

Not Learning from Others*

John J. Conlon[†] Malavika Mani Gautam Rao
Matthew Ridley Frank Schilbach

March 4, 2026

Abstract

We study social learning using experiments where two people independently learn relevant information and can share it to make accurate private decisions. Across three experiments, people are less sensitive to information others discover than to equally relevant information they discover themselves. This is not due to a failure to elicit information from others. Instead, it reflects a systematic tendency to underweight others' information relative to one's own. This holds even when others' information is perfectly communicated or directly observed. Our findings illustrate a powerful barrier to social learning.

Keywords: information aggregation, learning, social learning, experience effects

*We thank the editor and four anonymous referees for their helpful comments and suggestions. We also thank Sandro Ambuehl, Nava Ashraf, Ned Augenblick, Oriana Bandiera, Abhijit Banerjee, Rafael Batista, Leonardo Bursztyn, Arun Chandrasekhar, Katie Coffman, Stefano DellaVigna, Esther Duflo, Ben Enke, Christine Exley, Ben Golub, Emir Kamenica, David Laibson, Shengwu Li, Ulrike Malmendier, Madeline McKelway, Michel Maréchal, Sendhil Mullainathan, Tommy O'Donnell, Ricardo Perez-Truglia, Matthew Rabin, Chris Roth, Dmitry Taubinsky, David Yanagizawa-Drott, and numerous seminar and conference participants for helpful suggestions and comments. We thank Sangeetha Ramanathan and the entire team at the Behavioral Development Lab in Chennai for excellent research assistance. We thank all our study participants for their time and patience. This project was funded by the Wellspring Philanthropic Fund. We received IRB approval from MIT, protocol #1810538700. Each experiment was pre-registered on the AEA registry, number AEARCTR-0004253.

[†]Conlon: Carnegie Mellon University (jconlon@andrew.cmu.edu); Mani: University of Warwick (Malavika.Mani@warwick.ac.uk); Rao: UC Berkeley and NBER (grao@berkeley.edu); Ridley: University of Warwick (Matthew.Ridley@warwick.ac.uk); Schilbach: MIT and NBER (fschilb@mit.edu)

1 Introduction

We often learn new information through our own actions or experiences: experimenting with a new technology, trying out a new restaurant, or personally visiting several schools before choosing one. But the actions and experiences of *others* also present us with a vast trove of useful data. We can talk to others who have used the technology, read restaurant reviews, or ask other students or parents for their opinions. Efficient social learning requires us to elicit information from others and correctly aggregate it with our own private information. How well do people do this, and what frictions might prevent social learning?

We study these questions experimentally. In our first experiment in Chennai, India, 500 participants face a simple learning task: making incentivized guesses of the fraction of red balls in an urn. To inform their guesses, participants have access to two sets of draws with replacement from the urn. In a control condition, participants make both sets of draws themselves and then make a guess. In the main treatment condition, the information is instead split between two participants. Each gets one set of private draws and the chance to learn each others' draws in an unstructured face-to-face discussion. After the discussion, each participant makes a private guess about the contents of the urn. Participants face aligned incentives to share information since the earnings from one partner's randomly-selected guess are equally split between the two of them.

Absent frictions to social learning, participants should be equally sensitive to information they uncover themselves and information their partner uncovered. In our reduced-form analysis, participants' guesses instead respond 54% less ($p < 0.01$) to signals their partners uncover—and which they can access through the discussion—than to signals they uncover themselves. This lower sensitivity to their partner's information is a mistake that reduces earnings: given the incentives for accurate guesses, an additional draw by one's partner improves one's earnings 73% less than receiving an additional own draw.

What drives this failure of social learning? First, it could reflect imperfect *communication*: participants might be reluctant to elicit or share information, or may mistrust others' competence or memory, preventing them from reliably learning what their partner uncovered. Second, there might be a failure of *information aggregation*: participants might not appropriately combine others' information with their own when forming guesses. To test between these possibilities, we run a treatment in which communication

frictions are shut down by design: the experimenter directly informs each participant about their partner’s draws.

Removing communication frictions does not increase sensitivity to others’ information, implying that our finding is driven by a systematic bias in information aggregation. When the experimenter perfectly communicates each partner’s draws, participants’ guesses are 86% less sensitive to others’ information than to own signals ($p < 0.01$). Discussion with one’s partner after being informed of their draws by the experimenter still leads to 46% less sensitivity to the partner’s information ($p < 0.05$). People thus appear to simply treat others’ signals as inherently less informative, even when communication is clear.

We replicate these findings and rule out several potential mechanisms in a second experiment with 292 adults in the same setting. Most strikingly, participants underweight their partner’s information by 41% relative to their own signal even when they sit beside their partner and can observe them drawing balls from the urn *with their own eyes* ($p = 0.04$). This result eliminates or diminishes any role for (i) distrust of the information communicated by the partner and/or experimenter, since the information is directly observed with one’s own eyes; and (ii) the mode of presentation of the information, including both its visual salience and whether the information is learned draw-by-draw or communicated in summary form.

Our third experiment demonstrates the external validity of our findings among a more educated population. In a simpler between-subjects experiment with 4,489 participants from the UK and US on Prolific, we randomize the order of learning one’s own signals versus a partner’s signals. Again, participants are less sensitive to others’ information than to their own, by 18% ($p < 0.01$), despite it being perfectly communicated to them. Presenting own and others’ signals using visually identical animations does not reduce underweighting, nor does doubling the stakes of the experiment or reducing any sense of competition by barring the partner from making any guess.

Participants appear largely unaware of their underweighting of others’ information, and we provide suggestive evidence that this bias is not fully explained by differential later recall. In a debriefing survey at the end of the third experiment, 77% of participants reported that they treated their own and their partner’s information the same. Yet these same participants are 15% less sensitive to their partner’s information than to their own ($p < 0.01$). In the survey, we also ask participants to recall their own and their partner’s

draws from the last round of the experiment. Among the majority of participants who perfectly recall theirs and their partner’s draws, others’ information is still underweighted by 9% ($p = 0.06$).

In our experiments and, arguably, in many natural environments, generating one’s own information requires taking action or experiencing something oneself, whereas learning from others is often more passive. Our third experiment provides evidence that this difference can at least in part explain the bias against others’ information. Specifically, in our online experiment, we vary whether participants click a button to generate their own draws or passively observe draws appearing on the screen with a label identifying each draw as ‘Your’ or ‘Partner’s’. When clicking to generate their “own” draws, they are 18 to 19% less sensitive to their partner’s information than to their own ($p < 0.01$). In contrast, when they must take no action to generate their own information, this under-sensitivity falls to only 4%, significantly smaller than with active involvement ($p < 0.05$) and not distinguishable from zero ($p = 0.27$). This finding suggests that when people take active efforts to uncover others’ information, social learning might be more effective.

We view the main contribution of this paper as being to the literature on social learning that investigates how agents learn from others (see Mobius and Rosenblat 2014 for a review).¹ We provide evidence for a novel, potentially far-reaching bias in information aggregation that may hinder social learning whenever people have to aggregate their own and others’ information. This phenomenon may underlie other documented cases of incomplete social learning. For example, lab studies of observational learning (largely following Anderson and Holt 1997) find that people put more weight on private information than on what can rationally be inferred from the actions of others (Weizsäcker 2010). Though these findings are consistent with our results, they are also consistent with other (sometimes rational) explanations, such as mistrust of others’ ability (De Filippis et al., 2017), noisy choice (Goeree et al. 2007), overconfidence (Angrisani et al., 2021), altruism (March and Ziegelmeyer, 2020), base-rate neglect (Benjamin et al., 2019), or other behavioral biases (Guarino and Jehiel, 2013). Our finding that agents down-weight others’ *signals*, not just their actions or beliefs, along with other features of our experiment, rules out these explanations.²

¹Existing research finds that people sometimes also react very differently to information depending on the identity of the sender, e.g., they may react more to information coming from celebrities (Alatas et al., 2021) or from people who are socially or economically similar to them or who are of a particular gender (BenYishay and Mobarak, 2019; BenYishay et al., 2020).

²Drehmann et al. (2005) include in their online experiment a treatment similar to Anderson and Holt

Our paper also relates to literatures on experience effects and reinforcement learning. Using observational data, previous work has shown that people’s beliefs and economic decisions are powerfully shaped by their personal experiences, even when much more complete data are readily available (Malmendier and Nagel, 2016; Malmendier et al., 2021; D’Acunto et al., 2021; Malmendier and Shen, 2024). Consistent with this, experiments on reinforcement learning show that people’s beliefs and actions react more to events that personally affect them (Merlo and Schotter 2003, Simonsohn et al. 2008, Miller and Maniadis 2012). Intuitively, burning one’s hand on a hot surface is more impactful than watching someone else do it or being told about it. Our findings echo this idea, with an important difference: even *before* (or without) any consequences or feedback, people are biased against information from others. Our findings might thus apply to decisions where any payoff realization lies in the distant future, such as education or career choices.

Our paper also builds on recent work on “ownership effects” in attention to information. In their experiments, Hartzmark et al. (2021) show that owning an asset draws attention to information about it, even when there is no instrumental reason for this differential attention.³ In our experiments, who uncovers information (oneself or others) affects how much a person reacts to it, even though it bears no relevance for the informativeness of the signals. In both cases, these instrumentally irrelevant factors may affect bottom-up attention to information. To the extent that generating information creates a sense of ownership, we extend the notion of ownership in Hartzmark et al. (2021) to information itself.

More broadly, this paper adds to a growing literature on the drivers of under- and

(1997) but where subjects can see previous decision-makers’ signals (a or b) as well as their choice of urns (A or B). Using their data, we find (analyses not shown) that, even conditional on the total number of a and b signals available to the agent (her signal and all previous players’ signals), she is much more likely to choose urn A if “her” signal was a than if “her” signal was b . These results, while suggestive of and consistent with our finding of intrinsic under-weighting of “others” compared to “own” information, cannot rule out that subjects put more weight on their own information because they receive it last, for example due to base-rate neglect. Though Drehmann et al. (2005) note that people in this treatment choose the urn associated with their own signal more than theory would predict, their paper primarily focuses on the effect (in other experimental treatments) of adding asset prices to the Anderson and Holt (1997) paradigm.

³Hartzmark et al. (2021) ask experimental participants to update their beliefs about experimental assets’ fundamentals after observing their prices change. Incentives for accurate beliefs are identical across owned and non-owned assets, and there are no buy-or-sell decisions to be made. Thus, by design, there is no instrumental reason to pay extra attention to owned assets. Yet, Hartzmark et al. (2021) find strong evidence of greater responsiveness to information about owned assets.

overreaction to information. Recent work mostly focuses on people failing to learn effectively on their own, including neglecting the strength of signals (Augenblick et al., 2024), being influenced by salient features (Bordalo et al., 2025), complexity and cognitive uncertainty (Enke and Graeber, 2023), selective attention (Ba et al., 2025; Schwartzstein, 2014), and the role of associative memory (Enke et al., 2024). We focus on inefficiencies in learning that may arise when some of the signals reach people from others.

Insufficient learning from others could arise in many real-world situations. For example, business analysts can either learn information from co-workers or do the research themselves, educators can impart facts passively or invite students to discover them by trial and error, farmers can use new agricultural technologies themselves or learn from their neighbors about them (Kondylis et al., 2024), and medical experts may supplement their own judgment with information from AI tools (Agarwal et al., 2024). Our results suggest that people will be systematically less sensitive to information originating from others, even when that information is perfectly communicated.

The remainder of this paper is organized as follows. Section 2 presents the broad aspects of the design shared by the different experiments. Section 3 presents the empirical framework. Sections 4, 5, and 6 present the detailed designs and results of the three experiments. Section 7 concludes.

2 Overview of Design

In all three experiments, participants play multiple rounds of the same basic statistical learning exercise: a balls-and-urns task based on a large literature studying individual learning (Benjamin, 2019). Here, we describe the task and features of the design common to all experiments. We defer the discussion of treatment variations and details specific to each experiment to the corresponding sections below. Full experimental scripts are shown in Online Appendix C.

The goal in the experimental task is to guess the number of red balls in an urn containing 20 balls. Participants are informed that the number of red balls is drawn uniformly from 4 to 16 in each round, as explained with the help of the illustration in Appendix Figure A.I(a) in the in-person experiments. In the online experiment, we explain that “the computer will randomly choose the exact number of red marbles [in the urn], where every number between 4 and 16 was equally likely to be chosen.”

In each round, participants receive two independent, noisy signals about the composition of the urn, by privately drawing a number of balls from the urn with replacement.⁴ The number of draws in each ‘signal’ is randomized—either 1, 5, or 9 draws—creating variation in how informed each participant is.⁵ We provided training in the task to participants in Experiments 1 and 2 before the first round. These participants individually played two unincentivized practice rounds with two guesses in each, and during these rounds received two ‘tips’ on making good guesses.⁶ The vast majority understood the tasks, as measured by excellent performance on comprehension checks (Table A.I).

Depending on the treatment condition, participants either play the game entirely on their own—the *Individual* treatment—drawing two sets of balls themselves, or else draw one set of balls themselves and have access to another set of balls that a partner (another participant) drew. The different treatments vary how the information obtained by one’s partner can be learned: via open-ended discussion, directly communicated by the experimenter, and/or observed with one’s own eyes. Guesses are made after making each set of draws (or potentially learning them via one’s partner). We test for frictions in social learning by comparing the sensitivity of guesses to draws across conditions.

Participants have incentives to make accurate guesses. The incentives provided were chosen to be easy for participants to understand: a penalty per ball away from the truth. Formally, each guess is incentivized by a piece-wise linear loss function.⁷ In Experiments 1 and 2, a perfectly accurate guess earns each member of the pair Rs. 105 and the payment decreases by Rs. 15 per ball the guess deviates from the truth. This incentive scheme was explained to participants in Experiments 1 and 2 using the illustration shown in Appendix Figure A.I(b). These incentives are sizable. Rs. 105 is about \$1.50 and Rs. 15 is about \$0.20, while average daily earnings in our Chennai sample are about Rs. 350 (\$5).

These incentives imply that participants’ guesses cannot be interpreted as the mean

⁴In Experiments 1 and 2, participants physically drew balls from an urn in our lab, while in Experiment 3 (the online experiment), the drawing was simulated using an animation of an urn. In each case, participants were informed that both partners were drawing from the same urn.

⁵For each pair, we randomly choose the number of draws in the two signals with uniform probability from $\{(1, 1), (1, 5), (5, 1), (5, 5), (1, 9), (9, 1)\}$. This excludes cases with more than 10 draws total.

⁶The first tip explains that it makes sense to guess there are more red than white balls if you draw more red than white, and vice-versa. The second tip is that “the more balls you draw, the more confident you can be in your guess”.

⁷On top of their participation fee, each person receives a payment equal to $\max\{(A - B \times |g - r|), 0\}$, where g is the guess, r the true number of red balls for the randomly-selected guess, and A and B are constants.

(or median or mode) of their beliefs about the color composition of the urn. We chose this payment rule to be transparent and easily comprehensible to participants, as none of our analyses require identification of any particular statistic of participants' beliefs (which a more complex proper scoring rule could theoretically elicit). Guesses should therefore be thought of as actions that participants have an incentive to tailor to the signals they receive. Our non-parametric and reduced-form results simply test whether these guesses are equally sensitive to one's own vs one's partner's signals. However, as a benchmark, we also compute what a risk-neutral Bayesian seeking to maximize expected payoffs would guess given the signals and our incentive structure. In addition, our structural model accounts for the incentive structure faced by participants.

Participants make multiple guesses throughout the experiment, and we randomly select one guess to score and pay participants for its accuracy. In Experiments 1 and 2, we select one guess among all the guesses that either partner made (including intermediate guesses). We then pay the two participants equally, irrespective of who made the guess, in separate envelopes at the end of the experiment. Thus, each person has an incentive to increase the accuracy of each guess from their pair. Neglecting to ask one's partner for information, withholding information from them, or ignoring their information reduces one's own expected payoff. In Experiment 3, the online experiment, participants never need to (and, indeed, cannot) communicate with each other, and information is shared by design. Each participant is rewarded for a randomly selected one of their own guesses: i.e., we do not split incentives between partners in Experiment 3.

3 Empirical Framework

Our goal is to test whether individuals' guesses are equally sensitive to signals drawn by themselves versus by others. We further examine how this depends upon the precise mode of social learning, such as whether the partner's information must be learned through a discussion, is communicated by a third party (the experimenter), and/or is directly observed. We present three types of empirical analyses—non-parametric, reduced form, and structural—to answer these questions. These three approaches impose different assumptions and have different strengths, but ultimately lead to similar conclusions.

3.1 Non-parametric Approach

In the non-parametric approach, we use minimal assumptions and simply plot average guesses in each treatment against the signals drawn. For simplicity, we summarize each signal by the net number of red draws (i.e., the number of red minus the number of white draws). That is, if a participant saw 4 red draws and 1 white draw, we would classify the signal as being 3 net red draws.⁸ To enable a transparent comparison of the sensitivity of guesses to own versus others' signals, we plot the guesses separately against the signals drawn by oneself versus those drawn by one's partner.

3.2 Reduced-form Approach

In our second empirical approach, we impose a linear relationship between signals and the resulting guesses and test for differences in this relationship across treatments. The starting point for our analysis is estimation of the following equation by OLS, separately by treatment:

$$Guess_i = \alpha + \beta^o \cdot Own\ Info_i + \beta^p \cdot Partner's\ Info_i + \epsilon_i, \quad (1)$$

where $Guess_i$ is i 's guess of the number of red balls (after having a chance to learn both signals), and $Own\ Info_i$ and $Partner's\ Info_i$ are the net number of red draws (i.e., red minus white draws) drawn oneself and by one's partner respectively. β^o and β^p capture the sensitivity of participants' guesses to signals drawn themselves and by others. If participants learn their partner's signals and treat them the same as their own signals, then we should expect that $\beta^o = \beta^p$. If instead $\beta^p < \beta^o$, participants in that treatment are less sensitive to their partner's draws than to their own.

When estimating equation (1), we add controls for the order in which participants complete the different treatments. In all three experiments, participants play multiple rounds of the game, doing the different treatments in randomized order. Although they receive no feedback after each round, and thus the scope for learning is limited, we control for treatment order effects by including dummies for round number interacted

⁸This simplification loses some information, e.g., it does not capture the total number of draws. A signal with 1 net red could come from a single draw of a red ball or from 9 draws with 5 red and 4 blue. A Bayesian should react differently to these two signals. The structural model does not share this weakness.

with *Own Info_i* and *Partner’s Info_i*.⁹

We also control for the possibility that differences between β^o and β^p may reflect ‘information order’ effects: i.e., the order in which one receives information may affect the weight placed on it. For a Bayesian, the order of receiving information should not matter. Nonetheless, ex ante it is possible that participants put more weight on signals they saw first (‘first impressions matter’) or on signals they saw last (‘recency effects’). In Experiment 3, the order of learning one’s own and one’s partner’s signals is randomized with equal probabilities, so the comparison between β^o and β^p is unbiased by information order effects. In Experiment 1, instead, participants learn their partner’s signals only after they have received their own signals. Therefore, we compare β^p in different treatments with the coefficient on one’s own *second* set of draws in the *Individual* treatment, estimated as β_2^o in the following regression:

$$Guess_i = \alpha + \beta_1^o \cdot Own\ First\ Info_i + \beta_2^o \cdot Own\ Second\ Info_i + \epsilon_i \quad (2)$$

In practice, we find that participants tend to put more weight on the signals they receive *second*, so treatments that provide partners’ information last would tend to bias us *against* finding under-sensitivity to others’ information.¹⁰

3.3 Structural Approach

In our third empirical approach, we estimate a simple model of quasi-Bayesian updating (Grether, 1980, 1992). This approach has several strengths relative to the reduced-form analysis. First, it exploits the full information content of the signals, including the number of draws, rather than the simplified ‘net red draws’ employed in the reduced form. Second, it accounts for the incentive structure faced by participants, modeling them as risk-neutral agents trying to maximize expected payoffs given their beliefs. Third, by taking the form of a standard learning model, it allows us to estimate interpretable weights placed on one’s own and others’ signals, with a clear Bayesian benchmark. Finally, it also accounts for noisy choice together with censoring in guesses at 4 and 16, which might otherwise cause guesses to appear less sensitive than those of a risk-neutral

⁹To include these controls, we stack the regressions for all treatment conditions in a given experiment and estimate them jointly in one regression, allowing the coefficients α , β^o and β^p to vary by treatment.

¹⁰Experiment 2 has aspects of the design of both Experiment 1 and Experiment 3. Some comparisons involve a randomized order of receiving information, as in Experiment 3. Others are similar to Experiment 1 in that one’s partner’s information is received after one’s own information.

Bayesian. On the other hand, the structural model makes more assumptions than the non-parametric and reduced-form analysis, including imposing risk-neutrality.

Let d_1 be the participant’s own signal and let d_2 be her partner’s signal, e.g., d_1 might equal {Red, Red, White, Red, White} and d_2 might equal {Red}. We then assume that the participant updates her beliefs about the state of the world s (the number of red balls in the urn) according to a modified version of Bayes’ Rule:

$$Posterior(s|d_1, d_2) \propto Prior(s) * P(d_1|s)^{\omega_{1rt}} * P(d_2|s)^{\omega_{2rt}} \quad (3)$$

where $Prior(s)$ is the participant’s prior about the probability of state s , and $P(d_i|s)$ is the (objective) probability of observing a set of draws d_i conditional on state s . Recall that participants are told each state is equally likely, and there are 13 possible states $s \in \{4, 5, \dots, 16\}$, so $Prior(s) = \frac{1}{13}$. Next, ω_{1rt} and ω_{2rt} are the weights that the participant puts, respectively, on her own and her partner’s signals in treatment t when that round occurs in chronological order r . For $\omega_{1rt} = \omega_{2rt} = 1$, Equation (3) reduces to Bayes’ Rule.

We allow ω_{1rt} and ω_{2rt} to differ from the Bayesian benchmark depending on both the treatment condition and the chronological order of the round. In particular, we assume the following functional form to mirror the reduced-form analysis described above:

$$\begin{aligned} \omega_{1rt} &= \beta_{1t} + \mu_{1r} \\ \omega_{2rt} &= \beta_{2t} + \mu_{2r} \end{aligned}$$

where β_{1t} and β_{2t} are, respectively, the weight the participant puts on her own and her partner’s signal (or, in Experiment 1’s *Individual* treatment, on her own second signal) in treatment t , and μ_{1r} and μ_{2r} are the additional weight she puts on each signal when that treatment occurs in chronological order r .

Just as with the reduced-form analysis, we control for the order in which information arrives in one of two ways. In Experiment 1, the partner’s information is always conveyed second, so we use a control condition (the *Individual* treatment) where both signals are drawn by the participant herself. We then compare ω_{2rt} , the weight on the second set of draws across treatments to identify the effect of drawing information oneself net of any information order effects. In Experiment 3, we randomize whether participants’ own information or their partner’s information comes first, so ω_{1rt} and ω_{2rt} will not be biased by differential treatment of earlier or later signals.

In addition to systematically biased updating, we allow for noisy choice. Doing so allows us to account for heterogeneity in guesses conditional on signals (i.e., not everyone with the same signals makes the same guess). We assume that agents are risk-neutral but calculate the expected payoff of each possible guess with noise. In particular, let $Earnings(g, s)$ be the earnings that a participant would earn if they made guess g and the true state was s . Given the experimental incentives, $Earnings(g, s) = \max\{0, 105 - 15 * |g - s|\}$. We assume that the agent calculates the expected payoff of each guess g using the (potentially biased) updating rule given by Equation 3 plus a random additive error term. That is, we assume the perceived expected payoff from making guess g given draws d_1 and d_2 is given by

$$EP(g|d_1, d_2) = \sum_{s=4}^{16} Posterior(s|d_1, d_2)Earnings(g, s) + \alpha\epsilon_{i,g}. \quad (4)$$

The agent then chooses the guess that maximizes this perceived expected payoff. For simplicity, we assume $\epsilon_{i,g}$ is i.i.d. Type 1 extreme value. The parameter α then governs the extent of noisy choice (Goeree et al., 2007). We estimate the model by maximum likelihood.¹¹

4 Experiment 1: Establishing the Main Results

4.1 Recruitment and Sample

Experiment 1 was conducted in person at the Behavioral Development Lab in Chennai, India, between July and December 2019. Participants were recruited on a rolling basis, with about 4 to 10 individuals completing the experiment on a given day. We recruited individuals—not pairs—residing in low- to middle-income neighborhoods within a reasonable travel time of the lab. Surveyors went door-to-door to advertise an academic

¹¹In particular, given the assumptions above, the probability that an agent with signals d_1 and d_2 will choose guess g is $P(i \text{ guesses } g|d_1, d_2) \propto \exp\left(\frac{1}{\alpha} \left[\sum_{s=4}^{16} Posterior(s|d_1, d_2)Earnings(g, s) \right]\right)$. We then choose parameters that maximize the joint likelihood of observing all the choices in our data. We calculate standard errors by bootstrapping the data, drawing pairs with replacement from the data. Throughout, we report bootstrapped standard errors for legibility but denote significance using bootstrapped confidence intervals (e.g., an estimate is significant at the 5% level if the center 95% of bootstrapped estimates do not include zero).

study on ‘your choices and how you aggregate information’ which would ‘help us understand how you make decisions’. No more specific study details were provided at this stage. Potential participants were informed that they would spend 2 to 3 hours at the study office and could expect to earn Rs. 150 to 280 (\$2 to \$3.90) per person, plus a payment of Rs. 100 (\$1.40) to cover travel expenses. Recruitment stopped when we reached our pre-specified target of 500 individuals. Participants were randomly assigned to pairs within an experimental session.¹²

Column 1 of Table 1 reports demographic characteristics of our sample. 50% of the participants are male. Participants are on average 35 years old and have a bit less than 8 years of education. Participants answered about 80 percent of comprehension questions correctly on the first attempt, indicating fairly high levels of attention and comprehension for a task that was unusual and somewhat complex given the local context.

4.2 Experiment 1: Design

Participants play five rounds of the task, as illustrated in Figure 1, with no feedback between rounds. Participants first play, in randomized order, an *Individual* round and a *Discussion* round. In each round, participants have access to two sets of draws with 1, 5, or 9 draws each.

***Individual* round.** In the *Individual* round, the participant first draws a set of balls from the urn with replacement, then guesses how many red balls are in the urn. Then, they draw a second set of balls from the urn and make a second (and final) guess. All drawing and guessing is done privately, without any need to share information. This round serves as a control condition—a benchmark against which we compare the other treatments.

***Discussion* round.** The *Discussion* round models a common mode of social learning, where we learn from others’ experiences through direct communication with them. Instead of drawing two sets of draws oneself as in the *Individual* round, each participant’s *partner’s* draws—accessible through a discussion—serve as their second set of

¹²Each participant plays four of their five rounds with one randomly-assigned partner of a different gender, and one round with a randomly-assigned partner of the same gender. Participants were introduced to their partner at the start of each round. This variation was induced to study the effect of gender composition on learning and to contrast these findings with a study of learning between spouses. These results are reported in a companion paper (Conlon et al., 2024). Here, we pool results across gender and both types of pairs.

draws. Each person first makes one set of draws followed by a private guess, exactly as in the *Individual* round. Next, the pair are asked to hold a face-to-face discussion and enter a joint guess.¹³ After their discussion, the teammates are separated and each person makes one final, private guess.

Participants can take as long as they like for the unstructured, face-to-face discussion with their partner. While we do not require communication to be truthful, participants have an incentive to share information accurately since one guess per team is randomly chosen to be paid for accuracy at the end of the experiment, with the payment split between the two partners. Participants also have an incentive to help their partner deliberate and make better guesses conditional on information, as in Cooper and Kagel (2005). We record the audio of the discussion (with participants' consent) and present summary statistics of the transcripts in Table A.II. Discussions last an average of 48 seconds and the vast majority (83%) involve some information sharing, with participants about as likely to share their signals as their guesses.

Comparing each participant's final guesses in the *Individual* and *Discussion* rounds reveals whether they learn as much through a discussion with a partner as from information they uncovered themselves. By design, participants have access to the exact same number of draws to inform their final guess in these two rounds, provided they share information. If participants are instead less sensitive to information collected by their partner, this implies either a failure of communication or a failure to aggregate information provided by one's partner.

Participants next play three more rounds, in randomized order, consisting of a second *Discussion* round and two additional treatments in which the experimenter informs the participant of their partner's draws or guesses.

***Informed of Partner's Draws* round.** This round (which we abbreviate as the '*Informed*' round) is designed to shut down any communication frictions between the partners. It is identical to the *Discussion* round except that after participants receive their first set of draws and enter their first guess, they are told their partner's draws (both number and composition) directly by the experimenter, e.g., "Your partner had

¹³The joint guess was included as a comparison to joint guesses made by teams composed of married couples and is not the focus of this paper. Note that having to enter a joint guess might cause teammates to come closer to agreement about the optimal guess, which might be expected to reduce under-sensitivity to each others' information. Experiment 3 and most treatments in Experiment 2 do not include such a joint guess.

five draws, of which three were red and two were white.” Participants then make an additional private guess before moving on to the discussion and their final private guess.

Comparing the guess made after the experimenter informs the participant of their partner’s signal (but *before* discussion) with the second guess in the *Individual* round allows us to directly test whether participants use information they gathered themselves in the same way as information collected by others but perfectly shared with them by a third party. In each case, there is no possibility of joint deliberation.¹⁴ Finally, the *post*-discussion guesses in the *Informed* round reveal how participants use information shared by a third party when additionally given the chance to confirm the information with its source (their partner) and deliberate.

***Informed of Partner’s Guess* round.** This round is the same as the *Informed of Partner’s Draws* round except that the experimenter informs each person of their partner’s private guess (made based on their own draws only), rather than their partner’s draws. The experimenter also shares the number of draws this guess was based on, e.g., “Your partner had 5 draws and, after seeing these draws, they guessed that the urn contains 12 red balls.” Thus, while in the *Informed of Partner’s Draws* round we transmit the signal received by one’s partner, in the *Informed of Partner’s Guess* round we transmit the action (guess) taken based on that signal as well as a measure of the precision of the signal. This round parallels more closely the literature that investigates social learning based on observing others’ actions (Weizsäcker, 2010). In this treatment, less information is transmitted to the participant. Moreover, beliefs about others’ competence might affect how these actions are interpreted and how much is learned about the signals.

4.3 Experiment 1: Results

4.3.1 Non-parametric results

Before comparing participants’ guesses across treatment conditions, we show evidence of good comprehension of the task. Panel A of Figure 2 examines participants’ first guesses, made after drawing the first set of balls by themselves. Reassuringly, the average number of red balls guessed increases in the number of “net red” draws uncovered oneself (pooling

¹⁴Note that this comparison requires controlling for round order effects, since the *Individual* round is always in the first two rounds, while the *Informed* round falls in rounds 3-5.

across all treatments), implying that participants respond sensibly to the information they receive. We can compare this sensitivity to a normative benchmark by computing, for each guess that participants make, what a risk-neutral Bayesian seeking to maximize expected payoffs would guess given the same signals and faced with our incentive structure. On average, participants' individual guesses (blue dots and lines) are fairly close to this benchmark (pink dashed lines), though somewhat less sensitive to signals than a risk-neutral Bayesian would be. Pooling across all data, the average participant guess increases by 0.61 percentage point for every 1 percentage point increase in what the risk-neutral Bayesian would guess.

Figure 3 contrasts the sensitivity of participants' guesses to their second set of draws in the *Discussion* and *Informed* rounds, comparing each to the *Individual* round. The blue curve representing the *Discussion* round (Panel A) is distinctly flatter than the grey curve representing the *Individual* round, revealing that participants' guesses are less sensitive to information gathered by their partner compared to information they collected themselves. This difference is statistically significant: we can reject (F -test, $p = 0.001$) that the differences in average guesses across treatments for each 'net red' value are all zero (i.e., that for each pair of dots in Figure 3 the true values lie on top of each other).

Strikingly, the curve is even flatter in the *Informed* round (Panel B), in which we plot participants' guesses after their partners' information is *directly* communicated to them by the experimenter (and before any joint deliberation with their partner). Despite having been given *all* decision-relevant information about their partner's draws directly, participants react to this information much less than they do to information they collected themselves. We can again reject that average guesses conditional on each 'net red' value are always equal across treatments (F -test, $p < 0.001$).

Panel C of Figure 3 shows a similar, though somewhat less pronounced, under-sensitivity to partner's information in the *Informed* round's post-discussion guesses compared to the pre-discussion guesses shown in Panel B. Since participants can confirm the information given to them by the experimenter directly with its source (their partner), this suggests that low sensitivity to others' information is not primarily driven by mistrust in the experimenter (though we cannot rule out some role of mistrust in the pre-discussion guesses of the *Informed* rounds).

4.3.2 Reduced-form and structural results

The reduced-form and structural models provide quantitative estimates of sensitivity to own and others' information. Figure 4 plots participants' average sensitivity to the second set of signals, separately within each treatment, using reduced-form estimates of Equation (1). In their final private guesses in the *Discussion* round, participants are less than half as sensitive to their partner's signals (second bar) compared to the corresponding signals in the *Individual* round (first bar, $p < 0.01$).¹⁵ This implies they respond less to information their partner collected than to their 'own' information. Even more starkly, participants put close to zero weight on their partner's information in the *Informed* round, right after it is *directly* shared with them (third bar, $p < 0.01$). Adding a face-to-face discussion with their partner after being informed of their draws somewhat increases participants' sensitivity to their partner's signals. Still, it remains significantly below the sensitivity to their own signals (fourth bar, $p = 0.02$).

The corresponding regression estimates are presented in Table 2 Panel A (columns 1 to 4). Comparing the coefficient β_2 on the second set of information by treatment condition shows a clear result. Participants are 54 percent (0.28/0.52) less sensitive to information collected by their partner in the *Discussion* round relative to information they collected themselves in the *Individual* round ($p < 0.01$). Even more strikingly, they are 86 percent (0.45/0.52) less sensitive to their partner's draws in the pre-discussion *Informed* guess compared to in the *Individual* round ($p < 0.01$). Put differently, participants are seven times more sensitive to their own information than to their partner's, and we cannot reject that participants put no weight on their partner's information at all. The face-to-face discussion increases sensitivity to the partner's information relative to the pre-discussion guess, perhaps through joint deliberation regarding the right answer or increased trust in the information about the partner's signal.¹⁶

The structural estimates in Panel B of Table 2 mirror the reduced-form results.

¹⁵For a risk-neutral Bayesian, this sensitivity would be around 0.72: that is, for every additional net red draw, the expected-payoff maximizing guess on average increases by 0.72. Participants' sensitivity to signals in the *Individual* round is 0.52, so discounting of partners' signals moves participants further away from the risk-neutral rational benchmark.

¹⁶The pattern of results in the *Informed of Partner's Guess* round are similar or more extreme than those that we find in the *Discussion* and *Informed of Partner's Draws* rounds. In the reduced-form estimates, participants are 85% less sensitive to their partner's information in this treatment. Because lower sensitivity to others' information in this treatment can be explained by additional factors such as guesses containing less information than draws or players' beliefs about their partners' ability to make good guesses, we defer the analysis of the *Informed of Partner's Guess* round to Appendix B.1.

Column 1 shows that participants put close to the Bayesian weight ($\beta_1 = 0.92$ vs. the Bayesian benchmark of 1) on their own first signal in the *Individual* treatment, and somewhat greater weight ($\beta_2 = 1.50$) on their second signal in that round. In contrast, participants put much less weight on their partner’s information in the *Discussion* and *Informed* rounds. Most strikingly, participants put no weight at all on their partner’s signals in the (pre-discussion) *Informed* round. The weight on others’ information is somewhat higher in the rounds involving discussion but still 69% to 74% lower than in the *Individual* round.¹⁷

Earnings implications. The expected earnings from guesses are a direct measure of performance in the experiment. Table 3 estimates average expected earnings from guesses as a function of the number of draws in each set of signals. As expected, more draws in the second set of signals in the *Individual* round significantly increases earnings, by Rs. 3.25 per extra draw. However, participants earn only Rs. 0.86, 73% less ($p = 0.02$), for each extra draw their partner makes in the *Discussion* round and Rs. 0.27, 92% less ($p = 0.01$), for each draw their partner makes in the pre-discussion guess in the *Informed* round. In the post-discussion guess in the *Informed* round, additional draws by the partner also earn less than own draws in the *Individual* round (Rs. 1.87 vs Rs 3.25), though this difference is not statistically significant ($p = 0.25$).¹⁸

5 Experiment 2: Exploring Mechanisms and Confounds

Why do participants discount their partner’s information even when it is directly communicated to them? Experiment 2 is designed to isolate potential mechanisms, rule out potential confounds, and evaluate the robustness of our findings.

¹⁷The parameters of the quasi-Bayesian model have a different scale and interpretation than the reduced-form results discussed above. But frictionless social learning implies $\beta_2^o = \beta_2^p$ in both cases, where o and p refer to own and partner’s draws respectively. Appendix B.2 shows that the reduced-form and structural estimates are consistent with each other: data simulated using the structural model produces the same reduced-form results as the empirical data.

¹⁸Appendix Table A.III shows similar regressions but where the dependent variable is the absolute difference between participants’ guesses and the true number of red balls in the urn. Mirroring the results in Table 3, additional draws by the participant reduce this error on average, but this improvement is 70 to 100% smaller for draws that come from participants’ partners and in fact we cannot reject that such draws do not reduce the guessing error at all.

5.1 Recruitment and Sample

Experiment 2 was run at the Behavioral Development Lab in Chennai, India between February and March 2020, after observing the results of Experiment 1. We recruited new participants following a similar procedure as Experiment 1. Data collection ended in March 2020 due to the Covid-19 pandemic, with a sample size of 292 participants (out of an intended sample of 800).¹⁹ Compared to Experiment 1, participants have a similar average age (38 versus 35) and years of education (9 versus 8), but are less likely to be female (31% versus 50%), as reported in column 2 of Table 1.

5.2 Experiment 2: Design

Participants played six rounds corresponding to different treatment conditions, with no feedback between rounds. They first played a *Discussion* round, exactly as in our first experiment, to replicate our earlier findings and provide a baseline to compare other treatments with.²⁰ They then played five rounds in randomized order, consisting of an *Informed* round just as in Experiment 1 and four additional variations of the *Informed* round, described below and in Appendix Figure A.II, Panel A.

***Observe Partner’s Draws* round.** In this round (which we abbreviate to ‘*Observe*’), both participants are in the same booth, so they can each watch their partner drawing balls from the urn with their own eyes.²¹ After both participants have drawn their signals they are separated and each makes a private guess. There is no discussion between partners and no need for the experimenter to share draws. Nor is there any scope for distrust of the experimenter or partner. Both one’s own and one’s partner’s signals are perfectly observable and revealed draw-by-draw in randomized order across individuals. The *only* difference between the two sets of draws is who physically drew the balls from the urn. We designed this to be an extreme treatment, where we anticipated

¹⁹The pre-registered sample size was chosen to be powered to test for gender differences in treatment effects (which we explore in Conlon et al. 2024). Thus, even though the final sample size is smaller than intended, we remain tolerably well-powered to estimate the treatment effects described here. For example, the minimum detectable effect size for the *Informed* treatment is around 50% lower sensitivity to partner’s information relative to own information. This is close to the estimated effect in Experiment 1.

²⁰Participants always played the *Discussion* round first in order to avoid the possibility that our new treatments primed participants to behave differently in the *Discussion* round, and thus provide a cleaner test of whether the *Discussion* results replicated.

²¹We would like to thank Christine Exley for suggesting this treatment.

equal sensitivity to one’s own and others’ information. The remaining treatments are subtler and largely subsumed under this treatment.

***Draw-by-Draw* round.** In the *Informed* treatment, participants draw their own balls one at a time from the urn, while their partner’s information is communicated in summary form (‘2 red and 3 white balls’). Certain updating biases (e.g., base-rate neglect) could cause participants to respond differently to summary information than to learning information draw-by-draw. To test for this channel, the *Draw-by-Draw* round proceeds identically to the *Informed* round, except that the experimenter shares their partner’s draws with each participant one draw at a time, e.g., by saying, ‘Your partner first drew a red ball’, then after a brief pause, ‘Your partner then drew a white ball, ...’ and so on.

***Reverse-Order* round.** In this round, one participant learns their partner’s signal before making any draws themselves. They then make a guess, make their own draws, and make another private guess. Since this treatment is only possible for one person in each pair, we only include guesses from the treated person while analyzing this round.

***No-First-Guess* round.** This round was identical to the *Informed* round except that participants do not make a guess directly after making their own set of draws. We implemented this change to test whether, for example, people are more open to others’ information when they have not yet taken an action or stated a belief based on their own information.

***Higher-Stakes* treatment.** We increased the incentives for accurate guessing by 50% in a randomly-chosen 3 out of 6 rounds. The maximum amount each individual could earn from a guess and their loss in earnings per ball away from the truth were both increased by 50%, to Rs. 158 (\$2.25) and Rs. 22.5, respectively. Participants were informed about the stakes at the beginning of each round.

5.3 Experiment 2: Results

Figure 5 shows the results from Experiment 2. Since this experiment does not include an *Individual* round, we simply compare the sensitivity to own information (β^o) and the partner’s information (β^p) within round, estimating Equation (1) by OLS.²² The

²²We did not include an *Individual* round since the previous experiment established that, if anything, participants are more sensitive to their most recent signals compared to earlier signals. Thus, if the

corresponding regression coefficients are presented in Panel A of Table 4.²³

We first replicate the main finding from Experiment 1: Figure 5 shows that participants are 87% and 58% less sensitive to their partner’s information in the *Discussion* and *Informed* rounds, respectively, and we can reject $\beta^o = \beta^p$ with $p < 0.01$.²⁴ In addition, though not all differences are statistically significant, participants underweight their partner’s information in *every* other treatment.

Most strikingly, participants are less sensitive to their partner’s signals even in the *Observe* treatment, in which they see their partner drawing balls from the urn with their own eyes while sitting beside them. Participants are 41% less sensitive to their partner’s information than to their own in this treatment ($p = 0.04$). This result rules out a large set of confounds, including distrust of information communicated by others, perceived skill differences, order effects, and subtle differences in how information is communicated. Instead, it suggests that the act of producing information (i.e., physically drawing balls from the urn) or associating one piece of information with oneself as opposed to with one’s partner may be driving factors. We explore these mechanisms further in Experiment 3.

The remaining treatments provide further evidence of our core result that participants put less weight on their partner’s information than on their own. When participants learn their partner’s signals before drawing their own signals, in the *Reverse Order* treatment, they are still 54% less sensitive to their partner’s information ($p = 0.04$). While still sizable, the effects in the *Draw-by-Draw* and *No First Guess* treatments are somewhat less pronounced at 38% ($p = 0.18$) and 41% ($p = 0.12$), respectively. The latter two estimates are not statistically significant, perhaps due to the lower-than-intended sample size, but the difference in point estimates is roughly comparable across all six treatments, and we cannot reject that it is the same in all treatments ($p = 0.82$).

The under-sensitivity to others’ information is also not meaningfully affected by the size of the incentives for accurate guessing (Appendix Figure A.V and Table A.IV). In

partner’s information is learned last, this biases us *against* finding less sensitivity to others’ information.

²³Appendix Figure A.III, Panel A shows non-parametric estimates, plotting participants’ guesses first against their own signal and then against their partner’s signal in the *Discussion*, *Informed*, and *Observe* treatments. Like the reduced-form and structural results we discuss below, the non-parametric results indicate underweighting of others’ information: the slope of guesses against one’s own signal is steeper.

²⁴To conserve space, in all rounds except the *Discussion* round, we focus in the main text on only the pre-discussion guesses, after the participant is informed of their partner’s draws. Figure A.IV shows reduced-form results for the post-discussion guesses (except for the *Observe* round in which there was no discussion and thus no post-discussion guess), which look broadly similar to those for the pre-discussion guesses.

particular, we see no significant change in underweighting in rounds that were randomly assigned to have 50% higher stakes, and in fact relative underweighting is somewhat higher with high stakes (52% versus 42%).

The structural estimates (Table 4 Panel B) paint a similar picture, with participants putting significantly less weight on their partner’s information in every round. The weight participants put on their own information tends to be at or above the Bayesian benchmark of $\beta^p = 1$, while they tend to underweight their partner’s information (though, given the imprecision of these estimates, we typically cannot reject equality with the Bayesian benchmark). In particular, in the *Observe* round, participants place 60% less weight (0.91/1.51) on their partner’s information than their own ($p < 0.01$).

6 Experiment 3: External Validity and Mechanisms

Experiment 3 is a large-scale, between-subjects online experiment with three goals. First, it has a simpler, between-subjects design and an online format that other researchers can easily adopt. Second, we further investigate mechanisms, and particularly the importance of taking some action to generate one’s own signals rather than passively receiving information. Third, we test the external validity of our findings with higher-education participants from a different cultural context.

6.1 Recruitment and Sample

We recruited 4,489 participants from the US and UK on the online survey platform Prolific in February 2022, asking participants to complete a “short decision-making experiment” that involved a 15-minute survey. We required participants to have completed at least 50 previous surveys on Prolific with an approval rating above 95%. Participants were paid \$2.50 for completing the survey, plus up to \$2.80 as a bonus for accurate guessing. The resulting sample is similar in age and gender to our Experiment 1 and 2 samples (column 3 of Table 1). A key difference is that the sample is more highly educated, though participants’ task comprehension and performance are comparable across

the three experiments.²⁵

6.2 Experiment 3: Design

Participants recruited on Prolific were directed to a Qualtrics survey that embedded the experiment. Each participant was randomly matched to a partner.²⁶ The experiment had a purely between-subjects design, with participants randomized to one of the treatments—variants of the *Informed* condition—described below. Each participant played five identical rounds of the same treatment without feedback. We randomized across participants whether they drew their own signals first or instead first learned about their partner’s signals.²⁷

Informed of Partner’s Draws treatment ($N = 1,008$). This treatment sought to emulate the *Informed* round from Experiments 1 and 2 as closely as possible in an online format. Participants saw a virtual urn and clicked to draw balls from it one at a time. The drawing and replacement of the balls from the urn was animated. Participants were shown a summary of their partner’s draws, as in the previous *Informed* treatments (e.g., “Your partner got to draw 5 marbles out of **the same** jar. They drew 1 red marble and 4 blue marbles.”).

Observe Partner’s Draws treatment ($N = 1,497$). This treatment (which we abbreviate to ‘*Observe*’) differed from the *Informed* treatment in that participants saw their partner’s draws being revealed using the same ball-by-ball animations as their own draws. The goal was to make the mode of presentation of the two sets of draws as similar as possible. Comparing the *Observe* and *Informed* treatments isolates the role of the presentation of others’ information, including its visual presentation and whether the information is delivered in summary form or signal-by-signal.

Labels Only treatment ($N = 1,487$). This treatment was the same as the *Observe*

²⁵In Experiment 3, we included eight multiple-choice comprehension questions asking participants to explain aspects of the instructions. Participants had to answer each question correctly before they could proceed. The average participant answered 92% of these questions on the first attempt, and more than 80% did so for all eight questions. The results are unchanged if we include only those who answered all questions correctly.

²⁶Since the experiment did not include any direct communication between partners (as there was no *Discussion* round), it was not necessary for partners to be playing the game at the same time. Instead, we pre-generated the signals for each partner from the same ‘urn’.

²⁷The script and a link to the online experiment are provided in Appendix C. Appendix Figure A.II Panel B illustrates the design of the different treatments. Selected screenshots showing how draws were presented to participants are reproduced in Appendix Figure A.VI.

treatment, except that participants no longer had to click a button to generate each of their own draws. The only difference between one’s own and one’s partner’s signals was one word in the text that appeared below the animation (e.g., ‘Your first marble’ versus ‘Partner’s first marble’). If participants are less sensitive to their partner’s draws even in this minimal treatment, it implies that a subtle label is enough to generate a sense of ownership. In turn, comparing this treatment with the *Observe* treatment isolates the effect of taking an action to generate your own information. Taking action to generate information might be necessary to create a sense of ownership or to make that information more salient or vivid.

Non-Rivalry treatment. This treatment aimed to reduce any sense of competition with one’s partner. A randomly-selected half of the participants in the *Informed* treatment were truthfully informed that their partner would not be guessing the contents of the urn. Instead, the partner would only draw signals and be asked to remember them.

Higher-Stakes treatment. We randomized across participants the size of the incentives for accurate guessing. Half of those in each treatment earned a \$1.40 bonus minus \$0.20 cents times the absolute difference between their guess and the true number of red balls in the urn. For the other half of participants, the incentives were doubled.

Survey. After completing the five rounds of the experiment, participants completed a short survey. In the survey, without prior warning, we collected unincentivized measures of recall of their own and their partner’s draws from the last round as a measure of attention and memory. We also elicited participants’ perceptions of whether they used their own and their partner’s signals equally in informing their guesses.

6.3 Experiment 3: Results

To test whether participants are less sensitive to their partner’s information, we simply estimate equation (1) within each treatment condition and test $\beta^o = \beta^p$. Since the order of learning one’s own and one’s partner’s signals was randomized with equal probabilities, order effects do not confound this comparison. Figure 6 and Table 5 Panel A report the reduced-form results.²⁸ In the *Informed* treatment, we see that participants are 18% less

²⁸Appendix Figure A.III, Panel B, reports non-parametric results for Experiment 3, plotting participants’ guesses against their own signal and then against their partner’s. The pattern of results is similar to our reduced-form results, showing a greater responsiveness to own signal in the *Informed* and *Observe* treatments, but a more equal responsiveness in *Labels Only*.

sensitive to their partner’s information than their own ($p < 0.01$). We thus qualitatively replicate our previous findings from Experiments 1 and 2, despite the large differences in presentation, instructions, and samples.

The design of Experiment 3 also permits an even simpler test of sensitivity to own and others’ information. Since participants were randomized to receive their own or their partner’s signals first, we can examine the *first* guess they make—after seeing only the first set of draws—and test whether this guess was less sensitive to draws made by their partner. Appendix Table A.V column 1 reports these results. Once again, we find that participants’ guesses are 17% less sensitive to their partner’s signal than to their own ($p < 0.01$).²⁹ Overall, these results provide strong evidence of lower sensitivity to others’ information even with a very different sample and experimental format.

In the *Observe* treatment, participants continue to be significantly less sensitive to their partner’s information ($p < 0.01$). Indeed, the magnitudes are nearly identical to the *Informed* treatment (19% vs. 17%, $p=0.74$). Consistent with the findings of Experiment 2, this suggests that differences in the presentation of own and others’ information do not explain the lower sensitivity to others’ information. Appendix Table A.V column 2 shows that this also holds for the first guess (11% lower sensitivity to others’ information, $p < 0.01$).

By contrast, participants in the *Labels Only* treatment were only 4% less sensitive to their partner’s information, a difference that was not statistically significant in the reduced-form estimates ($p = 0.27$). The difference in sensitivity to own and partner’s information ($\beta^o - \beta^p$) is significantly lower in the *Labels Only* treatment than in the *Informed* ($p = 0.02$) and *Observe* ($p < 0.01$) treatments. We find a similar pattern in the first guess (Appendix Table A.V column 3). We interpret this result as showing that taking an action to gather information—which, plausibly, increases its salience—plays a role in producing greater sensitivity to it. Merely labeling information as ‘own’ versus ‘partner’s’ when the participant receives the information passively might not create this differential attention.

The structural estimates presented in Panel B of Table 5 again show clear evidence of underweighting of others’ information in the *Informed* and *Observe* treatments, by 31% and 33%, respectively (each $p < 0.01$). Again, the difference in weights is significantly

²⁹Appendix Table A.VI shows that participants’ first guess also underweights others’ signals relative to their own in the first round of the experiment, before participants have experienced any signals of the opposite type.

smaller in the *Labels Only* treatment compared to the *Observe* treatment ($p=0.04$), implying that taking an action to generate one’s own draws increases the weight on own relative to others’ information. However, in contrast to the reduced-form estimates, the structural estimates show significant underweighting of partners’ information even in the *Labels Only* round ($p<0.01$).³⁰

The results of Experiment 3 reinforce our findings from the *Observe* treatment in Experiment 2 that mechanisms such as visual presentation of the information, or distrust of the experimenter, do not drive underweighting. In Experiment 3, both one’s own and one’s partner’s draws are simply displayed on the computer screen, such that it is not clear why one would trust draws assigned to oneself more. The visual presentation of own and others’ draws is also identical in the *Observe* treatment in Experiment 3.

Our additional treatment variations in Experiment 3 enable us to rule out two further possibilities: first, that participants *consciously* underweight others’ information and second, that some form of competitiveness or rivalry drives our main result. Finally, we provide correlational evidence on the role of memory.

Stakes and awareness. As in Experiment 2, the relative sensitivity to own vs. others’ information is not affected by the size of the incentives that participants faced for accurate guesses (Figure A.V, Panel B). For all three treatments, the differential sensitivity to own and partner’s signals is very similar (and statistically indistinguishable) between the low- and high-stakes groups. This finding suggests that participants are either unaware that they are less sensitive to others’ information or that they mistakenly believe it is optimal to discount others’ information. Consistent with the former interpretation, 77% of participants reported in the debriefing survey that they treated both pieces of information the same. Yet these same participants were 15% less sensitive to their partner’s information in the *Informed* treatment, nearly identical to the result in the full sample (Table A.VII Panel C).³¹

³⁰The structural estimates account for the full information content of the signal. For example, a person should place more weight on seeing 9 draws (5 red and 4 white) than on seeing just 1 red draw, whereas the reduced-form analysis treats these as identical (1 net red draw). The structural model also accounts for noisy choices and the fact that guesses are constrained to be between 4 and 16. These differences may explain the discrepancy in the reduced-form and structural analysis of the *Labels Only* treatment.

³¹This also provides evidence against the possibility that participants felt experimenter demand to use their own information more. In all experiments, our instructions emphasized that participants’ goal was just to guess the number of red balls in the urn, and that their and their partner’s draws were coming from the same urn. Among participants who did not report treating both pieces of information equally, 14% reported using their own information more while 8% reported using their partner’s information more.

Competition. Despite the incentives to make accurate guesses, one concern could be that participants underweight their partner’s information out of a sense of competitiveness; e.g., they may enjoy ‘winning’ by making good guesses precisely when their partner guesses poorly. This might plausibly lead to a strategy of ignoring the partner’s draws. However, the *Non-Rivalry* treatment—a sub-treatment of *Informed* in which the partner does not make any guesses—does not increase sensitivity to the partner’s signals (Appendix Figure A.VII and Table A.VIII).

Memory. Our survey at the end of the experiment allows us to provide suggestive evidence on the role of memory (Appendix Figure A.VIII and Table A.IX). Participants were asked to recall their own and their partner’s draws in the final round. Recall of one’s own draws was slightly higher on average (60% vs. 55%, $p < 0.01$).³² However, even when restricting the sample to those who perfectly remember *both* sets of draws, participants still place (marginally) significantly less weight on other’s information, by 9% (Appendix Figure A.IX and Table A.VII Panel C), and we cannot reject that the difference in weights is the same in this subsample as in the overall sample ($p = 0.75$). This suggests that even participants who can recall others’ information (at least when specifically asked to) fail to use it in the same way as their own information. Of course, participants who remember others’ draws are a selected subsample. One limitation of our study is that we did not exogenously manipulate memory of own or others’ information, as would be necessary to quantitatively determine how much of the gap in weights it can explain. Such experiments are a promising direction for future work.

6.4 Experiment 3: Discussion of Mechanisms

What do we learn from Experiment 3 about the potential mechanisms underpinning asymmetric learning from own vs others’ information? We view several aspects of our results as broadly consistent with a bottom-up attention story whereby information gathered oneself simply looms larger in people’s later decisions. First, the asymmetry is

In open-ended responses, participants who reported using their own information more often explained that this was because they (randomly) received more draws than their partners. A 10 percentage point increase in the share of draws received oneself is associated with a 12 percentage point increase in the likelihood of reporting using one’s own information more ($p < 0.01$).

³²The difference in recall is significant in all treatments, although it is smaller in the *Labels Only* treatment. Specifically, recall of own vs. partner’s draws in the different treatments are 64% vs. 56% for the *Informed* treatment, 60% vs. 55% for *Observe*, and 58% vs. 55% for *Labels Only*. We can reject equal memory gaps in the *Labels Only* and the *Informed* treatments ($p=0.01$).

greatly reduced (or even eliminated) when participants need not take action to generate their own signals. Future work could fruitfully explore what types of actions are sufficient to generate these effects, especially in more naturalistic settings. This could also shed light on whether taking similar actions to uncover others' information would improve social learning.

Second, the relative under-sensitivity to others' signals in Experiment 3 is notably smaller than in Experiments 1 and 2, largely driven by the fact that participants put a higher weight on their partner's information in Experiment 3 than in Experiments 1 and 2.³³ One possibility is that this reflects a difference in samples,³⁴ but another is that the online interface reduces the difference between own and partner's information in a way that may limit the scope for own information to draw more attention. For example, in Experiment 3, participants do not see, learn any details about, or interact with their partner, and thus it might not truly feel like information coming from others. They uncover their own signals only by clicking a button rather than physically drawing from a urn, which might be a less engaging way to discover information. Future work could aim to understand these differences better. For instance, one could experimentally vary participants' knowledge of the 'other', or whether participants interact with their partner in person, via (video) chat, or not at all, or vary information about whether others have seen the same signal. Gathering independent data on what participants attend to or what is salient (Krajbich and Rangel 2011, Li and Camerer 2022, Bohren et al. 2024, Bordalo et al. 2025) would help to further understand the mechanisms involved.

7 Conclusion

This paper presents evidence of a powerful and potentially far-reaching barrier to social learning: people place more weight on information gathered themselves than on information gathered by others. Our results rule out failures of communication as the primary source of this asymmetry, as participants underweight others' signals even when

³³For ease of comparison, Figure A.X repeats the non-parametric results from the *Informed* round in Experiments 1, 2 and 3 side by side, with the average guesses of a risk-neutral Bayesian displayed as a benchmark.

³⁴For example, it could reflect a documented tendency for lower sensitivity to treatment conditions in online compared to in-person experiments (Gupta et al., 2021). Another potential explanation is that our sample in Experiment 3 is has higher education levels, though we do not find stronger effects among lower-education participants within any of the three experiments (Table A.VII).

they are frictionlessly communicated or even physically observed, and instead point toward failures of information *aggregation* as the primary mechanism: people simply focus more on information that they themselves gathered. This phenomenon appears across three experiments with very different study populations, cultural contexts, experimental formats, and treatment variations.

A limitation of our lab setting is that, while precise and closely controlled, it is fairly abstract and with at best moderate stakes (up to about half a day’s income in Experiments 1 and 2). An open question is to what extent similar findings will appear in field settings and with higher stakes. We speculate that the tendency we identify may underlie documented failures of social learning, whether in information cascade experiments (Weizsäcker, 2010), people reacting to information interventions (Haaland et al., 2023), farmers learning more from their own plots than from neighbors (Duflo et al., 2020; Chandrasekhar et al., 2022), decision-makers under-using AI-based recommendations (Agarwal et al., 2024), or central bankers being sensitive to their own personal economic experiences beyond aggregate data (Malmendier et al., 2021). But underweighting of others’ information could play a role in numerous other settings where social learning is possible.

We document this phenomenon in teams of strangers. In Conlon et al. (2024), we find that the marital context—learning from one’s spouse—appears to counteract the discounting of others’ information for women but not for men. Future work should study the underlying mechanisms behind these differences and, more generally, what types of social or work relationships and contexts shape how effectively people learn from each other. For example, do people learn better from friends or colleagues? How do social status hierarchies affect the weight placed on a person’s independent information? Finally, additional work is required to understand how sensitivity to information coming from others may be increased, where appropriate, particularly in naturalistic settings.

References

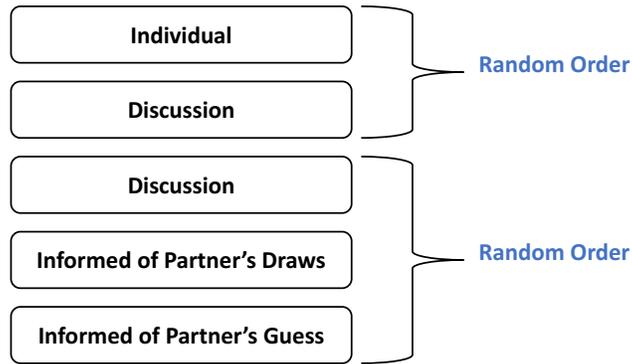
- Agarwal, Nikhil, Alex Moehring, Pranav Rajpurkar, and Tobias Salz**, “Combining Human Expertise with Artificial Intelligence: Experimental Evidence from Radiology,” *NBER Working Paper Number 31422*, 2024.
- Alatas, Vivi, Arun G Chandrasekhar, Markus Mobius, Benjamin A Olken, and Cindy Paladines**, “Designing Effective Celebrity Public Health Messaging: Results from A Nationwide Twitter Experiment in Indonesia,” 2021.
- Anderson, Lisa R and Charles A Holt**, “Information Cascades in the Laboratory,” *American Economic Review*, 1997, 87 (5), 847–862.
- Angrisani, Marco, Antonio Guarino, Philippe Jehiel, and Toru Kitagawa**, “Information redundancy neglect versus overconfidence: a social learning experiment,” *American Economic Journal: Microeconomics*, 2021, 13 (3), 163–197.
- Augenblick, Ned, Eben Lazarus, and Michael Thaler**, “Overinference from weak signals and underinference from strong signals,” *The Quarterly Journal of Economics*, 2024, p. qjae032.
- Ba, Cuimin, J Aislinn Bohren, and Alex Imas**, “Over-and Underreaction to Information: Belief Updating with Cognitive Constraints,” 2025.
- Benjamin, Dan, Aaron Bodoh-Creed, and Matthew Rabin**, “Base-Rate Neglect: Foundations and Implications,” *Working Paper*, 2019.
- Benjamin, Daniel J.**, “Chapter 2 - Errors in Probabilistic Reasoning and Judgment Biases,” *Handbook of Behavioral Economics: Applications and Foundations 1*, 2019, 2, 69–186.
- BenYishay, Ariel and A. Mushfiq Mobarak**, “Social learning and incentives for experimentation and communication,” *The Review of Economic Studies*, 2019, 86 (3), 976–1009.
- , **Maria Jones, Florence Kondylis, and Ahmed Mushfiq Mobarak**, “Gender gaps in technology diffusion,” *Journal of Development Economics*, 2020, 143.
- Bohren, J Aislinn, Josh Hascher, Alex Imas, Michael Ungeheuer, and Martin Weber**, “A cognitive foundation for perceiving uncertainty,” Technical Report, National Bureau of Economic Research 2024.
- Bordalo, Pedro, John Conlon, Nicola Gennaioli, Spencer Kwon, and Andrei Shleifer**, “How people use statistics,” *Review of Economic Studies*, 2025, p. rdaf022.
- Chandrasekhar, Arun G, Esther Duflo, Michael Kremer, João F Pugliese, Jonathan Robinson, and Frank Schilbach**, “Blue spoons: Sparking communication about appropriate technology use,” Technical Report, National Bureau of Economic Research 2022.

- Conlon, John J., Malavika Mani, Gautam Rao, Matthew Ridley, and Frank Schilbach**, “Learning in the Household,” *Working Paper*, 2024.
- Cooper, David J. and John H. Kagel**, “Are two heads better than one? Team versus individual play in signaling games,” *American Economic Review*, 2005, *95* (3), 477–509.
- D’Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber**, “Exposure to Grocery Prices and Inflation Expectations,” *Journal of Political Economy*, 2021, *129* (5), 1615–1639.
- De Filippis, Roberta, Antonio Guarino, Jehiel Philippe, and Toru Kitagawa**, “Updating Ambiguous Beliefs in a Social Learning Experiment,” *Working Paper*, 2017.
- Drehmann, Mathias, Jörg Oechssler, and Andreas Roider**, “Herding and Contrarian Behavior in Financial Markets: An Internet Experiment,” *American Economic Review*, 2005, *95* (5), 1403–1426.
- Dufo, Esther, Daniel Keniston, Tavneet Suri, and Céline Zipfel**, “Chat Over Coffee? Technology Diffusion Through Social and Geographic Networks in Rwanda,” *Working paper*, 2020.
- Enke, Benjamin and Thomas Graeber**, “Cognitive uncertainty,” *The Quarterly Journal of Economics*, 2023, *138* (4), 2021–2067.
- , **Frederik Schwerter, and Florian Zimmermann**, “Associative memory, beliefs and market interactions,” *Journal of Financial Economics*, 2024, *157*, 103853.
- Goeree, Jacob K., Thomas R. Palfrey, Brian W. Rogers, and Richard D. McKelvey**, “Self-Correcting Information Cascades,” *The Review of Economic Studies*, 2007, *74* (3), 733–762.
- Grether, David M.**, “Bayes rule as a descriptive model: The representativeness heuristic,” *The Quarterly journal of economics*, 1980, *95* (3), 537–557.
- , “Testing Bayes rule and the representativeness heuristic: Some experimental evidence,” *Journal of Economic Behavior & Organization*, 1992, *17* (1), 31–57.
- Guarino, Antonio and Philippe Jehiel**, “Social Learning with Coarse Inference,” *American Economic Journal: Microeconomics*, 2013, *5* (1), 147–74.
- Gupta, Neeraja, Luca Rigotti, and Alistair Wilson**, “The Experimenters’ Dilemma: Inferential Preferences over Populations,” *arXiv preprint arXiv:2107.05064*, 2021.
- Haaland, Ingar, Christopher Roth, and Johannes Wohlfart**, “Designing information provision experiments,” *Journal of economic literature*, 2023, *61* (1), 3–40.
- Hartzmark, Samuel M, Samuel Hirshman, and Alex Imas**, “Ownership, Learning, and Beliefs,” *The Quarterly Journal of Economics*, 2021, *136* (3), 1665–1717.

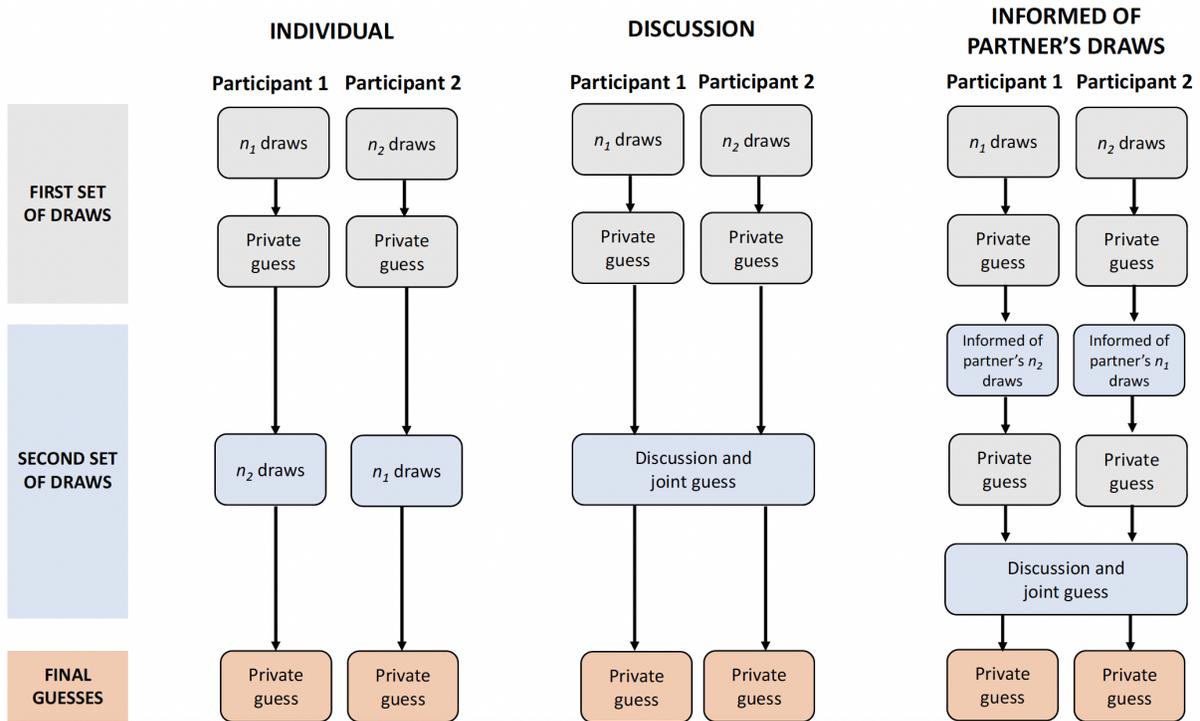
- Kondylis, Florence, John Loeser, Mushfiq Mobarak, Maria Jones, and Daniel Stein**, “Decentralizing Agricultural Demonstration to Accelerate Social Learning,” 2024.
- Krajbich, Ian and Antonio Rangel**, “Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions,” *Proceedings of the National Academy of Sciences*, 2011, *108* (33), 13852–13857.
- Li, Xiaomin and Colin F Camerer**, “Predictable effects of visual salience in experimental decisions and games,” *The Quarterly Journal of Economics*, 2022, *137* (3), 1849–1900.
- Malmendier, Ulrike and Leslie Sheng Shen**, “Scarred consumption,” *American Economic Journal: Macroeconomics*, 2024, *16* (1), 322–355.
- and **Stefan Nagel**, “Learning from inflation experiences,” *The Quarterly Journal of Economics*, 2016, *131* (1), 53–87.
- , – , and **Zhen Yan**, “The making of hawks and doves,” *Journal of Monetary Economics*, 2021, *117*, 19–42.
- March, Christoph and Anthony Ziegelmeyer**, “Altruistic Observational Learning,” *Journal of Economic Theory*, 2020, *190*, 105–123.
- Merlo, Antonio and Andrew Schotter**, “Learning by not doing: an experimental investigation of observational learning,” *Games and Economic Behavior*, 2003, *42* (1), 116–136.
- Miller, Joshua B and Zacharias Maniadis**, “The weight of personal experience: An experimental measurement,” Technical Report, Working paper 2012.
- Mobius, Markus and Tanya Rosenblat**, “Social Learning in Economics,” *Annual Review of Economics*, 2014, *6* (1), 827–847.
- Schwartzstein, Joshua**, “Selective attention and learning,” *Journal of the European Economic Association*, 2014, *12* (6), 1423–1452.
- Simonsohn, Uri, Niklas Karlsson, George Loewenstein, and Dan Ariely**, “The tree of experience in the forest of information: Overweighing experienced relative to observed information,” *Games and Economic Behavior*, 2008, *62* (1), 263–286.
- Weizsäcker, Georg**, “Do We Follow Others When We Should? A Simple Test of Rational Expectations,” *American Economic Review*, 2010, *100* (5), 2340–60.

Figure 1: Experimental Design

Panel A: Randomization of Rounds



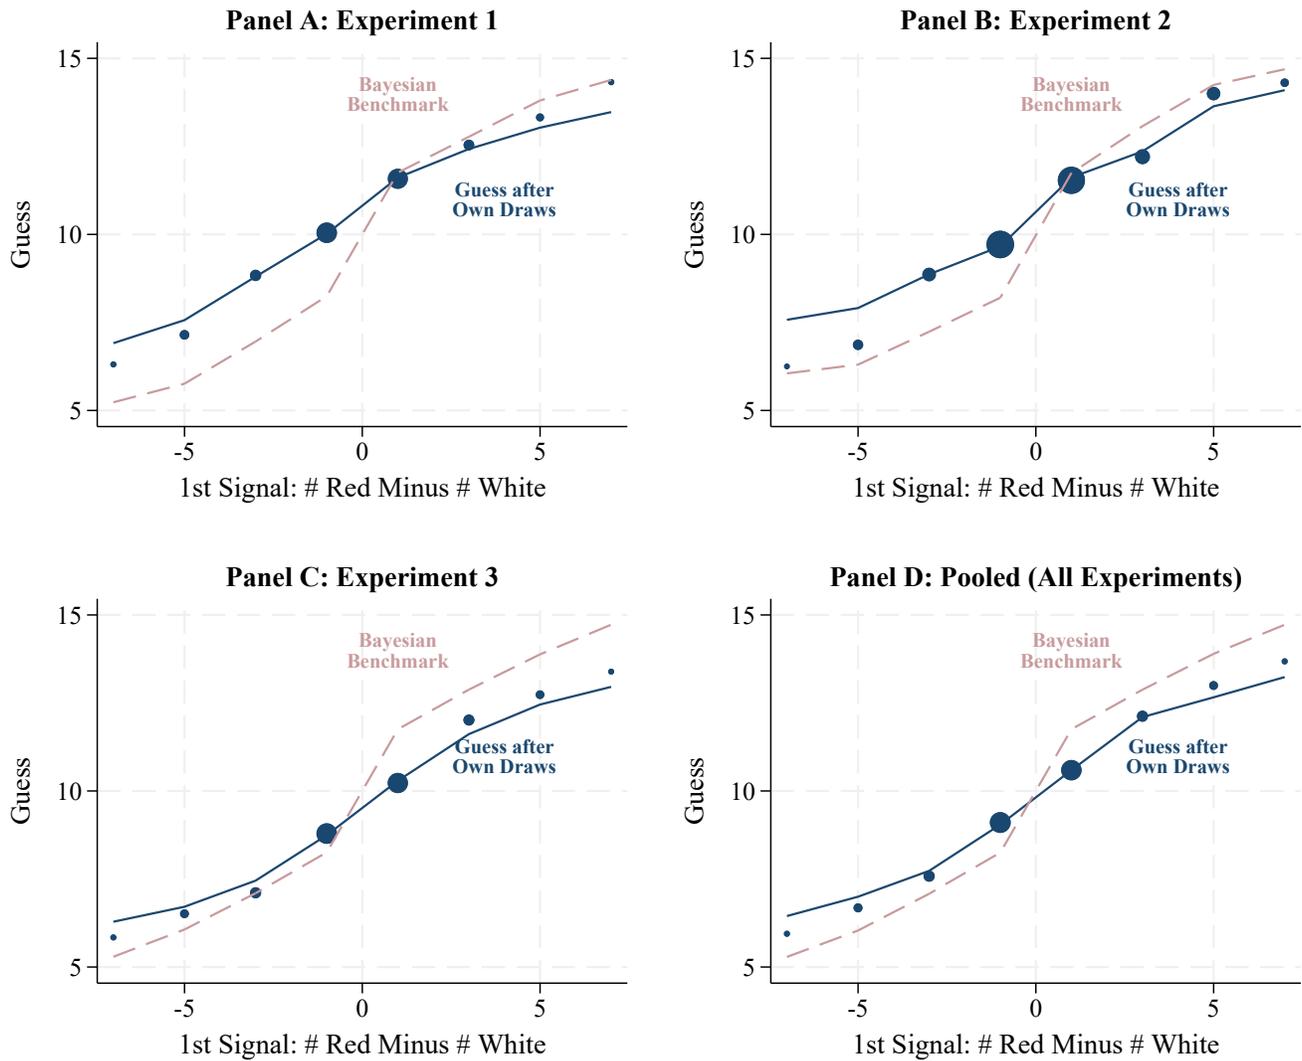
Panel B: Structure of Individual, Discussion, and Informed Rounds



Panel A shows the five rounds of Experiment 1. All participants get matched to a previously unknown partner and complete all five rounds with this partner (with the exception that in one randomly-selected *Discussion* round, participants were re-matched for that round only to generate variation in the relative gender of the partners. We do not exploit this variation in our paper). We randomized the order of the first two rounds (*Individual*, *Discussion*) and the order of the following three rounds: *Discussion*, *Informed of Partner's Guess*, and *Informed of Partner's Draws*.

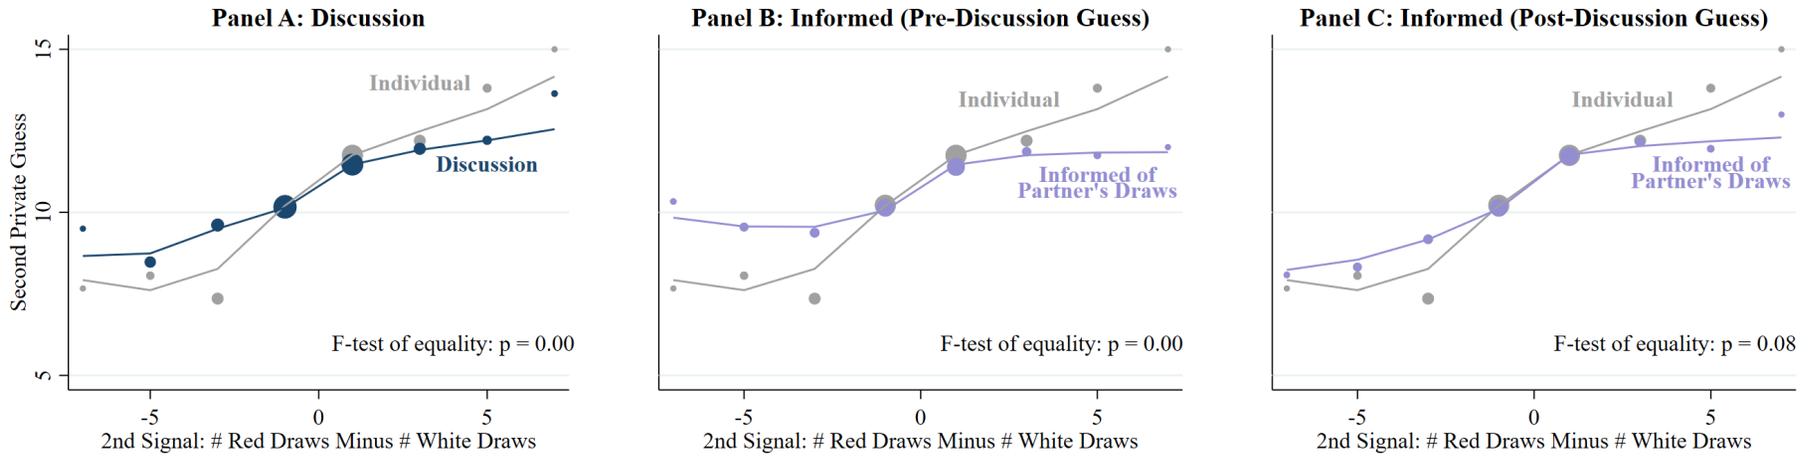
Panel B describes the structure of the different rounds. In the *Individual* round, each participant gets two sets of private draws from the urn and makes a private guess after each set of draws. In the *Discussion* round, each participant makes one set of draws followed by a private guess. The two participants are then asked to discuss and make a joint guess before each makes a final private guess. The *Informed of Partner's Draws* round is identical to the *Discussion* round, except that pre-discussion, each participant is informed about their partner's first set of draws and then asked to make a private guess. In the *Informed of Partner's Guess* round (Appendix B.1), each participant is instead informed pre-discussion about their partner's first private guess and then asked to make a private guess.

Figure 2: Individual Performance vs. Risk-Neutral Bayesian



Notes: This figure plots participants' first private guess against the net number of red draws (red draws minus white draws) in participants' own first (private) signal. We only include observations where participants saw their own signal first (in Experiment 1, this is all observations). The blue solid curve shows locally weighted means (lowess). The pink dotted lines show the average of a risk-neutral Bayesian's guesses given the same signals. Dot size indicates number of observations for each net number of red draws. Panels A through C show data from each of the three experiments separately. Panel D shows pooled data from all three experiments.

Figure 3: Experiment 1: Non-Parametric Results



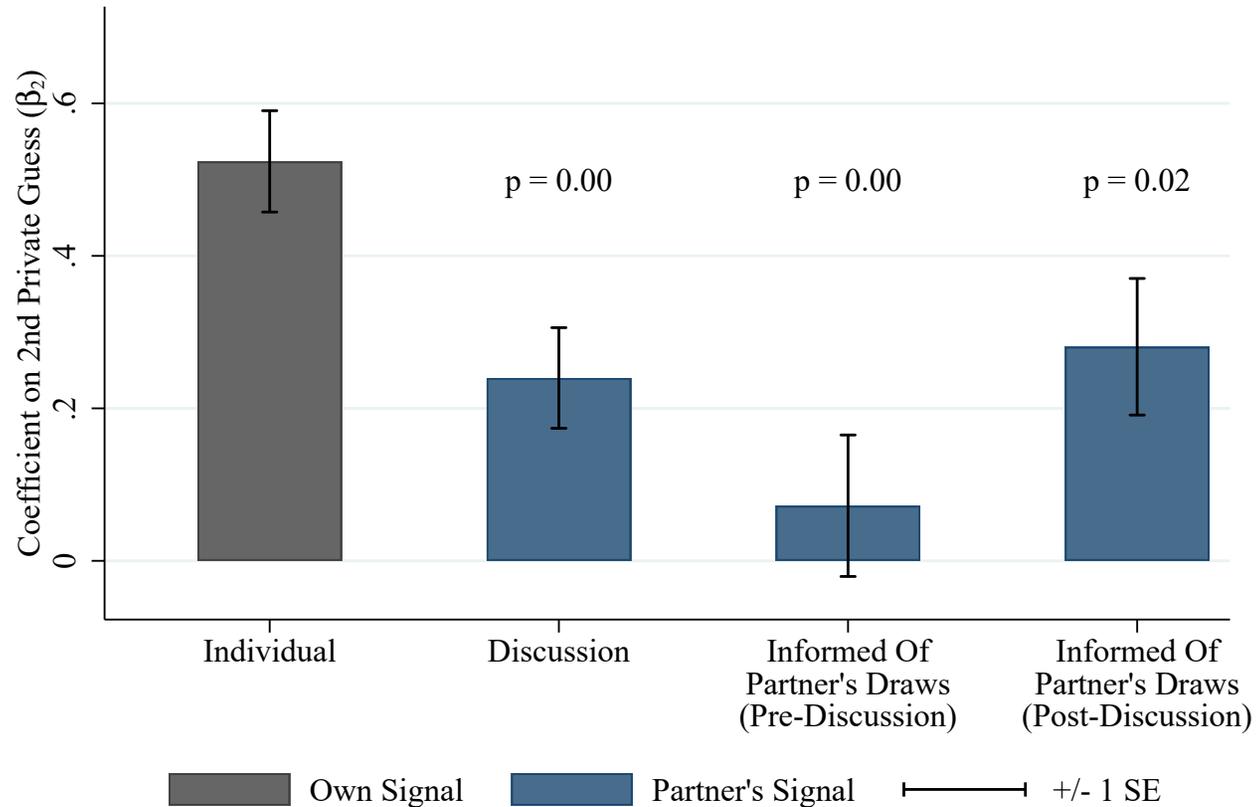
Notes: This figure shows average second private guess of participants in Experiment 1. The x-axis shows the net number of red draws (i.e., red draws minus white draws) in the second signal of the round. Dot size indicates number of observations for each net number of red draws. Lines show locally weighted means (lowess).

- In **Panel A**, the gray dots indicate average guesses in the *Individual* Round, where participants made the second set of draws themselves. The dark-blue dots in the graphs on the left show guesses in the *Discussion* Round, where the second set of draws had to be communicated to the participant via discussion.
- In **Panel B**, the lavender dots show average guesses in the *Informed of Partner's Draws* round, after the respondent is told of his/her partner's draws by the experimenter (but before the joint discussion).
- In **Panel C**, the lavender dots show average guesses in the *Informed of Partner's Draws* round after the joint discussion.

'F-test of equality' in the bottom right shows the p -value of a test of the joint hypothesis that the mean guess is equal across the two rounds at every value of net red draws.

Figure 4: Experiment 1: Reduced-Form Estimates

Weights on Own vs. Others' Signals

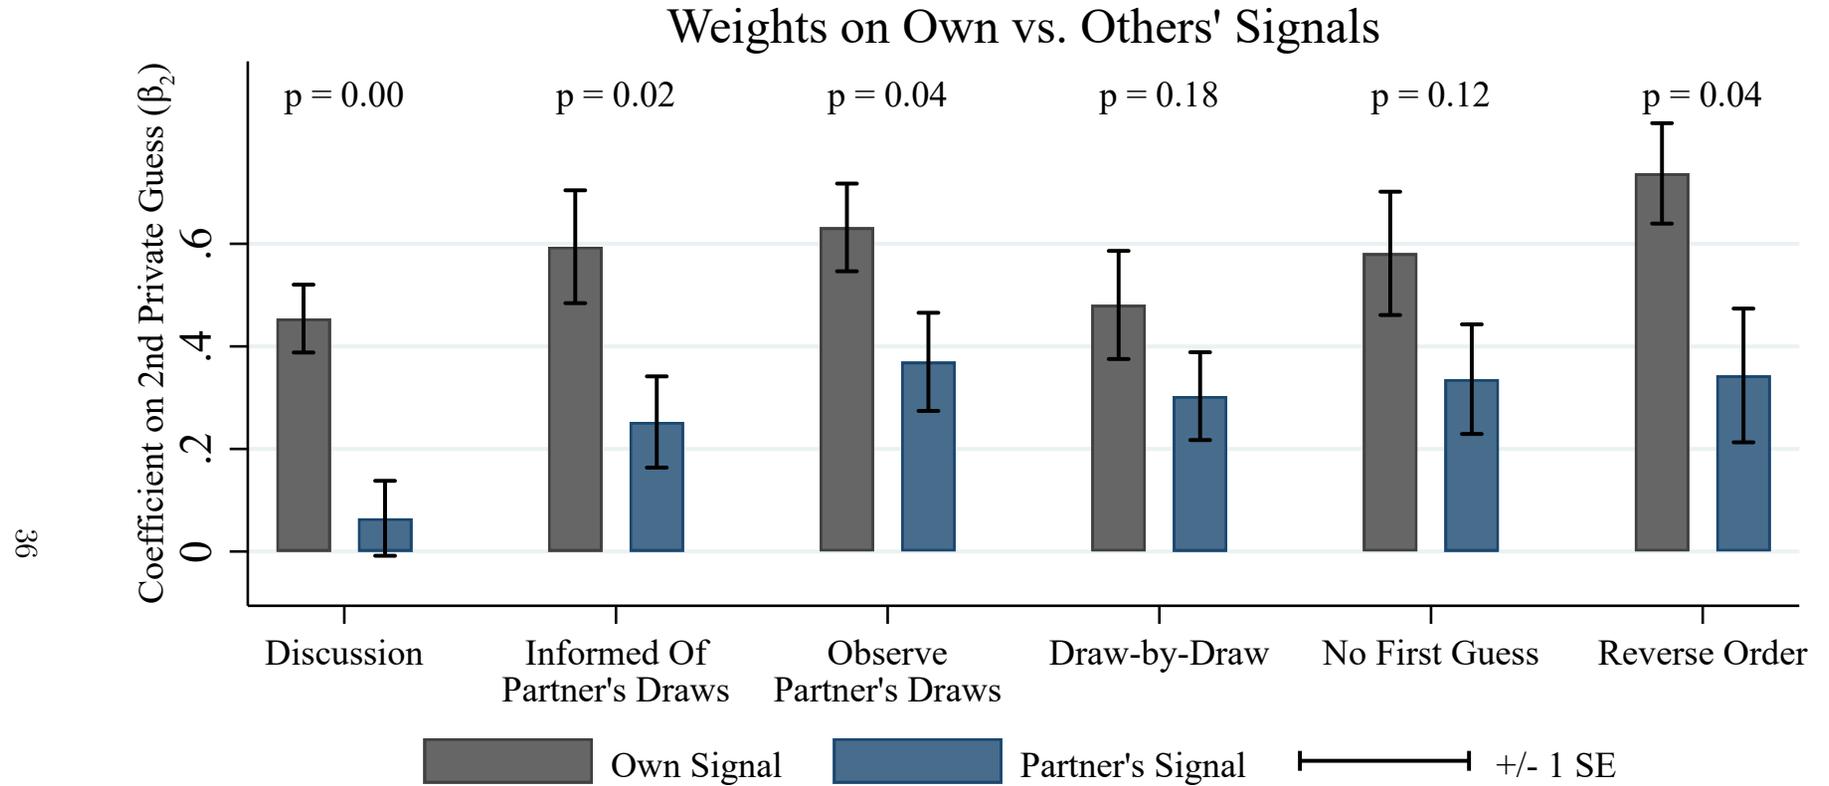


Notes: This figure shows the weights participants put on different signals in Experiment 1. We estimate Equation (1) and then display β_2 for each of the following four types of private guesses:

- (a) *Individual*, in which participants collect all information on their own. For this round, we report the coefficient on the net red draws in the participant's second set of draws, which replaces *Partner's Signal* in Equation (1);
- (b) *Discussion*, in which participants collect the first signal on their own and the second signal (their partner's) is only accessible via discussion;
- (c) *Informed of Partner's Draws (pre-discussion)*, where participants receive the second set of information directly from the experimenter but before any discussion with their partner;
- (d) *Informed of Partner's Draws (post-discussion)*, in which participants receive the second set of information directly *and* have the chance to discuss it with their partner.

For each of the dark-blue bars, we show the p -value of testing whether the weight in that round equals the corresponding weight in the *Individual* round (gray bar).

Figure 5: Experiment 2: Reduced-Form Estimates



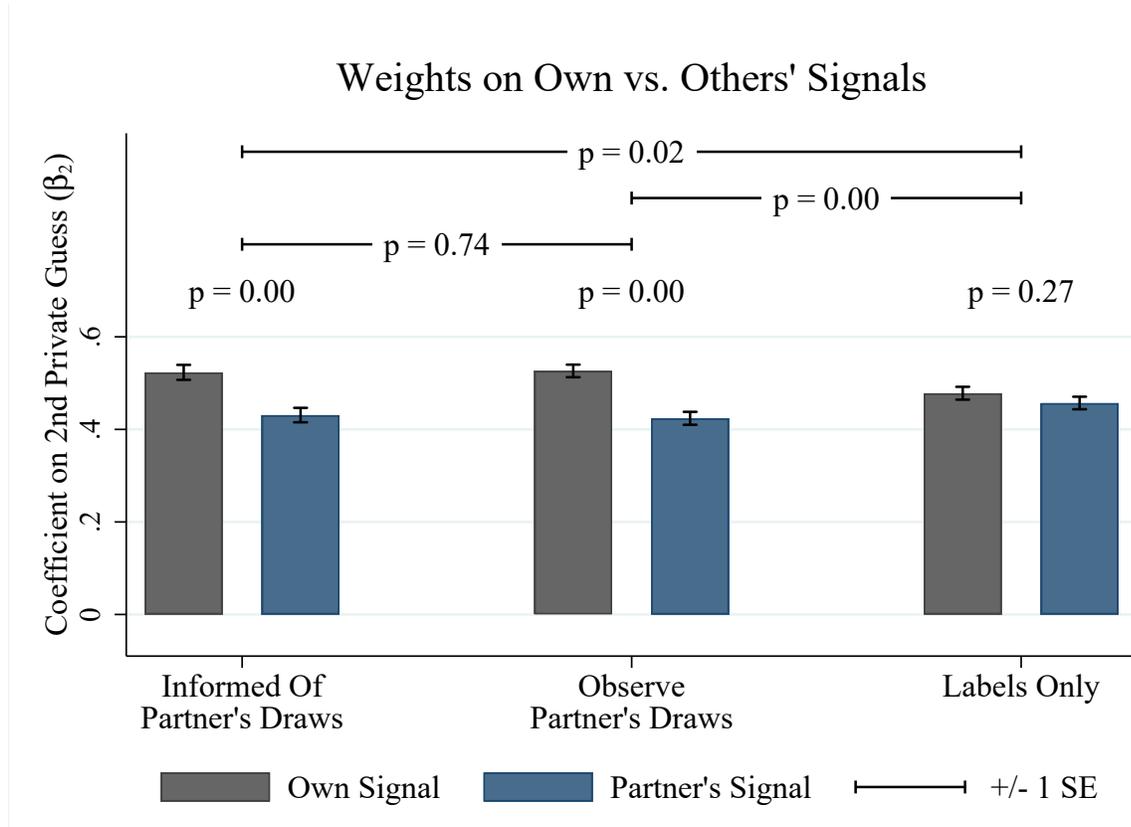
36

Notes: This figure shows the weights participants put on different signals in Experiment 2. We estimate Equation (1) and then display β_1 in gray and β_2 in dark blue for each treatment. The dependent variable is participants' pre-discussion guess, except in the *Discussion* round. In the *Discussion* round, it is the post-discussion guess as there was no pre-discussion guess. In addition to the *Discussion* and *Informed of Partner's Draws* rounds, we look at the following treatments:

- (a) *Observe Partner's Draws*, in which each participant directly observes their partner's draws (as well as making their own);
- (b) *Draw-by-Draw*, in which participants receive the second set of signals directly one draw at a time;
- (c) *No First Guess*, in which participants receive their partner's signals (and their own) before making their first and only private guess;
- (d) *Reverse Order*, in which one participant receives their partner's signals first and makes their first private guess, and then receives their own signals and makes their second private guess.

For each round, we show the p -value of testing whether the weight on their signal (β_1) equals the corresponding weight on their partner's signal (β_2) in that round.

Figure 6: Experiment 3: Reduced-Form Estimates



Notes: This figure shows the weights participants put on different signals in Experiment 3. We estimate Equation (1) and then display β_1 in gray and β_2 in dark blue for each treatment.

- The first set of bars shows the weights participants put on signals in the *Informed of Partner's Draws* treatment, in which participants clicked to draw their own balls one at a time and were told their partner's number of red and white draws.
- The second set of bars represents the *Observe Partner's Draws* treatment, in which participants clicked to draw their own balls one at a time and directly observed their partner's draws appearing from the urn one at a time.
- The third set of bars corresponds to the *Labels Only* treatment, in which participants did not take any actions and instead passively observed their own and their partner's labeled draws one by one in the exact same format.

For each round, we show the p -value of testing whether the weight on their signal (β_1) equals the corresponding weight on their partner's signal (β_2) in that round.

Table 1: Sample Characteristics

	Experiment 1	Experiment 2	Experiment 3
Female	0.50 (0.50)	0.31 (0.46)	0.57 (0.50)
Age	34.66 (8.58)	38.40 (7.31)	37.70 (13.87)
Years of Education	7.86 (3.94)	9.02 (3.49)	15.04 (2.03)
Expected Earnings (Relative to Bayesian)	0.82 (0.12)	0.84 (0.12)	0.89 (0.10)
Fraction of Comprehension Questions Correct	0.79 (0.14)	0.79 (0.13)	0.92 (0.13)
Observations	500	293	4489

Notes: This table shows averages of key background characteristics for individuals in each of our three experiments. Standard deviations are in brackets. “Expected Earnings (Relative to Bayesian)” is calculated as the expected payoff of the participant’s guess given the draws they observed, divided by the expected payoff that the Bayesian risk-neutral guess (i.e., expected payoff-maximizing guess) would make given those same draws. “Fraction of Comprehension Questions Correct” shows the proportion of participants who correctly answer questions about the task (summary of questions in Table A.I).

Table 2: Experiment 1: Reduced-Form and Structural Estimates

	Individual (1)	Discussion (2)	Informed (Pre) (3)	Informed (Post) (4)
<i>Panel A: Reduced-Form Estimates</i>				
β_1 : Own First Signal	0.43 (0.06)	0.56 (0.06)	0.56 (0.09)	0.36 (0.09)
β_2^o : Own Second Signal	0.52 (0.07)			
β_2^p : Partner's Signal		0.24 (0.07)	0.07 (0.09)	0.28 (0.09)
Constant	10.71 (0.16)	10.73 (0.16)	10.64 (0.23)	10.66 (0.23)
$\beta_2^p - \beta_2^o$		-0.28*** (0.08)	-0.45*** (0.11)	-0.24** (0.10)
<i>Panel B: Structural Estimates</i>				
β_1 : Own First Signal	0.92 (0.43)	0.87 (0.17)	1.02 (0.32)	0.57 (0.31)
β_2^o : Own Second Signal	1.50 (0.50)			
β_2^p : Partner's Signal		0.40 (0.14)	-0.01 (0.21)	0.46 (0.27)
$\beta_2^p - \beta_2^o$		-1.11*** (0.47)	-1.51*** (0.53)	-1.04*** (0.51)
N	500	1000	500	500

Notes: This table shows reduced-form and structural estimates for the weights on signals in Experiment 1. The dependent variable is participants' private guess. "Informed (Pre)" means the second private guess from the *Informed of Partner's Draws* round, after the participant was directly told their partner's signal but before the joint discussion. "Informed (Post)" means the third private guess, after the discussion. All standard errors are clustered at the pair (of two participants) level. Standard errors of the structural estimates are bootstrapped. *, **, and *** indicate significance at the $p < 0.10$, 0.05, and 0.01 levels of the difference $\beta_2^p - \beta_2^o$.

Reduced-form coefficients: Panel A shows reduced-form results, estimating Equation 1 by OLS. "Own First Signal" is the net number of red draws (i.e., red draws minus white draws) in the participant's first set of draws, which they drew themselves in all rounds. "Own Second Signal" is the net number of red draws in the participant's second set of draws in the individual round. "Partner's Signal" is the net number of red draws in the set of draws by the participant's partner, which was the second signal available to the participant in the *Discussion* and *Informed of Partner's Draws* rounds. All regressions include order fixed effects interacted with the participants' first and second signal.

Structural parameters: Panel B shows estimates of the structural model described in Section 3.3. "Own First Signal", "Own Second Signal" and "Partner's Signal" indicate the weights placed on the first set of signals, second set in the *Individual* round, and second (partner's) set in each other round in the agents' quasi-Bayesian updating rule.

Table 3: Experiment 1: Expected Earnings by Type of Guess and Number of Draws

	Individual (1)	Discussion (2)	Informed (Pre) (3)	Informed (Post) (4)
γ_1 : # Own First Draws	1.53 (0.83)	2.52 (0.56)	3.25 (0.76)	2.58 (0.78)
γ_2^o : # Own Second Draws	3.25 (0.86)			
γ_2^p : # Partner's Draws		0.86 (0.65)	0.27 (0.82)	1.87 (0.82)
Constant	102.51 (5.13)	103.36 (4.45)	97.94 (6.18)	97.12 (5.91)
$\gamma_2^p - \gamma_2^o$		-2.39** (1.04)	-2.98*** (1.13)	-1.38 (1.19)
N	500	1000	500	500

Notes: This table compares participants' expected earnings in the *Discussion* and *Informed of Partner's Draws* rounds to their earnings in the *Individual* round. The table shows OLS estimates of the following equation for the *Discussion* and *Informed of Partner's Draws* rounds:

$$Expected\ Earnings_i = \alpha + \gamma_1 \# Own\ First\ Draws_i + \gamma_2^p \# Partner's\ Draws_i + \epsilon_i \quad (5)$$

and OLS estimates of the following equation for the *Individual* round:

$$Expected\ Earnings_i = \alpha + \gamma_1 \# Own\ First\ Draws_i + \gamma_2^o \# Own\ Second\ Draws_i + \epsilon_i \quad (6)$$

where $Expected\ Earnings_{irt}$ is the expected earnings from i 's guess in the round in question, given the signals, and $\# Own\ First\ Draws_i$ indicates the number of draws in the first set of signals, drawn oneself. $\# Own\ Second\ Draws$ is the number of draws in the participant's second set in the *Individual* round and $\# Partner's\ Draws$ is the participant's partner's number of draws, in the *Discussion* and *Informed of Partner's Draws* rounds. In estimation, we stack the estimating equations for all treatments and estimate them jointly, including controls for round order fixed effects. Standard errors are clustered at the pair level. *, **, and *** indicate significance at the $p < 0.10$, 0.05, and 0.01 levels of the difference $\gamma_2^p - \gamma_2^o$.

Table 4: Experiment 2: Reduced-Form and Structural Estimates

	Discussion (1)	Informed (2)	Observe (3)	Draw-by-Draw (4)	No First Guess (5)	Reverse Order (6)
<i>Panel A. Reduced Form Estimates</i>						
β_1 : Own Signal	0.45 (0.07)	0.59 (0.11)	0.63 (0.09)	0.48 (0.11)	0.58 (0.12)	0.74 (0.10)
β_2^p : Partner's Signal	0.06 (0.07)	0.25 (0.09)	0.37 (0.10)	0.30 (0.09)	0.34 (0.11)	0.34 (0.13)
Constant	10.66 (0.19)	10.64 (0.22)	10.63 (0.22)	10.45 (0.21)	10.51 (0.24)	10.38 (0.26)
$\beta_2^p - \beta_1$	-0.39*** (0.09)	-0.34** (0.14)	-0.26** (0.13)	-0.18 (0.13)	-0.25 (0.16)	-0.39** (0.19)
<i>Panel B. Structural Estimates</i>						
β_1 : Own Signal	0.48 (0.09)	1.32 (0.33)	1.51 (0.43)	1.01 (0.32)	1.25 (0.36)	1.29 (0.47)
β_2^p : Partner's Signal	0.07 (0.07)	0.23 (0.17)	0.60 (0.27)	0.33 (0.16)	0.48 (0.25)	0.52 (0.41)
$\beta_2^p - \beta_1$	-0.41*** (0.09)	-1.09*** (0.34)	-0.91*** (0.30)	-0.68** (0.33)	-0.77** (0.32)	-0.77* (0.44)
hline N	288	292	292	292	292	146

Notes: This table shows reduced-form and structural estimates for rounds in Experiment 2 (our second lab experiment).

Reduced-form coefficients: Panel A shows reduced-form results, estimating Equation 1 by OLS. The dependent variable is participants' private guess. "Informed" refers to the *Informed of Partner's Draws* round and "Observe" to the *Observe Partner's Draws* round. "Own Signal" indicates the net number of red draws (i.e., red draws minus white draws) in the participant's own set of draws. Similarly, "Partner's Signal" indicates the net number of red draws in their partner's set of draws. In estimation, we stack the estimating equations for all treatment and estimate them jointly. The joint regression also includes fixed effects for the order in which participants played treatment conditions, interacted with "Own Signal" and "Partner's Signal." Standard errors are clustered at the pair level. *, **, and *** indicate significance at the $p < 0.10$, 0.05, and 0.01 levels of the difference $\beta_2^p - \beta_1$.

Structural parameters: Panel B shows estimates of the structural model described in Section 3.3. "Own Signal" and "Partner's Signal" indicate the weights placed on their own and their partner's set of draws in the agents' quasi-Bayesian updating rule. Bootstrapped standard errors (clustered at the pair level) in parentheses. *, **, and *** indicate significance at the $p < 0.10$, 0.05, and 0.01 levels of the difference $\beta_2^p - \beta_1$.

Table 5: Experiment 3: Reduced-Form and Structural Estimates

	Informed (1)	Observe (2)	Labels Only (3)
<i>Panel A. Reduced Form Estimates</i>			
β_1 : Own Signal	0.52 (0.02)	0.53 (0.01)	0.48 (0.01)
β_2^p : Partner's Signal	0.43 (0.02)	0.42 (0.01)	0.46 (0.01)
Constant	9.56 (0.04)	9.55 (0.03)	9.61 (0.03)
N	5040	7485	7435
$\beta_2^p - \beta_1$	-0.09*** (0.02)	-0.10*** (0.02)	-0.02 (0.02)
p -value: $\beta_2^p - \beta_1$ same as in Informed		0.74	0.02
p -value: $\beta_2^p - \beta_1$ same as in Observe			0.00
<i>Panel B. Structural Estimates</i>			
β_1 : Own Signal	0.49 (0.04)	0.51 (0.03)	0.46 (0.03)
β_2^p : Partner's Signal	0.34 (0.03)	0.34 (0.03)	0.36 (0.03)
$\beta_2^p - \beta_1$	-0.15*** (0.04)	-0.17*** (0.03)	-0.10*** (0.03)
p -value: $\beta_2^p - \beta_1$ same as in Informed		0.57	0.27
p -value: $\beta_2^p - \beta_1$ same as in Observe			0.04
N	5040	7485	7435

Notes: This table shows reduced-form and structural estimates for rounds in Experiment 3 (the online experiment).

Reduced-form coefficients Panel A shows reduced-form results, estimating Equation 1 by OLS. The dependent variable is participants' private guess. "Informed" refers to the *Informed of Partner's Draws* round and "Observe" to the *Observe Partner's Draws* round. "Own Signal" indicates the net number of red draws (i.e., red draws minus white draws) in the participant's own set of draws. Similarly, "Partner's Signal" indicates the net number of red draws in their partner's set of draws. Randomization was between participants in this experiment so we estimate the equation separately for each treatment condition. Standard errors are clustered at the pair level. *, **, and *** indicate significance at the $p < 0.10$, 0.05, and 0.01 levels of the difference $\beta_2^p - \beta_1$.

Structural parameters: Panel B shows estimates of the structural model described in Section 3.3. "Own Signal" and "Partner's Signal" indicate the weights placed on their own and their partner's set of draws in the agents' quasi-Bayesian updating rule. Bootstrapped standard errors (clustered at the pair level) in parentheses. *, **, and *** indicate significance at the $p < 0.10$, 0.05, and 0.01 levels of the difference $\beta_2^p - \beta_1$.