)	-0.089*** (0.029)	-0.038** (0.016)	-0.051*** (0.016)
SIGNALS (Effect) × Late Rounds		-0.007 (0.017)		-0.011 (0.012)		0.026* (0.015)
Game Round Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	44,460	44,460	5,499	5,499	5,499	5,499
Adjusted R ²	0.021	0.021	0.135	0.135	0.062	0.062

Notes: ** p<0.05; *** p<0.01. Clustered standard errors by session in parentheses. Late rounds are 11 through 20.

Table A.2: Treatment Effects for NO INFO, ACTIONS, and SIGNALS

	<i>Dependent variable:</i>	
	Correct Actions	
	(1)	(2)
ACTIONS (Baseline)	0.583*** (0.017)	0.572*** (0.015)
NO INFO (Effect)	-0.076*** (0.019)	-0.053*** (0.016)
NO INFO (Effect) × Late Rounds		-0.043* (0.026)
SIGNALS (Treatment)	0.078*** (0.029)	0.096*** (0.028)
SIGNALS (Treatment) × Late Rounds		-0.034** (0.013)
Game Round Fixed Effects	Yes	Yes
Observations	56,430	56,430
Adjusted R ²	0.028	0.028


Notes: **p<0.05; ***p<0.01. Clustered standard errors by session in parentheses. Late rounds are 11 through 20.

Table A.3: Treatment Effects for Group Size

	<i>Dependent variable:</i>			
	Correct Actions			
	(1)	(2)	(3)	(4)
ALL4 (Baseline)	0.614*** (0.017)	0.603*** (0.016)		
ACTIONS4 (Baseline)			0.546*** (0.016)	0.539*** (0.012)
ALL (Effect)	0.065*** (0.023)	0.082*** (0.022)		
ALL (Effect) × Late Rounds		-0.033 (0.031)		
ACTIONS (Effect)			0.008 (0.022)	0.019 (0.020)
ACTIONS (Effect) × Late Rounds				-0.021 (0.024)
Game Round Fixed Effects	Yes	Yes	Yes	Yes
Observations	44,080	44,080	40,090	40,090
Adjusted R ²	0.029	0.029	0.020	0.020

Notes: **p<0.05; ***p<0.01. Clustered standard errors by session in parentheses. Late rounds are 11 through 20.

Figure A.3: Frequency of correct actions, by information structure and group size



Notes: Both panels present the average frequency of correct actions in each treatment per each round, averaged across games. Shaded regions represent confidence intervals from 50% (darkest) to 95% (faintest) probability levels. Confidence intervals are constructed with a variance-covariance matrix clustered by session.

B Open-ended Questions

We quantify the answers to the open-ended questions designed to elicit participants' strategies throughout the game and analyze their relationship with participants' characteristics, specifically players' IQ.

For each participant in the ALL4 and ALL treatments, we examine the three open-ended survey questions asked after all rounds have been played:

- (i). What strategy did you use in the game (if any)? Please elaborate.
- (ii). Did you look at the balls drawn for other players in your group? Did you find them useful/not useful? Please elaborate.
- (iii). Did you look at the bets made by other players in your group? Did you find them useful/not useful? Please elaborate.

We implement a machine learning algorithm for a probabilistic topic model known as the structural topic model (STM) (Roberts et al., 2013). Under this framework, a *topic* is defined as a probability distribution over words and a participant's response in our data is modeled as a distribution over topics. Thus, each participant's response in the data can

belong to multiple topics with a probability distribution estimated from the data. We are interested in the proportion of an answer spent covering each estimated topic.

Each participant’s response to question $q \in \{(i), (ii), (iii)\}$, r_q , has its own distribution over topics, θ_{r_q} . We label this parameter topic prevalence and interpret it as the proportion of each topic $k = 1, \dots, K$ in response r_q . We can think of each topic k as drawn from a multinomial distribution with parameter θ_{r_q} . Conditional on the topic selected, word $w_{r_q, n}$ included in response r_q is drawn from a multinomial distribution over the vocabulary $n = 1, \dots, N$ with parameter $\beta_{k, n}$. This is a probability vector over the N words in the vocabulary.

In probabilistic topic models such as the STM, the number of topics K needs to be selected a priori. In doing so, there is a trade-off between interpretability (consistent with a lower K) and goodness-of-fit (consistent with a higher K). We favor the former and choose the low value of $K = 4$ for the three open-ended questions. We find that $K = 4$ gives us topics that are very straightforward to interpret and with a higher semantic coherence (i.e., the most probable words in a given topic co-occur together) than models with more topics.

The structural topic model allows for the inclusion of participants’ characteristics to inform the topic prevalence. Specifically, $\theta_{r_q} \sim \text{LogisticNormal}(X_{r_q}\gamma, \Sigma)$, where X_d is a vector of participant characteristics. In addition to our binary measure of IQ, we include several participant characteristics: **female**, which is an indicator variable that takes the value of one if the participant identifies as female. **stem**, which is an indicator variable that takes the value of one if the participant’s major is STEM. **overconfidence** measures the extent of a participant’s over-estimation of her IQ, which is given by the difference between the number of questions a participant believes she solved correctly and the actual number of correct answers. **risk** is measured by the number of points invested in a risky asset as specified in a risky investment task solved at the end of the game.

The topic model for each open-ended question is estimated using a variational Expectation-Maximization (EM) algorithm, as implemented in the `stm` package in R (Roberts et al., 2013). Prior to estimation, we pre-processed the raw responses using standard conventions: we stem words (i.e., reduce words to their root form), drop punctuation, as well as common stop-words, and remove words that were used less than 0.5% over all responses.

Figure B.1 shows the labels of three estimated topics for the strategy question along with the actual responses that are estimated to be highly associated with each topic. The topics we labeled *Other Strategies* encompass participants who either choose actions at random, ignoring signals and actions of others, or used heuristics that deviate from Bayesian updating. The topic we labeled *Signals + Actions of Others* encompasses answers where participants emphasized using both signals and others’ actions to make their choices. The topic labeled *Signals* describes answers where participants said making their choices based

on the aggregate number of green and red balls drawn throughout a game. Figures B.3 and B.5 show the labels and response examples for questions (ii) and (iii), respectively. For these questions, the estimated topics capture the range from negative to positive reliance on signals and others’ actions, respectively.

Figure B.2 shows for each estimated topic in the strategy question (i), the relationship between our measure of IQ and the estimates of topic prevalence, θ_{rq} . We capture each topic with a word cloud of the top words associated with it. Figures B.4 and B.6 present these estimates for questions (ii) and (iii), respectively.

We find that participants’ description of their strategies is consistent with their actual behavior in the game according to their IQ type. In particular, high-IQ participants, who appear closer to Bayesian behavior in the game (see Figure 4), described using strategies based on signals to a larger extent than low-IQ ones (14% larger with $p < 0.001$, see panel (a) in Figure B.2) who, in turn, relied significantly more on heuristics which, in many cases, disregarded raw signals as well as actions of others (see panel (b) in Figure B.2). Moreover, when directly asked “Did you look at the balls drawn for other players in your group? Did you find them useful/not useful?”, low-IQ participants were more skeptical about the information provided by other players’ signals, compared to high-IQ players (5% more with $p < 0.001$)

In terms of explicitly looking at the actions of others, the topic prevalence from question (i), which we labeled “Signals + Actions of Others” shows that high-IQ participants describe other players’ actions as useful in a similar proportion to low-IQ ones (37% for high-IQ participants versus 42% for low-IQ ones ($p = 0.055$), see panel (c) in Figure B.2). Moreover, the analysis of the answers to “Did you look at the bets made by other players in your group? Did you find them useful/not useful?” shows that high-IQ participants describe other players’ actions as useful similarly to low-IQ ones (36% vs 29% with $p = 0.07$, see panel (a) of Figure B.6), while low-IQ players described the social signal as sometimes/somewhat useful to a larger extent than high-IQ ones (4% and 30% more ($p < 0.05$) for topics “Sometimes/Somewhat Useful”, respectively). Overall, these answers are consistent with a similar responsiveness to others’ actions by IQ types found in the game data (see Figure 4). At the same time, more than two thirds of those participants who find the information in others’ actions redundant come from the group of high-IQ players (see panel (b) in Figure B.6), which is consistent with the behavior of a subset of high-IQ participants who take decisions consistent with Bayesian updating.

Other strategies	Signals + Actions of Others	Signals
<p>based my decision off of the amount of drawn colored balls. I would wait six turns to see which balls were being drawn the most and then selecting the urn with the highest amount. I would continue with the same urn for the remainder of the game unless there was an obvious discrepancy.</p> <p>-----</p> <p>No strategy. I randomly chose based on patterns I selected in my head</p> <p>-----</p> <p>I was actively trying to come up with a strategy but since each urn had a 50/50 chance of being chosen no matter what balls were pulled, I didn't think a strategy would make any difference.</p>	<p>I looked at the set of balls drawn and the generally asked bets and see if theirs matched up against my own balls and bets.</p> <p>-----</p> <p>I tried to pick one color first, and try to see the probability of color of draw URN, if more 3 same color out of 4 balls then I bet that color</p> <p>-----</p> <p>When betting on the colored balls I would compare the likelihood of the color popping up on each of the other players picks.</p>	<p>I kept track of the number of green and red balls throughout each game and after Round 19 but before Round 20, I divided each total by 152 (8*19). As this was a decently large sample size of 160, I used these percentages to figure out which urn it was as they tended to go toward either the 60% or 40% mark.</p> <p>-----</p> <p>I kept track of the ratio of red to green across the game and guessed whichever was higher.</p> <p>-----</p> <p>I counted the total number of red balls throughout each game. If the number I have got till the current round is larger than #of round*4/2, I will pick Red Urn in the next round. Otherwise, I will pick Green Urn. If the numbers are equal, I will continue my old guessing</p>

Figure B.1: Example Responses to “What strategy did you use in the game (if any)?” by topic (ALL4 and ALL treatments)

Notes: Each panel presents the three most highly associated open-ended responses with each of the four estimated topics for the question “What strategy did you used in the game (if any)?”

Panel (a): Topic “Signals”

Panel (b): Topic “No Strategy”

Panel (c): Topic “Signals + Actions”





Figure B.2: Responses to “What strategy did you use in the game (if any)?” of low-IQ and high-IQ Participants (ALL4 and ALL treatments)

Notes: The upper row of each panel presents the word cloud of the top 100 words with the highest likelihood of being drawn from each topic, with the size of the word representing the magnitude of this likelihood. The lower row of each panel presents the estimated topic prevalence for the subgroups of low-IQ and high-IQ participants with 95% confidence intervals.

Sometimes/Somewhat Useful	Yes, useful	Yes, more information
<p>I looked, but I didn't find it useful, since we didn't draw the same balls.</p> <p>-----</p>	<p>Yes, it's useful to find the trend</p> <p>-----</p>	<p>es, because the larger the sample size, the closer the average mean will be to the population mean.</p> <p>-----</p>
<p>es, I found them useful since I assumed more of one color of balls correlates to the urn</p> <p>-----</p>	<p>Yes I used this to figure out the majority of what color ball was chosen in each round</p> <p>-----</p>	<p>es. I thought these were extremely useful as instead of just a sample size of 19 balls (if you only saw one drawn for you), you now had a sample size of 15 balls (8*19).</p> <p>-----</p>
<p>It was useful in looking at the probability of which urn is being used. The different balls other players got sometimes had an unlucky round that threw me off but most of the time it helped determine the right ball.</p> <p>-----</p>	<p>Yeah, I think they are helpful.</p> <p>-----</p>	<p>es because it expanded the size of my reference data set, gave me more information to make decisions based on.</p> <p>-----</p>

Figure B.3: Example Responses to “Did you look at the balls drawn for other players in your group?” by topic (ALL4 and ALL treatments)


Notes: Each panel presents the three most highly associated open-ended responses with each of the four estimated topics for the question “Did you look at the balls drawn for other players in your group? Did you find them useful/not useful? Please elaborate.”

Panel (a): Topic “Sometimes/Somewhat Useful”

Panel (b): Yes, more information



Topic: Sometimes/Somewhat Useful



Topic: Yes, more information




Figure B.4: Responses to “Did you look at the balls drawn for other players in your group?” of low-IQ and high-IQ Participants (ALL4 and ALL treatments)

Notes: The upper row of each panel presents the word cloud of the top 100 words with the highest likelihood of being drawn from each topic, with the size of the word representing the magnitude of this likelihood. The lower row of each panel presents the estimated topic prevalence for the subgroups of low-IQ and high-IQ participants with 95% confidence intervals.

Sometimes/Somewhat Useful	Yes, useful	No, no extra information
<p>Sometimes when I found myself not really sure and rushed I automatically looked at the other guesses and would pick the same or check to make sure we were following the same pattern. Usually I stuck to my own pattern and guesses though.</p> <p>-----</p> <p>Yes, at times. I did find them useful and most of the time, most people also went with the majority color of draws so I felt like I was also guessing right. But my bet wasn't dependent on others' bets.</p> <p>-----</p> <p>I did not look at the bets from other players. They were guessing based on the same information I was able to have confidence in myself so changing my guess based on them would have made it too complicated.</p>	<p>s. to count how many red urns and green urns in each run. Probably useful?</p> <p>-----</p> <p>Yes, I found them sort of useful, because it would tell me whether or not they were firm with their decision or if they had uncertainty as to which urn was picked.</p> <p>-----</p> <p>yes. useful in me feeling confident.</p>	<p>no, i didn't look at them because they don't know the actual answer either</p> <p>-----</p> <p>I didn't really look at. I don't think it will be helpful, because computer selected didn't balls for every one. There isn't really points to look at other players</p> <p>-----</p> <p>I didn't. I only looked at the balls.</p>

Figure B.5: Example Responses to “Did you look at the bets made by other players in your group?” by topic (ALL4 and ALL treatments)

Notes: Each panel presents the three most highly associated open-ended responses with each of the four estimated topics for the question “Did you look at the bets made by other players in your group? Did you find them useful/not useful? Please elaborate.”

C Participants' Beliefs

We analyze the precision of subjects' incentivized beliefs, as elicited at the end of the experiment (see section E.3 for details). For the analysis, we pool all participants' answers across treatments.³⁴ In the beliefs section, each participant provides an estimate of the fraction of correct actions (in the last round of the game) in both her own treatment and other treatments. We combine these answers and define participant i 's precision as her mean squared error across all treatments: $MSE_i = \sum_{treat} (y_{treat} - \hat{y}_{i,treat})^2$, where y_{treat} denotes the actual fraction of correct actions in treatment $treat \in \{\text{NO INFO}, \text{ACTIONS4}, \text{ACTIONS}, \text{SIGNALS}, \text{ALL4}, \text{ALL}\}$ and $\hat{y}_{i,treat}$ denotes participant i 's estimate about the fraction of correct actions in treatment $treat$. Both actual values and beliefs are discretized in 10 equally spaced bins from 0% to 100%, consistent with subjects' answers.

To assess the difference between low-IQ and high-IQ subjects on their beliefs' precision, we estimate a regression of the form:

$$MSE_i = \delta_{treat,i} + \beta_1 \cdot \text{low-IQ} + X_i' \gamma + \epsilon_i,$$

where **low-IQ** is a dummy variable that takes a value of one if a subject is low-IQ according to our IQ auxiliary measure and 0 if she is high-IQ. $\delta_{treat,i}$ is a fixed effect of the treatment assigned to subject i . X_i is a vector of participant i 's characteristics as defined in Appendix B including: **female**, **stem**, **overconfidence** and **risk**. ϵ_i denotes the error term with a variance-covariance matrix clustered by session.

The results of estimating the difference between low-IQ and high-IQ subjects on their beliefs' precision are presented in Table C.1 and visually in Figure C.1. Across all specifications, we find that low-IQ subjects are significantly less precise than their high-IQ counterparts. In particular, controlling for both treatment effects and participant characteristics (column (3) of Table C.1), this effect translates into low-IQ participants miscalculating the actual fraction of correct actions by more than 10% (i.e., $\sqrt{1.58} = 1.26$ bins of 10% increments each).

³⁴We do not have data for subjects in the first session of each treatment, because it was used to collect data and calibrate subjects' payments in subsequent sessions.

Table C.1: Precision of Beliefs (MSE) by IQ Type (ALL treatments)

	<i>Dependent variable:</i>		
	MSE		
	(1)	(2)	(3)
Intercept	5.281*** (0.414)	3.428*** (0.283)	3.770*** (0.818)
low-IQ	1.097** (0.537)	1.088** (0.545)	1.580*** (0.602)
Treatment Fixed Effects	No	Yes	Yes
Participant Covariates	No	No	Yes
Observations	476	476	476
Adjusted R ²	0.005	0.012	0.013

Notes: *p<0.1; **p<0.05; ***p<0.01. Clustered standard errors by session in parentheses




Figure C.1: Beliefs' Accuracy and IQ Types

Notes: The figure shows predicted mean squared errors for the beliefs of low-IQ and high-IQ participants along with 95% confidence intervals. For these estimated effects we use the estimated coefficients from column (3) in Table C.1. We include as controls: `female`, `stem`, `overconfidence` and `risk` and set them to their median values in the data.

Online Appendix

D Structure of the ALL treatment and payments

Each participant earns \$10 for completing the session (participation fee). In addition, subjects earn money for other parts as described below.

- Part I: Main game
 - 10 games with 20 rounds in each game
 - random re-matching into groups of 8 subjects at the beginning of a game
 - random round from a random game determines participants' payments for this part: correct guess pays \$20, wrong guess pays \$5
- Part II: Beliefs
 - one questions is randomly selected for payment,³⁵ \$5 for a correct answer
 - questions appear in a random order for each subject
 - in the first session of each treatment, this part was not present because we were collecting the data used in the next sessions for subjects' payments.
- Part III: Risk Attitudes
 - one of the two investment tasks is randomly selected for payment
 - points earned are converted into dollars using the rate 1 point = 1 cent
- Part IV: IQ and Overconfidence
 - six matrices for IQ measure, each correctly solved matrix earns 50 cents
 - overconfidence is measured using two related dimensions: over-estimation and over-placement
 - correct prediction for over-estimation task earns 50 cents
 - in the first session of each treatment (those without beliefs questions) over-placement question is not incentivized, while in the remaining sessions it is and correct prediction is rewarded by 50 cents³⁶
- Part V: Open-ended Questionnaire for Strategies

³⁵For the NO INFO, ACTIONS, and ALL treatment we had 3 beliefs questions corresponding to these three information structures. In the SIGNALS treatment, we had 4 beliefs questions each corresponding to one of the all four information structures.

³⁶The first session is used to collect data for comparing actual rank of students to that reported for the follow-up sessions.

E Instructions for ALL treatment

Welcome. You are about to participate in an experiment on decision-making. You will be paid for your participation in cash privately at the end of the session. Please turn off all electronic devices, especially phones. During the experiment you are not allowed to open or use any other applications on these laboratory computers, except for the interface of the experiment.

The experiment consists of four parts: Part I of the experiment is the main and the longest part. The other parts (Part II, III and IV) are short. You will receive the instructions for each part of the experiment before that part begins. You have already earned \$7 for coming to the lab. In addition, you can earn money in each part of the experiment. The instructions for each part of the experiment will be very precise about that.

Part I

Part I of the experiment consists of 10 games. Each game consists of 20 rounds. Before the beginning of each game, you will be randomly assigned to a group of 8 players. You will play all 20 rounds of the game with the same group of people. At the end of each game, you will again be randomly assigned to a new group of 8 players and will play 20 rounds with them, and so on.

At the beginning of each game, the computer randomly selects one of the two Urns for your group for this game (independent of the urns selected in previous games):

- RED URN contains 6 RED balls and 4 GREEN balls
- GREEN URN contains 4 RED balls and 6 GREEN balls

That is, there is 50% chance that the RED URN is selected, in which case the urn has 6 RED balls and 4 GREEN balls and 50% chance that the GREEN URN is selected, in which case the URN has 6 GREEN balls and 4 RED balls. We will refer to the selected urn as the URN.

The composition of the URN is determined once at the beginning of each game (before round 1) and stays the same throughout the game (in all 20 rounds). All players that were assigned to the same group share the same composition of the URN. At the beginning of a new game, after players are assigned to new groups but before round 1 begins, the computer again determines the composition of the URN for each group separately using the rule described above. Thus, the composition of the URN in your group stays the same in all 20 rounds of the same game but is not related to its composition in later games.

In every round of a game, your task is to guess (bet on) the URN selected for your group at the beginning of this game. If you choose to bet on RED URN, it means that you are

betting that the RED URN was selected for your group at the beginning of this game. If you choose to bet on GREEN URN, it means that you are betting that the GREEN URN was selected for your group at the beginning of this game. At the end of each round, one ball will be randomly drawn from the selected URN and its color revealed to you. The same thing happens with all other players in your group: each player bets on the URN selected for your group and then observes one randomly drawn ball from the selected URN. It is important that different players observe different draws, but that all drawn are made from the same URN.

Your payment in Part I. To determine your payment in Part I, at the end of the experiment, the computer will select one game from the 10 games played. Each game is equally likely to be chosen for payment. Then the computer will select one of the 20 rounds in the selected game. If you guessed correctly the URN in the selected round of the selected game, then you will receive \$20. If your guess (bet) was wrong then you will receive \$5.

End-of-Round Information. At the end of each round, you will be reminded of the URN you betted on and you will observe the color of the randomly drawn ball from the URN. Moreover, at the end of each round you will also observe the bets that other players in your group made in this round regarding the selected URN and the balls that were drawn from the selected URN for other players. This information will be summarized in the table on your screen. This table will keep track of all decision that you have made in this game (highlighted in yellow), as well as the decisions made by other players in your group.

Note that the computer draws one ball for each member of your group in each round from the URN with replacement. That is, every ball that is drawn from the URN is placed back in the URN before the next draw. Therefore, if the URN contains, say, 6 RED balls and 4 GREEN balls then for each player in your group there is exactly 60% chance that the drawn ball is RED and 40% chance that it is GREEN.

Are there any questions?

E.1 Screenshots for ALL treatment

This screenshot presented in Figure [E.1](#) shows round 1 of a game. On the top of the screen, subjects are reminded about the compositions of two urns. The game history table keeps track of all what has transpired in the current game, and the bottom of the screen is where subjects make their bets about which urn was selected for their group for this game.

The game history table starts filling up at the beginning of round 2 (as shown in Figure [E.2](#)), and it displays own bet and own draw as well as bets and draws of all group members

Figure E.1: Screenshot 1 from All treatment

Game 1

Round 1 of 20

RED URN: ●●●●●●●●

GREEN URN: ●●●●●●●●

One of these two urns was randomly selected at the beginning of this game.

Game history

Round	Me bet/draw	Player 2 bet/draw	Player 3 bet/draw	Player 4 bet/draw	Player 5 bet/draw	Player 6 bet/draw	Player 7 bet/draw	Player 8 bet/draw
1	--	--	--	--	--	--	--	--

YOUR DECISION IN ROUND 1:

I bet on **RED URN** **GREEN URN** Next

(as seen on the second screenshot). Bets are always displayed as rectangles, i.e., lottery tickets, and signals are displayed as circles, i.e., balls drawn from the selected urn.

E.2 Strategy questions in ALL treatment

Think about the main game in the experiment (betting on the urn selected for your group).

- What strategy did you use in the game (if any)? Please elaborate.
- Did you look at the bets made by other players in your group? Did you find them useful/not useful? Please elaborate.
- Did you look at the balls drawn for other players in your group? Did you find them useful/not useful? Please elaborate.
- Was anything unclear about the game?
- What is your gender?
- What is your major?

E.3 Beliefs questions in ALL treatment

Below we present the three beliefs questions that subjects in the ALL treatment were asked to answer. The formulation for the beliefs questions in the other treatments is adapted to

Figure E.2: Screenshot 2 from All treatment

Game 1

Round 2 of 20

RED URN: ● ● ● ● ● ● ● ● ● ●

GREEN URN: ● ● ● ● ● ● ● ● ● ●

One of these two urns was randomly selected at the beginning of this game.

Game history

Round	Me bet/draw	Player 2 bet/draw	Player 3 bet/draw	Player 4 bet/draw	Player 5 bet/draw	Player 6 bet/draw	Player 7 bet/draw	Player 8 bet/draw
1	● ● ● ● ● ● ● ●	● ● ● ● ● ● ● ●	● ● ● ● ● ● ● ●	● ● ● ● ● ● ● ●	● ● ● ● ● ● ● ●	● ● ● ● ● ● ● ●	● ● ● ● ● ● ● ●	● ● ● ● ● ● ● ●
2	--	--	--	--	--	--	--	--

YOUR DECISION IN ROUND 2:

I bet on **RED URN** **GREEN URN** Next

the information structure subjects experienced in the treatment, but follow the same idea. The questions were presented in the random order across subjects.

Question 1. In the past session, groups of 8 subjects played the same game that you just played. Namely, at the beginning of each game, the computer randomly selected one of the two urns: with probability 50% the RED URN was selected, which contained 6 RED balls and 4 GREEN balls, and with probability 50% the GREEN URN was selected, which contained 4 RED balls and 6 GREEN balls. Each game lasted for 20 rounds. In each round group members submitted their guesses (bets) about the URN selected for their group for this game. After making their choices, each group member observed the colors of the 8-randomly drawn balls from the selected urn (with replacement) and also guesses made by members of their group regarding the selected urn for this game.

Please answer the following question: What fraction of last round bets (bets submitted by subjects in round 20) were correct? The correct guess means that a subject guessed correctly which URN was selected for her group for this game.

[[Radio buttons with ten options: 0%-9%, 10%-19%, ..., 90% - 100%]]

If this question is selected for payment, you will get paid \$5 if you chose correctly.

Question 2. In the past session, groups of 8 subjects played game that was slightly different from the game you just played. Just like in your game, at the beginning of each

game, the computer randomly selected one of the two urns: with probability 50% the RED URN was selected, which contained 6 RED balls and 4 GREEN balls, and with probability 50% the GREEN URN was selected, which contained 4 RED balls and 6 GREEN balls. Each game lasted for 20 rounds. In each round group members submitted their guesses (bets) about the URN selected for their group for this game. After making their choices, each group member observed the colors of the one randomly drawn ball from the selected urn and the guesses made by members of their group regarding the selected urn.

However, contrary to the game you just played, participants in these previous sessions did not observe the colors of the balls that were randomly drawn for other participants in their group. These past participants only observed the guesses that their group members made in each round and one randomly drawn ball from the selected urn.

Please answer the following question: What fraction of last round bets (bets submitted by subjects in round 20) were correct? The correct guess means that a subject guessed correctly which URN was selected for her group for this game.

[[Radio buttons with ten options: 0%-9%, 10%-19%, ..., 90% - 100%]]

If this question is selected for payment, you will get paid \$5 if you chose correctly.

Question 3. In the past session, groups of n (parameter in the program) subjects played game that was slightly different from the game you just played. Just like in your game, at the beginning of each game, the computer randomly selected one of the two urns: with probability 50% the RED URN was selected, which contained 6 RED balls and 4 GREEN balls, and with probability 50% the GREEN URN was selected, which contained 4 RED balls and 6 GREEN balls. Each game lasted for 20 rounds. In each round group members submitted their guesses (bets) about the URN selected for their group for this game. After making their choices, each group member observed the colors of the one randomly drawn ball from the selected urn.

However, contrary to the game you just played, participants in this previous session were not provided with any additional information about the guesses made by other members of their group or the colors of the randomly drawn balls from the selected urn for other participants. The only information these past participants had was the color of one randomly drawn ball from the selected urn.

Please answer the following question: What fraction of last round bets (bets submitted by subjects in round 20) were correct? The correct guess means that a subject guessed correctly which URN was selected for her group for this game.

[[Radio buttons with ten options: 0%-9%, 10%-19%, ..., 90% - 100%]]

If this question is selected for payment, you will get paid \$5 if you chose correctly.

E.4 Other Control Tasks

Risk Attitudes. Risk attitudes were measured using two investment tasks, in each of which subjects were endowed with 200 points (worth a total of \$2), any portion of which they could choose to invest in a risky project. In the first investment task, the risky project was successful 50% of the time and had a return of 2.5 points for each point invested in it, while in the second investment task the risky project was successful 40% of the times and returned 3 points for each point invested in it. Points not invested in the risky project had a return of 1 to 1 point. One of these two tasks was randomly selected for payment. This is one of the standard methods used in the experimental literature to elicit subjects' attitudes towards risk (see Gneezy and Potters (1997) and Charness, Gneezy, and Imas (2013)). Administering this task twice with two sets of parameters allows to reduce measurement error (see ORIV technique developed by Gillen, Snowberg, and Yariv (2018)).

IQ and Overconfidence. Subjects were asked to solve six matrices from the ICAR database (see ICAR, Condon and Revelle (2014)). Subjects earned 50 cents for each correctly solved matrix. So, the IQ of a subject is measured by the number of correctly solved matrices with the smallest number being 0 and the largest number being six. After solving these matrices, we asked subjects two questions:

- (i). How many of the six puzzles do you think you correctly answered? You will receive 50 cents if you answer correctly.
- (ii). Now, think about 100 UCSD students. Where do you think you rank in terms of how many correct rotation cubes puzzles you got? For example, if you think you got the most correct, you should answer 1, while if you think you got the least correct, you should answer 100.


The last two questions are used to measure overconfidence of subjects. We chose to measure over-estimation and over-placement of subjects. Specifically, overestimation is the difference between how many ICAR questions a subject thinks she solved correctly minus how many she actually solved correctly, while the over-placement is the reported rank minus actual rank in a sample of the 100 randomly selected students. These measures are quite standard in the literature (see Chapman et al. 2019b).

F Physical *vs* Virtual Lab

We compare outcomes and behavior observed in the ALL treatment across sessions conducted in the physical lab at UCSD and in the virtual lab at OSU. For each location we have four

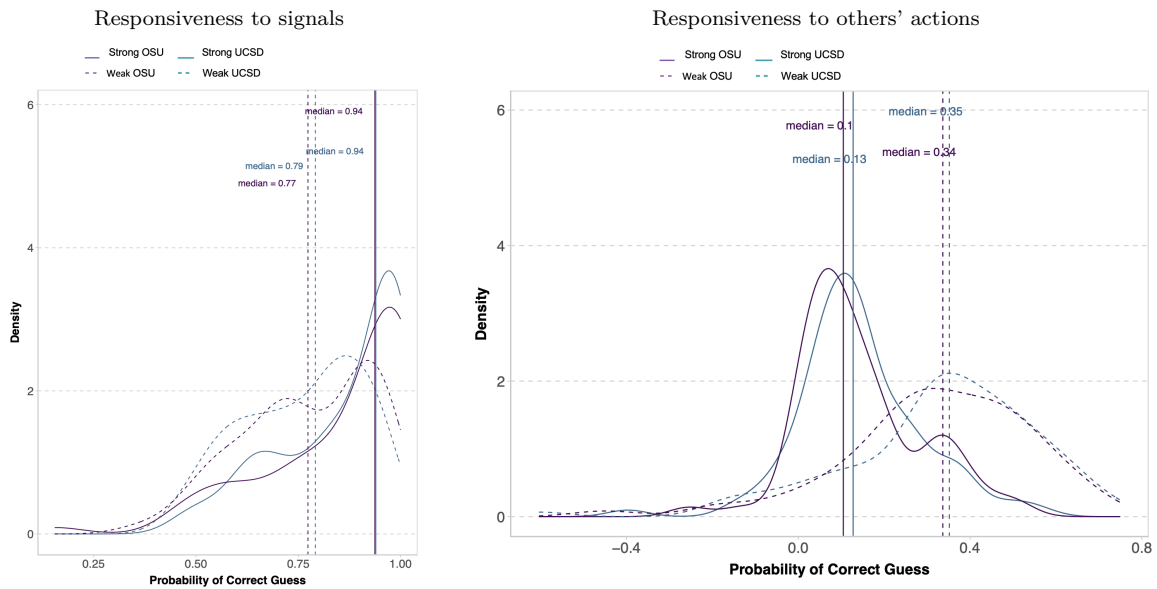
sessions with 64 subjects all together at UCSD and 88 subjects at OSU. In both locations, we used the standard subjects' pool of undergraduate students.

Figure F.1: Aggregate statistics in the ALL treatment



Notes: Panel (a) presents the average frequency of correct guesses in the ALL treatment in each round, averaged across games. Panel (b) depicts the evolution of consensus in each round, i.e., the relative size of the majority, averaged across games.

Figure F.2: Individual responsiveness to signals and others' actions in ALL treatment



Notes: Panel (a) presents Kernel distributions of participants' responsiveness to signals separately for weak and strong signals. The vertical lines depict median responsiveness for each group. Responsiveness to signals is given by the probability that a participant's action matches signal majority, when actions of others group members in the previous round are split equally between green and red. Panel (b) presents Kernel distributions of participants' responsiveness to others' actions separately for weak and strong signals. The vertical lines depict median responsiveness for each group. Responsiveness to others' actions is measured by the change in the probability of choosing the action of the majority of signals when all versus none of the other group members choose the majority of signals in the last round.

Figure F.1 presents two main outcomes of interest in each location: fraction of correct guesses in Panel (a) and consensus rates in Panel (b). The evolution and the levels of both outcomes are extremely similar in the two locations, which is confirmed by the statistical analysis. Regression analysis detects no significant differences between two locations with $p > 0.10$ in all comparisons. Figure F.2 complements aggregate results by depicting individual responsiveness to signals and to others' actions in the two locations. We find no significant differences between the distribution of individual behavior of subjects in the two locations, as measured by these two statistics.

We conclude by noting that we detect no significant differences in neither aggregate results nor in the individual level results between sessions conducted in person at UCSD and sessions conducted online at OSU, which is why the analysis in the main text pools together these sessions.

G Learning Across Games

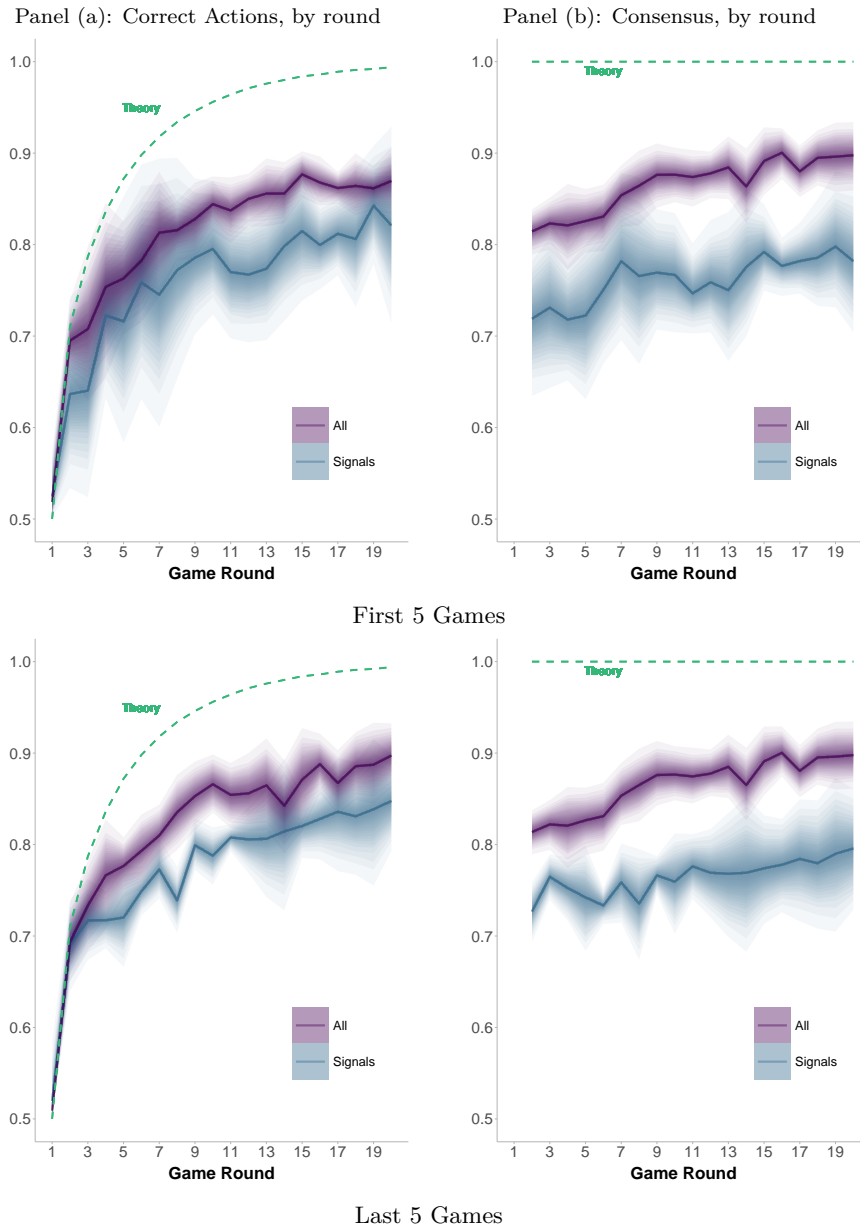


Figure G.1: ALL and SIGNALS treatments (early versus late games)

Notes: Panel (a) presents the average frequency of correct actions in each treatment in each round, averaged across the first 5 games (top figure) and the last 5 games (bottom figure). Panel (b) depicts the evolution of consensus in each round, i.e., the relative size of the majority, averaged across the first 5 games (top figure) and the last 5 games (bottom figure). For Panel (b) we exclude cases with equal number of green and red signals. Shaded regions represent confidence intervals from 50% (darkest) to 95% (faintest) probability levels. Confidence intervals are constructed with a variance-covariance matrix clustered by session.

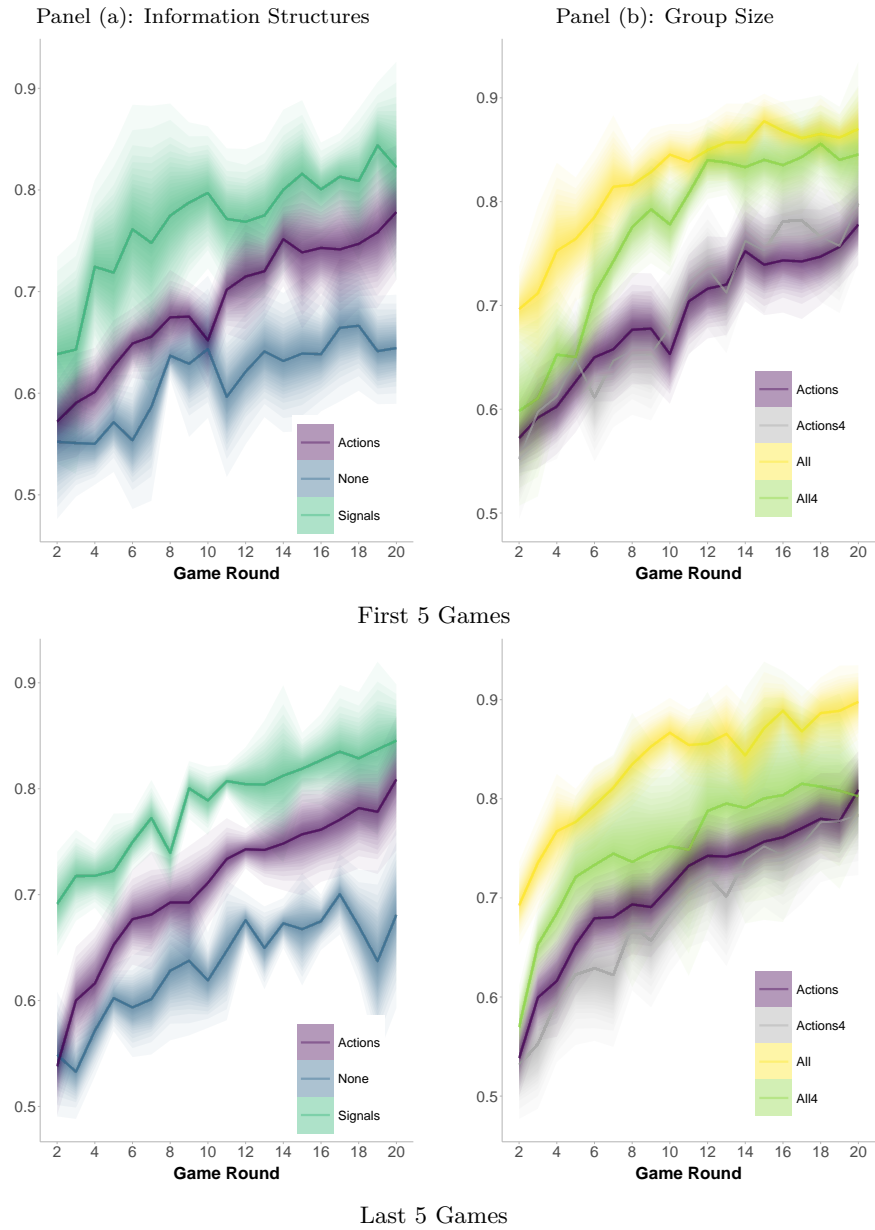


Figure G.2: Frequency of correct actions (early versus late games), by information structure and group size

Notes: Both panels present the average frequency of correct actions in each treatment per each round, averaged across the first 5 games (top figure) and the last 5 games (bottom figure). Shaded regions represent confidence intervals from 50% (darkest) to 95% (faintest) probability levels. Confidence intervals are constructed with a variance-covariance matrix clustered by session.

H Learning From Others' Actions

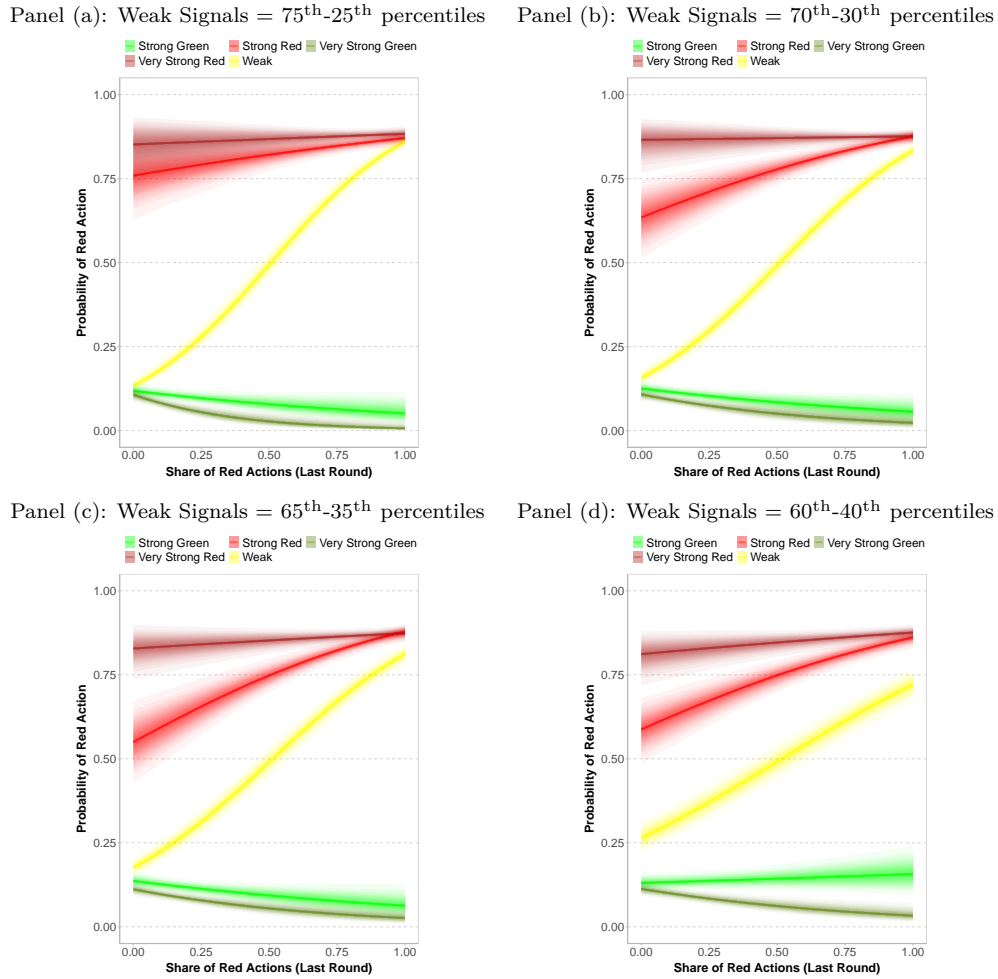


Figure H.1: Different Cutoffs for Signal Strength in ALL treatment

Notes: Each panel depicts the probability of guessing red as a function of the share of red actions of other group members in the previous round. The estimates are obtained from a Bayesian logistic regression of subjects' actions on the share of others' actions in the previous round conditional on signal strength. For panel (a) we use the classification in the main text: *Very Strong Green* with *percentile* = (0, 0.1], *Strong Green* with *percentile* = (0.1, 0.25], *Weak* with *percentile* = (0.25, 0.75], *Strong Red* with *percentile* = (0.75, 0.9] and *Very Strong Red* with *percentile* = (0.9, 1). For the cutoffs of panels (b), (c) and (d) we reduce the *Weak* signals category and increase the "strong" categories proportionally: *Very Strong Green* with *percentile* = (0, 0.1 + x], *Strong Green* with *percentile* = (0.1 + x , 0.25 + x], *Weak* with *percentile* = (0.25 + x , 0.75 - x], *Strong Red* with *percentile* = (0.75 - x , 0.9 - x] and *Very Strong Red* with *percentile* = (0.9 - x , 1), where $x \in \{0.05, 0.1, 0.15\}$.

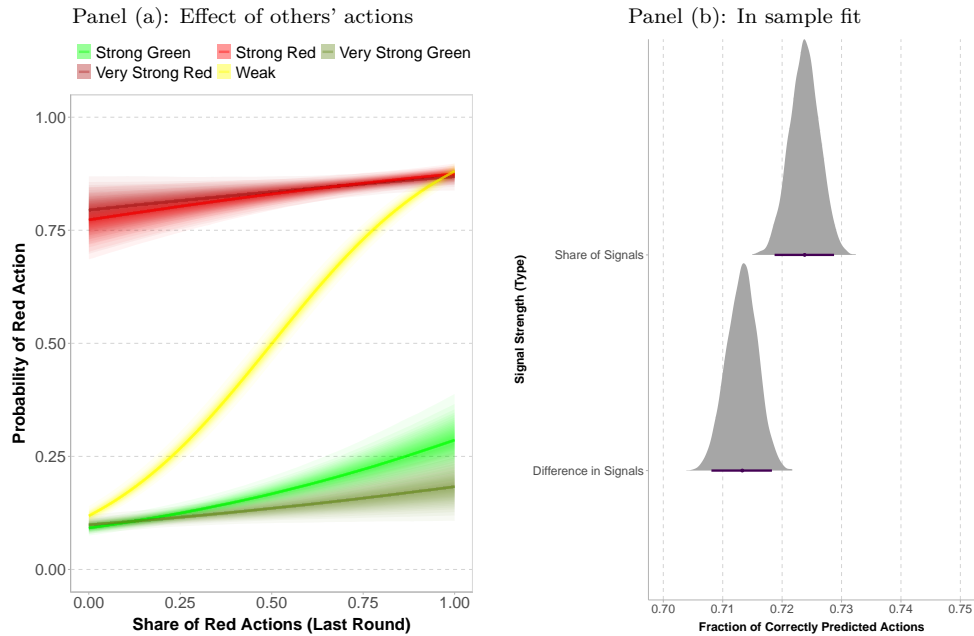


Figure H.2: Alternative Signal Strength Measures (Difference in Signals *vs* Share of Signals)

Notes: Panel (a) depicts the probability of choosing red as a function of the share of red actions of other group members conditional on signal strength as constructed from the share of red signals. Shaded regions represent 95% credible intervals from 50% (darkest) to 95% (faintest) probability levels. Panel (b) depicts the posterior distributions, along with 95% credible intervals, of the fraction of correctly predicted guesses using either the difference in red and green signals or the share of red signals to construct signal strength. For both measures of signal strength, we estimate a Bayesian logistic regression of subjects' actions on the share of others' actions in the previous round conditional on signal strength and session random effects.


I Additional Aggregate Results

Table I.1: Treatment Effects by Group Size

	<i>Dependent variable:</i>			
	Consensus Rate			
	(1)	(2)	(3)	(4)
ACTIONS (Baseline)	0.686*** (0.011)	0.681*** (0.010)		
ALL (Baseline)			0.820*** (0.010)	0.812*** (0.011)
ACTIONS4 (Effect)	0.090*** (0.010)	0.101*** (0.009)		
ACTIONS4 (Effect) \times Late Rounds		-0.020* (0.010)		
ALL4 (Effect)			0.049*** (0.009)	0.064*** (0.011)
ALL4 (Effect) \times Late Rounds				-0.028** (0.014)
Game Round Fixed Effects	Yes	Yes	Yes	Yes
Observations	5,550	5,550	6,841	6,841
Adjusted R ²	0.160	0.161	0.058	0.061

Notes: *p<0.1; **p<0.05; ***p<0.01. Clustered standard errors by session in parentheses. Late rounds are 11-20.

Figure I.1: Consensus Rates, by Group Size



Notes: Both panels present the consensus rates in each treatment per each round, averaged across games within a session. Shaded regions represent confidence intervals from 50% (darkest) to 95% (faintest) probability levels. Confidence intervals are constructed with a variance-covariance matrix clustered by session.

Table I.2: Probability of Red Bet by Signal Strength (ALL versus SIGNALS)

	<i>Dependent variable:</i>			
	Red Bet			
	(1)	(2)	(3)	(4)
Constant	0.088*** (0.011)	0.199*** (0.023)	0.174*** (0.056)	-0.078 (0.060)
SIGNALS	0.032 (0.032)	0.032 (0.032)	0.046* (0.026)	0.300*** (0.026)
<i>Strong Green</i>	0.020** (0.009)	0.015* (0.009)	0.032** (0.016)	0.031* (0.016)
<i>Weak</i>	0.369*** (0.019)	0.315*** (0.026)	0.321*** (0.024)	0.306*** (0.022)
<i>Strong Red</i>	0.765*** (0.021)	0.753*** (0.022)	0.747*** (0.020)	0.740*** (0.017)
<i>Very Strong Red</i>	0.790*** (0.022)	0.790*** (0.022)	0.786*** (0.024)	0.780*** (0.021)
<i>Strong Green</i> × SIGNALS	0.028*** (0.011)	0.025** (0.011)	0.006 (0.024)	0.003 (0.022)
<i>Weak</i> × SIGNALS	-0.111*** (0.030)	-0.108*** (0.030)	-0.120*** (0.031)	-0.123*** (0.030)
<i>Strong Red</i> × SIGNALS	-0.089 (0.076)	-0.090 (0.074)	-0.116* (0.067)	-0.122* (0.063)
<i>Very Strong Red</i> × SIGNALS	-0.066 (0.068)	-0.065 (0.068)	-0.099 (0.069)	-0.114* (0.067)
Game Round Fixed Effects	No	Yes	Yes	Yes
Game Fixed Effects	No	No	Yes	Yes
Participant Fixed Effects	No	No	No	Yes
Observations	44,460	44,460	44,460	44,460
Adjusted R ²	0.271	0.274	0.285	0.321

Notes: *p<0.1; **p<0.05; ***p<0.01. Clustered standard errors by session in parentheses.

J Additional Individual Results

For participant covariates we include: **female**, which is an indicator variable that takes the value of one if the participant identifies as female and zero, otherwise. **stem**, which is an indicator variable that takes the value of one if the participant's major is STEM and zero, otherwise. **overconfidence** measures the extent of a participant's over-estimation of her IQ, which is given by the difference between the number of questions a participant believes she solved correctly and the actual number of correct answers. **risk** is measured by the number of points invested in a risky asset as specified in a risky investment task solved at the end of the game.

Table J.1: Probability of Correct Guess by IQ Level (ALL versus SIGNALS)

	<i>Dependent variable:</i>			
	Correct Guess			
	(1)	(2)	(3)	(4)
Constant	0.740*** (0.021)	0.614*** (0.019)	0.577*** (0.022)	0.506*** (0.038)
SIGNALS	-0.055** (0.026)	-0.055** (0.026)	-0.055** (0.026)	-0.041 (0.029)
IQ	0.024*** (0.004)	0.024*** (0.004)	0.024*** (0.004)	0.031*** (0.006)
female				0.023 (0.022)
stem				0.018 (0.021)
overconfidence				0.015** (0.007)
risk				0.0002 (0.0001)
IQ × SIGNALS	0.003 (0.011)	0.003 (0.011)	0.003 (0.011)	0.001 (0.010)
Game Round Fixed Effects	No	Yes	Yes	Yes
Game Fixed Effects	No	No	Yes	Yes
Participant Covariates	No	No	No	Yes
Observations	44,460	44,460	44,460	44,460
Adjusted R ²	0.014	0.032	0.034	0.037

Notes: *p<0.1; **p<0.05; ***p<0.01. Clustered standard errors by session in parentheses.

Table J.2: Probability of Correct Guess by low-IQ/high-IQ Subjects (ALL versus SIGNALS)

	<i>Dependent variable:</i>			
	Correct Guess			
	(1)	(2)	(3)	(4)
Constant	0.863*** (0.015)	0.737*** (0.014)	0.701*** (0.015)	0.653*** (0.033)
SIGNALS	-0.047 (0.045)	-0.047 (0.045)	-0.047 (0.045)	-0.044 (0.042)
<i>low-IQ</i>	-0.074*** (0.016)	-0.074*** (0.016)	-0.074*** (0.016)	-0.088*** (0.026)
female				0.023 (0.024)
stem				0.023 (0.020)
overconfidence				0.009 (0.007)
risk				0.0003* (0.0001)
<i>low-IQ</i> × SIGNALS	0.010 (0.031)	0.010 (0.031)	0.010 (0.031)	0.017 (0.031)
Game Round Fixed Effects	No	Yes	Yes	Yes
Game Fixed Effects	No	No	Yes	Yes
Participant Covariates	No	No	No	Yes
Observations	44,460	44,460	44,460	44,460
Adjusted R ²	0.012	0.029	0.031	0.035

Notes: *p<0.1; **p<0.05; ***p<0.01. Clustered standard errors by session in parentheses.




Figure J.1: Learning from others' actions in ALL treatments

Notes: Panel (a) depicts the probability of choosing red as a function of the share of red actions of other group members in the ALL4 treatment, obtained from a Bayesian logistic regression of the subject's action on the share of others' actions in the previous round conditional on the difference between red and green signals. Solid lines depict the median of the posterior distribution and dashed lines depict 95% confidence intervals. Panel (b) presents the same exercise for the ALL treatment. Panel (c) shows the difference in the probability of a red bet between the ALL and ALL4 treatments for a *Weak* signal strength.




Figure J.2: Responsiveness to signals and actions, individual level data (ACTIONS)

Notes: Panel (a) shows the kernel distributions of participants' responsiveness to signals for both weak and strong signals. The vertical lines and the numbers next to them depict median responsiveness for each group. Responsiveness to signals is calculated based on equation (3) in the main text for ACTIONS. Responsiveness is given by the probability that a participant's action matches signal majority, i.e., the probability of an optimal action. Panel (b) shows the kernel distributions of participants' responsiveness to others' actions for weak and strong signals. The vertical lines and the numbers next to them depict median responsiveness for each group. Responsiveness to others' actions is measured by the change in the probability of choosing the action of the majority of signals when all versus none of the other group members choose the majority of signals in the last round.

K Informational Value of Others' Actions

One of the main findings from comparing the SIGNALS and ALL treatments is that participants in the ALL information treatment outperform those in the SIGNALS treatment. In this Appendix we decompose this improvement in participants' performance across treatments into the contributions of private information from signals and social information from others' actions. We quantify the contribution of each source relative to the performance implied by Bayesian updating. Specifically, we compute the root mean squared error (RMSE) between the Bayesian benchmark probabilities of correct actions (as shown in Table 2) and model-predicted probabilities under two specifications: one in which participants' behavior depends only on signals, and another in which it depends on both signals and others' actions.

As before, let v_{it}^g denote the correct action of participant i in round t of game g , where $v_{it}^g \sim \text{Bernoulli}(p_{it}^g)$. We model the probability of a correct action as a flexible function of the information contained in signals and in others' actions, while controlling for treatment differences and round effects.

$$p_{it}^g = \text{logit}^{-1}\left(\alpha_{ig} + \lambda_t + \delta_{treat[i]} + \beta_1 \text{log-odds}_{it}^{\text{correct}} + \beta_2 \text{share}_{v_{-i,t-1}}\right), \quad (7)$$

where, α_{ig} is a participant–game-specific random intercept, λ_t is a round-specific fixed effect, and $\delta_{treat[i]}$ is a fixed effect for the treatment assigned to participant i .

We capture the information contained in signals using the posterior log-odds in favor of the true state ω relative to the alternative state ω^c :

$$\text{log-odds}_{it}^{\text{correct}} = \log\left(\frac{\mathbb{P}(\omega \mid I_{i,t})}{\mathbb{P}(\omega^c \mid I_{i,t})}\right) = \begin{cases} \text{log-odds}_{it}, & \text{if } \omega = R, \\ -\text{log-odds}_{it}, & \text{if } \omega = G, \end{cases}$$

where log-odds_{it} , defined in equation (8), measures the evidence in favor of $\omega = R$.

The information conveyed by others' actions is summarized by the share of others' correct choices in the previous round, $\text{share}_{v_{-i,t-1}}$.

We estimate the parameters in equation (7) using a Bayesian logistic regression that pools participants' choices from the SIGNALS and ALL treatments.

Figure K.3 summarizes the main results from the estimation of equation (7). The left panel reports the posterior distribution of the probability of a correct action under the SIGNALS and ALL treatments, averaging across participants, games, and rounds, while holding fixed the information from signals (log-odds_{it}) and others' actions ($\text{share}_{v_{-i,t-1}}$) at their median values in the data. We find a significant increase in performance of 6 percentage points when participants are assigned to the ALL treatment relative to the SIGNALS treatment, that cannot be directly attributed to the informational value of signals or others' actions.

The middle panel reports the effect of signals on performance by treatment. We show the probability of a correct action as $\log\text{-odds}_{it}^{\text{correct}}$ varies from its minimum to its maximum in the data, while holding others' actions ($share_{v-i,t-1}$) fixed at their median values. For participants in both the SIGNALS and ALL treatments, performance improves substantially as signals become more informative about the true state. This effect is stronger for participants in the ALL treatment, particularly when signals are less informative.

The right panel isolates the effect of others' actions on performance by reporting the probability of a correct action as $share_{v-i,t-1}$ varies from zero to one, while holding the informativeness of signals ($\log\text{-odds}_{it}^{\text{correct}}$) fixed at its median value. Consistent with the results in the main text (Observation 2), participants in the ALL treatment improve their performance by relying on the redundant information conveyed by others' actions. This improvement remains substantial even after accounting for the informativeness of observed signals, increasing the probability of a correct action by more than 10 percentage points as the share of others' correct actions in the previous round rises from zero to one.

Using the estimated parameters, we construct a *baseline* predicted probability of a correct action driven only by variation across participants and games (α_{ig}), rounds (λ_t), and treatments ($\delta_{treat[i]}$), setting $\beta_1 = \beta_2 = 0$ and thereby shutting down the contribution of both signals and others' actions.

We then construct predicted probabilities driven by signals by adding to the *baseline* specification the information provided by signals, $\beta_1 \log\text{-odds}_{it}^{\text{correct}}$, while keeping $\beta_2 = 0$.

Finally, we construct predicted probabilities driven by *signals + actions* by augmenting the *signals* specification with the effect of others' actions, $\beta_2 share_{v-i,t-1}$.

Figure K.4 reports, for each game round, the additional informational value of others' actions beyond signals by comparing the posterior RMSE of predicted probabilities relative to the Bayesian benchmark under the *signals* and *signals + actions* specifications. Averaging across rounds, incorporating others' actions reduces RMSE by approximately 1 percentage point. This reduction is driven by improved accuracy from round seven onward, where the RMSE decreases by 1.5 percentage points.

Table K.3 reports the median RMSE estimates across scenarios. Overall, the largest improvement in accuracy comes from incorporating signals, which reduces RMSE by approximately 3 percentage points relative to the *baseline* scenario, with a particularly large effect in later rounds (up to 6 percentage points). Beyond the information provided by signals, incorporating the redundant information contained in others' actions further improves participants' performance relative to the Bayesian benchmark, with an additional reduction in RMSE of 1 percentage point, equivalent to one-third of the informational value of signals.

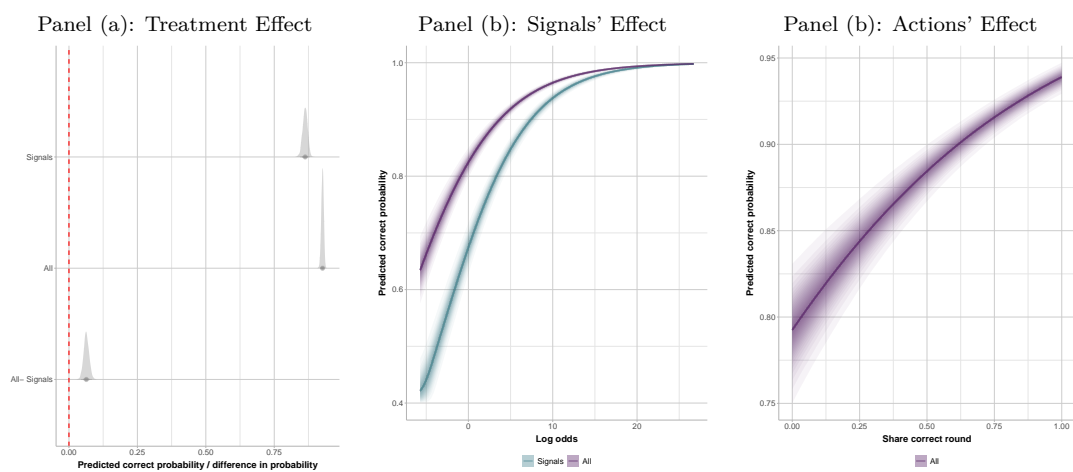


Figure K.3: Effect of Information on Signals and Actions

Notes: Panel (a) presents the posterior probability of a correct choice for the SIGNALS and ALL treatments, holding fixed both the signal information and the information from others' actions. Panel (b) presents the posterior probability of a correct choice as log-odds vary from their minimum to maximum values in the data, for both the SIGNALS and ALL treatments, while holding the information from others' actions fixed. Panel (c) presents the posterior probability of a correct choice as the share of others' correct choices varies from its minimum to maximum values in the data for the ALL treatment, while holding the signal information fixed.

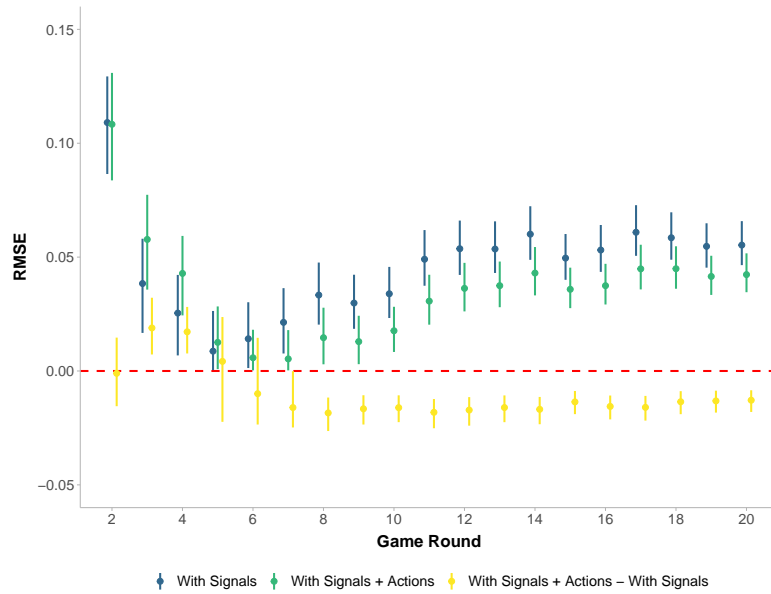


Figure K.4: RMSE relative to Optimal Behavior

Notes: This figure reports the posterior median root mean squared error (RMSE) between the Bayesian benchmark probabilities of correct actions and the model-predicted probabilities under two specifications. Predicted probabilities based on *signals* include treatment fixed effects, game-round fixed effects, and signal information measured using posterior log-odds. Predicted probabilities based on *signals + actions* additionally incorporate social information from others' actions, measured as the share of others' correct choices in the previous round.

Table K.3: The Effect of Signals and information on RMSE.

Round	Baseline	Signals	S + A	(S-B)	(S+A - Signals)
Round 2	0.216	0.109	0.108	-0.106	-0.001
Round 3	0.135	0.038	0.058	-0.096	0.019
Round 4	0.098	0.025	0.043	-0.072	0.017
Round 5	0.057	0.009	0.013	-0.049	0.004
Round 6	0.036	0.014	0.006	-0.022	-0.010
Round 7	0.017	0.021	0.005	0.004	-0.016
Round 8	0.004	0.033	0.015	0.029	-0.018
Round 9	0.010	0.030	0.013	0.020	-0.017
Round 10	0.022	0.034	0.018	0.012	-0.016
Round 11	0.043	0.049	0.031	0.006	-0.018
Round 12	0.054	0.054	0.036	-0.000	-0.017
Round 13	0.062	0.054	0.037	-0.008	-0.016
Round 14	0.077	0.060	0.043	-0.017	-0.017
Round 15	0.071	0.050	0.036	-0.022	-0.014
Round 16	0.083	0.053	0.037	-0.030	-0.016
Round 17	0.101	0.061	0.045	-0.040	-0.016
Round 18	0.108	0.058	0.045	-0.049	-0.014
Round 19	0.113	0.055	0.042	-0.058	-0.013
Round 20	0.124	0.055	0.042	-0.069	-0.013
Average	0.075	0.045	0.035	-0.030	-0.010

Notes: This table reports posterior median RMSE between the probabilities of correct actions under the Bayesian Benchmark and predicted probabilities for three different scenarios: (i) *baseline* includes a treatment fixed effect and game-round fixed effects. (ii) *signals* additionally incorporate signal information measured using log-odds. (iii) *signals + actions* additionally incorporate others' actions, measured by the share of others' correct choices in the previous round.

L Heterogeneity in the Effect of Redundant Information

The main finding that participants perform better in the ALL treatment relative to the SIGNALS treatment, as shown in Figure 1, reflects an average treatment effect that masks heterogeneity across participants in their likelihood of making a correct bet. In this Appendix, we show that the performance advantage of the ALL treatment over the SIGNALS treatment is robust to conditioning on the distribution of prediction quality across participants.

For this purpose, we model the probability of making a correct bet at the individual level using a flexible binary-response specification with individual-specific random effects.

For a given treatment (ALL or SIGNALS), the data generating process for the correct action of participant i in round t of game g , v_{it}^g , is given by $v_{it}^g \sim \text{Bernoulli}(p_{it}^g)$, where

$$p_{it}^g = \text{logit}^{-1}(\alpha_i + \beta_{it} + v_g),$$

where α_i is a participant-specific random intercept, β_{it} is a participant-specific random effect for round t . We allow a flexible correlation pattern across individual effects:

$$\begin{pmatrix} \alpha_i \\ \beta_{it} \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Sigma_u\right).$$

Finally, $v_g \sim \mathcal{N}(0, \sigma_g^2)$ denotes the game-specific random intercept.

We leverage variation across rounds and games within participants to recover subject-level response probabilities via Bayesian logistic regression. Using the posterior estimates of individual probabilities, we replicate the analysis in Section 5.1, comparing how often subjects correctly guess the state in each game round across treatments by quantiles of these individual probabilities.

Figure L.1 summarizes these results. Consistent with the aggregate findings, subjects across quantiles perform better in the ALL treatment relative to SIGNALS. Moreover, the effect of redundant information is significantly larger for participants in the lower quantiles of the distribution (i.e., the 5th, 10th, and 25th percentiles). Although the gap between treatments is smaller at higher quantiles, it remains statistically significant and economically meaningful (i.e., around 5 p.p.), particularly in earlier rounds.

Figure L.2 shows the full distribution of individual-level correct probabilities across treatments and game rounds. Consistent with the aggregate findings, as the game progresses, individuals in both SIGNALS and ALL treatments improve their performance. However, we find substantial heterogeneity across participants, with individual probabilities ranging from below 0.25 to above 0.95. Comparing treatments, we find that in later rounds (i.e., rounds 10, 15, and 20) a larger share of individuals performs better in the ALL treatment than in

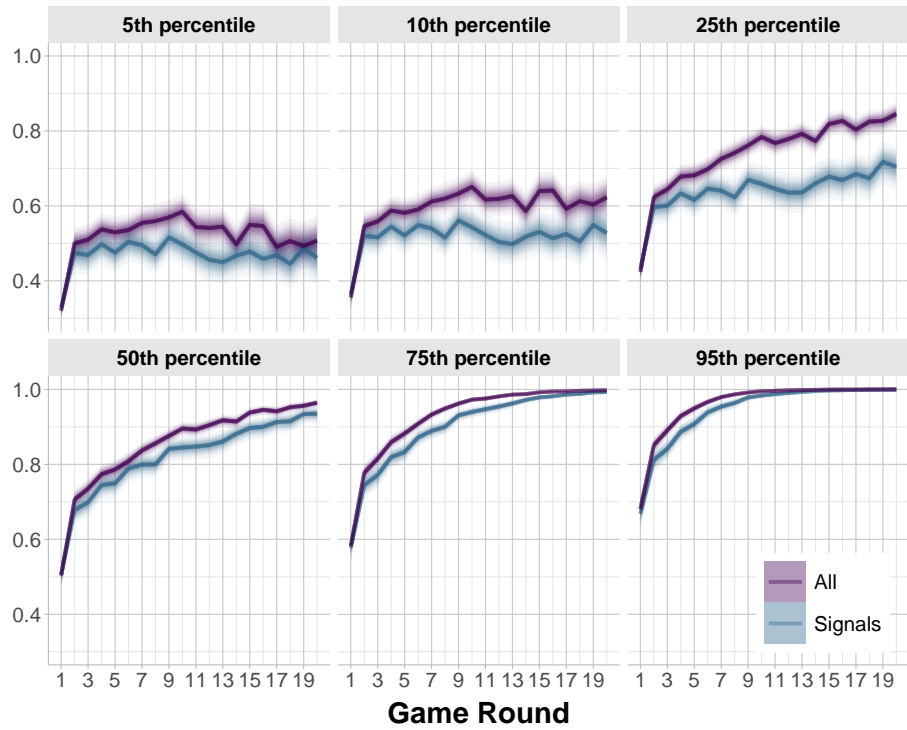


Figure L.1: Statistics in the ALL and SIGNALS treatments by participant heterogeneity

Notes: The figure presents the average probability of correct actions in each round, averaged across games. Shaded regions represent credible intervals from 50% (darkest) to 95% (faintest) levels.

the SIGNALS treatment.




Figure L.2: Statistics in the ALL and SIGNALS treatments by participant heterogeneity

Notes: The figure presents the average probability of correct actions in each round, averaged across games. Shaded regions represent credible intervals from 50% (darkest) to 95% (faintest) levels.

M Covariate Balance in Participant Characteristics

In Appendix F, we show that aggregate outcomes and individual-level behavior are very similar and statistically indistinguishable across experimental sessions conducted in the physical lab at UCSD and in the virtual lab at OSU. However, because half of the experimental sessions in the ALL and SIGNALS treatments were conducted at different locations, with distinct pools of college students (UCSD and OSU), we assess in this Appendix potential imbalances in participants' characteristics across treatments and their implications for inference on individual behavior.

As a measure of covariate balance, Figure M.1 reports absolute standardized mean differences between the ALL and SIGNALS treatments. We find almost no imbalance for the constructed covariates *overconfidence* and *risk* (less than 0.1), moderate imbalance for *female* and the constructed *iq* (less than 0.25), and substantial imbalance for *STEM*, with 31% of participants in the ALL treatment having a STEM major, compared to 21% in the SIGNALS treatment.

We summarize the extent of imbalance across covariates by estimating the propensity score, defined in our setting as

$$\mathbb{P}(\{T_i = \text{ALL} \mid X_i\},$$

where X_i is the vector of participant covariates (`overconfidence`, `risk`, `female`, `iq`, `STEM`). We estimate the propensity score using a logistic regression of T_i on X_i .

Figure M.2 compares the distribution of the estimated propensity score between the SIGNALS and ALL treatments by contrasting their quantiles across groups. The purple dots lie consistently above the 45-degree line, indicating that participants in the ALL treatment differ systematically from those in the SIGNALS treatment in their observable characteristics.

Based on the estimated propensity score, we implement one-to-one optimal matching. Each of the 82 participants in the SIGNALS treatment is matched to one of the 152 participants in the ALL treatment, yielding a total of 164 matched observations (82 matched pairs), with 70 unmatched participants in the ALL treatment.

The yellow triangles in Figure M.2, which lie close to the 45-degree line, indicate that matching yields nearly identical propensity score distributions across the ALL and SIGNALS treatments at most quantiles.

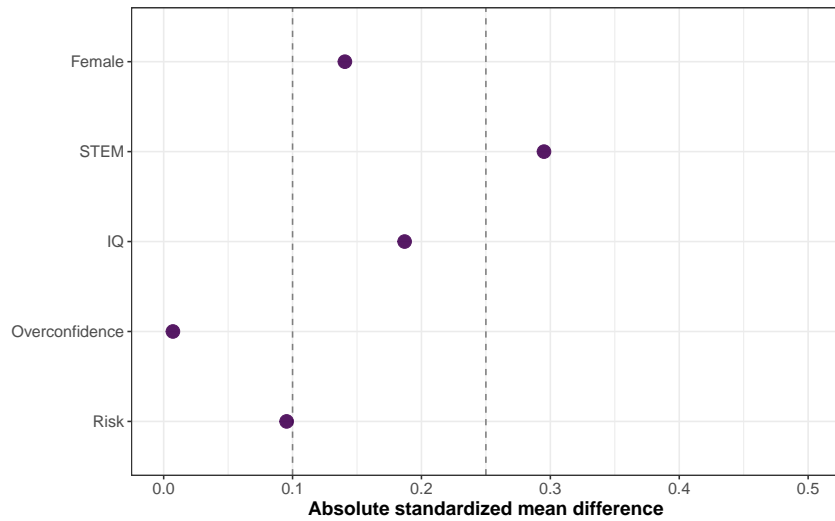


Figure M.1: Covariate Balance

Notes: The figure presents absolute standardized mean differences in covariate value between ALL and SIGNALS treatments.

Using the matched sample, we replicate the main analysis comparing the fraction of correct actions between the SIGNALS and ALL treatments. Figure M.3 reports the results of this exercise.

First, in the balanced subsample, the main aggregate finding remains unchanged: participants perform better in the ALL treatment than in the SIGNALS treatment. Moreover, relative to the full sample, participants in the matched sample perform significantly better in both treatments, with outcomes approaching the theoretical probabilities of correct actions under

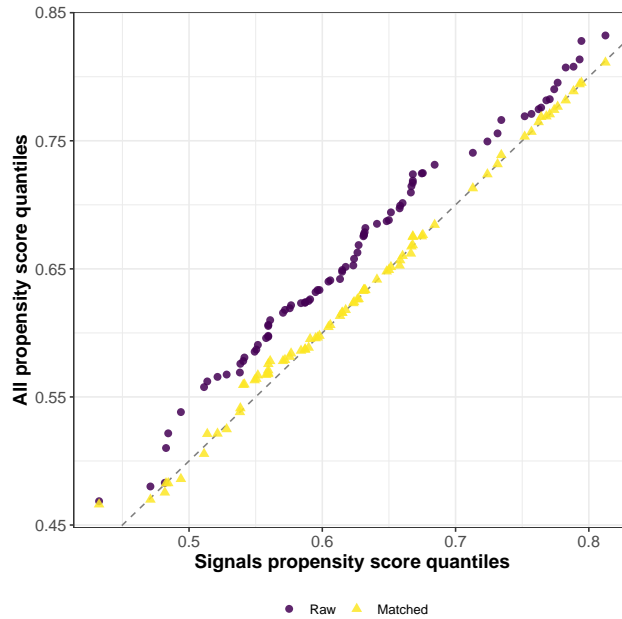


Figure M.2: QQ Plot of Propensity Score for Treatment Assignment (ALL and SIGNALS)

Notes: The purple dots represent empirical QQ estimates for the raw data. The yellow dots represent QQ estimates for the matched data. The 45-degree line indicates identical distributions. The propensity score is estimated with a logistic regression of treatment (ALL and SIGNALS) on **female**, **stem**, **iq**, **overconfidence** and **risk**.

Bayesian updating. Moreover, compared to the raw data, individuals in the matched sample perform significantly better in both treatments, almost matching the theoretical probabilities of correct actions under Bayesian updating and common knowledge of rationality.

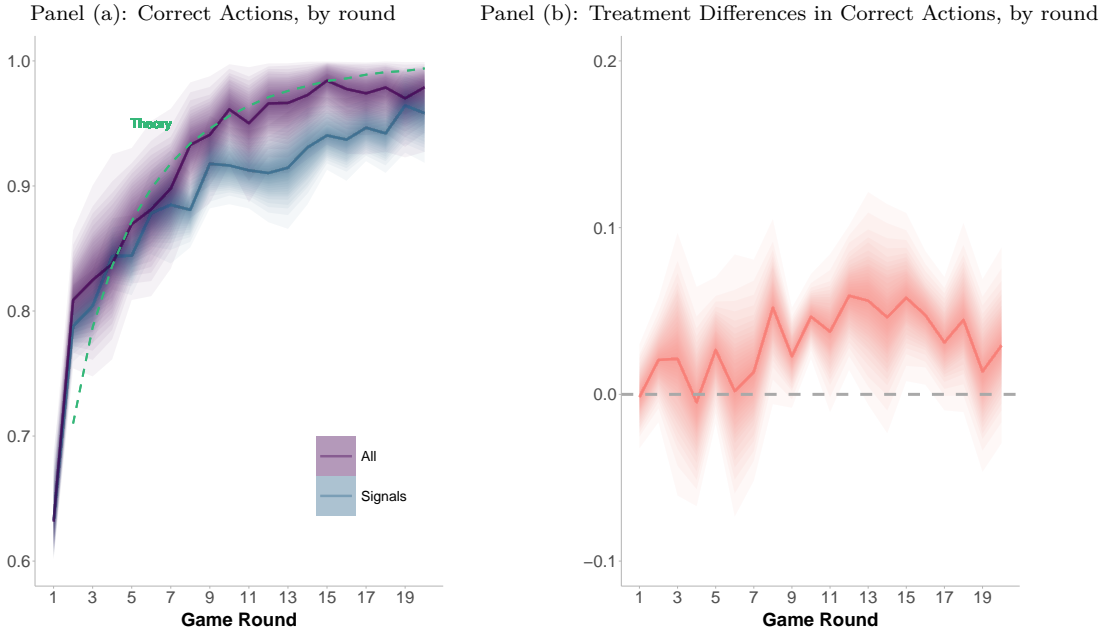


Figure M.3: Aggregate statistics in the ALL and SIGNALS treatments for the Exact Matched Sample

Notes: Panel (a) presents the average frequency of correct actions in each round, averaged across games. Panel (b) presents the treatment difference in the average frequency of correct actions in each round, averaged across games. Shaded regions represent 95% confidence intervals from 50% (darkest) to 95% (faintest) probability levels. Confidence intervals are constructed with a variance-covariance matrix clustered by session.

N Random Actions in the Data

Our main analysis examines how individuals learn from others' signals and actions, and the extent to which their behavior deviates from the Bayesian optimum. In this Appendix, we assess the extent to which such deviations are driven by participants who behave as if random, ignoring available information. We focus on the NONE and SIGNALS treatments, where relevant information is directly provided through observed signals.

To quantify random behavior, we estimate the effect of information $I_{i,t}$ on individual choices. We measure the information from signals with the posterior log-odds,

$$\text{log-odds}_{it} = \log \left(\frac{\Pr(\omega = R \mid I_{i,t})}{\Pr(\omega = G \mid I_{i,t})} \right), \quad (8)$$

which captures the relative evidence in favor of $\omega = R$ relative to $\omega = G$.³⁷ In particular, under the NONE treatment, $I_{i,t} = (s_{i,\tau})_{\tau < t}$, whereas under the SIGNALS treatment, $I_{i,t} =$

³⁷Similar results obtain when using the difference between the number of red and green signals.

$(s_{j,\tau})_{j \in N, \tau < t}$.

Given this measure of signals' informativeness, we model participant i 's choice in round t of game g , denoted a_{it}^g , as

$$a_{it}^g \sim \text{Bernoulli}(p_{it}^g),$$

where

$$p_{it}^g = \text{logit}^{-1}(\alpha_i + v_g + \lambda_t + \beta_i \log\text{-odds}_{it}).$$

The parameter of interest is β_i , which captures participant i 's sensitivity to the posterior log-odds. The term α_i is a participant-specific random intercept, v_g is a game-specific random intercept, and λ_t is a round-specific random intercept. We assume

$$\alpha_i \sim \mathcal{N}(0, \sigma_\alpha^2), \quad \beta_i \sim \mathcal{N}(0, \sigma_\beta^2), \quad v_g \sim \mathcal{N}(0, \sigma_v^2), \quad \lambda_t \sim \mathcal{N}(0, \sigma_\lambda^2).$$

We estimate responsiveness to information using a Bayesian logistic regression. Within this framework, we evaluate the posterior distribution of β_i , focusing on whether it places substantial mass near zero, consistent with random behavior or instead indicates learning, i.e., $\beta_i > 0$.

Figures N.4 and N.5 report the results of this estimation for the SIGNALS and NONE treatments, respectively. Panel (a) in these figures reports the posterior change in the probability of betting Red following a one-standard-deviation increase in the log-odds (in favor of $\omega = R$) across participants, corresponding to a homogeneous effect of information.

For both the SIGNALS and NONE treatments, we find a sizable effect of signals on behavior, with an increase of approximately 19 percentage points in the probability of betting red. However, this average effect masks substantial heterogeneity across participants. For more than half of participants in both treatments, the effect of a one-standard-deviation increase in log-odds favoring $\omega = R$ exceeds 25 percentage points.

On the other hand, around a quarter of participants in both treatments have posterior 95% credible intervals for β_i that include zero, suggesting weak or no responsiveness to information, consistent with random choices.




Figure N.4: Effect of Information on Actions (ALL Treatment)

Notes: Panel (a) presents the posterior change in bet probability after a one-SD increase in the log-odds ratio, where the slope parameter is common across participants ($\beta_i = \beta$, for all i). Panel (b) presents the posterior change in bet probability after a one-SD increase in the log-odds ratio, where the slope parameter, β_i is individual-specific




Figure N.5: Effect of Information on Actions (NONE Treatment)

Notes: Panel (a) presents the posterior change in bet probability after a one-SD increase in the log-odds ratio, where the slope parameter is common across participants ($\beta_i = \beta$, for all i). Panel (b) presents the posterior change in bet probability after a one-SD increase in the log-odds ratio, where the slope parameter, β_i is individual-specific

References

- M. Agranov, B. Gillen, and D. Persitz. Behavioral and structural barriers to information aggregation in networks. *working paper*, 2025.
- L. R. Anderson and C. A. Holt. Information cascades in the laboratory. *The American economic review*, pages 847–862, 1997.
- M. Angrisani, A. Guarino, P. Jehiel, and T. Kitagawa. Herd behavior in a laboratory financial market. *American Economic Review*, 95:1427–1443, 2005.
- M. Angrisani, A. Guarino, P. Jehiel, and T. Kitagawa. Information redundancy neglect versus overconfidence: a social learning experiment. *American Economic Journal: Microeconomics*, 13(3):163–97, 2021.
- Y. Azrieli, C. Chambers, and P. Healy. Incentives in experiments: A theoretical analysis. *Journal of Political Economy*, 126:1472–1503, 2018.
- A. V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, pages 797–817, 1992.
- D. J. Benjamin. Errors in probabilistic reasoning and judgment biases. *Handbook of Behavioral Economics*, edited by Doug Bernheim, Stefano Della Vigna, and David Laibson, 2019.
- S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, pages 992–1026, 1992.
- S. Bikhchandani, D. Hirshleifer, O. Tamuz, and I. Welch. Information cascades and social learning. Technical report, National Bureau of Economic Research, 2021.
- A. Caplin, M. Dean, and D. Martin. Search and satisficing. *American Economic Review*, pages 2899–2922, 2011.
- B. Çelen and S. Kariv. Distinguishing informational cascades from herd behavior in the laboratory. *American Economic Review*, 94(3):484–498, 2004.
- A. Chandrasekhar, H. Larreguy, and J. Xandri. Testing models of social learning on networks: Evidence from two experiments. *Econometrica*, pages 1–32, 2020.
- J. Chapman, M. Dean, P. Ortoleva, E. Snowberg, and C. Camerer. Econographics. *NBER working paper #24931*, 2019.

- D. L. Chen, M. Schonger, and C. Wickens. otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97, 2016.
- S. Choi, D. Gale, and S. Kariv. Social learning in networks: a quantal response equilibrium analysis of experimental data. *Review of Economic Design*, pages 135–157, 2012.
- D. M. Condon and W. Revelle. The international cognitive ability resource: Development and initial validation of a public-domain measure. *Intelligence*, 43:52–64, 2014.
- K. Dasaratha and K. He. An experiment on network density and sequential learning. *Games and Economic Behavior*, 128:182–192, 2021.
- R. De Filippis, A. Guarino, P. Jehiel, and T. Kitagawa. Non-bayesian updating in a social learning experiment. *Journal of Economic Theory*, 199:105188, 2022.
- J. Duffy, E. Hopkins, T. Kornienko, and M. Ma. Information choice in a social learning experiment. *Games and Economic Behavior*, 118:295–315, 2019.
- J. Duffy, E. Hopkins, and T. Kornienko. Lone wolf or herd animal? information choice and learning from others. *working paper*, 2020.
- B. Enke and F. Zimmermann. Correlation neglect in belief formation. *Review of Economic Studies*, pages 313–332, 2019.
- P. Evdokimov and U. Garfagnini. Individual vs. social learning: An experiment. *working paper*, 2020.
- E. Eyster and M. Rabin. Extensive imitation is irrational and harmful. *The Quarterly Journal of Economics*, 129(4):1861–1898, 2014.
- E. Eyster, M. Rabin, and G. Weizsacker. An experiment on social mislearning. *Available at SSRN 2704746*, 2018.
- D. Gale and S. Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 45(2):329–346, 2003.
- U. Gneezy and J. Potters. An experiment on risk taking and evaluation periods. *Quarterly Journal of Economics*, pages 631–645, 1997.
- J. K. Goeree, T. Palfrey, B. Rogers, and R. McKelvey. Self-correcting information cascades. *Review of Economic Studies*, 74(3):733–762, 2007.

- D. Griffin and A. Tversky. The weighing of evidence and the determinants of confidence. *Cognitive psychology*, 24(3):411–435, 1992.
- V. Grimm and F. Mengel. Experiments on belief formation in networks. *Journal of the European Economic Association*, 18(1):49–82, 2020.
- M. Harel, E. Mossel, P. Strack, and O. Tamuz. Rational groupthink. *The Quarterly Journal of Economics*, 136(1):621–668, 2021.
- W. Huang, P. Strack, and O. Tamuz. Learning in repeated interactions on networks. *arXiv preprint arXiv:2112.14265*, 2021.
- A. Hung and C. Plott. Information cascades: Replication and an extension to majority rule and conformity-rewarding institutions. *American Economic Review*, 91(5):1508–1520, 2001.
- D. Kubler and G. Weiszacker. Limited depth of reasoning and failure of cascade formation in the laboratory. *Review of Economic Studies*, 71(2):425–441, 2004.
- R. D. Luce. A probabilistic theory of utility. *Econometrica: Journal of the Econometric Society*, pages 193–224, 1958.
- E. Mossel, A. Sly, and O. Tamuz. Asymptotic learning on bayesian social networks. *Probability Theory and Related Fields*, 158(1-2):127–157, 2014.
- M. Mueller-Frank and C. Neri. A general model of boundedly rational observational learning: Theory and evidence. *working paper*, 2015.
- R. Parikh and P. Krasucki. Communication, consensus, and knowledge. *Journal of Economic Theory*, 52(1):178–189, 1990.
- M. E. Roberts, B. M. Stewart, D. Tingley, and E. M. Airoldi. The structural topic model and applied social science. In *Advances in neural information processing systems workshop on topic models: computation, application, and evaluation*, volume 4, pages 1–20. Harrahs and Harveys, Lake Tahoe, 2013.
- B. Scheibehenne and R. Greifeneder. Can there ever be too many options? a meta-analytic review of choice overload. *Journal of Consumer Research*, pages 409–425, 2010.
- L. Smith and P. Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.
- X. Vives. How fast do rational agents learn? *The Review of Economic Studies*, 60(2): 329–347, 1993.

- G. Weizsacker. Do we follow others when we should? a simple test of rational expectations. *American Economic Review*, pages 2340–236, 2010.
- A. Ziegelmeyer, F. Koessler, B. J., and E. Winter. Fragility of information cascades: An experimental study using elicited beliefs. *Experimental Economics*, 13(2):121–145, 2010.
- A. Ziegelmeyer, C. March, and S. Krugel. Do we follow others when we should? a simple test of rational expectations: Comment. *Experimental Economics*, 103(6):2633–42, 2013.