UC Irvine

UC Irvine Previously Published Works

Title

Lone wolf or herd animal? Information choice and learning from others

Permalink

<https://escholarship.org/uc/item/1sb8g2d7>

Authors

Duffy, J
Hopkins, E
Kornienko, T

Publication Date

2021-05-01

DOI

10.1016/j.euroecorev.2021.103690

Peer reviewed

Lone wolf or herd animal? Information choice and learning from others[☆]

John Duffy^a, Ed Hopkins^b, Tatiana Kornienko^{b,*}

^a UC Irvine, USA

^b University of Edinburgh, Scotland, United Kingdom

ARTICLE INFO

Article history:

Received 6 March 2020

Revised 7 February 2021

Accepted 12 February 2021

JEL classification:

C72

C92

D83

Keywords:

Social learning

Information

Experiments

Conformity

Social influence

Information design

ABSTRACT

We report on an experiment that distinguishes between rational social learning and behavioral information source bias. Subjects are asked to correctly guess the current binary state of the world. Differently from other social learning studies, subjects must choose between receiving a private, noisy signal about the current state or observing the past guesses of other subjects in the prior period. Our design varies the persistence of the state across time, which affects whether private or social information is optimal. Thus our design enables us to separate subjects who choose information optimally from those who excessively use either social information (“herd animals”) or private information (“lone wolves”). We find sizable proportions of both behavioral types.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Humans are highly social. We seem to have a keen interest in the actions chosen by others. Does this interest reflect an *intrinsic*, “herd animal,” desire to follow and imitate the actions of others even when relying on such information is suboptimal?¹ By contrast, according to the social learning literature, it can be rational to observe others *instrumentally* because there is information contained in their actions. What has often been neglected in the debate between conformity and rationality is a *third* possibility that individuals could have a “lone wolf” bias *against* social information or following others.

[☆] We thank Philipp Albert, Jose Apesteguia, Juergen Bracht, Steven Broomell, Ignacio Esponda, Maia Guell, Steffen Huck, Chad Kendall, Gilat Levy, Mingye Ma, Dimitra Politi, Emanuel Vespa, Georg Weizsäcker, and Leeat Yariv for helpful comments, and Qixin He and Fiona Ross for research assistance. We also are grateful to the audiences at BEAT, Berlin Behavioral Economics Seminar, Chapman University, Higher Economic School, New Economic School, UC Irvine, UC San Diego, UC Santa Barbara, University of Aberdeen, University of Innsbruck, University of Manchester, and WZB for thought-provoking questions and suggestions. Funding for this project was provided by the Dietrich School of Arts and Sciences.

^{*} Corresponding author.

E-mail addresses: duffy@uci.edu (J. Duffy), E.Hopkins@ed.ac.uk (E. Hopkins), Tatiana.Kornienko@ed.ac.uk (T. Kornienko).

URL: <http://www.socsci.uci.edu/~duffy/> (J. Duffy), <http://homepages.econ.ed.ac.uk/hopkinse> (E. Hopkins), <http://homepages.econ.ed.ac.uk/~tatiana/> (T. Kornienko)

¹ Grant et al. (1998) consider agents with purely *intrinsic*, non-instrumental preferences for more information (or information avoidance) and make a connection between attitudes toward information and attitudes toward risk.

In this paper, we test whether individuals' interest in social or private information is rational, and if not, whether they exhibit persistent *information source biases*, with some individuals always preferring private information while others always preferring social information, even if that information source leads to a lower payoff.² We conduct an experiment that has several novelties compared with the classic social learning experiment of [Anderson and Holt \(1997\)](#),³ which is based on the seminal information herding theory of [Banerjee \(1992\)](#) and [Bikhchandani et al. \(1992\)](#). First, we use a direct elicitation approach, requiring subjects to *choose* between receiving either private or social information. Second, in our design, the binary state of the world that all subjects are seeking to guess can change over time, and depending on the persistence of this state of the world, social information is the optimal choice in some environments, while private information is the optimal choice in other environments. Third, our within-subjects design exposes subjects to a pair of these contrasting environments.⁴ Fourth, choices in our setup are simultaneous, so the decision problem is symmetric across subjects.

We find clear and sizeable heterogeneity in information choices. While a little more than 63% of subjects choose information optimally more than three quarters of the time, the remaining nearly 37% of subjects can be roughly divided up equally between lone wolves and herd animals. Specifically, more than 19% of subjects can be broadly classified as lone wolves, as they chose private information in more than three quarters of their decisions, while the remainder of more than 17% of subjects can be broadly classified as herd animals as they chose social information in more than three quarters of their decisions. Thus, one should not simply contrast rationality with conformism; it is really a *three-way* contest.

Our model predicts that the past choices made by others can be informative about the state of the world if that state is sufficiently persistent. By contrast, in a volatile environment, where the state of the world is likely to change, information about the past choices of others would not be as useful as a new private signal draw. For example, compare restaurant reviews that praise a restaurant's chef and hotel reviews that praise the hotel's central location. If both reviews are years old, we might doubt whether the chef is still working at the restaurant, but may trust a hotel review as it is unlikely the hotel has changed location.⁵ The question we explore in this experiment is the degree to which subjects can solve such relatively simple inference problems.

Most prior experiments on social learning, including [Anderson and Holt \(1997\)](#); [Celen and Kariv \(2004\)](#); [Goeree et al. \(2007\)](#); [Ziegelmeyer et al. \(2010\)](#) and [De Filippis et al. \(2018\)](#) involve a *permanent* state of the world that subjects are seeking to determine. Further, when it is a subject's turn to make a choice in these experiments, they are given *both* a private signal and social information on the prior choices of others in the sequence; there is no information choice.

Thus, our model differs from the classic social learning design, where subjects move in sequence and are required to combine statistical inference with strategic reasoning and higher order beliefs. However, from the above examples, our model's predictions are also important for making everyday social inferences. Further, our revealed preference approach identifies more clearly the relative value that subjects place on private versus social information. Given the many repetitions and full feedback we provide on the information that was not chosen, subjects in our design are as well placed as in the standard design to judge the relative merits of social versus private information. Importantly, our design allows for analysis of both the *choice* of private or social information as well as how that information is subsequently *used* in guessing the state of the world, which is not possible in the standard design. Perhaps surprisingly, we find that some subjects choose not to follow the information that they choose in making their guess about the state. By contrast, the standard social learning design has to back out the weighting of private versus social information sources as well as errors, all from a single choice, i.e. the subjects' guess about the state.

[Kübler and Weizsäcker \(2004\)](#) is closer to our model in that they make private information optional and only visible if a subject chooses to pay a small fee. [Cavatorta et al. \(2018\)](#) make one action unobservable and then vary how the information about the unobserved action is presented to subjects. Experiments examining social learning in finance, [Cipriani and Guarino \(2005\)](#), [Drehmann et al. \(2005\)](#), allow for the cost of different options to vary but not a choice between different sources of information. Thus, a novelty of our approach is that we ask subject to choose the type of information (private or social) they wish to receive.

Meta-study analysis of the standard, sequential move social learning experimental design conducted by [Weizsäcker \(2010\)](#) finds that subjects follow their own private information more frequently than is empirically optimal. [Ziegelmeyer et al. \(2013\)](#) analyse an augmented meta-dataset, and, using an updated methodology, confirm that subjects overweight their private signals though to a lesser degree than was previously found.

Yet suboptimal actions can be due to a variety of mistakes, whether in Bayesian updating, statistical inference, strategic reasoning, second order beliefs, or something else. Such mistakes are difficult to disentangle in the standard binary choice sequential design, as one mistake (say, due to mistrust in the rationality of a predecessor) might offset a second mistake

² Such persistent biases could be taste or belief-based - for example, one could have a desire to observe others or a strong belief that social information is uninformative because other people are unreliable.

³ The classic experimental design of [Anderson and Holt \(1997\)](#) involves subjects taking turn in a *sequence* guessing the *fixed* state of the world. Subjects are exogenously provided with *both* a private signal *and* the (social) history of their predecessors' guesses. In equilibrium, it is optimal to follow private signals relatively early in the sequence, and switch to following others later on.

⁴ [Charness et al. \(2019\)](#) similarly use pairs of contrasting environments to analyze choices over information structures.

⁵ According to surveys ([Sterling, 2018](#)), consumers are aware of such issues, with 85% claiming to disregard online reviews of local businesses that are more than three months old. This may be because positive reviews for a business typically are not erased even when that business changes management or ownership.

(say, in Bayesian reasoning), so that a subject making two mistakes might end up behaving in a manner that appears to be perfectly rational. (Most of the existing literature deals with this problem by ruling out one type of mistake altogether.) Even more, the behavior of someone who is biased towards a particular information source might be observationally equivalent to rational behavior, depending on their position in the sequence. For example, a subject who is biased towards private information and who happens to act early in the sequence, might be misidentified as acting rationally. If, instead, the same player made the decision later in the sequence, she might correctly be identified as someone who relies too much on private signals. Equally, another subject's bias towards social information might only be uncovered if she made her decisions early in the sequence. Since in a sequential design it is optimal to follow the herd in most rounds beyond the earliest, such design might be prone to uncovering a built-in information source bias largely in one direction, over-reliance on private information, and conflate biases towards social information with rational behavior.

In contrast, some more recent and complex experiments, Goeree and Yariv (2015), Eyster et al. (2018), March and Ziegelmeyer (2018), and Duffy et al. (2019) modify the standard social learning design so that excessive, suboptimal over-weighting of social information is possible, and all of these studies find evidence in this direction.

Most similar to this paper, Duffy et al. (2019) also examine information choice but in the context of the standard, sequential-move social learning design. There, at each point in the sequence, each subject must choose whether to receive a new private signal or to look at the history of guesses made by other subjects (social information) about the unchanging state of the world. While they document both herd animal and lone wolf-like behaviour in this sequential design, they find a tendency towards favoring social information relative to equilibrium predictions. They attribute this bias to particular features of the classic sequential design, which are not an issue in the present study.⁶ Indeed, in the model considered here, the proper choice of social versus private information depends only on the persistence of the state of the world. Thus for the purpose of identifying the frequency of herd animals versus lone wolves versus fully rational types in the population, we believe the symmetric design of the present experimental study provides the better framework.

While we urge caution about whether the bias we observe toward one information source or the other is driven by a persistent preference for or against social information versus simple errors, nonetheless, there are indications that the behavioral types we identify in our symmetric, information choice-based design are a more general phenomenon, worthy of further study. Overall, we see our design as complementary to the classic sequential one in understanding social learning and inference.

2. A simple model of social learning

We first develop a simple model of social learning, which we will use to form hypotheses that we test in our experiment.⁷

There are two periods. In each period, the state of the world is either X or Y . In period 1, the state is X or Y with equal probability. The state of the world in period 2 is the same as in period 1 with probability p and changes to the other state with probability $1 - p$. This probability, p , is the main treatment variable in our experiment.

There are n agents/subjects. In each period, all agents must choose an action X or Y . The payoff to choosing X when the state is X is $k > 0$ dollars and similarly the payoff to choosing Y when the state is Y is also k dollars. The payoff to choosing X when the state is Y or to choosing Y when the state is X is zero. In period 1, each agent receives a noisy but informative private signal, x or y , with commonly known precision $\Pr(x|X) = \Pr(y|Y) = q \in (\frac{1}{2}, 1)$. Each agent's signal is independent of the signals of others. At the end of period 1, **no** feedback or payoff information is given. In period 2, each agent must choose between receiving another informative, independent, private signal ("private information") having the same precision, q , as in period 1, but based on the (possibly different) period 2 state of the world **or** seeing all actions taken by the other agents in period 1 ("social information") when the state of the world was possibly different. Once the chosen information is received, the agent makes her period 2 choice of X or Y . Importantly, the number of agents n , the persistence parameter, p , and the precision of the private signal, q , are all common knowledge.

A strategy for an agent consists of three decisions: (i) a decision about whether to follow (or comply with) her signal in period 1, (ii) a decision about which information to receive at the beginning of period 2, and finally (iii) a final period 2 decision about which state to guess, conditional on the type of information received. This decision problem is formally a game as the payoff to the information choice in period 2 depends on whether agents follow their signal in period 1, and thus we write the result below in terms of perfect Bayesian equilibrium. However, the strategic aspects of this game are minimal, so weaker concepts such as rationalizability generate the same outcome.

In period 1, following one's signal is a dominant strategy (in expected payoffs) because, as $q > \frac{1}{2}$, the signal is informative. In period 2, if an agent chooses to see an independent signal, then as before, that signal will be correct with probability q . It will be optimal for the agent to follow this signal in period 2 and consequently the agent's expected probability of guessing

⁶ In particular, when actions and signals are binary, choosing social information in the third position of the sequence is, in theory, no more informative than choosing it in the second position, because the first two agents can contradict each other. However, in practice, subjects in the third position chose social information more frequently than those in the second position.

⁷ Our model draws inspiration from Samuelson (2004) who shows that in a sufficiently stable world, it can be optimal to observe the past actions of others as these convey information about the underlying state of the world. Moscarini et al. (1998) also study social learning with a changing state of the world.

correctly from choosing private information in period 2 will also be q .⁸ If instead an agent chooses social information in period 2, that agent can recall her own action and now sees the actions chosen by the $n - 1$ others in period 1. Let n be odd so that there is always a majority for one of the two actions. Assume that all took the dominant action in period 1 of following one's signal, then the optimal action, given an agent's choice of social information, is to copy the majority (or median) action in period 1 if $p > 0.5$ and to take the opposite action if $p < 0.5$.

The probability of success with this strategy depends on the key parameters n , p and q in the following way. The probability that the majority (or median) action was the correct action in period 1 for $n = 3$ is $Q_3 = q^3 + 3q^2(1 - q)$ and for $n = 5$ it is $Q_5 = q^5 + 5q^4(1 - q) + 10q^3(1 - q)^2$ and so on. That is,

$$Q_n(q) = \sum_{i=1}^m \binom{n}{i-1} q^{n+1-i} (1-q)^{i-1} \quad (1)$$

where $m = (n + 1)/2$.

The probability that the majority (or median) action is still the correct action in period 2 is equal to the probability that it was the correct action in period 1 multiplied by p . The probability that the majority action was incorrect in period 1 is $1 - Q_n(q)$. The probability that it was both incorrect in period 1 but it would be correct to follow it (because meanwhile the state of the world has changed) is $(1 - p)(1 - Q_n(q))$. Thus, the overall accuracy, A , of social information, that is, the probability of correctly predicting the state in period 2 by following the majority (or median) action in period 1, is

$$A(n, p, q) = pQ_n(q) + (1 - p)(1 - Q_n(q)). \quad (2)$$

The value of social information is increasing in all three variables n , p and q . This allows us to prove the following proposition regarding when the choice of social or private information is optimal.

Proposition 1. *In the unique Perfect Bayesian Equilibrium, in period 1 all follow their signals. In period 2, there exists a $p^*(n, q) \in (\frac{1}{2}, 1)$ such that the equilibrium strategy is to select social information and to follow it when the persistence of the state, p , is greater than p^* . Further, if $p < 1 - p^*$, then in equilibrium all select social information and go against the majority. Finally, for p in $(1 - p^*, p^*)$, the equilibrium strategy is to choose private information and to follow it.*

Proof. In period 1, the returns to following and not following one's signal are q and $1 - q$, respectively. As this is independent of the actions of others, following one's signal is a dominant action. Thus, any equilibrium strategy will place probability one that the accuracy of social information is equal to $A(n, p, q)$ as given in (2). Thus any such strategy will select social information when A is greater than q , the accuracy of private information. Because $Q_n(q) > q > \frac{1}{2}$, then clearly $A(n, p, q)$ is strictly increasing in p with $A(n, \frac{1}{2}, q) = \frac{1}{2}$ and $A(n, 1, q) = Q_n(q) > q$. So there is a unique $p^*(n, q) \in (\frac{1}{2}, 1)$ such that $A(n, p, q) = q$, and the result follows. Next, note that the accuracy of going against the majority is $1 - A(n, p, q)$ and by a similar argument this will be greater than q when $p < 1 - p^*$. It similarly follows that A is less than q in $(1 - p^*, p^*)$. $\square$

2.1. Experimental parametrization

In our experiment, the values of the three parameters n , p and q were chosen with the following considerations. The group size, n , was chosen to be 9, as this number is sufficiently large for social information to be an attractive choice and an odd number of subjects facilitates analysis of majority actions. Our choice for the signal precision, $q = 0.7$, was influenced by parameter values used in previous social learning experiments (see Weizsäcker, 2010).⁹

While n and q were fixed, we varied the persistence, p , across treatments, having three different environments, labeled *Persistent*, *Erratic* and *Anti-Persistent*. We intended the *Persistent* environment to have high persistence, p , and so $p = 0.9$ is an obvious choice. We then specified an *Anti-Persistent* environment to test for blind conformism and by symmetry, $p = 0.1$ is a natural choice. Finally, we chose $p = 0.6$ for the *Erratic* environment (where choosing private information is optimal) so that the *Persistent* and *Erratic* environments are almost exactly reverse symmetric in terms of the strength of incentives. Given $n = 9$ subjects per group, Eq. (1) results in the probability that the majority is correct being $Q_9(0.7) = 0.901$. Thus, it follows from Eq. (2) that the accuracy of social information in the *Persistent* environment, where $p = 0.9$, is $A(9, 0.9, 0.7) = 0.821$, so that in the *Persistent* environment, social information is more accurate than one's own private signal (of precision 0.7) by $0.821 - 0.7 = 0.121$. In contrast, in the *Erratic* environment, where $p = 0.6$, $A(9, 0.6, 0.7) = 0.580$, and thus private information is $0.7 - 0.58 = 0.12$ more accurate than social information.

In the *Anti-Persistent* environment, where $p = 0.1$, using Eqs. (1) and (2), one can calculate that in period 2 the accuracy from following the majority action of period 1 is $A(9, 0.1, 0.7) = 0.179$. Thus, choosing the action *opposite* to the majority action in period 1 yields the optimal period 2 action with probability $1 - 0.179 = 0.821$. As this probability is greater than drawing a private signal in period 2 with precision $q = 0.7$, the equilibrium strategy in the *Anti-Persistent* environment is to choose social information in period 2 but to guess the choice opposite to the majority's choice in period 1. By design, this

⁸ As Appendix A shows, the expected payoff from choosing private information in period 2 is q , even though the agent already has a private signal from period 1. This is because these two private signals may disagree.

⁹ Furthermore, this choice for q is very close to 0.697, which is the value of q that maximizes the difference between social and private information, $A(n, p, q) - q$, when $n = 9$ and $p = .9$.

Table 1

Summary of environments and parameterizations. Note: p denotes the persistence of the state of the world, q denotes the accuracy of private information, n is the number of subjects, and A is the (theoretical) accuracy of social information.

Environment	Persistence p	Accuracy of Private Info q	Number of Subjects n	Accuracy of Social Info A	Equilibrium Strategy (Period 2)
<i>Persistent</i>	0.9	0.7	9	0.821	Social (S), Follow (F)
<i>Erratic</i>	0.6	0.7	9	0.580	Private (P), Follow (F)
<i>Anti-Persistent</i>	0.1	0.7	9	0.821	Social (S), Not Follow (N)

strategy has the same expected success rate (0.821) as the equilibrium strategy (choose the social information and follow the majority choice) in the *Persistent* environment.

2.2. Experimental predictions

Given our parameterization of the model, the equilibrium strategy involves following the private signal in period 1, while in period 2 it varies with the persistence environment. In the *Erratic* environment, the equilibrium strategy is to choose private information and to follow the private signal received. In the other two environments, the equilibrium strategy is to choose social information, and to follow the majority's action choice from the previous period in the *Persistent* environment, and to do the opposite of the majority in the *Anti-Persistent* environment.¹⁰

The actual accuracy of social information in period 2 depends on other subjects having played optimally in period 1, namely guessing the state corresponding to their period 1 private signal. It would not be optimal to choose social information if these period 1 actions were sufficiently noisy. Let the *compliance rate*, γ , be the probability that an agent follows her signal in period 1. Then $\tilde{q} = 0.7\gamma + 0.3(1 - \gamma)$ is the *realized* accuracy of the period 1 guess. We show that so long as the compliance rate, γ , is greater than a certain critical level, the main results of the theory remain unchanged.

Proposition 2. For fixed n , q , and $p > p^*$, there exists a compliance rate $\gamma^* \in (\frac{1}{2}, 1)$ such that choosing social information is optimal for all $\gamma > \gamma^*$. For $n = 9$, $q = 0.7$ and $p = 0.9$, then $\gamma^* = 0.77$.

Proof. Given the possibility of error in the period 1 behavior, the *realized* social accuracy $\tilde{A} = A(n, p, \tilde{q})$ is clearly strictly increasing in $\tilde{q}$ with $A(n, p, \frac{1}{2}) = \frac{1}{2}$ and $A(n, p, 1) = 1$, so that there is unique $\tilde{q}$ such that $\tilde{A} = q$. Given $\tilde{q}$ is strictly increasing in γ , the result follows. The final specific value for γ^* is calculated numerically. $\square$

In contrast, if the compliance rate $\gamma < 0.77$, then choosing private information and following it is always the optimal choice in period 2 of all three persistence environments. However, as we will see later in [Section 4.1](#), the overall realized compliance rate is very high, averaging 0.978, which is far above the cutoff level of 0.77, so the strategy described in [Proposition 1](#) is indeed optimal.

3. Experimental design

Our experiment involved 144 inexperienced subjects recruited from the undergraduate population of the University of Pittsburgh. Each subject participated in a single experimental session involving 18 subjects, who were divided up into two groups of size $n = 9$ and remained in the same group for the duration of the experiment. Thus each session yielded two independent groups, and we have a total of 16 such independent groups of size 9.

The experiment consists of two main parts. Subjects were initially given written instructions only for the first part that were read aloud in an effort to make these instructions common knowledge. Subjects had to answer some control questions to verify their understanding of these experimental instructions and they then completed the first part of the experiment. After the first part was completed, the experiment was paused. Subjects were handed out new written instructions for the second part which emphasized that the only change from part 1 was to the persistence parameter, p .

One of the two parts of each experimental session always had a persistence of $p = 0.6$, referred to as the “Erratic” environment. The other part had a persistence of either $p = 0.9$, the “Persistent” environment, or $p = 0.1$, the “Anti-Persistent” environment. Thus each subject went through 48 rounds of two different persistence levels, for a total of 96 rounds (see [Table 2](#)). We refer to the different “within subjects” treatments by number labels. For example, “69” (*Erratic* then *Persistent*) refers to the treatment where the $p = 0.6$ environment was played for 48 rounds followed by 48 rounds of play of the $p = 0.9$ environment; “61” (*Erratic* then *Anti-Persistent*) to $p = 0.6$ followed by $p = 0.1$. We controlled for possible order

¹⁰ Risk aversion is unlikely to affect the equilibrium. In each period, the strategy with the highest expected payoff also has the lowest variance. Potentially, subjects could reduce the variance of their payoffs across the two periods by hedging, i.e., in period 2 guessing the opposite of their period 1 guess, irrespective of period 2 information, but it comes at high cost. In the *Persistent* environment, someone defecting to this strategy would reduce her expected return by almost 50% in order to reduce the standard deviation by 10%.

Table 2Experimental parameters. Group size $n = 9$ and accuracy of private signal $q = 0.7$ was fixed in all sessions.

Treatment	Part 1		Part 2		Total Rounds	Sessions	Groups	Subjects
	Persistence	Rounds	Persistence	Rounds				
96	$p = 0.9$	48	$p = 0.6$	48	96	2	4	36
69	$p = 0.6$	48	$p = 0.9$	48	96	2	4	36
16	$p = 0.1$	48	$p = 0.6$	48	96	2	4	36
61	$p = 0.6$	48	$p = 0.1$	48	96	2	4	36
Total						8	16	144

Table 3Summary statistics for the Signal Compliance Rate, (i.e., the subject-specific proportion of period 1 signals followed by each subject in each persistence environment, out of 48 rounds), and the ensuing realized accuracy, $\bar{A}$, of social information (based on the pooled compliance rate), for each persistence environment. Part 1 (2) means that the p environment was in the first (second) part of the session.

Persistence	Period 1 Signal Compliance Rate per Environment									Social Info
p	Part 1			Part 2			Both Parts Pooled			Realized
	Mean	(SD)	No. Obs.	Mean	(SD)	No. Obs.	Mean	(SD)	No. Obs.	Accuracy $\bar{A}$
$p = 0.9$	0.992	(0.021)	36	0.976	(0.100)	36	0.984	(0.072)	72	0.815
$p = 0.6$	0.970	(0.081)	72	0.979	(0.061)	72	0.975	(0.072)	144	0.578
$p = 0.1$	0.964	(0.071)	36	0.992	(0.027)	36	0.978	(0.055)	72	0.812
Pooled	0.974	(0.069)	144	0.982	(0.067)	144	0.978	(0.068)	288	

effects by varying the order of the two environments faced across sessions, so the other treatments are thus, “96” (*Persistent then Erratic*) and “16” (*Anti-Persistent then Erratic*).¹¹ Different pairs of *contrasting* persistence environments were run in different sessions and the order of each persistence pair was varied (see Table 2).

In each part, subjects were repeatedly confronted with the “main task” of: (i) guessing the state of the world in period 1 followed by (ii) an information choice and finally, (iii) guessing the state of the world in period 2.¹² In each part, this main task, consisting of three decisions, was repeated 48 times or “rounds” under a part-specific constant value for the persistence parameter, p . We chose to have a large number of rounds for each environment to allow time for subject learning. The number of group members $n = 9$, the precision of the private signal, $q = 0.9$, and the number of rounds in each part $N = 48$, were held constant across all treatments/sessions, were always at least public knowledge; they were explained in the written instructions which were then read aloud to all subjects at the start of each part. The persistence parameter, p , which varied across parts/treatments, was always announced in the instructions and aloud at the beginning of each part.

At the end of each 2-period round of play, subjects were presented with a complete, updated history of individually-relevant outcomes from all prior rounds of play, including the state of the world, and both subject’s own private signal and the other group members’ guesses, whether or not they chose to see that. Our design prevented subjects from imitating other group members’ behavior, as the only social information that is available is other subjects’ period 1 guesses, which subjects have to choose whether or not to view. Since our focus is on the choice between private and social information, we wanted to rule out any other channels of social influence that might affect that information choice.

At the end of the experiment, one round was randomly chosen from each of the two parts of the experiment. Since subjects could earn up to 2 points for each round (one point per period), they could earn up to a maximum of 4 points total for parts one and two. Points were converted into money payments at the fixed and known rate of 1 point=\$6. In addition, all subjects earned a fixed show-up payment of \$6 that required completion of an ex-post experimental survey. Thus, subjects could earn between \$6 and \$30 for a session lasting between 1.5 and 2 h. Average total subject earnings were \$21.80.

4. Experimental results

Each of our 144 subjects participated in a two-part “within subjects” treatment, with each part consisting of 48 rounds of a single persistence environment (thus 96 rounds in total across the two parts). While some variables exhibit small differences across the two parts, we do not find evidence for significant order effects in our primary outcome variables of interest, namely signal compliance and information choice rates. (See Appendix C for additional details including order effects, as well as Appendix D for dynamical aspects of our experimental data.)

¹¹ Copies of the written instructions used in the “96” treatment ($p = 0.9$ in part 1 and $p = 0.6$ in part 2) are found in the Appendix. Instructions for the other treatments are similar.

¹² The details of the experimental implementation of the main task are described in Appendix B.

4.1. Period 1 signal compliance rate

In period 1, subjects receive a private signal with a commonly known signal precision $q = 0.7$, and have to decide whether to follow it or not. Since the signal is informative (i.e., $q > 0.5$), the period 1 optimal action is to follow it. This first decision is important as the optimality of choosing social information in period 2 of the *Persistent* ($p = 0.9$) and *Anti-Persistent* ($p = 0.1$) environments depends on subjects being sufficiently “compliant”, i.e., correctly following their signal in period 1. We find that in fact subjects are highly compliant.

In Table 3, the signal compliance rates are broken down according to whether the p -treatment value was in place for the first (part 1) or the second (part 2) 48 rounds of a session. By design, all subjects faced the *Erratic* ($p = 0.6$) environment, as well as either the *Persistent* ($p = 0.9$) or the *Anti-Persistent* ($p = 0.1$) environment. Thus, there are twice as many subjects facing the *Erratic* ($p = 0.6$) environment (72 in each part, 144 total) than the other two environments (36 in each part, 72 total for each).

Overall, as Table 3 demonstrates, 97.8 percent of all period 1 decisions were to follow the signal.¹³ Thus, the period 1 compliance rates are significantly above the cut-off value, $\gamma^* = 0.77$, from Proposition 2 (one-sided $t = 29.915$, p -value = 0.000), and so the equilibrium strategy in period 2 is always to choose social information in the *Persistent* and *Anti-Persistent* environments. We note further that even for the lowest observed mean compliance rate (0.964 in the $p = 0.1$ environment for part 1), the realized accuracy $\hat{A}$ of social information is 0.806. This is not as high as the theoretical benchmark of 0.821, but it is still clearly above the accuracy of private information which is 0.7.

Result 1. The average period 1 signal compliance rate, 0.978, is significantly greater than the cut-off value, $\gamma^* = 0.77$. Thus, choosing social information is the optimal strategy in the *Persistent* ($p = 0.9$) and *Anti-Persistent* ($p = 0.1$) environments.

4.2. Period 2 information choice

We now turn to evaluation of the theoretical predictions of Section 2.2, and explore whether subjects choose information optimally in period 2 of each round of the different persistence environments. Given our Result 1, it is optimal to choose social information in the *Persistent* ($p = 0.9$) and *Anti-Persistent* ($p = 0.1$) environments, and choose private information about the period 2 state in the *Erratic* ($p = 0.6$) environment. We introduce the Information Optimality Index (IOI_i), which is the overall average proportion of the time that each subject i made optimal information choices across the 48 rounds of each of the two persistence environments that they faced:

$$IOI_i = \frac{1}{2} \frac{\sum_{t=1}^{N_{1,i}} I_{i,t}^{IO}(\text{if private optimal})}{N_{1,i}} + \frac{1}{2} \frac{\sum_{t=1}^{N_{2,i}} I_{i,t}^{IO}(\text{if social optimal})}{N_{2,i}} \quad (3)$$

$$= \frac{1}{2} \frac{\sum_{t=1}^{48} I_{i,t}^{IO}(p = 0.6)}{48} + \frac{1}{2} \frac{\sum_{t=1}^{48} I_{i,t}^{IO}(p \neq 0.6)}{48}$$

Here, $I_{i,t}^{IO}(\cdot)$ is a binary indicator variable equal to 1 if subject i chose optimal information in round t in the corresponding environment (i.e., private information if $p = 0.6$ and social information if $p \neq 0.6$), and zero if subject i chose the suboptimal information instead, in $N_{1,i} = N_{2,i} = 48$ rounds. Note that $IOI_i \in [0, 1]$, and here it can be regarded as the overall rate of optimal information choice across both of the persistence environments faced, 96 rounds in total.¹⁴ If subject i chose social information when it is optimal to do so and private information when that is optimal, then her IOI_i is 1. If she always made the wrong choice, then her IOI_i is 0. If she always chose either social information or private information in all 96 rounds, then her IOI_i is 0.5.

Across all treatments, the mean (st. dev.) IOI is: 0.787 (0.217) (see Table C.6 in the Appendix C.6), which means that an average subject made optimal choices 78.7% of the time. However, as the population shares in Fig. 1 (left panel) show, the distribution of this Information Optimality Index is bimodal. First, there is a large spike at perfect optimality ($IOI = 1$), with 21.5% of subjects being perfectly optimal in their information choice across both persistence environments. Second, there is another large spike at 0.5, with 13.2% of subjects making exactly half optimal choices and half suboptimal choices.

Result 2. Only about one fifth of subjects are perfectly optimal in their information choice across our two different persistence environments. About one-eighth of subjects make exactly half optimal choices and half suboptimal choices.

Note that if the subject-specific rates of optimal information choice across the two persistence environments, were instead calculated separately for a single persistence environment p , we would get a very different picture of subjects' individual rationality. As reported in Appendix C.3, such a “between-subjects” slicing of our data leads to an overstatement of the proportion of optimal choices since it is more difficult to distinguish information source biases from optimal behavior in a single environment, as a bias towards private information may look like an optimal choice in *Erratic* environment, while a

¹³ While period 1 compliance rates varied over rounds, the lowest compliance rate in any single round was 93.1%. We do not find any systematic or robust order effects in signal compliance rates across the two parts, and in any case, order effects for the compliance rate are not important as long as the signal compliance rate is sufficiently high relative to the cut-off value, $\gamma^* = 0.77$.

¹⁴ Observation counts $N_{1,i}$ and $N_{2,i}$ in each environment may vary across subjects (e.g., Appendix E).

Fig. 1. The distributions of period 2 subject-specific indices (pooled over all treatments, calculated over all 96 rounds, 144 individual observations). Left: Information Optimality Index IOI_i , right: Payoff Optimality Index POI_i (or period 2 expected payoff). The vertical axes report population shares, and the horizontal axes report the subject-specific indices.

Table 4

Table of information use strategies, corresponding theoretical payoffs, and empirical frequencies for each persistence environment. Roman numerals denote ranking of strategies in terms of payoff: I is best (in bold), II is second best (in italics), III is second worst, and IV is worst. QRE represents Quantal Response Equilibria (QRE) fitted frequencies (see Section 5).

Strategy	$p = 0.9$			$p = 0.6$			$p = 0.1$			Pooled
	Payoff π (Rank)	Empir. Freq.	QRE Freq.	Payoff π (Rank)	Empir. Freq.	QRE Freq.	Payoff π (Rank)	Empir. Freq.	QRE Freq.	Empir. Freq.
PF (Private, Follow)	0.700 (II)	0.172	0.284	0.700 (I)	0.741	0.654	0.700 (II)	0.186	0.284	0.460
PN (Private, Not Follow)	0.300 (III)	0.039	0.012	0.300 (IV)	0.045	0.027	0.300 (III)	0.028	0.012	0.039
SF (Social, Follow)	0.821 (I)	0.761	0.699	0.580 (II)	0.199	0.247	0.179 (IV)	0.030	0.005	0.297
SN (Social, Not Follow)	0.179 (IV)	0.028	0.005	0.420 (III)	0.015	0.072	0.821 (I)	0.755	0.699	0.203
No. Obs.			3456			6912			3456	13,824

bias towards social information may look like an optimal choice in the other two environments. Using the between subjects optimality measure would lead us to conclude that 47.9% of our subjects were perfectly optimal, more than twice the population share of rational types found using the overall, within-subjects, information optimality measure. For this reason, in what follows, we conduct our analysis using only our overall, within-subjects IOI measure, which we believe to be a better indicator of individual rationality.

4.3. Period 2 information use and expected payoff

We now examine whether subjects optimally used the information they chose to receive before making their period 2 choices as predicted in Section 2.2, as well as the payoff consequences of their choices.

Since there are only two types of information, private (P) and social (S), and the state of the world is binary, a subject can choose either to follow (F) the private signal or, in case of social information, the majority (or median) guess; or not to follow (N) this same information, i.e., do the opposite. This results in 4 strategies denoted as: PF , PN , SF , and SN - e.g., the strategy SF is to choose social information and to follow it. The strategies can be ordered in terms of their theoretical expected payoff π , with rank I being the equilibrium strategy in that environment and rank IV being the worst, and this ordering will depend on the persistence environment.

We find that there is a high degree of optimal information use, but it is still far from 100%. As Table 4 (and also Table C.3 in Appendix C.4) shows, the frequencies of strategies chosen are largely ordered in terms of their relative payoffs. The frequency of the optimal information use strategy, ranked I in Table 4, averages 0.75 across both parts of the three persistence environments. Further, by summing up the frequencies of the strategies ranked I and II, one can see from Table 4 that the overall frequency of the correct use of chosen information was high (at 0.933, 0.940, and 0.942 in environments $p = 0.9$, $p = 0.6$, and $p = 0.1$, respectively), with an overall average of 0.939.

Result 3. In the aggregate, across all three environments, information-use strategies in period 2 are highly rational, with three quarters of all observations corresponding to the optimal (best-ranked) strategy.

We now turn to examining the overall payoff consequences of subjects' choices. Recall that all subjects participated in two of the three environments, one of which was always the *Erratic* environment, and the other one was either the

Persistent or the Anti-Persistent environment. As predicted in Section 2.2, the latter two have the same theoretical expected payoff. We thus further calculate an overall subject-specific Payoff Optimality Index, POI_i , which is the expected period 2 payoff, averaged over both parts of the experiment.¹⁵ It is defined as the sum of payoffs to each of the four strategies, PF, PN, SF and SN, defined earlier and weighted by the proportion of times that a subject used those strategies in period 2 of each part (and thus in each persistence environment):

$$POI_i = \frac{1}{96} \sum_{t=1}^{48} [0.7 \cdot I_{i,t}^{PF}(p=0.6) + 0.3 \cdot I_{i,t}^{PN}(p=0.6) + 0.580 \cdot I_{i,t}^{SF}(p=0.6) + 0.420 \cdot I_{i,t}^{SN}(p=0.6)] \\ + \frac{1}{96} \sum_{t=1}^{48} \begin{cases} 0.7 \cdot I_{i,t}^{PF}(p) + 0.3 \cdot I_{i,t}^{PN}(p) + 0.821 \cdot I_{i,t}^{SF}(p) + 0.179 \cdot I_{i,t}^{SN}(p) & \text{if } p = 0.9 \\ 0.7 \cdot I_{i,t}^{PF}(p) + 0.3 \cdot I_{i,t}^{PN}(p) + 0.179 \cdot I_{i,t}^{SF}(p) + 0.821 \cdot I_{i,t}^{SN}(p) & \text{if } p = 0.1 \end{cases}$$

The theoretical minimum and maximum values for the overall average expected payoff are $POI_{\min} = 0.2395$ and $POI_{\max} = 0.7605$, so that $POI_i \in [0.2395, 0.7605]$.

As Fig. 1 (right panel) shows, the distribution of period 2 expected payoffs is skewed, with a mean expected payoff (st. dev.) of 0.709 (0.051) and a median of 0.720. There is a clear and distinct mode around the maximum expected payoff of $POI_{\max} = 0.7605$ from following the (best) optimal strategy in each environment, however only 15 subjects out of 144 (10.4%) achieved this maximum payoff.¹⁶ As we will explore in the next section, there is significant payoff heterogeneity, with a long leftward tail in Fig. 1 (right panel) representing subjects who persistently chose suboptimally. The worst total expected payoff was $POI_i = 0.503$, which is close to a pure random guess.

Result 4. On average, subjects achieved an expected payoff of 0.709. However, only 10.4% of all subjects obtained the maximum period 2 overall expected payoff, $POI_{\max} = 0.7605$, by both choosing and using information correctly.

4.4. Information source bias and its consequences

We will now turn to our main question of interest, whether subjects exhibit any bias toward a particular source of information, and develop a methodology which allows one to explore the consequences of such a bias.

4.4.1. Identifying information source bias

To quantify subjects' tendency to choose private information, we construct a subject-specific "Lone Wolf Index" (LWI_i) by adding the proportions of time that each subject i made private information choices across 48 rounds of each of the two persistence environments they faced, and subtracting 1:

$$LWI_i = \frac{\sum_{t=1}^{N_{1,i}} I_{i,t}^{PI}(\text{if private optimal})}{N_{1,i}} + \frac{\sum_{t=1}^{N_{2,i}} I_{i,t}^{PI}(\text{if social optimal})}{N_{2,i}} - 1 \quad (4) \\ = \frac{\sum_{t=1}^{48} I_{i,t}^{PI}(p=0.6)}{48} + \frac{\sum_{t=1}^{48} I_{i,t}^{PI}(p \neq 0.6)}{48} - 1$$

Here, $I_{i,t}^{PI}(\cdot)$ is a binary indicator variable equal to 1 if subject i chose private information in round t in the relevant environment, and zero if she chose social information instead, in $N_{1,i} = N_{2,i} = 48$ rounds. As $LWI_i \in [-1, 1]$, here it is simply a rescaled overall rate of private information choice across both persistence environments, 96 rounds in total. For example, if a subject chooses optimally in the 69 treatment, she chooses private information 100% of the time (or 48 times) in the first part, and after the environment switches from Erratic to Persistent, she chooses social information 100% of the time in the second part. Thus, her LWI_i is 0, as she is unbiased. If, however, she always chooses social information, then she will have a LWI_i of -1, as a fully prosocial, pure "herd animal". If she always chooses private information, then her LWI_i would be 1, as a fully antisocial, pure "lone wolf".

Fig. 2 depicts population shares of subject LWI_i and shows that the distribution of the Lone Wolf Index is broadly unimodal and symmetric around zero. In the aggregate, the population of subjects is effectively unbiased, with the mean (st.dev.) LWI of $-0.001(0.558)$ and median of 0, and a 95% confidence interval of $[-0.093, 0.091]$. The mean LWI is not different from zero according to a two-tailed, one-sample t -test, p -value=0.975). A test for skewness and kurtosis gives $\Pr(\text{Skewness}) = 0.784$ and $\Pr(\text{Kurtosis}) = 0.544$, with adjusted $\chi^2(2) = 0.45$ (p -value = 0.799), suggesting that there is no systematic bias for or against a particular source of information.

We find that scores for the Lone Wolf Index vary widely across subjects, allowing us to classify them by their choice rates of a particular type of information (see Fig. 2 and Table 5). We find that 28 (19.4%) of subjects are "broad lone wolves"

¹⁵ Given that the theoretical range of the period 2 expected payoff varies with the persistence parameter, p , there is no obvious way to normalize a combination of the expected payoffs from the two parts. As we are interested in understanding differences in subjects' behavior, we see the "raw" average expected payoff as a meaningful payoff metric.

¹⁶ The spike at the maximum value of period 2 expected payoff in Fig. 1 (right panel) includes extra three subjects who earned slightly less than the maximum payoff.

Fig. 2. The distribution of period 2 subject-specific Lone Wolf Index LWI_i (pooled over all treatments, 144 individual observations). The vertical axis reports population shares, and the horizontal axis reports the rescaled subject-specific private information choice rates.

Table 5

Summary statistics on the Payoff Optimality Index POI_i (or the overall expected payoff in period 2, and its normalization, $NPOI_i = \frac{POI_i - POI_{\min}}{POI_{\max} - POI_{\min}}$, averaged over all 96 rounds), by broad subject type as defined by their LWI_i .

Broad Type	LWI_i	#Subj. (%)	Expected Payoff in Period 2, POI_i				Normalized Expected Payoff, $NPOI_i$			
			Mean	(St.Dev.)	Min	Max	Mean	(St.Dev.)	Min	Max
Herd Animal	$[-1, -0.5]$	25 (17.4%)	0.687	(0.016)	0.652	0.721	0.859	(0.031)	0.792	0.924
Unbiased	$[-0.5, 0.5]$	91 (63.2%)	0.731	(0.041)	0.503	0.7605	0.943	(0.078)	0.506	1.000
Lone Wolf	$(0.5, 1]$	28 (19.4%)	0.655	(0.054)	0.548	0.728	0.797	(0.104)	0.591	0.937

who choose private information in at least 75% of their information choices. As we will see later in Fig. 3, a typical “broad lone wolf” would always choose private information in the Erratic ($p = 0.6$) persistence environment, and more than half of the time in the other (non-erratic) persistence environment (either the Persistent ($p = 0.9$) or the Anti-Persistent ($p = 0.1$) environment, depending on the session). Approximately the same number, 25 (17.4%) of subjects are “broad herd animals” who chose social information at least 75% of the time, typically choosing social information in all rounds of the non-erratic environment they faced, as well as more than half of the time in the other, Erratic ($p = 0.6$), persistence environment.

Moreover, as Fig. 2 shows, some subjects demonstrated an extreme bias toward a particular type of information. Surprisingly, 8 out of 144 subjects (5.6%) were pure lone wolves, that is, they chose private information 96 out of 96 times so their LWI score was 1. Almost the same number, 9 out of 144 subjects (6.3%), were pure herd animals, that is, they chose social information 96 out of 96 times giving them an LWI score of -1 . Thus we find that 17 out of the 144 subjects (11.8%) were 100% optimal in one persistence environment but 100% suboptimal in the other environment. As Appendix D.2 demonstrates, the behavior of these pure types cannot be explained by individual histories of signal realizations.

We classify the remaining 91 (63.2% of all 144 subjects) of subjects as “broadly unbiased”, because they have no strong bias in either direction, among whom 34 (23.6%) subjects have a LWI score of exactly 0 (with 31 (21.5%) subjects always choosing optimal information and 3 subjects making suboptimal choices of private information that were exactly offset by suboptimal choices of social information).

Result 5. There is no aggregate bias in information choice. 19.4% of subjects are “broad lone wolves” who choose private information in at least 75% of their information choices, and 17.4% of subjects are “broad herd animals” who chose social information in at least 75% of their information choices. The remaining 63.2% of subjects are “broadly unbiased” in their information choices which includes 21.5% of subjects who make perfectly optimal information choices in both environments.

4.4.2. Information source bias vs. information optimality

We will now combine the Information Optimality and Lone Wolf Indices and explore the interaction between optimality and information source bias that neither index by itself can address. Fig. 3 plots the Information Optimality Index IOI_i

Fig. 3. The “Diamond of Rationality” plots each of the 144 subjects’ optimal information choice rate, as measured by their Information Optimality Index, IOI_i (horizontal axis) against their information bias as measured by the Lone Wolf Index, LWI_i (vertical axis). By construction, LWI and IOI are constrained by the dashed lines. Circle sizes are proportional to the number of subjects with a given pair of indices. Shaded areas represent classification into broad types.

(horizontal axis) against the corresponding Lone Wolf Index LWI_i (vertical axis) for each subject i , creating a “diamond shape” of possible values. Full or pure optimality, i.e., always choosing the correct information without bias, results in the point (1,0). A pure lone wolf subject corresponds to the point (0.5, 1), as she always chooses private information and thus is always correct only in the *Erratic* environment and always wrong in the *Persistent* or *Anti-Persistent* environments. Pairs of indices located on the line between these two points represent subjects who always choose private information in the *Erratic* environment, and who choose a mixture of private and social information in the other two environments. Similarly, the point (0.5,-1) represents pure herd animal subjects who always choose social information in the two environments. The final vertex of the diamond, point (0,0) represents consistently incorrect behavior, only choosing social information when private information is optimal and vice-versa.

Fig. 3 shows how the population of subjects is distributed within the Diamond of Rationality, with the size of each circle reflecting the frequency count of subjects. A majority of subjects are close to the far right vertex of the diamond, corresponding to optimal behavior and no bias, yet only 31 (21.5%) subjects are completely rational with scores of (1, 0). There are also significant numbers of subjects at two of the other vertices, with 9 (6.3%) subjects (pure lone wolves) at (0.5, 1), and 8 (5.6%) subjects (pure herd animals) at (0.5,-1).

Result 6. There is considerable heterogeneity with herd animals and lone wolves coexisting with unbiased/optimal types. The distribution is symmetric, with herd animals and lone wolves appearing in approximately equal numbers.

The distribution of types reported on in Result 6 is not much affected if we allow for learning over time, specifically if we classify subjects based on their behavior in the first 6 versus the last 6 rounds of each part (see Appendix D.3 for a detailed analysis).

4.4.3. Information source bias vs. payoff optimality

In Fig. 4, we map subjects’ Lone Wolf Index LWI_i scores against their total expected payoffs POI_i earned given their strategy choices, with the dashed lines representing the maximum payoffs (obtainable by always using chosen information correctly), conditional on a particular LWI score. For example, pure lone wolf subjects ($LWI=1$) who always followed their chosen private information, would guess the state of the world correctly with probability 0.7 which is the precision of private information. Indeed there is a small cluster around the point (0.7,1) in Fig. 4, which can be contrasted with another small cluster around the point (0.7005,-1) representing pure herd animal subjects ($LWI=-1$) who always followed their chosen social information. There is a larger cluster around the equilibrium strategy at $LWI=0$ with an expected payoff of

Fig. 4. The “Arrow of Rationality” plots each of the 144 subjects’ Payoff Optimality Index POI_i (horizontal axis) against their information bias as measured by the Lone Wolf Index LWI_i (vertical axis) with dashed lines outlining the envelope of possible combinations. Circle sizes are proportional to the number of subjects with a given pair of indices, and point D depicts the average of the data. Shaded areas represent classifications of subjects into broad types as in Fig. 3. The solid curve, labeled logit QRE, depicts the set of logit QRE for all values of the QRE precision parameter, β , and the point Q shows the aggregate QRE prediction with 99% confidence intervals around it (see Section 5 for a discussion of the QRE estimates).

0.7605 (derived from 0.7 being the highest accuracy in the *Erratic* environment and 0.821 in the *Persistent* and *Anti-Persistent* environments).¹⁷

Interestingly, subjects’ biases, as captured by their Lone Wolf Indices (LWI_i), result in marked reductions in their payoffs. As Table 5 shows, the 91 “broadly unbiased” subjects (63.2% of the subject population) lost, on average, 5.7% of the normalized payoff across all period 2 decisions. By contrast, the 25 (17.4%) “broad herd animals” lost, on average, 14.1% of the normalized payoff. Interestingly, the 28 (19.4%) “broad lone wolves” stand out, as on average they lost the most, 20.3% of the normalized payoff, at least in part because these biased subjects were more prone not to use their information correctly (see Appendix C.4).

Result 7. Information source biases have payoff consequences across all period 2 decisions, with broadly unbiased subjects losing on average only 5.7% of the normalized overall payoff, while broad lone wolves and broad herd animals lose on average 20.3% and 14.1% of the normalized overall payoff, respectively.

5. Alternative explanations

We now explore whether prominent behavioral models can explain our findings. We start with the logit Quantal Response Equilibrium (QRE) model (see Goeree et al., 2016 for details), which assumes that all subjects play noisy best responses to the play of others, and which has been successful in explaining deviations from Nash equilibrium. We calculate the payoffs to choosing social information using the empirical average period 1 signal compliance rates as given in Table 3. The payoff precision parameter, β , that maximizes the log-likelihood for the logit QRE model, is found to be 7.955. One can see that the associated QRE frequency estimates shown in Table 4 are close to the experimental data frequencies, but that QRE slightly underestimates the frequencies of both the highest-paying and the lowest-paying strategies. Importantly, the standard, symmetric QRE model does not predict any particular bias in favor of or against following a specific source of information.

To examine the fit of a logit QRE to our data, we first calculate the logit QRE for different possible values of the logit precision parameter, β . In Fig. 4, the set of QRE is the curve running from (0.5, 0) to (0.7605, 0). When β is zero, in QRE agents choose at random, giving a payoff of 0.5. As β becomes very large, subjects choose the correct information with a

¹⁷ Theoretically, pure lone wolves ($LWI=1$) and pure herd animals ($LWI=-1$) are expected to earn at most 0.7 and 0.7005, respectively, thus losing money at an approximately symmetric rate relative to those following equilibrium strategy.

probability close to one, giving them the highest possible payoff. These calculations take into account the noisy behavior in period 1. Apart from a small non-monotonicity for relatively small payoffs, for the most part the predicted QRE frequencies are very close to being symmetric (see Fig. 4). The estimated logit QRE (as given in Table 4) is represented by the point labeled Q with coordinates (0.703, -0.023) in Fig. 4, while the point D with coordinates (0.709, -0.001) shows the average from the experimental data pooled across all sessions. As one can see, the two points Q and D are remarkably close.

The good fit of QRE at the *aggregate* level masks a considerable mismatch at the *individual* level. Subjects who are very strongly biased either for or against private information cannot easily be explained by a symmetric QRE, in which the choice probabilities are the same for all subjects. Compare the point Q with the data point D also plotted in Fig. 4. If the data had been generated by a symmetric QRE then there should be a unique mode around the estimated QRE, point Q. Indeed, on Fig. 4, we have drawn a two-dimensional interval such that if a subject chose information according to the estimated QRE frequencies across 48 repetitions, her choices would be within both intervals with 99% probability. However, as Fig. 4 clearly reveals, the experimental data look very different. There are many subjects that are far more accurate than is predicted by QRE, so the mode is not around (0.703, -0.023) but is instead at (0.7605, 0), which lies well outside the 99% confidence intervals for the QRE. There are further modes in two other corners representing subjects who appear to be biased for or against social information. Thus, a symmetric QRE cannot explain the simultaneous existence of lone wolves and herd animals.¹⁸

Result 8. Symmetric logit QRE estimates are a good fit to the aggregate empirical data, slightly underestimating the frequencies of both the highest and the lowest paying strategies. However, the QRE fails to predict the heterogeneity in information source bias found in our data.

Another prominent behavioral model is the Level-*k* model of differing levels of strategic sophistication (see, for example, Crawford et al. (2013)). In that model, the lowest “Level 0” subjects are assumed to randomize uniformly over their choice of the two types of information and would thus not contribute to any overall bias. However, the next, Level 1 subjects, believing that all others are Level 0, would reason that period 1 actions of others are uninformative and thus would always choose the private signal. Thus, the Level-*k* model predicts an overall bias towards lone-wolf behavior, as long as there is some positive measure of Level 1 types as is typically the case. However, as we saw in Section 4.4, overall subjects are unbiased in their information choices.¹⁹

6. Conclusions

We have conducted an experiment on social learning using a novel experimental design that enables us to distinguish between rational behavior and information source bias. Subjects have to choose whether to observe a private signal or the previous period choices of fellow subjects (social information). By altering the persistence of the state, we alter the optimal information source to choose. Mistakes can therefore run in both directions: subjects can choose private information when social information is optimal and vice versa. This allows for a clearer identification of biases in subject behavior.

Most importantly, we find that there is considerable subject heterogeneity with what we call lone wolves and herd animals both being present alongside rational individuals. We argue that these deviations from optimal behavior are driven at least in part by persistent behavioral types rather than error. First, in this experiment subjects make many decisions with detailed feedback and so have plenty of opportunity to learn if they desire. Second, standard behavioral models based on bounded rationality and error-making such as quantal response equilibrium and Level-*k* do not predict the observed pattern of behavior.

Our finding of no bias in favor of private information is another main take-away result from our paper and stands in contrast to previous social learning experiments. In the sequential structure of these previous studies, both social and private information are given to subjects, and it is optimal to follow social information unless one's position is very early in the sequence. Thus, it is more likely that bias is in terms of excessive use of private information.²⁰ In contrast, our current design is more symmetric, and thus provides cleaner identification of the frequencies of lone wolves and herd animals in the population.²¹

We realize that some readers may view the observed heterogeneity in our experiment as reflecting preferences with respect to social interaction rather than “rules of thumb”, habit or heuristics. In settings such as the one we study, it may be difficult to separate these different explanations. For example, when trying to find one's way in an unknown city, some individuals may opt to use a map, while others may ask passers-by for directions. Do these different choices express a

¹⁸ Even allowing the logit precision parameter β to vary at the individual level cannot explain data in which there are subjects who always choose social (private) information when it is optimal, implying a very high precision level, and who also always choose same type of information when it is not optimal, implying a low or even negative precision level.

¹⁹ Another behavioural theory, ambiguity aversion, also suggests a bias towards private information that is not observed.

²⁰ Nevertheless, using the data from previous standard sequential binary choice social learning experiments compiled in the meta-studies of Weizsäcker (2010) and Ziegelmeyer et al. (2013), a within-subjects analysis (see Appendix E) finds evidence for both lone wolf and herd animal player types.

²¹ In future research, it would be of interest to implement a modified within-subjects design to determine whether the lone wolves, herd animals and rational types identified here act in a similar manner in the standard social learning setting.

preference for private versus social interaction or for different rules of thumb? Equally here, having a preference for private over social information could be considered as a preference for different ways of approaching decision problems. We are open to either way of describing the observed behavior in our experiment.

In conclusion, we find evidence in support of the notion that social influence is an important aspect of human behavior. At the same time we also find that its reverse, an aversion to social influence, also exists, and that both herd animals and lone wolves coexist in approximately equal numbers – a finding that is new to the literature. Further, we also find that many subjects in our experiment appear to have no intrinsic interest in, nor an aversion to learning from others' behavior, but simply choose to rely on such information when it is most useful. Our findings parallel the literature on social preferences where inequality averse and reciprocal agents have been identified to co-exist alongside the entirely self-interested. We hope that the menagerie of social types introduced in this paper will also find widespread applications.

Appendices

A. Accuracy of period 2 private signal

One might think that if an individual chooses private information, then the accuracy with which that individual can guess the state in period 2 is greater than q as she already has a signal from period 1. However, this previous signal is dominated by the new signal in that it can never be optimal to follow the first signal over the second signal when the two signals disagree. Hence, the first signal has no effect on the individual's accuracy in period 2 if she opts for private information.

Formally, if an individual follows the equilibrium strategy and guesses the state is X in period 2 if and only if the signal is x in period 2 (and the period 1 signal has no effect on the guess she makes in period 2), then the probability that this guess is correct is the probability that the state is X when the signal is x , that is $\Pr(X_2|x_2) = \Pr(x_2|X_2) = q$, where the subscript indicates period.

It is true that if the signals from periods 1 and 2 agree, then the probability of being correct, conditional on agreement is greater than q . But it can also happen that the two signals disagree. In this case, it is still optimal to follow the more recent signal, but the conditional probability of being correct is lower. It turns out that the overall expected accuracy is exactly equal to the signal precision, q . To see this, note that the probability of the state being X in period 2, given that the signals are x in both periods, is

$$\Pr(X_2|x_1x_2) = \frac{\Pr(x_1x_2|X_2) \Pr(X_2)}{\Pr(x_1x_2|X_2) \Pr(X_2) + \Pr(x_1x_2|Y_2) \Pr(Y_2)}$$

where $\Pr(x_1x_2|X_2) = q(pq + (1-p)(1-q))$ and so on. Then, for example, in the *Persistent* environment where $p = 0.9$, $\Pr(X_2|x_1x_2) = 0.819$ which is indeed greater than 0.7. Thus, if an individual sees the same signal in both periods, she would correctly infer that the probability of the state truly being X is high. But equally one has that

$$\Pr(X_1|y_1x_2) = \frac{\Pr(y_1x_2|X_2) \Pr(X_2)}{\Pr(y_1x_2|X_2) \Pr(X_2) + \Pr(y_1x_2|Y_2) \Pr(Y_2)}$$

where y_1x_2 is the event of having signal y in period 1 and x in period 2. For example, in the *Persistent* environment where $p = 0.9$, $\Pr(X_2|y_1x_2) = 0.546$. Thus, the presence of the period 1 signal can also reduce the predicted accuracy of the period 2 signal.

The overall accuracy is

$$\begin{aligned} \Pr(X_2|x_2) &= \frac{\Pr(x_1x_2) \Pr(x_1x_2|X_2) \Pr(X_2)}{\Pr(x_1x_2|X_2) \Pr(X_2) + \Pr(x_1x_2|Y_2) \Pr(Y_2)} + \frac{\Pr(y_1x_2) \Pr(y_1x_2|X_2) \Pr(X_2)}{\Pr(y_1x_2|X_2) \Pr(X_2) + \Pr(y_1x_2|Y_2) \Pr(Y_2)} \\ &= \Pr(x_1x_2|X_2) \Pr(X_2) + \Pr(y_1x_2|X_2) \Pr(X_2) = q, \end{aligned}$$

as $\Pr(y_1x_2|X_2) = \Pr(x_1x_2|X_2) = q$ and $\Pr(X_2) = \frac{1}{2}$. So the overall accuracy is q , as claimed.

B. Experimental design: implementation of the main task

In the main task, each round consists of two periods, period 1 and period 2. Subjects were instructed to imagine that there exist two urns, a “black” urn containing 7 black balls and 3 red balls, and a “red” urn containing 3 black balls and 7 red balls. These distributions of balls in the two urns reflect our signal precision choice of $q = 0.7$, which was fixed across all treatments. For all members of each 9-member group, one urn was randomly chosen at the start of each new period 1 (in a two period round) with an equal (0.5) probability of either urn. Subjects were instructed that: “it is as though a coin flip determines which of the two urns is chosen in each round”.

We used a paired-group design involving paired groups of $n = 9$ subjects (18 subjects per session). This allowed for only one sequence of random numbers to determine the sequence of urn draws for both paired groups. Thus, when one of the paired groups faced a “black” urn, the other faced a “red” urn, and vice versa. This was done to ensure that the dataset contained the same number of red and black urns. The random urn choice draws were “live” for the first session of 18 subjects, but thereafter we used the same sequence of random draws in all subsequent sessions. We did this so that subjects in the different sessions faced the same empirical frequencies of urn colors.

For period 2, the urn color remains the same as in period 1 with probability p , or changes to the other colored urn with probability $1 - p$. The paired-group design ensures that the urn draws in period 2 follow the same switching pattern for

both paired groups. That is, if in period 2 there was a switch from the black to the red urn in the first paired group, then the second paired group would have a switch in period 2 from the red to the black urn.

Given the urn that is in place for a given period, ball colors drawn corresponded to the color of the urn with probability $q = 0.7$ or not with probability $1 - q = 0.3$. The individual ball draws were made randomly, independently and with replacement for each subject viewing a colored ball; the latter draws were live (i.e., real-time) in all sessions; only the urn sequence was pre-determined following the first experimental session.

In period 1, subjects are shown the color of a ball selected from the unknown urn and must guess the color of that unknown urn. The decision screen at the end of period 1 shows the subject's choice of Red or Black for the urn color, but subjects do not immediately learn the color of the urn that was chosen. Instead, we move to a second decision screen, where subjects are reminded of the color of their period 1 ball draw (their signal) and their period 1 choice of urn. They are told that for period 2, there is a p percent chance that the (still unknown) urn will be the same color urn as in period 1 and a $1 - p$ percent chance that for period 2, the urn will be the opposite of the one used in period 1. Note that the urn in each period is common to all $n = 9$ members of a group. Following period 1 but prior to the start of period 2, each subject is asked whether s/he would prefer to draw a ball from the urn chosen for period 2, or would prefer to look at the actual urn choices (guesses) made by the other 8 subjects in his/her matching group for period 1 (and *not* the 8 signals the other 8 players received in period 1).

After making this information choice, play proceeds to period 2. If the subject chose to draw a new ball (i.e., private information), then a ball is drawn randomly from the urn that is in place for period 2 and the color of that ball is revealed to the subject. The subject then chooses the color of the period 2 urn that s/he thinks the ball was drawn from. On the other hand, if the subject chose to look at the urn choices made by the other 8 subjects in period 1 (i.e., social information), then the subject is shown the numbers m and $8 - m$ of the other 8 subjects who chose the Black and the Red urns, respectively. The subject was reminded of her own choice for period 1 and is asked to make an urn choice for period 2.

After all period 2 urn choices were submitted, the round was over and subjects received feedback on the outcomes of that round. Specifically, subjects were reminded of the color of the ball they had drawn for period 1, their guess of the urn for period 1 and the actual color of the urn in period 1. They were further reminded of their information choice prior to period 2 (i.e., new ball draw or group information from period 1), and the contents of *both* their chosen and foregone information (i.e., both the group information and the random ball draw from the period 2 urn) - so that subjects had an opportunity to assess whether their choice of information was optimal or not, without experimenting with different sources of information. Finally, subjects were also informed of their guess of the urn color for period 2 and the actual color of the urn in period 2, as well as their payoffs for the round. For each period in which they correctly guessed the true color of the urn, they received 1 point and 0 points otherwise. Thus, for each round, subjects could earn 0, 1, or 2 points, depending on the accuracy of their guesses for the urn colors in periods 1 and 2.

At the end of each 2-period round of play, a complete, updated history of outcomes from all prior rounds of play was reported at the bottom of subjects' decision screens. This history included (i) the color of the ball drawn in the period 1 of each prior round, (ii) the subject's own guess of which urn was selected in each period, (iii) the information the subject chose to view (New Draw /Group), (iv) the *other* piece of information the subject did not choose to view (to allow for learning), (v) the actual group urn that was selected for periods 1 and 2, and (vi) the subject's points earned for the round.

Note that our design prevents subjects from imitating other group members' information choices or period 2 guesses as we do not reveal any information about those choices. The only social information that is available is other subjects' period 1 guesses, which subjects have to choose whether or not to view. Since our focus is on the choice between private and social information, we wanted to rule out any other channels of social influence that might affect that information choice.

C. Further results

Before presenting additional results, we first check for order effects, i.e., whether those subjects who faced a particular environment in part 1 made different decisions from subjects who faced that same environment in part 2. In our design, no subject experienced the same persistence environment in both parts, permitting the use of tests for two independent samples.

C1. Period 1 signal compliance rate, order effects

We define the period 1 signal compliance rate $SCR_i(p) \in [0, 1]$ per environment p as the proportion out of 48 rounds when a given subject i followed his/her period 1 signal when facing a specific persistence p , or

$$SCR_i(p) = \frac{\sum_{t=1}^{48} I_{i,t}^{SC}(p)}{48}$$

where $I_{i,t}^{SC}(p)$ is a binary index equal to 1 if subject i 's guess of the state in period 1 of round t coincided with their period 1 signal, and zero if subject's guess was the opposite of their signal, in an environment p .²²

²² Note also that a subject's expected payoff in period 1 is linear in her compliance rate.

Table C.1

Subject-specific signal compliance (SCR_i). Left: random effects tobit censored at upper limit of 1 (with bootstrapped errors), right: random effects OLS (with robust errors clustered on groups), without and with session dummies. (Robust) standard errors in parentheses, Erratic environment ($p = .6$) as base. (* p -value < 0.10, ** p -value < 0.05, *** p -value < 0.01.).

SCR_i	Random-Effects Tobit				Random-Effects OLS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$p = 0.9$	0.065*** (0.025)	0.061** (0.029)	0.064** (0.027)	0.061* (0.031)	0.010* (0.005)	0.009* (0.005)	0.010** (0.005)	0.009* (0.005)
$p = 0.1$	−0.003 (0.022)	−0.009 (0.022)	0.002 (0.027)	−0.006 (0.026)	0.002 (0.005)	0.001 (0.005)	0.002 (0.006)	0.001 (0.006)
Part2	0.059*** (0.017)	0.310 (0.254)	0.059*** (0.022)	0.305 (0.298)	0.008** (0.004)	0.075 (0.073)	0.008** (0.004)	0.073 (0.073)
q_i^{emp}		0.584** (0.258)		0.576* (0.300)		0.117* (0.068)		0.113 (0.070)
$Part2 \times q_i^{emp}$		−0.378 (0.371)		−0.370 (0.422)		−0.098 (0.104)		−0.095 (0.104)
_cons	1.100*** (0.032)	0.703*** (0.179)	1.088*** (0.067)	0.700*** (0.209)	0.971*** (0.005)	0.891*** (0.047)	0.979*** (0.009)	0.901*** (0.050)
Session Controls	No	No	Yes	Yes	No	No	Yes	Yes
chi2	14.24	21.61	25.28	31.33	10.04	24.67	112.82	657.33
p	0.00	0.00	0.00	0.00	0.02	0.00	0.00	0.00
N	288	288	288	288	288	288	288	288

Table C.2

Summary statistics for “between” period 2 optimal information choice rate $IOR_i(p)$ (i.e., subject-specific period 2 optimal information choice rates by each subject per each environment, out of 48 rounds). The numbers of observations are as in Table 3.

p	Persistence	Information Optimality Rate per environment, $IOR_i(p)$					
		Part 1		Part 2		Both Parts Pooled	
		Mean	(SD)	Mean	(SD)	Mean	(SD)
$p = 0.9$		0.774	(0.360)	0.804	(0.351)	0.789	(0.354)
$p = 0.6$		0.790	(0.319)	0.782	(0.370)	0.786	(0.344)
$p = 0.1$		0.752	(0.395)	0.819	(0.355)	0.786	(0.374)
Pooled		0.776	(0.347)	0.797	(0.360)	0.787	(0.353)

As mentioned earlier, order effects in signal compliance rates are not important as long as the compliance rate exceeds the cutoff value, which would be particularly important in non-Erratic environments. We find that subjects' signal compliance rate SCR_i responds to the empirical period 1 signal accuracy q_i^{emp} in random effects regressions, both using tobit regressions with censoring at the upper limit of 1, and OLS regressions (see Table C.1), and that controlling for this accuracy in each part eliminates any order effect, namely whether the Erratic environment occurs in the first part (baseline) or the second part (dummy variable labeled “Part 2.”)²³ Relative to the (baseline) Erratic ($p = 0.6$) environment, subjects appear to follow their period 1 signal marginally more often in Persistent ($p = 0.9$) environment, which is important for accuracy of social information. Session controls further weaken the above effects. Overall, while the order effects and learning may exist for the signal compliance rate, they are neither systematic nor robust. More importantly, order effects in signal compliance are not important in our setup.

Result C1. There is no robust evidence for any order effects in signal compliance rates across treatments.

C2. Period 2 information choice, order effects

We now define the Individual Optimality Rate $IOR_i(p) \in [0, 1]$ for a given persistence environment, p , which is a subject-specific optimal information choice rate across 48 rounds of that environment:

$$IOR_i(p) = \frac{\sum_{t=1}^{48} I_{i,t}^O(p)}{48}$$

Using this “between-subjects” measure, we find no systematic order effect in the main task of information selection (see summary statistics in Table C.2).

Fig. C.1. Cumulative distributions (population shares) of period 2 optimal information choice rate $IOR_i(p)$ for all persistence/treatment combinations.

Fig. C.1 presents the cumulative distributions (population shares) of optimal information chosen in each persistence environment, in the relevant treatments. We find a statistically significant difference only for the pair of treatments $Tr=61$ and $Tr=96$, equal to 0.361, with one-sided exact Kolmogorov-Smirnov p -value of 0.009, though the difference in the two means is only marginally different (one-tailed t -test $t=1.367$, p -value=0.088). In all other pairwise comparisons, the differences are not statistically significant.

Result C2. There is no systematic order effect in the period 2 optimal information choice rate.

C3. Period 2 information choice, between- vs. within-subjects optimality measures

As noted in the main text, using the overall, within-subjects optimality rate over 96 periods, we find a surprisingly low degree of rationality, for instance, just 21.5% of subjects make perfectly optimal information choices. By contrast, the “between-subjects” optimal information choice rates, $IOR_i(p)$ in [Table C.2](#) suggest that there is a high level of optimal information choice, with rates all exceeding 75%. That is, if we looked only at a between-subjects measure of optimal information choice, we obtain an artificially inflated measure of the frequency of rational, period 2 information choices.

Further, there does not appear to be any systematic differences in optimal information choice rates among the three persistence environments. Indeed, when comparing cumulative distributions within the same part but across different environments, the Kolmogorov-Smirnov two-tailed test for two independent samples fails to find any significant difference. If we ignore the within-subjects design (and thus potential non-independence) and pool all observations for each environment, this test again fails to find any pairwise differences among the three environments. This leads us to the following important result.

Result C3. Subject-specific period 2 optimal information choice rates are distributed similarly across the three persistence environments.

Furthermore, as [Fig. C.2](#) (left panel) shows, the modal IOR_i is at 1. That is, out of 72 “between” subjects (or subject-part observations) in the *Erratic* ($p = 0.6$) environment, 44% involve choices of private information 100% of the time; while out of 36 subject-part observations in the *Persistent* ($p = 0.9$) and *Anti-Persistent* ($p = 0.1$) environments 50% and 54%, respectively, involve choices of social information 100% of the time. Across the three environments, 138 out of 288 (47.9%) “between” subjects make optimal information choice 100% of the time, with 74.6% of subjects making at least 75% of optimal choices.

Interestingly, despite the high prevalence of perfectly optimal behavior, there is notable heterogeneity in optimal information choice rates. Specifically, 8% of “between” subjects made 100% *suboptimal* choices - that is, 100% choices of social information in the *Erratic* ($p = 0.6$) environment, and 100% choices of private information in the *Persistent* ($p = 0.9$) and *Anti-Persistent* ($p = 0.1$) environments. Overall 45 out of 288 (15.6%) “between” subjects made at least 75% of suboptimal choices. These tendencies for optimal and suboptimal information choices can be further observed when one explores the cumulative frequency distributions (population shares) as in the right panel of [Fig. C.2](#).

Result C4. Across the three environments, almost half of the “between” subjects made perfectly optimal information choices, with almost three quarters of these subjects making optimal information choices at least 75% of the time. However, more than 15% of the “between” subjects made suboptimal information choices at least 75% of the time including 8% of subjects who made perfectly suboptimal information choices.

Thus, our “between-subjects” analysis suggests an interesting pattern. There appears to be a remarkably high average level of optimal information choice accompanied by a substantial rate of suboptimal choices by some subjects. Yet, as we pointed out earlier in [Section 4.2](#), in our “between” subject analysis, errors can run only in one direction. Instead, as we showed in our “within” subjects analysis in the text, some subjects follow (suboptimal) payoff-independent strategies, but

²³ The subject-specific empirical accuracy of the period 1 signal q_i^{emp} has a mean (st.dev.) of 0.696 (0.065) with a range of [0.479, 0.875]. Two subjects experienced uninformative signals, with q_i^{emp} of 0.479 and 0.5, but their compliance rates SCR_i were 1 and 0.896, respectively.

Fig. C.2. Frequencies (left) and cumulative distributions (right) of “between” subjects’ optimal information choice rates $OIR_i(p)$ in each of the three p environments, overlayed (as proportions out of 48 rounds). The vertical axes report (cumulative) frequencies, and the horizontal axes - rates of “between”-subject-specific information choices (out of 48 rounds)..

this suboptimality cannot be detected by the “between” subjects analysis simply because their payoff-independent choices happened to coincide with optimal choices.

As the above results together with Result 2 suggest, that there could be important differences between the “within-subjects” and “between-subjects” estimates of the optimality of information choices. In the above between-subjects analysis, we explored how often each individual subject chose information optimally in a *single environment*. This approach, however, pools together those subjects who select information optimally with those subjects who made optimal information choices simply by accident as their bias happened to match a particular persistence environment. In contrast, in our within-subjects analysis, we look at how often each individual subject chose information optimally in a *pair of contrasting environments*, where, by construction, the benefits to any given bias in a “matching” persistence environment are canceled by out by losses from a “mismatching” persistence environment, thus separating choices driven by subjects’ rationality from those driven by their biases. If each individual chose equally optimally across the two environments, there would be no difference between the two measures. Thus, the fact that a difference is observed indicates that a substantial number of subjects perform consistently better in one environment than in the other.

Fig. C.3 compares the cumulative frequency distribution (population shares) of “between” optimal information choice rates for each subject-part, pooled over all three environments with the cumulative frequency distribution of “within” optimal information choice rates for each subject for the pair of environments. Notice that according to the “within” optimality measure only 21.5% of subjects are perfectly optimal in their information choice across *both* persistence environments. By contrast, as reported earlier, using the between optimality measure, 47.9% of subjects are found to be perfectly optimal, more than twice bigger than the population share found using the within subject optimality measure. As noted in Section 4.2, 13.2% of subjects make exactly half optimal choices and half suboptimal choices. That is, more than half of the subjects who were 100% optimal in one environment made a significant proportion (up to 100%) of suboptimal information choices in the other environment, leading to the following result.

Result C5. The “between-subjects” measure of optimal information choice overstates optimal and understates suboptimal behavior relative to the “within-subjects” measure.

Support for Result C5 is immediate from Fig. C.3. As, by construction, the two measures are simply two different partitions of the same data set, the “between” and “within” mean optimal information choice rate is identical at 0.787. However, the “between” information optimality rate is “noisier” (or more dispersed) with the “between” measure having a standard deviation of 0.353, as compared with the 0.217 standard deviation of the “within” measure.

C4. Period 2 information use: by expected payoff rank and by information source

Table C.3 gives the overall frequencies of strategy choices, where the strategies are labeled S for social, P for private, F for follow the signal received and N for not following. For example, the strategy SF is to choose social information and to follow it, or go against it. The strategies are ordered in terms of their expected payoff, with I being the equilibrium strategy in that environment and IV being the worst.

Furthermore, in the aggregate, *conditional on choosing the optimal information*, that information was used correctly 95.3% of the time. Even in the *Anti-Persistent* ($p = 0.1$) environment, where the equilibrium strategy is to go against the previous period majority, subjects followed this strategy 96.1% of the time. Note, however, that in the *Erratic* ($p = 0.6$) environment where the equilibrium strategy was simply to choose private information and follow it, the frequency of following this

Fig. C.3. Cumulative distributions of “between-subjects” period 2 optimal information choice rates IOR_i (pooled over all 3 persistence environments and both parts, 288 individual observations) and “within-subjects” optimal information choice rate (IOI_i) (both parts combined, 144 unique subject observations). The vertical axes report cumulative frequencies, and the horizontal axes - subject-specific optimal information choice rates (as proportions out of 48 rounds for “between” measures, and out of all 96 rounds for “within” measures).

Table C.3

Table of summary statistics for choice of Period 2 payoff-ranked strategies by environment and by part. N is the number of subjects and mean (SD) is the percentage rate of the number of times (out of 48) that each strategy was chosen.

PersistenceStrategy		Part 1			Part 2			Total		
		N	Mean	(SD)	N	Mean	(SD)	N	Mean	(SD)
$p = .9$	SF (I)	36	0.741	(0.350)	36	0.781	(0.358)	72	0.761	(0.352)
	PF (II)	36	0.180	(0.285)	36	0.164	(0.307)	72	0.172	(0.294)
	PN (III)	36	0.046	(0.106)	36	0.031	(0.087)	72	0.039	(0.097)
	SN (IV)	36	0.032	(0.051)	36	0.023	(0.061)	72	0.028	(0.056)
$p = .6$	PF (I)	72	0.742	(0.328)	72	0.740	(0.364)	144	0.741	(0.345)
	SF (II)	72	0.191	(0.303)	72	0.207	(0.355)	144	0.199	(0.329)
	SN (III)	72	0.019	(0.029)	72	0.011	(0.025)	144	0.015	(0.027)
	PN (IV)	72	0.048	(0.084)	72	0.042	(0.087)	144	0.045	(0.086)
$p = .1$	SN (I)	36	0.719	(0.388)	36	0.792	(0.351)	72	0.755	(0.369)
	PF (II)	36	0.200	(0.326)	36	0.172	(0.348)	72	0.186	(0.335)
	PN (III)	36	0.047	(0.094)	36	0.009	(0.029)	72	0.028	(0.072)
	SF (IV)	36	0.034	(0.035)	36	0.027	(0.036)	72	0.030	(0.036)
Pooled	I (Best)	144	0.736	(0.347)	144	0.763	(0.357)	288	0.750	(0.352)
	II	144	0.191	(0.302)	144	0.188	(0.340)	288	0.189	(0.321)
	III	144	0.033	(0.075)	144	0.016	(0.049)	288	0.024	(0.064)
	IV (Worst)	144	0.041	(0.067)	144	0.034	(0.071)	288	0.037	(0.069)

optimal information was 94.3%, which is still very high but slightly lower than in the other environments, and also lower than the frequency with which subjects followed their period 1 signal, the “between” Signal Compliance Rate, SCR_i (see Table 3).

Interestingly, as Fig. C.4 shows, the rates of correct use of private information by 135 subjects who chose it at least once is both more variable across subjects, and lower than the similar figure for 136 choosers of social information, but there are far more subjects who always followed private, rather than social, information. Specifically, the mean (st. dev.) and median of optimal use rate of private information are 0.916 (0.162) and 1, while the corresponding numbers for social information use are 0.933 (0.144) and 0.975. Interestingly, optimal use of social information is negatively correlated with Lone Wolf Index (LWI_i) ($r = -0.3317$, p -value = 0.0001), indicating that herd animals are more likely to comply with the social information they chose to observe. In contrast, such correlation is absent for optimal use of private information ($r = 0.0677$, p -value = 0.4354).

Fig. C.4. Overall frequencies of correct use of private and social information, pooled across all environments. Left: Cumulative distributions. Right: Scatter plots (with jitter) of information use rates versus Lone Wolf Index. 9 and 8 subjects who never chose private and social information, respectively, are excluded.

Table C.4

Rates of correct use of optimal and suboptimal information, for each persistence environment. Numbers of subjects reflect those subjects who chose optimal or suboptimal information at least once.

Persistence	Use of Optimal Info				Use of Suboptimal Info			
	Definition	#Subj.	Mean	St.Dev.	Definition	#Subj.	Mean	St.Dev.
$p = 0.9$	SF_{+SN}	66	0.947	(0.141)	PF_{+PN}	36	0.857	(0.224)
$p = 0.6$	PF_{+PN}	132	0.927	(0.144)	SF_{+SN}	81	0.856	(0.261)
$p = 0.1$	SN_{+SF}	67	0.923	(0.157)	PF_{+PN}	33	0.788	(0.305)
Pooled	SF_{+SN}	144	0.943	(0.091)	PF_{+PN}	113	0.859	(0.209)

Result C6. Compliance use of chosen information depends on the type of information (private or social), with herd animals more likely to follow chosen social information.

Furthermore, as Table C.4 demonstrates, the rates of correct use of optimal information exceed the rates of correct use of suboptimal information. That is, those subjects who choose suboptimal information, also tend to use it incorrectly. As we will see in Section 4.3, both the incorrect choice of information and incorrect use of information contribute to subjects' expected payoffs.²⁴

Result C7. For each of the three environments, the rates of optimal use of optimal information source exceed the rates of optimal use of suboptimal information source.

C5. "Between-Subjects" expected payoff

Here we examine the subject-specific "between" expected period 2 payoff per environment, denoted by $\pi_i(p)$, which is defined as the sum of payoffs to each of the four strategies defined in Table 4, weighted by the proportion of times that a subject used that strategy in period 2 in each part (and thus in each persistence environment):

$$\begin{aligned}\pi_i(p = 0.9) &= \frac{\sum_{t=1}^{48} [0.7I_{i,t}^{PF}(p = 0.9) + 0.3I_{i,t}^{PN}(p = 0.9) + 0.821I_{i,t}^{SF}(p = 0.9) + 0.179I_{i,t}^{SN}(p = 0.9)]}{48} \\ \pi_i(p = 0.6) &= \frac{\sum_{t=1}^{48} [0.7I_{i,t}^{PF}(p = 0.6) + 0.3I_{i,t}^{PN}(p = 0.6) + 0.580I_{i,t}^{SF}(p = 0.6) + 0.420I_{i,t}^{SN}(p = 0.6)]}{48} \\ \pi_i(p = 0.1) &= \frac{\sum_{t=1}^{48} [0.7I_{i,t}^{PF}(p = 0.1) + 0.3I_{i,t}^{PN}(p = 0.1) + 0.179I_{i,t}^{SF}(p = 0.1) + 0.821I_{i,t}^{SN}(p = 0.1)]}{48}\end{aligned}$$

where $I_{i,t}^j(p)$ is a binary index equal to 1 if subject i followed strategy $j \in \{PF, PN, SF, SN\}$ in round t given the environment p and zero otherwise. The maximum "between" expected payoff $\pi^{\max}(p)$ thus varies across environments, and thus $\pi_i(p) \in [1 - \pi^{\max}(p), \pi^{\max}(p)]$, where $\pi^{\max}(p)$ is the theoretical maximum listed in Table C.5.

²⁴ Our finding that those subjects who rely on suboptimal information also more likely not to follow it underscores the difficulties of identifying information source biases in standard binary choice designs, as these subjects might generate false positives.

Table C.5

Summary statistics for “between” subject-specific period 2 expected payoff $\pi_i(p)$ (averaged over 48 rounds), and the share of the corresponding maximum theoretical payoff $\pi(p)_{\max}$, in each environment p . The numbers of observations are as in Table 3.

Persistence	“Between” Expected Payoff in Period 2 per Environment $\pi_i(p)$								
	$\pi(p)_{\max}$	Part 1			Part 2			Both Parts Pooled	
		Mean	(SD)	$\% \pi(p)_{\max}$	Mean	(SD)	$\% \pi(p)_{\max}$	Mean	(SD)
$p = 0.9$	0.821	0.754	(0.079)	91.8%	0.770	(0.085)	93.8%	0.762	(0.082)
$p = 0.6$	0.7	0.653	(0.053)	93.3%	0.655	(0.053)	93.6%	0.654	(0.053)
$p = 0.1$	0.821	0.750	(0.079)	91.4%	0.778	(0.053)	94.8%	0.764	(0.068)
Pooled	0.7605	0.702	(0.084)	92.4%	0.715	(0.086)	93.9%	0.709	(0.085)

Table C.6

Summary statistics for overall “within” variables, Information Optimality Index IOI_i and Pay-off Optimality Index POI_i , and Lone Wolf Index LWI_i , in each of the four treatments.

Treatment	IOI_i		POI_i		LWI_i		N
	Mean	(S.D.)	Mean	(S.D.)	Mean	(S.D.)	
69	0.805	(0.215)	0.712	(0.062)	0.001	(0.529)	36
96	0.822	(0.201)	0.713	(0.051)	0.097	(0.514)	36
61	0.796	(0.201)	0.714	(0.032)	−0.045	(0.534)	36
16	0.723	(0.245)	0.695	(0.054)	−0.058	(0.655)	36
Total	0.787	(0.217)	0.709	(0.051)	−0.001	(0.558)	144

Fig. C.5. Cumulative distributions (population shares) of the overall “within” variables, Information Optimality Index IOI_i , Payoff Optimality Index POI_i , and Lone Wolf Index LWI_i , in each of the four treatments..

Recall that, as Table 4 shows, in the *Erratic* ($p = .6$) environment, the theoretical maximum payoff is lower and the theoretical minimum is higher than in the other two environments. This is reflected in the empirical findings of Table C.5, which shows that, on average, subjects earned similar proportions of the maximum payoff across the three environments, losing, on average, just under 7% of the theoretical maxima. However, as we show in the next section, there is strong effect of subjects’ biases on their payoffs.

Result C8. In period 2, on average over both persistence environments, subjects achieved 93.2% of the maximum theoretical expected payoff.

One can calculate the overall subject-specific Payoff Optimality Index POI_i (defined earlier in Section 4.3) in an alternative way, as the expected period 2 “between” payoffs $\pi_i(p)$, averaged over both parts of the experiment:

$$POI_i = \frac{1}{2} \pi_i(p = 0.6) + \frac{1}{2} \pi_i(p \neq 0.6)$$

C6. “Within-Subjects” variables, order effects

In this section we show that there are no systematic differences across treatments, allowing us to pool together all overall (“within”) observations.

As one can see in Table C.6, the mean of information optimality IOI_i is significantly higher in only one pairwise comparison (96 vs 16, one-tailed t -test $t = 1.873$, p -value = 0.033). The cumulative distributions (population shares) of IOI_i in these two treatments also markedly different in Fig. C.5, with the largest distance of 0.306 significant according to the one-sided exact Kolmogorov-Smirnov p -value of 0.035. However, this difference cannot be explained by the order effect, as the *Erratic* ($p = 0.6$) environment appears in the second part. A further marginal difference in means (61 vs 16, one-tailed t -test $t = 1.384$, p -value = 0.085) is not captured by the exact Kolmogorov-Smirnov p -value of 0.105.

Table D.1

Random-effects logit regressions on a dummy variable $Private_{i,t}$, indicating whether subject i chose private information in round t , for all 288 subject-part “individuals”, without and with session controls. Robust standard errors (in parentheses) clustered at the group level. (* p -value < 0.10, ** p -value < 0.05, *** p -value < 0.01.)

$Private_{i,t}$	(1)	(2)	(3)	(4)	(5)	(6)
$p = 0.9$	−9.707*** (1.512)		−10.495*** (0.815)	−10.572*** (1.373)		−11.277*** (0.884)
$p = 0.1$	−9.961*** (1.233)		−10.860*** (0.800)	−8.705*** (1.039)		−9.432*** (0.966)
t	0.041*** (0.012)		0.041*** (0.004)	0.041*** (0.012)		0.041*** (0.004)
$p = 0.9 \times t$	−0.017 (0.018)		−0.017** (0.008)	−0.018 (0.018)		−0.017** (0.008)
$p = 0.1 \times t$	−0.023 (0.020)		−0.022** (0.009)	−0.023 (0.020)		−0.022** (0.009)
$Private_{i,t-1} > Social_{i,t-1}$		0.344*** (0.107)	0.370*** (0.110)		0.344*** (0.107)	0.369*** (0.110)
$Private_{i,t-1} < Social_{i,t-1}$		−0.255** (0.107)	−0.290*** (0.109)		−0.254** (0.107)	−0.290*** (0.109)
_cons	3.826*** (0.634)	−1.769*** (0.238)	4.167*** (0.467)	4.619*** (0.923)	−0.975 (0.845)	5.088*** (0.911)
Session Controls	No	No	No	Yes	Yes	Yes
chi2	81.70	21.41	379.57	114.51	48.96	406.30
p	0.00	0.00	0.00	0.00	0.00	0.00
N	13,824	13,536	13,536	13,824	13,536	13,536

Furthermore, we find no systematic differences in the Payoff Optimality Index POI_i across different treatments. The mean of this index is significantly higher only in one pairwise comparison (16 vs 61, one-tailed t -test $t = 1.849$, p -value = 0.034), but even that difference is not captured by the exact Kolmogorov-Smirnov p -value of 0.105. The largest difference between the cumulative functions in Fig. C.5 is 0.333 (between treatments $Tr = 69$ and $Tr = 16$), with one-sided exact Kolmogorov-Smirnov p -value of 0.036 - however this difference cannot be explained by the order effect. The difference in means between these two treatments is only marginally different (one-tailed t -test $t = 1.401$, p -value = 0.083).

Finally, we find no systematic differences in the Lone Wolf Index LWI_i across different treatments. Again, the largest difference between the cumulative functions in Fig. C.5 is 0.306 (96 vs 61), significant according to one-sided exact Kolmogorov-Smirnov p -value of 0.035. However, again, this difference cannot be explained by the order effect.

Result C9. There is no systematic treatment effect in the overall “within” variables, Information Optimality Index IOI_i , Payoff Optimality Index POI_i , and Lone Wolf Index LWI_i .

D. Dynamics of Period 2 behavior

Recall that at the end of each round, subjects were given feedback about both the information they chose, and the other piece of information that they did not choose to view, and were also informed about the true state of the world, so in principle they could weigh the accuracies of the two different types of information. We will now explore the effect of feedback, as well as other dynamical aspects of subjects behavior.

To preview the results, while some subjects adjust their choices, there is very little overall change. First, 48 out of 144 subjects never change their information choices within each part, including those who are lone wolves and herd animals. Second, among those who do change, the overall number of unbiased subjects does not rise over time (Result D3). Some subjects learn to choose optimal actions more frequently, but others become less optimal over time.

D1. Dynamics of period 2 strategies (“between-subjects”)

Regressions in Table D.1 reveal that subjects are affected by both introspection about the persistence environment and by the feedback they received about the two information sources they could choose between.²⁵ First, subjects’ choice of private information is clearly affected by the persistence environment in which they were placed, as evidenced by the significant coefficients on the persistence environment dummy variables, $p = 0.9$ and $p = 0.1$ (the baseline environment is $p = 0.6$). Second, their choice of private information is also affected by the predictive accuracy of private relative to social information in the immediate preceding period. When private (social) information strictly outperforms social (private) information in the

²⁵ The results are qualitatively similar if one excludes 161 out of 288 (55.9%) “committed” subject-part “individuals” (explored further in Section D.2), who always choose the same type of information in all 48 rounds, and consider only the 127 out of 288 (44.1%) “adjusting” subject-part “individuals”, (explored further in Section D.3), who changed their choice of information at least once (results available on request).

Fig. D.1. Dynamics of information choice and use in each persistence environment. Proportion of subjects who have settled on a strategy versus time period. SF=follow social info, PF=follow private info, SN=go against social info, and PN=go against private info..

preceding period ($Private_{i,t-1} > (<) Social_{i,t-1}$), then private information is more (less) likely to be chosen.²⁶ Finally, subjects are also more likely to choose the correct information source over time, as evidenced by the positive coefficient on round number t (for the baseline $p = 0.6$ environment) and the environment specific interactive time trend variables $p = 0.9 \times t$ and $p = 0.1 \times t$.

Result D1. Subjects respond both to the persistence structure of the environment, p , and to feedback given about the different types of information they could choose to view. In all three environments, there is a trend toward more optimal information choice over time, i.e., toward the choice of private information in the Erratic ($p = .6$) environment, and toward the choice of social information in the other two environments.

Fig. D.1 further shows the speed with which subjects “settle” into a particular information choice, “social” versus “private”, and information use strategy, in the form of the four strategies “SF”, “PF”, “SN” and “PN” as in Table 4. These graphs are cumulative over time and they give the proportion of subjects who, starting from round t , have settled on a strategy choice, never subsequently changing their strategy.²⁷

The left panel of Fig. D.1 shows that, when $p = 0.9$, already in the period 1 approximately half the subjects have settled on choosing social information. However, only about half of these have settled on the correct use of that social information. Both proportions rise over time, with the gap between the two decreasing. In the middle panel, when $p = 0.6$, there is a similar pattern of the choice and use of private information. There is also a similar pattern in the last panel, when $p = 0.1$, though it takes time before subjects settle into (correctly) going against the social information they chose to see. So it seems that which information to choose is clearer to subjects than whether to follow it. Alternatively, subjects might be more prepared to experiment on whether to follow information than which information to choose.

To further explore the effect of past experience, we consider whether subjects use simple, naive best-response strategies of the “Win-Stay, Lose-Shift” variety (though we recognize that they could employ more sophisticated strategies as well). Specifically, we ask: if subject i ’s guess of the period 2 state at time $t - 1$ was correct (“Win $_{i,t-1}$ ”), would this subject’s information use strategy at time t be the same (“Stay $_{i,t}$ ”) Table D.2 indicates that winning is indeed associated with keeping the same strategy in the following round. However, as the round t variable indicates, it also appears that subjects settle on a particular (optimal or suboptimal) strategy with time as well.

Result D2. The dynamics of subjects’ information use exhibits both “settling down” and “Win-Stay, Lose-Shift” behavior.

D2. Signal experiences of the pure types

Do subjects’ individual signal histories affect their self-selection into a particular player type? We note first that 48 of our 144 subjects, a third of the total, never switch their information choice *within* each of the two parts, i.e., they were “committed” to choosing a particular information in each part. 31 (65% of these committed 48 subjects) are fully rational – always choosing private information in the part with the Erratic ($p = 0.6$) environment and always choosing social information in the other part. By contrast, 8 (17% of the committed subjects) are apparent “pure” herd animals and always choose social information, while 9 (19%) are “pure” lone wolves. Thus, of the subjects who never adjust their information choice, most are rational and the distribution of bias is approximately symmetric (see Fig. 3).

We now check whether individual histories might be responsible for creating these “pure” player types. For example, pure “lone wolves” might have rationally opted to always choose private information simply because, by chance, their sequence of private information draws was more accurate than past social information in the non-Erratic environments. Similarly, pure “herd animals” might have rationally opted to always choose social information because the private information they drew happened to be less accurate in the Erratic environment. If there was such a history-dependent explanation for emergence of the pure types, then the group of 8 pure “lone wolves” would have more accurate private ball draws as compared with the

²⁶ We need more than one dummy variable for this distinction since there are instances where both sources of information are equally correct or equally incorrect.

²⁷ These graphs were inspired by the graphing of convergent behavior in Esponda and Vespa (2018).

Table D.2

Random-effects logit regression (without and with session dummies) of $Stay_{i,t}$, which is a dummy variable indicating whether subject i 's information use strategy at time t is the same as in the prior round, $t - 1$. Robust standard errors (in parentheses) have been clustered at the group level. (* p -value < 0.10, ** p -value < 0.05, *** p -value < 0.01.)

$Stay_{i,t}$	(1)	(2)	(3)	(4)	(5)	(6)
$p = 0.9$	0.107 (0.403)		0.064 (0.386)	-0.172 (0.405)		-0.229 (0.387)
$p = 0.1$	0.028 (0.330)		-0.064 (0.337)	0.323 (0.356)		0.245 (0.369)
t	0.027*** (0.007)		0.028*** (0.007)	0.027*** (0.007)		0.028*** (0.007)
$p = 0.9 \times t$	-0.006 (0.008)		-0.008 (0.009)	-0.006 (0.008)		-0.008 (0.009)
$p = 0.1 \times t$	0.001 (0.010)		-0.001 (0.010)	0.001 (0.010)		-0.001 (0.010)
$Win_{i,t-1}$		1.033*** (0.080)	1.030*** (0.081)		1.034*** (0.081)	1.031*** (0.081)
_cons	2.217*** (0.249)	2.185*** (0.167)	1.602*** (0.263)	2.213*** (0.405)	2.029*** (0.410)	1.609*** (0.399)
Session Controls	No	No	No	Yes	Yes	Yes
chi2	57.48	164.98	221.26	75.46	181.51	261.93
p	0.00	0.00	0.00	0.00	0.00	0.00
N	13,536	13,536	13,536	13,536	13,536	13,536

Fig. D.2. Evolution of the variable $CumPrivateBetter_{i,t}$, representing the cumulative experience with greater accuracy from private rather than social information, averaged over pure subject types.

group of 9 pure “herd animals” and the group of 31 pure “rational” subjects. Similarly, the group of 9 pure “herd animals” would have less accurate private ball draws as compared with the group of 8 pure “lone wolves” and the group of 31 pure “rational” subjects.

For each subject in each round, we constructed an index that had a value of 1 when private information outperformed social, a value of -1 when social information outperformed private, and a value of 0 when both pieces of information were either correct or wrong – similar to the procedure followed in Section D.1. We then calculated the *cumulative* value of this index as each part of the experiment (*Erratic* or *Persistent/Anti-Persistent*) proceeded. Finally, we constructed the average cumulative index of private information being the better predictor, averaged over a particular subject type, pure lone wolf, pure herd animal or pure rational type. We focus on pure types as these provide the least ambiguous cases for which the history of signal accuracies could have played a role.

The average cumulative indices of the superiority of private information in predicting the state of the world are plotted in Fig. D.2, representing the experiences of the pure types, as a group. These figures suggest that the experience of the 8 pure “herd animals” and the 31 pure “rationals” were similar. In the *Erratic* environment, the superiority of private information is increasing over time while in the other two environments it is decreasing over time. Notice that in the early rounds of the *Erratic* environment, 9 pure “lone wolves”, as a group, actually experienced a slightly worse performance from private draws than from a choice of social information, which works against the hypothesis that history matters for the creation of types. (Similar figures disaggregated by subject are available on request.)

The panel regression in Table D.3 confirms that private information becomes cumulatively better in the *Erratic* ($p = 0.6$) environment, and cumulatively worse in the non-*Erratic* ($p \in \{0.1, 0.9\}$) environments, as evidenced by the significant positive and negative coefficients, respectively, on the time trend, t . However, there are no differences in the experiences of the pure types, as evidenced by non-significant interaction of the time trend variable with each pure type index. Thus, the decisions of the 8 pure “herd animals” and, in particular, of the 9 pure “lone wolves” to always choose, respectively, social and private information in all environments cannot be explained by their signal experiences within our experiment – since these experiences are not statistically different from those of the 31 pure “rational” subjects.

Table D.3

Fixed-effects panel regression on a variable $CumPrivateBetter_{i,t}$, which is a cumulative experience of more accurate private information than social information, for subject i in round t , for 48 “pure type” subjects (pure “rational” subjects are a baseline). Robust standard errors (in parentheses). (* p -value < 0.10, ** p -value < 0.05, *** p -value < 0.01.)

$CumPrivateBetter_{i,t}$	$p = 0.6$	$p \in 0.1, 0.9$
t	0.082*** (0.015)	-0.101*** (0.016)
$LoneWolf \times t$	-0.052 (0.049)	-0.053 (0.033)
$HerdAnimal \times t$	0.002 (0.028)	0.000 (0.031)
_cons	-0.100 (0.326)	-0.148 (0.308)
F	14.16	27.28
p	0.00	0.00
N_g	48	48
N	2304	2304

Table D.4

Random-effects logit regressions on a dummy variable $Private_{i,t}$, indicating whether subject i chose private information in round t for all “non-pure” type 192 subject-part “individuals”, by their broad type, without and with session controls. Robust standard errors (in parentheses) clustered at the group level. (* p -value < 0.10, ** p -value < 0.05, *** p -value < 0.01.)

$Private_{i,t}$	(1) All	(2) UB	(3) LW	(4) HA	(5) All	(6) UB	(7) LW	(8) HA
$Private_{i,t-1} >$	0.346*** (0.106)	0.521*** (0.132)	-0.378 (0.241)	0.442* (0.250)	0.345*** (0.106)	0.521*** (0.132)	-0.380 (0.241)	0.448* (0.250)
$Private_{i,t-1} <$	-0.252** (0.106)	-0.222* (0.132)	-0.211 (0.237)	-0.542* (0.319)	-0.252** (0.107)	-0.222* (0.132)	-0.204 (0.237)	-0.541* (0.319)
_cons	-0.401 (0.386)	-0.852* (0.446)	3.743*** (0.394)	-4.126*** (0.553)	1.005 (0.992)	-0.659 (1.102)	4.282*** (0.845)	-2.601 (1.716)
Session Controls	No	No	No	No	Yes	Yes	Yes	Yes
chi2	21.73	23.73	2.70	7.88	34.65	28.68	6.27	11.57
p	0.00	0.00	0.26	0.02	0.00	0.00	0.71	0.17
N	9024	5640	1880	1504	9024	5640	1880	1504

Result D3. The self-selection of some subjects to “pure” information choice types cannot be explained by idiosyncratic draws of private signals.

D3. Dynamic adjustment of non-pure type subjects (“within-subjects”)

Let us now turn to the remaining 96 “adjusting” subjects (two thirds of our sample) who changed their choice of information at least once within at least one part of the experiment. Do these adjusting subjects learn to make optimal information choices over time?

To capture possible learning from the provided feedback by these adjusting types we conducted a regression analysis analogous to specifications (2) and (5) of Table D.1 for all of these adjusting subjects as well as those broadly classified as unbiased (UB), lone wolves (LW) or heard animals (HA). As the regressions in Table D.4 demonstrate, broadly unbiased “adjusting” subjects strongly react to history, choosing private information when it outperforms social information and vice versa (see specifications (2) and (6)). In contrast, broad lone wolves have a strong bias towards private information, which is manifested by a significantly positive intercept, and insignificant coefficients on lagged performance (specifications (3) and (7)). The patterns for broad herd animals are more complex, as they both have marginally significant coefficients on lagged performance, and the negative intercept is significant only without session controls (see specifications (4) and (8)). As we will show below, some, but not all subjects who behaved as broad herd animals in the early rounds of each part, turn out to learn over time, which might explain these regression results.

These regression results suggest that learning from past signals largely occurs among broadly unbiased subjects, as well as possibly among a subset of broad herd animals. Note that, because subjects’ types are determined by their information choices rather than being exogenous, this regression analysis is only illustrative, as it utilizes the same choice data twice - first, subjects’ choices of information source were aggregated to determine their broad type, and then these same information choices (though disaggregated) were explored to see whether subjects’ unique signal performance affected their information choices.

Thus, to better understand the evolution of the 96 “adjusting” subjects’ information choices over time, we compare their behavior in the first 6 and the last 6 rounds of each of the two parts. First, we construct two analogues of the Lone Wolf Index LWI_i for the first 6 and the last 6 rounds of each of the two parts, which we denote by $PI6_i^{initial}$ and by $PI6_i^{final}$ (for

Table D.5

Number of the adjusting subjects (who changed their information choice at least once, out of 96 subjects) by Private Information type classification (combined across both parts) using the $PI6_i^{initial}$ index for the first 6 rounds and the $PI6_i^{final}$ index for the final 6 rounds.

		$PI6_i^{initial}$		$PI6_i^{final}$										Total	
		6-Broad Herd Animal				6-Broad Unbiased				6-Broad Lone Wolf				in Initial	
		-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6	6 Rounds
6-Broad Herd Animal	-6	3	0	0	0	0	1	3	0	0	0	0	0	0	7
	-5	4	0	0	0	0	0	0	1	0	0	0	0	0	5
	-4	2	0	0	0	0	1	4	0	0	0	0	0	0	7
	-3	1	0	0	0	1	0	8	0	0	0	1	0	0	11
6-Broad Unbiased	-2	1	1	1	0	0	1	6	0	0	0	0	0	1	11
	-1	1	1	0	0	1	0	7	1	0	0	0	0	0	11
	0	0	0	0	0	0	0	14	0	0	1	0	1	3	19
	1	0	0	0	0	0	0	4	0	0	0	0	0	2	6
6-Broad Lone Wolf	2	0	0	0	0	0	0	4	0	0	0	0	1	4	9
	3	0	0	0	0	0	0	1	0	0	0	0	0	0	1
	4	0	0	0	0	0	0	0	0	0	0	1	1	1	3
	5	0	0	0	0	0	0	0	0	0	0	0	0	5	5
Total in Final 6 Rounds		12	2	1	0	2	3	51	2	0	1	2	3	17	96

Table D.6

Number and percentages of the adjusting subjects (out of 96 subjects) by aggregated 6-broad Private Information type classifications (combined across both parts) using the $PI6_i^{initial}$ index for the first 6 rounds and the $PI6_i^{final}$ index for the final 6 rounds.

Initial 6-Broad Type	Final 6-Broad Type					Total in Initial		
	6-Broad Lone Wolf		6-Broad Unbiased		6-Broad Herd Animal	6 Rounds		
6-Broad Lone Wolf	9	(9.4%)	1	(1.0%)	0	(0%)	10	(10.4%)
6-Broad Unbiased	13	(13.5%)	38	(39.6%)	5	(5.2%)	56	(58.3%)
6-Broad Herd Animal	1	(1.0%)	19	(19.8%)	10	(10.4%)	30	(31.1%)
Total in Final 6 Rounds	23	(24.0%)	58	(60.4%)	15	(15.6%)	96	(100%)

the initial and final number of private information choices):

$$PI6_i^{initial} = \sum_{t=1}^6 I_{i,t}^{PI}(p=0.6) + \sum_{t=1}^6 I_{i,t}^{PI}(p \neq 0.6) - 6$$

$$PI6_i^{final} = \sum_{t=43}^{48} I_{i,t}^{PI}(p=0.6) + \sum_{t=43}^{48} I_{i,t}^{PI}(p \neq 0.6) - 6$$

Subjects with $PI6_i$ values of $\{-6, -5, -4, -3\}$ are classified as “within-6-round-block” broad (“6-broad”) herd animals (since at least 75% of their choices in the corresponding 6-round blocks of each part were for social information); those with $PI6_i$ values of $\{3, 4, 5, 6\}$ are classified as “6-broad” lone wolves (since at least 75% of their choices were for private information), and the remaining subjects with $PI6_i$ values of $\{-2, -1, 0, 1, 2\}$ are classified as “6-broad” unbiased (since at least 75% of their information choices were unbiased).

As Table D.5 shows, while the number of adjusting subjects who made fully unbiased information choices in a 6-round block (i.e. $PI6_i = 0$) increased from 19 (19.8%) to 51 (53.1%), most of this net increase comes from those 21 subjects (21.9%) who started out (i.e., in the initial 6-round block) as broadly but not fully unbiased ($6 + 7 + 4 + 4$), as well as from those 15 subjects (15.6%) who started out as “6-broad” herd animals ($3 + 0 + 4 + 8$). By contrast, 5 subjects (5.2%) who started as fully unbiased moved to becoming “6-broad” lone wolves ($1 + 0 + 1 + 3$), with only one subject transiting in reverse.

Aggregating the number of adjusting subjects by their “6-broad” subject type, Table D.6 further shows that there is only a slight net increase in the number of “6-broad” unbiased subjects over time (from 56 in the initial 6 rounds of each part to 58 in the final six rounds). This is because the net inflow of $19 - 5 = 14$ (14.6%) subjects who moved from initially being “6-broad” herd animals to in the end being “6-broad” unbiased is almost fully offset by the net outflow of $13 - 1 = 12$ (12.5%) subjects who moved from being “6-broad” unbiased to being “6-broad” lone wolves. While in the first 6 rounds there were three times more “6-broad” herd animals (30) than “6-broad” lone wolves (10), by the final 6 rounds there were more broad lone wolves (23) than broad herd animals (15).

Result D4. Comparing behavior in the initial and final 6 rounds of each part, a majority of the 96 “adjusting” subjects remain in the same “6-broad” category, with a plurality of subjects (39.6% of the “adjusting” subjects) maintaining their

Table D.7

Number of the adjusting subjects (out of 96 subjects) making optimal information choices in the initial and final six rounds, (combined across both parts), as classified using the $OIG_i^{initial}$ and OIG_i^{final} indices.

$OIG_i^{initial}$	OIG_i^{final}												Total in Initial		
	0	1	2	3	4	5	6	7	8	9	10	11	12	6 Rounds	
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0	0	0	1	1	
3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
4	0	0	0	0	0	0	1	0	0	0	0	0	0	1	
5	0	0	0	0	0	0	3	0	1	0	0	0	1	5	
6	1	0	0	0	0	0	8	0	0	0	0	1	3	13	
7	0	0	0	0	0	0	6	1	0	0	0	0	2	9	
8	0	0	0	0	0	1	1	1	2	0	0	2	4	11	
9	0	0	0	0	0	0	1	0	1	0	1	0	7	10	
10	0	0	0	0	0	0	3	2	0	0	0	0	9	14	
11	0	0	0	0	0	0	3	1	0	0	0	1	10	15	
12	0	0	0	0	0	0	3	1	0	0	0	0	13	17	
Total in Final 6 Rounds	1	0	0	0	0	1	29	6	4	0	1	4	50	96	

“6-broad” unbiasedness, and almost equal numbers of subjects remaining in each of “6-broad” lone wolf and herd animal classifications (9.4% and 10.4%, respectively). The overall number of subjects who started out being “6-broad” unbiased is almost the same as the number who ended up being “6-broad” unbiased. However, this is because the net inflow of 14.6% of adjusting subjects who moved from being “6-broad” herd animals to being “6-broad” unbiased is similar to the net outflow of 12.5% subjects who moved from being “6-broad” unbiased to being “6-broad” lone wolves.

While most of the fully unbiased choices are fully rational, not all are. One could see this using two analogues of the Information Optimality Index IOI_i for the first 6 and the last 6 rounds of each of the two parts, which we denote by $OIG_i^{initial}$ and OIG_i^{final} , respectively:

$$OIG_i^{initial} = \sum_{t=1}^6 IOI_{i,t}(p = 0.6) + \sum_{t=1}^6 IOI_{i,t}(p \neq 0.6)$$

$$OIG_i^{final} = \sum_{t=43}^{48} IOI_{i,t}(p = 0.6) + \sum_{t=43}^{48} IOI_{i,t}(p \neq 0.6)$$

As Table D.7 shows, there is some improvement in terms of optimal information choice, as 46 subjects (47.9% of “adjusting” subjects) made more optimal decisions in the final 6 rounds relative to the initial 6 rounds, while 25 (27.1%) subjects made more suboptimal decisions - for a net of 21 (21.9%) “adjusting” subjects exhibiting improvement. However, the set of subjects who exhibited improvement in their choice of optimal information includes 27 (28.1%) subjects who made at least 75% optimal information choices (but not fully optimal) even in the first 6 rounds of each part (i.e., those with $OIG_i^{initial} \in \{9, 10, 11\}$), while among the group of those 16 (16.7%) subjects who initially made at least 75% of optimal choices, including fully optimal choices (i.e., those with $OIG_i^{initial} \in \{9, 10, 11, 12\}$) but who failed to improve in their choice of optimal information includes 10 (10.4%) subjects who closer to the end settled on making optimal choices only half of the time, which is consistent with choosing one type of information only.

Result D5. Comparing subjects’ behavior in the initial and final 6 rounds of each part, only a net of 21.9% of “adjusting” subjects exhibited an improvement in optimal information choice.

The dynamics of behavior of the 96 “adjusting” subjects, is summarized by Fig. D.3. In this figure, arrows represent the direction of movement from a point $\{OIG_i^{initial}, PIOI_i^{initial}\}$ in the initial 6 rounds of each part to a point $\{OIG_i^{final}, PIOI_i^{final}\}$ in the final 6 rounds of each part, and the area of the circles is proportional to the number of subjects with a given level of index counts in the final six rounds. Initial choices, as indicated by the beginning point of each of the arrows, are quite dispersed. In contrast, the arrow ends are more concentrated, suggesting that over time, subjects appear to experiment less, settling on a particular strategy. While many arrows point to the far right corner of the Diamond of Rationality (representing an increase in optimal choices over time), not all subjects settle on the correct strategy with experience, as many arrows point in other directions, with significant minorities locking onto lone-wolf-type behavior (top corner) or to herd-animal-type behavior (bottom corner). The lone arrow pointing to the far left corner represents one subject whose initial behavior was broadly unbiased and close to purely random, but whose final behavior was perfectly counter-optimal.

Fig. D.3. Evolution of information choices from the initial six rounds to the final six rounds, combined across both persistence environments (see Tables D.5 and D.7), for “adjusting” subjects (i.e., those who changed their choice of information at least once, total $N = 96$). Circle areas are proportional to the number of “adjusting” subjects with a given sum of frequency counts in the final 6 rounds in both environments (as given by the final row in each of the Tables D.5 and D.7). Arrows start at the sum of counts in the initial 6 rounds in both environments (frequencies of subjects at initial conditions are not shown to reduce clutter), so each arrow may represent more than one subject. Cases of identical initial and final frequency counts at points (6, 6), (6, 6), (8, 4), and (12, 0) are represented by downward pointing arrowheads without shafts.

E. Comparison to existing studies

Here, we ask whether our identification of persistent behavioral types in social learning environments has wider validity beyond this particular setting that we study.

E1. Identification issues with the standard sequential design

We first clarify the reasons why the standard sequential-move design of Anderson and Holt (1997) is not well suited to identify the kind of heterogeneity in social learning behaviour found in our simultaneous-move information choice design.

The main issue with the standard sequential design is its asymmetric nature. Apart from the very early rounds of the sequence, it is almost always optimal to follow the herd and ignore one’s own private signal. Thus, *excessive* herd behaviour, defined as following the herd when it is *not optimal* to do so (which we term here as “herd animal behavior”), is almost excluded *ex ante* by design. Instead, herd animal behaviour found in our experiment would be classified as rational behavior in the standard experiments.

A second problem with the standard design is that only a handful of previous experiments employed sufficient number of different sequential move games to allow for a meaningful within-subjects analysis, and so (as we will show below) the existing data sets are not well-suited for identification of behavioural types. Further, this design generates many realizations that are uninformative because a subject’s private information agrees with the previous history of guesses.

Third, the standard, sequential move, binary choice design cannot distinguish information source bias from irrational misuse of information. When a subject takes an off-equilibrium-path action, one cannot tell whether this is due to the subject’s bias in favor of a less accurate source of information, or this is instead due to the subject irrationally contradicting a more accurate source of information. In a similar vein, the standard design does not distinguish rational behavior from compounded mistakes. For example, in the standard binary setup, if a subject favors less accurate information and irrationally contradicts that less accurate information (thus making both types of mistakes simultaneously), their action would coincide with a rational choice! Thus, the irrational misuse of information that we found in our experimental design would instead be classified as bias – or even as rational behavior – in standard experiments.

The final difference between the standard sequential move setting and our own is, in our opinion, less important. The sequential-move game potentially requires second- or higher-order strategic reasoning. For example, how much information is contained in the action of player 2, given that she also observed player 1’s action prior to acting on her own? By contrast, our simultaneous move design involves a dominance solvable game and thus has little strategic uncertainty if all players

Table E.1

Streamlined subset of metaset of Ziegelmeyer et al. (2013) (ZMK): the list of exclusion criteria and the included observations.

Datasource	Original ZMK		Exclusions		Streamlined	
	NObs	SubjN	NObs	Reasons	NObs	SubjN
Alevy et al. (2007)	1647	109	1,005	$q_A \neq q_B$	642	42
Anderson (2001)	270	18			270	18
Anderson and Holt (1997)	810	54	540	$q_A \neq q_B$	270	18
Cipriani and Guarino (2005)	161	48			161	48
Dominitz and Hung (2009)	1760	90	560	"Replication rounds" retained	1,200	90
Drehmann et al. (2005)	2789	1,840	2,549	$Pr(A) \neq Pr(B) \neq 0.5$	240	240
Fahr and Irlenbusch (2011)	1080	72			1,080	72
Goeree et al. (2007)	8760	400			8,760	400
Hung and Plott (2001)	889	40			889	40
Kübler and Weizsäcker (2004)	482	35			482	35
Nöth and Weber (2003)	9834	126	9,834	q vary with positions	0	0
Oberhammer and Stiehler (2003)	840	36			840	36
Willinger and Ziegelmeyer (1998)	324	36			324	36
Ziegelmeyer et al. (2010)	1440	96	1,440	$Pr(A) \neq Pr(B) \neq 0.5$	0	0
Total	31,086	3,000			15,158	1075

are rational. However, even in our simpler setting, we observe substantial deviations from optimal behaviour and apparently persistent use of heuristics such as always choosing private over social information and vice versa. An obvious hypothesis is that the use of such heuristics would only be more frequent in more complex situations.

Further, a recent experiment by Kawamura and Vlaseros (2017) finds, in a slightly different context, that subjects' sub-optimal behaviour is largely driven by their failure to understand the combined accuracy of multiple independent signals rather than strategic complexity. Thus, we suspect that the use of suboptimal heuristics is widespread in all settings but may not always be observable due to reasons given above.

E2. Exploring the existing meta-data

Given the above mentioned difficulties, can we still utilize the existing data from the standard sequential move game to identify evidence for heterogeneous player types? To address this question, we employ the comparable existing sequential move experimental data compiled by Weizsäcker (2010) and extended by Ziegelmeyer et al. (2013). We follow Weizsäcker (2010) in focusing on the sequential move social learning games with symmetric signal structure (i.e., symmetric priors $Pr(A) = Pr(B) = 0.5$ and position-invariant symmetric precision of private signals $q_A = q_B$) as these designs are most comparable to our own. Using the resulting "streamlined" subset of the Ziegelmeyer et al. (2013) metaset (thereafter ZMK) for the sequential move games (see Table E.1 for details), we pursue the following identification strategy to attempt to detect the "lone wolf" and "herd animal" player types documented in our simultaneous move design.

First, we follow Weizsäcker (2010) and look only at players facing histories of unanimous prior action choices by others, i.e., when all predecessors choose the same action, and thus are in full agreement. This is because the sequential design does not allow a clear interpretation of non-unanimous histories. Indeed, there exist situations where it is optimal to go against the simple majority. For example, consider the sequence of choices *ABA* which is consistent with the first three subjects each optimally having followed their own signal. Then if the fourth subject gets a *b* signal, the equilibrium (with the standard tie-breaking rule) prescribes her to go with her signal thus contradicting the simple majority. Further, the lack of unanimity can be the result of other subjects having taken actions that are off the equilibrium path (e.g., *AAB*), which makes it difficult to determine the optimal action.

Second, we will look at situations where a subject's own private signal differs from the unanimous actions taken by all predecessors. Otherwise, if a subject's action coincides both with their own private signal and the unanimous actions of others, it is not clear whether one should classify that subject as being rational, or a herd animal, or a lone wolf.

Third, we will interpret subjects' choices as if they followed their preferred source of information. This is not an innocuous assumption, as subjects sometimes go against not only the most accurate available piece of information, but also – clearly irrationally – against *all* information available to them, even when all sources of information are in full agreement. Specifically, in the 1831 observations where the subjects' private signal coincided with the unanimous history in the "streamlined" dataset, subjects choose to contradict *both* sources of information in 65 (3.6%) of these observations (the highest proportion of such irrational behavior is 6.0% in position 4, and one subject went both against a 29-subject-long unanimous herd and against own signal which agreed with this very long herd).

Fourth, recall that in the symmetric sequential-move social learning model it is rational to follow one's own signal at early positions in the sequence. But at a certain point in the sequence it becomes optimal for the player to go with the herd and against their own private information. So, if instead a subject sufficiently far from the start of the sequence follows her private signal that differs from the herd choice, we identify that subject as exhibiting (suboptimal) "lone wolf" behaviour.

Fig. E.1. Rationality vs information source bias in the “streamlined” subset of Ziegelmeyer et al. (2013) (ZMK) metadata set of experimental data using the standard design, for 276 subjects who experienced both type identifying situations (as discussed in the text) at least once. Information optimality index (horizontal axis) against Lone Wolf Index (vertical axis). Diamonds represent the 62 subjects who experienced both of the type-identifying situations at least twice, and circles represent the other 214 subjects. The marker size is proportional to the number of subjects with a given pair of indices.

Instead, if a subject contradicts her own signal and follows the predecessors too early in the sequence, we identify that subject as exhibiting (suboptimal) “herd animal” behaviour.

Finally, to separate subject’s rationality from information source bias, we explore each subject’s actions in *both* types of situations – the one where it is optimal for the subject to contradict others and instead follow their own signal, and the other situation where it is optimal to do the opposite – to contradict one’s own signal and instead follow the others. We identify subjects as being a “rational” player type if they consistently switch their preferred source of information depending on the situation they are in. Similarly, we identify subjects as being “lone wolf” or “herd animal” types if, in both types of situations, they consistently follow their private signals or their predecessors, respectively.

Consistent with our identification strategy, we thus look at all cases where subjects face an unanimous history. As we have already argued, the standard sequential-move setting presents little opportunity for purely irrational herd behavior. Indeed, we can identify only one situation where following a herd is clearly irrational – when the private signal of the subject in position 2 contradicts the action of the first subject in the sequence (position 1). Under the refinement of equilibrium due to Banerjee (1992), the subject in position 2 should follow her own signal, which is also the empirically optimal choice, simply because in 8.5% of observations, subjects in position 1 irrationally do not follow their private signal. So following a “herd” consisting of a single subject is a very *strong* measure of suboptimal herd behavior. However, we cannot weaken our measure of excessive herd behavior – as following the unanimous choice made by others is optimal at *any* position beyond position 2.

Thus, all instances where subject i ’s signal contradicted the unanimous history, can be divided into two sets – those $N_{1,i}$ instances at position 2 where subject either optimally follows their private signal or behaves as a strong “herd animal”, and those $N_{2,i}$ instances starting with position 3 where a subject either optimally copies all of her predecessors, or behaves as a “lone wolf”. Using Eqs. (3) and (4) for each subject i with their subject-specific number of instances $N_{1,i}$ and $N_{2,i}$, we then construct the corresponding within-subjects Information Optimality Index (IOI_i^{ZMK}) and Lone Wolf Index (LWI_i^{ZMK}), by looking at the subject-specific fractions (or rates) of such clearly suboptimal lone wolf and herd animal choices among those instances where the private signal contradicts a uniform social history.

It is clear that the previous experiments were not designed to generate such data, as there are relatively few subjects who are observed in both types of the above-mentioned situations that permit clearly suboptimal choices. In our streamlined subset of the metadata, 244 (22.7%) out of the total of 1,075 subjects made only one decision, and thus are excluded by our identification strategy. We identify only 276 (25.7% of the total) subjects who experienced both types of “clear” identifying situations (where they could behave suboptimally) at least once. We also look at a smaller subset of these subjects – only 62 (5.8% of total) – who experienced both situations at least twice.

Fig. E.1 is a sequential-move counterpart of the earlier Figure 3. It presents the corresponding within-subjects measures for sequential move data, where the size of the markers represent the number of subjects with corresponding subject-

Fig. E.2. Rationality versus information source bias in the two-stage sequential move design of Duffy et al. (2019). Information Optimality Index (horizontal axis) against Lone Wolf Index (vertical axis), for each of 128 subjects. The marker size is proportional to the number of subjects with a given pair of indices.

specific indices. The diamonds represent the smaller subset of subjects who experienced each of the two kinds of type-identifying situations at least twice, while the circles represent the larger subset of subjects who experienced one type of situation only once, and the other type at least once. Note that there is heterogeneity across subjects with both lone wolves and herd animals present.

Furthermore, both pure biased types are less frequent in the smaller, more experienced, subset (represented by the diamond markers). In fact, pure herd animal types are absent in this smaller subset - which is hardly surprising, as repeatedly following a “group” consisting of a single subject would indicate a very strong herd animal bias that one would expect to happen only very rarely. Still, the two subjects in this smaller subset with the Lone Wolf Index values of -0.75 each (irrationally) followed “the herd” consisting of a single subject 3 times and also (rationally) followed larger herds at later positions another 3 times, to a total of 6 counts of following the herd out of 7 total instances (i.e., with an overall 85.7% rate of following social information). At the same time, only three subjects in the smaller subset turned out to be pure lone wolf types, choosing to always follow their own signals in fewer (only 4 and 5) total instances across all situations of our interest.

We are quite aware of the small numbers of observations per subject involved in this analysis, as well as very imperfect identification of information source biases, but these are some of the limitations of using data from the standard design. Nevertheless, we believe that this exercise provides indicative evidence for the existence of behavioral types even in standard sequential social learning experiments.

E3. More complex sequential move designs

Finally, we consider some other experiments where the design is in some way intermediate between sequential and simultaneous moves. Duffy et al. (2019) (DHKM) modify the standard sequential move design by introducing information choice similar to the way it is done here. That is, at each point in the sequence, each subject must choose whether to receive a new private signal or to look at the history of guesses by other subjects (in contrast to the standard sequential design where each subject sees both). The equilibrium prediction is to choose private information at positions 2 and 3 and social information at position 4. However, significant proportions of subjects choose social information too early, while others choose private information at position 4. There is thus evidence of both herd animal and lone wolf like behaviour in this sequential context. Overall, relative to equilibrium predictions, they found an overall (irrational) tendency towards favouring social information.

We can see this by again constructing Information Optimality and Lone Wolf Indices for the DHKM data using Eqs. (3) and (4), with $N_{1,i} = 18$ for choices at positions 2 and 3, where private information is optimal, and $N_{2,i} = 6$ for choices at position 4, where social information is theoretically (but not empirically) optimal. Fig. E.2 represents the resulting distribution of rationality vs bias.

Comparing the Figs. 3, E.1, and E.2, we see the following patterns. In the first, plotting data from the current symmetric design, there is a strong mode at optimality while many other subjects are symmetrically distributed between lone wolf and herd animal behavior. In Fig. E.1, which plots data from the standard sequential design (that as we argue above is highly asymmetric) errors that run in the direction of lone wolf behavior are much more common. The idea that this may simply be the result of the asymmetry of the design is supported by Fig. E.2, which plots data from DHKM whose design is asymmetric but in the opposite direction. In the DHKM setting it is mostly optimal to choose private information, so mistakes there generally run in the direction of choosing social information too early in the sequence. And, indeed, the biggest mode of subjects in Fig. E.2 is somewhat herd animal in behavior. Thus, we suggest that herd animals and lone wolves are common behavioral types. However, their apparent relative frequency is heavily influenced by the symmetry of the experimental design.

Eyster et al. (2018) examine social learning in a design that has both simultaneous and sequential moves as one treatment involves a sequence of groups, in which all members of the group move simultaneously. There, equilibrium behaviour involves ignoring or acting against the observed history. Instead, overall they found excessive imitation leading to substantial payoff losses relative to equilibrium. But, perhaps consistent with our finding of heterogeneity, they also observe a minority of subjects who do anti-imitate.

Goeree and Yariv (2015) also find evidence for excessive imitation. In their experiment, subjects choose between receiving a private signal and observing the previous choices of other subjects who themselves did not receive a private signal. Thus, the optimal policy in Goeree and Yariv's study is to always choose private information. Nonetheless, about one-third of subjects chose social information suggesting either confusion or conformism.

March and Ziegelmeyer (2018) employ a design with two parallel sequences of subjects. The second sequence can observe the first, but their own guesses are only observed by the experimenter. These unobserved subjects can have private signals of high, medium or low accuracy. When these signals contradict the public information, subjects with low or medium quality signals should follow the crowd, but those with high quality signals should follow the private signal. Thus, with this design, it is possible to make mistakes in both directions, and indeed, it is found that there is both over-reliance on private signals by some subjects and on social information by others.

In summary, in social learning experiments where excessive imitation (herd animal behavior) is possible, it is, in fact, observed. By design, in the standard sequential move social learning game, such herd animal behavior is difficult to observe as following the herd is largely optimal in that setting and thus it is difficult to separate bias toward social information from rational behavior. However, in our symmetric, two-stage within-subjects design it is possible to separate the information source bias from information misuse, as well as to identify simultaneously both lone wolf and herd animal behaviour. For this reason and the others identified above, we are broadly optimistic that these behavioural types are present in other settings, even though they are not easily detectable.

F. Quantifying individual differences

At the end of the experiment, subjects were offered a flat fee of \$6 for completing a non-incentivized post experiment survey. In all sessions, subjects were asked to report on their gender, age, and college major and were further asked to answer 20 "core" multiple-choice individual personality trait questions.²⁸ All 144 subjects completed the survey. The average and median age were 20.4 and 20 years, respectively, with an age range of 18–27; just 12 (8.3%) subjects were 23 years of age or older. The sample had 73 (50.7%) females and 71 (49.3%) males, almost perfectly gender balanced.

F1. Proxies for cognitive abilities

It is plausible that subjects' choices can be explained by their cognitive skills. College majors are widely seen as noisy signals of these skills. For example, if quantitative skills were important in our setup, engineering majors, on average, would make optimal choices more frequently than communications majors because engineering majors, on average, have higher quantitative ability. We do not take a position on the origin of any such differences. Instead, we simply attempt to explore the correlations between subjects' choices and their college majors as proxies for their cognitive abilities.

Specifically, we used subjects' self-reported college majors to construct a proxy for their cognitive abilities. There are two issues involved in using college majors as a cognitive proxy. First, college majors are categorical variables, which without additional information cannot be used to quantify or even rank subjects' cognitive abilities. Second, our subject pool consists of students from a large US university, many of whom study a combination of several subject disciplines, complicating the analysis.

We thus used a non-conventional (though novel) approach of using publicly available data on two standardized test scores, averaged by discipline (major), to quantify the expected differences (and similarities) across the different majors. First, we use the average scores from the quantitative section of the Graduate Record Examination (GRE) General Test (see

²⁸ The rest of the survey varied across sessions. In 5 sessions, subjects also faced additional multiple-choice questions: three cognitive reflection questions (to be described further), three probabilistic reasoning questions, and a single social interactions "vignette" question. In the remaining 3 sessions, there were additional 9 "individual traits" questions, and a simple calculation question, but neither cognitive reflection nor statistics questions. Among the additional questions, only the cognitive reflection scores were used as a supplementary variable, and other questions were discarded as uninformative.

Fig. F.1. Probability density distributions of GRE_Q , SAT_M , and CRT_{MCQ} .

Table F.1
Multiple-choice version of Cognitive Reflection Test of Frederick (2005).

1.	A bat and a ball cost \$1.10 in total. The bat costs \$1.00 more than the ball. How many cents does the ball cost?	(a) 0.10 (b) 10 (c) 100 (d) none of the above.
2.	If it takes 5 machines 5 minutes to make 5 widgets, how many minutes would it take 100 machines to make 100 widgets?	(a) 5 (b) 25 (c) 100 (d) none of the above.
3.	In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 32 days for the patch to cover the entire lake, how many days would it take for the patch to cover half of the lake?	(a) 4 (b) 16 (c) 31 (d) none of the above.

Service, 2014, Table 4), which is a standardized test used for most US graduate admissions. This test is typically taken by students prior to apply for graduate school. The association of GRE quantitative test scores with student majors is thus more clear, since test takers would almost certainly have a declared major to report. Second, we use the average scores for the mathematics section of the Scholastic Aptitude Test (SAT) (see Board, 2013, Table 25), which is a standardized test used for US college admissions.²⁹ The SAT test is taken prior to entering college and so the major associated with test scores is more speculative, based on student's pre-entry expectations.

We note that average SAT and GRE scores vary across disciplines in predictable ways. For example, the average scores of individuals who expect to major in mathematics ($GRE_Q = 162$ and $SAT_M = 604$), are higher than the average scores of individuals who expect to major in sociology ($GRE_Q = 149$ and $SAT_M = 553$), allowing us to rank disciplines using these two standardized tests. We thus ascribed these standardized test scores to each subject based on their self-reported major field of study. When a subject's reported major involved more than one field, we averaged the scores over the reported fields.³⁰

There are relative advantages to using each of the set of scores. Publicly available GRE scores offer finer classification, but are unavailable for some broad categories such as "Law" and "Undecided", - thus, we were able to ascribe GRE scores only to 136 subjects. In contrast, we were able to ascribe SAT scores to all 144 subjects, but publicly available SAT scores are aggregated into broad categories. As probability densities in Figure F.1 show, one can see that the SAT (middle panel) provides a coarser, and thus potentially noisier, measure relative to GRE (left panel). Average (st.d.) and median scores are 154.07 (4.23) and 153 for GRE_Q and 541.54 (28.32) and 547 for the SAT_M . Overall, the two sets of ascribed average tests scores are highly correlated, with $r(GRE_Q, SAT_M) = 0.77$, p -value = 0.00.

We emphasize that we do not have GRE or SAT scores for our experimental subjects. Furthermore, our subjects are by no means a random sample from the population who take either of the tests. Our subject pool involves both domestic and international students, attending a selective US university. Apart from a small number of postgraduate students, most of our subjects are undergraduate students, who may or may not intend to go to a graduate school. In comparison, the SAT exam is taken by both those who end up attending highly selective US universities and those who end up attending technical colleges, while the GRE exam is taken by individuals from all over the world who intend to study at US graduate schools.

To further check the validity of average quantitative GRE and mathematics SAT scores as noisy proxies for our subjects' cognitive skills, a subset of our 90 subjects were presented with a multiple-choice version of the free-form Cognitive Reflection Test (CRT) test of Frederick (2005) (see Table F.1). The CRT_{MCQ} score ranges from 0 to 3, based on the number of correct answers to the 3 questions of the test. Figure F.1 (right panel) presents a probability density of the CRT_{MCQ} scores of our subjects, who had average (st.d.) and median scores of 1.27 (1.09) and 1. These CRT_{MCQ} scores are significantly cor-

²⁹ See, for example, Leslie et al. (2015) for use of GRE General Test to account for potential differences across disciplines.
³⁰ When in doubt about the classification of major fields, we consulted the major field codes used in the corresponding test application forms administered by ETS and College Board, respectively.

Table F.2

20 “individual traits” questions employing frequency rating of “Always” (6), “Frequently” (5), “Sometimes” (4), “Occasionally” (3), “Rarely” (2), “Never” (1).

Question	Facet	α
Understand how others think	Understanding1	0.740
Able to predict others' behavior	Understanding2	
Feel that others don't know what they are doing	Mistrust1	0.740
Doubt others' abilities or intentions	Mistrust2	
Go my own way rather than following others	OwnWay1	0.748
Follow my judgment rather than what other people do	OwnWay2	
Am curious about what other people do	Curiosity1	0.766
Interested to see what others are up to	Curiosity2	
Feel that winning or losing doesn't matter to me	Rivalry1 (-)	0.810
Avoid situations involving competition	Rivalry2 (-)	
Drawn to compete with others	Rivalry3	(dropped)
Feel that I must win at everything	Rivalry4	
Feel I am no good at understanding others' behavior		
Feel that what other people do is irrelevant to me		
Take the opposite route from everyone else		
Am at ease when behaving differently from others		
Avoid being directed by others		
Avoid contradicting others		
Feel uncomfortable to do things differently from the group		
Ignore my own gut feeling and instead copy others' behavior		

Table F.3

Pairwise correlations between five “individual traits” scales. (* p -value < 0.10, ** p -value < 0.05, *** p -value < 0.01.)

	Understanding	Mistrust	OwnWay	Curiosity	Rivalry
Understanding	1				
Mistrust	0.098	1			
OwnWay	0.168*	0.324***	1		
Curiosity	0.271***	0.122	0.077	1	
Rivalry	0.206**	0.353***	0.275***	0.156*	1

related with both test scores assigned to majors, particularly the GRE, with $r(CRT_{MCQ}, GRE_Q) = 0.381$, p -value = 0.000 and $r(CRT_{MCQ}, SAT_M) = 0.256$, p -value = 0.015, cross-validating all three proxies for cognitive abilities.

F2. Proxies for behavioral biases (predispositions)

Since individuals tend to engage in social learning situations outside of the laboratory, it is plausible the they may bring these “home made” biases with them to the laboratory. Thus, our post-experiment questionnaire contained 20 “core” individual personality trait questions aimed at eliciting potential behavioral predispositions of subjects with regard to their preferences for social versus private interactions. The questions asked subjects to indicate the frequency of experiencing a feeling or taking an action on a 6-item frequency rating scale (“Always”, “Frequently”, “Sometimes”, “Occasionally”, “Rarely”, and “Never”).

All 144 subjects completed the 20 questions, but we discarded the individual trait responses of 14 subjects as being unreliable.³¹ Using responses of the remaining 130 subjects, we retained 12 out of 20 questions, and constructed five scales, which appear to capture the relevant aspects of social interaction, and have sufficient or close to sufficient internal consistency as captured by Cronbach's α (see Table F.2). The remaining 8 questions were omitted from the analysis as being internally inconsistent (see, e.g. Duffy and Kornienko (2010) for the relevant methodology).³² As Table F.3 shows, the five scales are heavily inter-correlated.

G. Individual differences

In this section we investigate the extent to which individual differences in subject characteristics (traits) described in Appendix F can explain the heterogeneity in subjects' behavior. As Table G.1 shows, the “Understanding (Others)” individual trait scale, age and gender affect the extent of bias toward private information, as measured by the Lone Wolf Index.

³¹ These subjects chose the same answer option on at least one questionnaire screen (i.e., they had zero within-screen variance in their answers). However, these subjects' answers to CRT_{MCQ} questions (where available) exhibited consistent variability, and thus were retained.

³² Our main results do not change if instead we use the common factor approach on all 20 question items. We also confirm that the Rivalry scale of Duffy and Kornienko (2010) is both internally consistent (with Cronbach's scale reliability coefficient α of 0.810) and factorable (with Kaiser–Meyer–Olkin (KMO) measure of “middling” sampling adequacy of 0.769).

Table G.1

Bias towards private information: censored (tobit) regressions, robust standard errors in brackets, clustered on unique subject group. Lone Wolf Index LWI_i left-censored at -1 and right-censored at 1 . Specifications (1)–(3) and (4)–(6) without and with inclusion of session dummies. (* p -value < 0.10 , ** p -value < 0.05 , *** p -value < 0.01 .)

LWI_i	(1)	(2)	(3)	(4)	(5)	(6)
female	−0.200** (0.089)		−0.186** (0.089)	−0.202** (0.080)		−0.186** (0.088)
age	−0.054 (0.033)		−0.065** (0.032)	−0.053 (0.034)		−0.067** (0.032)
Understanding		−0.133** (0.052)	−0.133*** (0.048)		−0.137** (0.053)	−0.143*** (0.052)
_cons	1.201* (0.700)	0.549*** (0.195)	1.975** (0.768)	1.275* (0.706)	0.747*** (0.271)	2.253*** (0.822)
Session Controls	No	No	No	Yes	Yes	Yes
F	3.22	6.54	3.80	34.99	10.73	170.23
p	0.04	0.01	0.01	0.00	0.00	0.00
N	144	130	130	144	130	130

Table G.2

Effect of cognitive proxy on payoff-relevant variables: censored (tobit) regressions, robust standard errors in brackets, clustered on unique subject group. Signal Compliance Index (SCI_i) right-censored at 1 , Total expected payoff POI_i right-censored at 0.7604 . Specifications (1)–(3) and (4)–(6) without and with inclusion of session dummies. (* p -value < 0.10 , ** p -value < 0.05 , *** p -value < 0.01 .)

Signal Compliance Index SCI_i	(1)	(2)	(3)	(4)	(5)	(6)
female	0.015 (0.029)		0.019 (0.035)	0.007 (0.027)		0.014 (0.033)
age	0.030*** (0.011)		0.026** (0.011)	0.030*** (0.011)		0.026** (0.012)
GRE_Q		0.006* (0.003)	0.006* (0.003)		0.005 (0.003)	0.005 (0.003)
_cons	0.456** (0.211)	0.116 (0.480)	−0.375 (0.584)	0.464** (0.217)	0.309 (0.500)	−0.215 (0.613)
Session Controls	No	No	No	Yes	Yes	Yes
F	4.48	3.85	2.60	2.31	16.08	5.62
p	0.01	0.05	0.05	0.02	0.00	0.00
N	144	136	136	144	136	136
Total Expected Payoff POI_i	(1)	(2)	(3)	(4)	(5)	(6)
female	0.007 (0.009)		0.008 (0.010)	0.005 (0.008)		0.008 (0.009)
age	0.007* (0.004)		0.005 (0.004)	0.006** (0.003)		0.005 (0.004)
GRE_Q		0.003** (0.001)	0.003** (0.001)		0.003** (0.001)	0.003** (0.001)
_cons	0.567*** (0.076)	0.254 (0.200)	0.159 (0.233)	0.579*** (0.063)	0.270 (0.202)	0.175 (0.239)
Session Controls	No	No	No	Yes	Yes	Yes
F	1.85	5.17	1.91	1.59	3.07	14.81
p	0.16	0.02	0.13	0.12	0.00	0.00
N	144	136	136	144	136	136

Result G1. A proxy for behavioral bias is correlated with the Lone Wolf Index. In addition, on average, young and male subjects display a greater tendency towards being Lone Wolves.

As Table G.2 shows, in addition to age, subjects' cognitive abilities (proxied by the average quantitative GRE score associated with their college major) indeed have an effect on both payoff-relevant indices. The effect of age is more pronounced for period 1 average payoff (which is linear in the Signal Compliance Index, SCI_i), while the effect of the cognitive proxy GRE_Q dominates for the period 2 payoff (given by total expected payoff, POI_i).

Result G2. Individual differences in payoff-relevant variables can be explained by age and a proxy for cognitive abilities.

We further double check whether our results are robust to alternative proxy specifications. As Table G.3 shows, when cognitive proxies are used as explanatory variables in addition to demographic characteristics, the cognitive measures have an effect on payoff-relevant variables, age has further effect on signal compliance, while information bias is correlated only with gender. The SAT_M scores (which, as Fig. F.1 shows, are coarser than GRE_Q scores) are worse at explaining subjects' payoff-relevant indices than GRE_Q , and thus is not used in the main analysis. The best measure is CRT_{MCQ} , however it is not available for a substantial number of subjects.

Table G.3

Censored (tobit) regressions, robust standard errors clustered on groups in brackets, without and with session dummies. Top panel: Signal Compliance Index SCI_i right-censored at 1. Middle panel: Total Expected Payoff POI_i right-censored at 0.7604; Bottom panel: Lone Wolf Index LWI_i left-censored at -1 and right-censored at 1. (* p -value < 0.10, ** p -value < 0.05, *** p -value < 0.01.)

(a) SCI_i	(1)	(2)	(3)	(4)	(5)	(6)
female	0.019 (0.035)	0.019 (0.032)	0.015 (0.023)	0.014 (0.033)	0.011 (0.030)	0.013 (0.023)
age	0.026** (0.011)	0.030*** (0.011)	0.012** (0.006)	0.026** (0.012)	0.030*** (0.011)	0.013** (0.006)
GRE_Q	0.006* (0.003)			0.005 (0.003)		
SAT_M		0.001 (0.001)			0.001 (0.001)	
CRT_{MCQ}			0.029** (0.013)			0.030** (0.013)
_cons	-0.375 (0.584)	0.106 (0.461)	0.740*** (0.109)	-0.215 (0.613)	0.132 (0.487)	0.717*** (0.132)
Session Controls	No	No	No	Yes	Yes	Yes
F	2.60	2.86	8.52	5.62	6.40	4.71
p	0.05	0.04	0.00	0.00	0.00	0.00
N	136	144	90	136	144	90
(b) POI_i	(1)	(2)	(3)	(4)	(5)	(6)
female	0.008 (0.010)	0.009 (0.011)	0.017 (0.012)	0.008 (0.009)	0.007 (0.010)	0.016 (0.012)
age	0.005 (0.004)	0.006* (0.004)	0.003 (0.004)	0.005 (0.004)	0.005* (0.003)	0.003 (0.004)
GRE_Q	0.003** (0.001)			0.003** (0.001)		
SAT_M		0.000* (0.000)			0.000 (0.000)	
CRT_{MCQ}			0.018*** (0.002)			0.018*** (0.002)
_cons	0.159 (0.233)	0.387** (0.151)	0.625*** (0.082)	0.175 (0.239)	0.398** (0.152)	0.625*** (0.081)
Session Controls	No	No	No	Yes	Yes	Yes
F	1.91	1.58	26.55	14.81	0.98	26.35
p	0.13	0.20	0.00	0.00	0.46	0.00
N	136	144	90	136	144	90
(c) LWI_i	(1)	(2)	(3)	(4)	(5)	(6)
female	-0.215** (0.091)	-0.207** (0.090)	-0.218** (0.098)	-0.217** (0.083)	-0.212** (0.082)	-0.224** (0.088)
age	-0.036 (0.033)	-0.052 (0.033)	-0.076 (0.046)	-0.035 (0.033)	-0.049 (0.034)	-0.072 (0.045)
GRE_Q	-0.007 (0.015)			-0.009 (0.017)		
SAT_M		-0.001 (0.002)			-0.002 (0.002)	
CRT_{MCQ}			-0.004 (0.072)			-0.004 (0.072)
_cons	1.841 (2.587)	1.793 (1.340)	1.659* (0.978)	2.230 (2.800)	2.154 (1.469)	1.686* (0.913)
Session Controls	No	No	No	Yes	Yes	Yes
F	2.00	2.28	1.92	45.82	25.70	84.04
p	0.12	0.08	0.13	0.00	0.00	0.00
N	136	144	90	136	144	90

As Table G.4 shows, when, instead, behavioral trait proxies are used as explanatory variables, both payoff relevant variables, Signal Compliance Index SCI_i and Pay Optimality Index POI_i , tend to correlate with age, as well as the “Mistrust” trait scale (though for POI_i , session effects are important, as indicated by the penultimate line reporting the significance of F statistics). In contrast, the information bias captured by the Lone Wolf Index LWI_i negatively correlates with age (weakly), and with the “Understanding” scale (strongly). Male subjects tend to have a stronger bias toward choosing private information. In what follows, we thus will only use “Understanding” and “Mistrust” scales as individual trait proxies. The combined effect of both cognitive and individual traits variables is presented in Table G.5.

Result G3. Payoff relevant variables tend to be correlated with age, cognitive proxies GRE_Q and CRT_{MCQ} , and the “Mistrust” trait scale. Information bias tends to be negatively correlated with age, being female, and with the “Understanding” trait scale.

Table G.4

Censored (tobit) regressions, robust standard errors clustered on groups in brackets, without and with session dummies. Top panel: Signal Compliance Index SCI_i right-censored at 1. Middle panel: Total Expected Payoff POI_i right-censored at 0.7604; Bottom panel: Lone Wolf Index LWI_i left-censored at -1 and right-censored at 1. (* p -value < 0.10, ** p -value < 0.05, *** p -value < 0.01.)

(a) SCI_i	(1)	(2)	(3)	(4)	(5)	(6)
female	−0.006 (0.018)	−0.004 (0.017)	−0.001 (0.019)	−0.014 (0.019)	−0.010 (0.018)	−0.007 (0.020)
age	0.020*** (0.007)	0.020*** (0.007)	0.022*** (0.007)	0.022*** (0.007)	0.022*** (0.007)	0.024*** (0.008)
Understanding	0.001 (0.011)	−0.003 (0.012)		0.004 (0.012)	−0.002 (0.014)	
Mistrust	0.023** (0.011)		0.017* (0.009)	0.023** (0.010)		0.016* (0.009)
OwnWay	−0.007 (0.012)			−0.006 (0.013)		
Curiosity	−0.011 (0.012)			−0.013 (0.012)		
Rivalry	−0.010 (0.015)			−0.015 (0.014)		
_cons	0.643*** (0.158)	0.650*** (0.137)	0.541*** (0.152)	0.600*** (0.184)	0.615*** (0.162)	0.512*** (0.164)
Session Controls	No	No	No	Yes	Yes	Yes
F	1.77	3.30	3.62	33.49	5.35	12.90
p	0.10	0.02	0.02	0.00	0.00	0.00
N	130	130	130	130	130	130
(b) POI_i	(1)	(2)	(3)	(4)	(5)	(6)
female	0.004 (0.008)	0.003 (0.009)	0.005 (0.009)	0.002 (0.007)	0.002 (0.008)	0.004 (0.008)
age	0.008** (0.004)	0.006* (0.004)	0.007* (0.004)	0.007** (0.003)	0.006* (0.003)	0.006* (0.003)
Understanding	0.002 (0.007)	0.002 (0.006)		0.004 (0.007)	0.002 (0.006)	
Mistrust	0.010* (0.005)		0.007 (0.005)	0.010** (0.004)		0.006 (0.004)
OwnWay	−0.002 (0.006)			−0.004 (0.006)		
Curiosity	0.005 (0.004)			0.003 (0.004)		
Rivalry	−0.008 (0.007)			−0.009 (0.007)		
_cons	0.508*** (0.094)	0.572*** (0.079)	0.541*** (0.086)	0.532*** (0.088)	0.579*** (0.078)	0.559*** (0.073)
Session Controls	No	No	No	Yes	Yes	Yes
F	1.76	1.08	1.50	83.84	1.22	1.76
p	0.10	0.36	0.22	0.00	0.29	0.08
N	130	130	130	130	130	130
(c) LWI_i	(1)	(2)	(3)	(4)	(5)	(6)
female	−0.143 (0.096)	−0.186** (0.089)	−0.186* (0.101)	−0.129 (0.092)	−0.186** (0.088)	−0.181* (0.096)
age	−0.058* (0.033)	−0.065** (0.032)	−0.050 (0.032)	−0.062* (0.033)	−0.067** (0.032)	−0.049 (0.031)
Understanding	−0.152** (0.060)	−0.133*** (0.048)		−0.170** (0.065)	−0.143*** (0.052)	
Mistrust	0.074 (0.083)		0.077 (0.080)	0.073 (0.080)		0.077 (0.080)
OwnWay	−0.003 (0.070)			−0.027 (0.066)		
Curiosity	−0.007 (0.060)			0.006 (0.063)		
Rivalry	0.056 (0.051)			0.079 (0.059)		
_cons	1.636* (0.896)	1.975** (0.768)	0.875 (0.807)	2.027** (0.953)	2.253*** (0.822)	1.002 (0.801)
Session Controls	No	No	No	Yes	Yes	Yes
F	2.31	3.80	2.64	131.05	170.23	109.44
p	0.03	0.01	0.05	0.00	0.00	0.00
N	130	130	130	130	130	130

Table G.5

Censored (tobit) regressions, robust standard errors clustered on groups in brackets, without and with session dummies. Signal Compliance Index SCI_i right-censored at 1; Total Expected Payoff POI_i right-censored at 0.7604; Lone Wolf Index LWI_i left-censored at -1 and right-censored at 1. Top panel: using GRE_Q ; bottom panel: using CRT_{MCQ} . (* p -value < 0.10, ** p -value < 0.05, *** p -value < 0.01.)

(a)	SCI_i		POI_i		LWI_i	
	(1)	(2)	(3)	(4)	(5)	(6)
female	0.004 (0.023)	0.001 (0.024)	0.006 (0.009)	0.005 (0.008)	-0.215** (0.096)	-0.204** (0.094)
age	0.021** (0.009)	0.023** (0.009)	0.006 (0.004)	0.005 (0.004)	-0.052* (0.031)	-0.054* (0.029)
GRE_Q	0.001 (0.003)	-0.001 (0.003)	0.003* (0.001)	0.002* (0.001)	-0.003 (0.015)	-0.006 (0.016)
Mistrust	0.015 (0.009)	0.018** (0.009)	0.004 (0.005)	0.002 (0.004)		
Understanding					-0.139*** (0.045)	-0.141*** (0.047)
_cons	0.380 (0.464)	0.656 (0.513)	0.184 (0.251)	0.234 (0.242)	2.261 (2.438)	2.794 (2.694)
Session Controls	No	Yes	No	Yes	No	Yes
F	1.98	50.61	1.24	2.30	4.02	108.42
p	0.10	0.00	0.30	0.01	0.00	0.00
N	124	124	124	124	124	124
(b)	SCI_i		POI_i		LWI_i	
	(1)	(2)	(3)	(4)	(5)	(6)
female	0.027 (0.027)	0.022 (0.029)	0.021 (0.013)	0.020 (0.013)	-0.241** (0.096)	-0.230*** (0.083)
age	0.013* (0.007)	0.015** (0.008)	0.003 (0.005)	0.003 (0.004)	-0.102** (0.043)	-0.094** (0.041)
CRT_{MCQ}	0.033** (0.015)	0.035** (0.016)	0.019*** (0.003)	0.018*** (0.003)	-0.005 (0.081)	-0.011 (0.078)
Mistrust	0.018* (0.009)	0.019** (0.009)	0.005 (0.003)	0.005 (0.003)		
Understanding					-0.106** (0.051)	-0.107* (0.054)
_cons	0.654*** (0.150)	0.599*** (0.166)	0.602*** (0.094)	0.605*** (0.094)	2.718** (1.066)	2.677** (1.063)
Session Controls	No	Yes	No	Yes	No	Yes
F	7.15	27.98	17.13	25.07	2.69	89.36
p	0.00	0.00	0.00	0.00	0.04	0.00
N	79	79	79	79	79	79

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.euroecorev.2021.103690](https://doi.org/10.1016/j.euroecorev.2021.103690)

References

- Alevy, J.E., Haigh, M.S., List, J.A., 2007. Information cascades: evidence from a field experiment with financial market professionals. *J. Financ.* 62, 151–180.
- Anderson, L.R., 2001. Payoff effects in information cascade experiments. *Econ. Inq.* 39, 609–615.
- Anderson, L.R., Holt, C.A., 1997. Information cascades in the laboratory. *Am. Econ. Rev.* 87 (5), 847–862.
- Banerjee, A.V., 1992. A simple model of herd behavior. *Q. J. Econ.* 107 (3), 797–817.
- Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* 100 (5), 992–1026.
- Cavatorta, E., Guarino, A., Huck, S., 2018. Learning from Aggregate and Partial Information. Working Paper.
- Celen, B., Kariv, S., 2004. Distinguishing informational cascades from herd behavior in the laboratory. *Am. Econ. Rev.* 94 (3), 484–498.
- Charness, G., Oprea, R., Yuksel, S., 2019. How do people choose between biased information sources? Evidence from a laboratory experiment. Working paper. https://economics.sas.upenn.edu/system/files/2018-10/CharnessOpreaYuksel_Aug2018.pdf.
- Cipriani, M., Guarino, A., 2005. Herd behavior in a laboratory financial market. *Am. Econ. Rev.* 95 (5), 1427–1443.
- College Board, 2013. 2013 College-Bound: Total group profile report. New York, NY: The College Board.
- Crawford, V.P., Costa-Gomes, Miguel, A., Iriberri, N., 2013. Structural models of non-equilibrium strategic thinking: theory, evidence, and applications. *J. Econ. Lit.* 51 (1), 5–62.
- De Filippis, D., Roberta, A.G., Jehiel, P., Kitagawa, T., 2018. Non-Bayesian Updating in a Social Learning Experiment. CEMMAP working paper CWP39/18.
- Dominitz, J., Hung, A.A., 2009. Empirical models of discrete choice and belief updating in observational learning experiments. *J. Econ. Behav. Org.* 69, 94–109.
- Drehmann, M., Oechssler, J., Roeder, A., 2005. Herding and contrarian behavior in financial markets: an internet experiment. *Am. Econ. Rev.* 95 (5), 1403–1426.
- Duffy, J., Hopkins, E., Kornienko, T., Ma, M., 2019. Information choice in a social learning experiment. *Games Econ. Behav.* 118, 295–315.
- Duffy, J., Kornienko, T., 2010. Does competition affect giving? *J. Econ. Behav. Org.* 74, 82–103.
- Educational Testing Service, 2014. GRE Guide to the Use of Scores. ETS, Princeton, NJ.

- Esponda, I., Vespa, E., 2018. Endogenous sample selection: a laboratory study. *Quant. Econom.* 9 (1), 183–216.
- Eyster, E., Rabin, M., Weizsäcker, G., 2018. An Experiment on Social Mislearning. Working paper.
- Fahr, R., Irlenbusch, B., 2011. Who follows the crowd - groups or individuals? *J. Econ. Behav. Org.* 80, 200–209.
- Frederick, S., 2005. Cognitive reflection and decision making. *J. Econ. Perspect.* 19 (4), 25–42.
- Goeree, J.K., Holt, C.A., Palfrey, T.R., 2016. *Quantal Response Equilibrium*. Princeton University Press, Princeton.
- Goeree, J.K., Palfrey, T.R., Rogers, B.W., McKelvey, R.D., 2007. Self-correcting information cascades. *Rev. Econ. Stud.* 74 (3), 733–762.
- Goeree, J.K., Yariv, L., 2015. Conformity in the lab. *J. Econ. Sci. Assoc.* 1 (1), 15–28.
- Grant, S., Kajii, A., Polak, B., 1998. Intrinsic preference for information. *J. Econ. Theory* 83, 233–259.
- Hung, A.A., Plott, C.R., 2001. Information cascades: replication and an extension to majority rule and conformity rewarding institutions. *Am. Econ. Rev.* 91, 1508–1520.
- Kawamura, K., Vlaseros, V., 2017. Expert information and majority decisions. *J. Public Econ.* 147, 77–88.
- Kübler, D., Weizsäcker, G., 2004. Limited depth of reasoning and failure of cascade formation in the laboratory. *Rev. Econ. Stud.* 71 (2), 425–441.
- Leslie, S.-J., Cimpian, A., Meyer, M., Freeland, E., 2015. Expectations of brilliance underlie gender distributions across academic disciplines. *Science* 347 (6219), 262–265. 16 Jan 2015
- March, C., Ziegelmeyer, A., 2018. Excessive herding in the laboratory: the role of intuitive judgments. CESifo Working Paper Series No. 6855, January 2018
- Moscarini, G., Ottaviani, M., Smith, L., 1998. Social learning in a changing world. *Econ. Theory* 11, 657–665.
- Nöth, M., Weber, M., 2003. Information aggregation with random ordering: cascades and overconfidence. *Econ. J.* 113, 166–189.
- Oberhammer, C., Stiehler, A., 2003. Does cascade behavior reflect Bayesian updating? an experimental study. Discussion Paper on Strategic Interaction 0301, Max Planck Institute of Economics, Jena.
- Samuelson, L., 2004. Information-based relative consumption effects. *Econometrica* 72 (1), 93–118.
- Sterling, G., 2018. 85% Of Consumers Disregard Local Reviews More than 3 Months Old. searchengineland.com.
- Weizsäcker, G., 2010. Do we follow others when we should? a simple test of rational expectations. *Am. Econ. Rev.* 100, 2340–2360.
- Willinger, M., Ziegelmeyer, A., 1998. Are more informed agents able to shatter information cascades in the lab? In: Cohendet, P., Llerena, P., Stahn, H., Umbhauer, G. (Eds.), *The Economics of Networks: Interaction and Behaviours*. Springer-Verlag.
- Ziegelmeyer, A., Koessler, F., Bracht, J., Winter, E., 2010. Fragility of information cascades: an experimental study using elicited beliefs. *Exp. Econ.* 13, 121–145.
- Ziegelmeyer, A., March, C., Krügel, S., 2013. Do we follow others when we should? a simple test of rational expectations: comment. *Am. Econ. Rev.* 103 (6), 2633–2642.