Lone Wolf or Herd Animal? An Experiment on Choice of Information and Social Learning*

John Duffy† UC Irvine

Ed Hopkins‡ University of Edinburgh

Tatiana Kornienko§ University of Edinburgh

July 2015

Abstract

We report on an experiment that uses revealed preference to distinguish between rational social learning and behavioral bias. Subjects must choose between receiving a private signal or observing the past guesses of other subjects before guessing the state of the world. The design varies the persistence of the state across time. This changes whether choosing social or private information is optimal. We can therefore separate subjects who choose optimally from both those who excessively use social information ("herd animals") and those with excessive use of private information ("lone wolves"). While aggregate behavior appears unbiased, this is because the numbers of lone wolves and herd animals are approximately equal.

Keywords: social learning, information, experiments, conformity, social influence.

JEL codes: C72, C92, D83.

*We thank Juergen Bracht, Steven Broomell, Steffen Huck, Chad Kendall, Dimitra Politi, Emanuel Vespa and Georg Weizsäcker for helpful comments.

†duffy@uci.edu, http://www.socsci.uci.edu/∼duffy/

‡E.Hopkins@ed.ac.uk, http://homepages.econ.ed.ac.uk/hopkinse

§Tatiana.Kornienko@ed.ac.uk, https://sites.google.com/site/tatianaedinburgh/

1 Introduction

Humans are highly social. We seem to have a keen interest in the actions and successes and failures of others. Does this reflect an intrinsic desire to conform or imitate? Individuals with such tastes for social information might be labeled "herd-animals" as they strongly desire to know what others are doing. Alternatively, the main idea in social learning is that it can be rational to observer others, because there is information contained in their actions. That is, responding to the behavior of others is a purely instrumental means toward the end of acquiring good information rather than an intrinsic taste for such information alone. Further, existing experimental research (Weizsäcker, 2010) finds that subjects are, on average, biased in the opposite direction, overweighting private information relative to social information. We label such individuals as "lone-wolves" according to their preference for private information over social information.

In this paper, we test whether individuals' interest in social information is rational and if not, whether this is better explained by taste or by error. We conduct experiments that have three major novelties. First, subjects must choose between receiving either private or social information. This is in contrast to the existing literature on observational learning, where subjects have access to both types of information. Second, our design allows the optimal choice of private or social information to differ across treatments; in certain environments, social information is the optimal choice while in others it is not. Third, in our within-subject treatments, subjects are exposed to two of these environments, which enables us to better separate rational subjects from being classified as lone wolves or herd animals. Finally, we attempt to minimize subjects' mistakes. We give subjects many repetitions of the main task with full feedback and, in our between-subjects design, they also go through a substantial number of training rounds to familiarize them with the induced random processes.

Our main finding is that, while there is no particular bias in favor or against the use of social information in the aggregate, there are clear individual differences in behavior. First, we find that about 70 percent of subjects choose correctly in two thirds or more of their decisions about which information to receive. Second, whether social or private information is optimal in a given environment has no significant effect on this success rate, so subjects on average are not biased for or against social information. This stands in strong contrast to previous studies, summarized in Weizsäcker (2010), on social learning that find a bias in favor of private information. Third, we find that across-subject measures overstate the degree of optimal behavior as compared with between-subject measures. There are some subjects who always choose social information and some who always choose private information irrespective of which type of information is optimal. Further, we have evidence that the incorrect choices are driven at least in part by taste and not just by error. Finally, we find that social science majors, including those studying economics and business, are more individualist than humanities or science majors. Thus, we provide evidence that there are significant numbers of all three types, rational agents, lone wolves and herd animals.

Furthermore, prominent alternative explanations for our findings based on error rather than taste are rejected by the data. A symmetric quantal response equilibrium (QRE) is broadly consistent with the aggregate deviations from optimal behavior. However, it is not consistent with the distribution of behavior across subjects, where the dispersion of behavior is far greater than a symmetric equilibrium would predict. An alternative, level-k model predicts an aggregate bias against social information, because unsophisticated subjects would not believe there is useful information in the actions of others (higher level subjects would not be biased). Yet, there is no aggregate bias in behavior.

How can we explain the absence of a strong bias in favor of private signals, as identified in previous experiments? While it is difficult to identify the exact cause as our experiment differs in more than one way from previous studies, we can make two hypotheses. First, subjects rely less on social information in traditional sequential move social learning experiments because it is more compled. The optimal action depends on one's position in the sequence and to what extent subjects early on in the sequence are following their private information rather than subjects in front of them. Second, here by giving subjects a choice between the two types of information (private or social), we avoid an experimenter demand effect that may have led to overweighting of private information in experiments where there was no choice of information. As Cooper and Rege (2011) suggest, subjects might feel obliged to respond to the private signals that the experimenters have given them.[1]

In our experiments we consider a simple situation where subjects are rewarded for identifying the true state of the world. Prior to making this choice, they can receive a noisy private signal about the true state of the world or consider the choices of others regarding the true state of the world in the previous period. Subjects could potentially improve their forecast of the true state of the world in the current period by inferring it from the past choices of others. Yet, as Samuelson (2004) pointed out, the relevance of such social information depends on the persistence of the environment. In a relatively volatile environment where the state of the world changes frequently, information about the past actions of others would not be as useful as it would be in a more stable world. Specifically, we vary the persistence of the environment, so that one's private signal about the current state of the world is more or less useful than the information about the past actions by others. Thus, when restricted to a choice between receiving a new private signal and the distribution of others' past actions, a rational agent should choose social information when the environment is persistent, and private information when it is not.

Our experimental setup has clear theoretical predictions. A group of experimental subjects face the same changing environment and are rewarded for correctly identifying the true, binary state of the world in each of two periods. In the first period, each subject receives a noisy but informative private signal. The optimal policy is clearly to follow this signal in

[1] The results of Miller and Maniadis (2012) also support the idea of bias driven by how signals are labelled. In their experiments, subjects respond more strongly to balls drawn from an urn designated as their own than to equally informative signals generated by drawing balls from other urns.

guessing the current state of the world. The next period, each subject has a choice between receiving a new private noisy signal, or, instead, observing the choices made by the other subjects in the group in the first period. In the PERSISTENT environment, the state of the world is sufficiently likely to be the same in both periods that the optimal policy is to choose social information rather than a new private draw, and, to guess the second period state they should imitate the past actions of the majority of subjects. In contrast, in the ERRATIC environment, the state of world is not very persistent across periods. In this case, the optimal policy is to draw a new private signal and to follow it. If subjects have an irrational taste for either social or private information, one would expect that the persistence of the environment would have little effect on subjects' choice over this type of information.

We are particularly interested whether there exists an irrational taste for social information. To see whether there is any scope for conformism, we further study a NON-CONFORMIST environment, where the state of the world is negatively persistent. In this treatment the optimal policy calls for subjects to select social information and then to choose the action opposite to that chosen by the majority of subjects in the prior period. By contrast, a purely conformist subject might simply imitate the past actions of the majority of his group or may seek to avoid the discomfort of behaving differently from the others, by choosing a less informative private signal instead.

Our experiment is a mix of both between-subject and within-subject designs. In the within-subject design, subjects were exposed successively to two different persistence environments. That is, all subjects had to make decisions in PERSISTENT environment as well as in one of the non-persistent environments - ERRATIC or NON-CONFORMIST. The order and composition of these two environments was varied across sessions. Subjects were not informed at the start of experiment that the persistence would be changed later. In the between subject design, subjects only faced one of the three environments, PERSISTENT, ERRATIC and NON-CONFORMIST. However, prior to completing this main task, subjects were incentivized to learn about both the precision of noisy private signals and the persistence of the state of the world using the same parameterization of the model environment used for the main task.

Deviations from rational behavior can happen for reasons other than taste, for example, subjects can make mistakes. Our design seeks to minimize the incidence of such mistakes in two ways. First, we allow subjects to learn about the usefulness of the two types of information, by confronting them with multiple information choice decisions in the same environment and by providing them with full feedback, including ex-post revelation of the content of information that they did not choose. Second as noted above, in the between-subject sessions, we included an opportunity for subjects to acquire "experience" with the precision of their noisy private signals and with the persistence of the environment, respectively, thus reducing the possibility that observed behavior might be due to a lack of experience with particular random events.[2]

[2]There is a growing literature on the differential effect of "stated" versus "experienced" probabilities -

Our environment differs substantially from previous experimental studies on social learning in two ways. First, our study does not involve purely sequential choice, subjects do not move one after another. Second, subjects must choose which type of information to view prior to making their choices. Prior experiments on social learning, including Anderson and Holt (1997), Celen and Kariv (2004), Goeree et al. (2007) and Ziegelmeyer et al. (2010) all involved purely sequential choice. Subjects in these experiments were given both a private signal and social information on the prior choices of others. These are also the assumptions of the classic theoretical papers of Banerjee (1992) and Bikhchandani et al. (1992). Kübler and Weizsäcker (2004) is closer to our model as it makes private information optional and only visible if a subject chooses to pay a small fee. These existing experiments on social learning, as summarized in Weizsäcker (2010), find that subjects follow their own private information more frequently than is optimal. Thus, prior social learning experiments suggest that if deviations from optimal behavior are due to preferences (and not errors) then subjects tend to be "lone wolves" rather than "herd animals".

Alternatively, there is another branch of the literature that considers social influence on what otherwise would seem to be single person decision problems. For example, Cooper and Rege (2011) find that subjects' choices over lotteries are affected by others' decisions (see also Linde and Sonnemans, 2012; Lahno and Serra Garcia, 2015). Rather than this being driven by conformity, Cooper and Rege identify this effect as being driven by "social regret", a desire not to be alone in suffering losses.

However, neither branch of the social learning literature allows subjects to choose which type of information (private or social) to receive. We think that this choice is important as it puts both types of information on an equal footing. The one exception and the paper that is closest to ours, is an experiment by Goeree and Yariv (2015) in which subjects choose between receiving a private signal and observing the previous choices of others who themselves did not receive a private signal. Thus, the optimal policy in Goeree and Yariv's study is always to choose private information. Nonetheless, about a third of subjects chose social information suggesting either confusion or conformism. The principal differences between Goeree and Yariv and the current study are, first, that their experiment involved sequential rather simultaneous choice. Second and more importantly, in their design, private information is always optimal and therefore errors always run in the same direction as conformism. Finally, because choice is sequential, it is possible to choose social information in order to copy exactly what others have already chosen and therefore ensure that one has the same payoff as other subjects (perhaps to avoid the social regret as identified by Cooper and Rege).

- see, for example, Hertwig et al. (2004). We did not include parts with "experienced probabilities" in the within subject sessions both because of time constraints and in order to test the effect of experienced probabilities on subject decisions.

2 A Simple Model of Social Learning

The environment is a simplified version of a model of social influence due to Samuelson (2004). The central idea is that it can be optimal to observe the actions of others as these convey information about the underlying state of the world. Here the model is modified in one important way. This is that individuals, rather than observing both a private signal and the actions of others, must choose which piece of information to observe. We show that the optimal policy depends on the level of persistence of the state of the world across time and set out how the optimal policy is different for the three parameter values implemented in the experiments.

There are two periods, time $t = 1$ and $t = 2$. 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 will change to the other state with probability $1 - p$. This probability p will vary across treatments.

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. The choice of X when the state is Y and choosing Y when the state is X have zero payoff.

In period 1, each agent receives a private signal x or y. The accuracy of the signal is such that $\Pr(x|X) = \Pr(y|Y) = q > \frac{1}{2}$. That is, as $q > \frac{1}{2}$, the signal is informative. Thus, the optimal action is to choose X if the signal is x and Y if the signal is y. 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"), which again will have accuracy q, **or** seeing all actions taken by the other subjects at time 1 ("social information"). Once the chosen information is received, the subject makes her period 2 choice of X or Y.

2.1 Optimal Policy

A policy or strategy for an individual is therefore a decision whether to follow her signal at $t = 1$, a decision about which information to receive at $t = 2$ and then a final 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 at $t = 2$ depends on whether agents follow their signal at $t = 1$, and thus the optimal policy can be thought of as a subgame perfect equilibrium. However, the strategic aspects of this game are close to trivial, as following one's signal at $t = 1$ is a dominant strategy (in expected payoffs).

If at $t = 2$ an individual chooses to see an independent signal, then as before its accuracy will be q. It will be optimal to follow the signal and consequently the accuracy of the individual's guess conditional on her choice of private information in period 2 will be q.[3] If instead an individual chooses social information at $t = 2$ that individual can recall her own action and now sees the actions chosen by $n - 1$ others. Let n be odd so that there is always a majority for one of the two actions. It is clear that the optimal action, given social information at $t = 2$, is to copy the majority action at time $t = 1$ if $p > 0.5$ and to take the opposite action if $p < 0.5$.[4]

The probability of success with this strategy depends on n, p and q in the following way. The probability that the majority action was the correct action at time $t = 1$ 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,

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

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

The probability that the majority action is still the correct action at $t = 2$ is equal to the probability that it was the correct action at time $t = 1$ multiplied by p. The probability that the majority action was incorrect at $t = 1$ is $1 - Q_n(q)$. The probability that it was both incorrect at $t = 1$ but it would be correct to follow it (because meanwhile the state changes) is $(1 - p)(1 - Q_n(q))$. Thus, the overall accuracy, A, of social information, that is, the probability of correctly predicting the state at $t = 2$ by following the majority action at $t = 1$, is

$$A(n, p, q) = pQ_n(q) + (1 - p)\left(1 - Q_n(q)\right). \tag{2}$$

The value of social information is increasing in all three variables n, p (if $p > 1/2$) and q.

The values of the three parameters n, p and q were chosen for the experiment with the following in mind. The value of the precision parameter, q, that maximizes the difference between social and private information, i.e., $A(n, p, q) - q$, is certainly below 1 and decreases toward 0.5 as n becomes large. At the same time, this maximal difference is increasing in n. Thus, n was chosen to be as large as practical, that is 9. Given that persistence, p, was meant to be high in the PERSISTENT environment, $p = 0.9$ was an obvious choice, and for symmetry, $p = 0.1$ is a natural choice for the NON-CONFORMIST treatment. Signal

[3] One might think that the accuracy with which an individual can guess the state in period 2 is greater than q as she already has an observation 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 when they disagree. Hence the first signal has no effect on the accuracy of the individual in the second period if she opts for private information. It is true that if the two signals agree then probability of being correct conditional on agreement is higher than q. But it can also happen that the two signals disagree. It is still optimal to follow the more recent signal, but the conditional probability of being correct is now lower. Overall expected accuracy is still exactly q.

[4] This assumes that all agents take their optimal action at time $t = 1$ and follow their private signals.

precision, q, was then chosen to maximize the advantage to social information in these two treatments given p and n.

Specifically we chose $q = .7$ for all of our treatments according to the following reasoning. As noted above, in the PERSISTENT environment, persistence, $p = 0.9$. Since we have $n = 9$ subjects per group, it follows from equation (1) that $Q_9(0.7) = 0.901$, and from equation (2) accuracy, $A(9, .9, .7) = 0.821$. Thus in the PERSISTENT treatment, the advantage to social information over following one's private signal is 0.121 (i.e. $0.821 - 0.7$). It follows that in the PERSISTENT environment, the optimal policy is to follow one's signal at $t = 1$, but to choose social information at $t = 2$ and to follow the majority's choice from $t = 1$ in period $t = 2$.[5] In the ERRATIC environment we chose $p = 0.6$, and thus from (1-2) we can compute that $A(9, .6, .7) = 0.580$. Thus in the ERRATIC treatment, private information, with accuracy 0.7, is 0.12 more accurate than social information which means that the two treatments PERSISTENT and ERRATIC are almost exactly reverse symmetric in terms of the strength of incentives. In the ERRATIC treatment the optimal policy is to choose private information and to follow the signal received.

In the NON-CONFORMIST environment, we set $p = 0.1$. Thus, by equations (1) and (2) the time $t = 2$ accuracy from following the majority action of $t = 1$, $A(9, .1, .7) = 0.179$. Thus, doing the opposite of the majority action at $t = 1$, gives one the optimal action at $t = 2$ with probability 0.821 $(1 - .179)$. As this probability is greater than drawing a private signal at $t = 2$ with accuracy $q = 0.7$, the optimal policy in the NON-CONFORMIST treatment is to choose social information at $t = 2$ and to guess the opposite to the majority. Note that, by design, this optimal policy has the same expected success rate (0.821) as the optimal policy (choose the social information and follow the majority choice) in the PERSISTENT environment.

Finally, the actual accuracy of social information in the second period depends on other subjects having played optimally in the first period, namely guessing the state corresponding to their first period signal. It would not be optimal to choose social information if these first period actions were sufficiently noisy. However, previewing one result from our experimental data, the realized frequency of period 1 optimal actions across all sessions of our experiment is very high, at 0.9779. We use this frequency to calculate a realized social accuracy $A(n, p, \tilde{q})$ where $\tilde{q}$ is not 0.7 but $0.7 \times 0.9779 + 0.3 \times 0.0221 = 0.691$. For example, in the PERSISTENT treatment, the theoretical accuracy is 0.821 while the experimental, realized accuracy is lower at 0.812. Nevertheless, the realized accuracy remains well above the 0.7 threshold that rationalizes the use of social information. Thus, the optimal policy conditional on actual first period behavior is unchanged. This same finding regarding the invariance of the optimal policy to the use of realized rather than theoretical accuracies, also applies to our other two treatments as well as shown in Table 1, which summarizes details of this section.

[5] Risk aversion does not affect the optimal policy here. The policy with the highest expected payoff also has the lowest variance.

Environment	**Persistence**	**A - Private**	**A - Social**	**A - Social Real**	**Opt. Policy**
PERSISTENT	0.9	0.7	0.821	0.805	Social
ERRATIC	0.6	0.7	0.58	0.576	Private
NON-CONFORM	0.1	0.7	0.821	0.805	Social - Rev

Table 1: Summary of Environments Used in Experiments

	Within	**Between**	
Signal Experience	-	24	
Persistence Experience	-	24	
Main Task	48	48	
Main Task 2	48	-	
Total Decisions	96	96	

Table 2: Outline of the Design by Numbers of Rounds

3 Experimental Design

We report results from two sets of experiments, which we refer to as the within-subjects experiment/design and the between subjects experiment/design In the within-subjects design, the experiment consists of two parts. In each part, subjects are repeatedly confronted with the "main task" of: 1) guessing the state of the world in period 1, followed by, 2) an information choice, and finally, 3) guessing the state of the world in period 2. In each part, this main task (consisting of three decisions) is repeated 48 rounds under a constant value for the treatment value of the persistence parameter, p. At the start of the second part of our within-subjects design, this persistence parameter was changed and remained constant for the remaining 48 rounds of that part; subjects were not informed of this change in the persistence parameter in advance. Specifically, one of the parts had a persistence, $p = 0.6$ and the other part had a persistence of either $p = 0.9$ or $p = 0.1$. Thus each subject went through 48 rounds of both persistence treatments, for a total of 96 rounds. We controlled for possible order effects by varying the order of two treatments faced across sessions.

Our between-subjects design involved three parts. Prior to performing the main task, subjects made choices that gave them "experience"with uncertain signals (part 1) and uncertain persistence (part 2). In the third part subjects had to make "main task"decisions in 48 rounds under a single type of persistence environment, i.e., a single, constant value for p. The outline of the two different designs in given in Table 2.

In either design, subjects were divided into groups, each with $n = 9$ subjects. Subjects remained in fixed matches with the same members of their matching group for all parts of the experiment, though they only interacted with matching group members during the main part (in the between-subjects design, the first and second parts were individual-choice experiments as detailed below).

3.1 The Main Task (Both Within and Between Designs)

In the main task, each round consists of two periods, period 1 and period 2. Subjects are instructed to imagine that there exist two urns, a "black" urn and a "red" urn. The black urn is so-named because it contains 7 black balls and 3 red balls and the red urn is so-named because it contains 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 matching group of size $n = 9$, one urn was randomly chosen at 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".

The random urn choice processes were "live" for the first session but thereafter, we used the same sequence of random draws in all subsequent sessions. We do this so all that the different sessions face the same empirical frequencies. 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 $1 - p$. Given the urn that is in place for a given period, ball colors drawn correspond to the color of the urn with probability $q = 0.7$ or not with probability $1 - q = 0.3$. The latter random draws are made randomly and independently for each subject viewing a colored ball; the latter draws were live (i.e. real-time) in all sessions; only the urn sequence was be pre-determined following the first session.

In period 1, subjects are shown the color of a ball selected from the unknown urn and must guess the color of the unknown urn. The decision screen at the end of period 1 shows the subject's choice of Red or Black 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 first period ball (signal) and their first period choice of urn. They are told that in period 2, there is a p percent chance that the (still unknown) urn will be the same urn from which the ball they observed was drawn for period 1, and a $1 - p$ percent chance that for period 2, the urn will the one that was not used in period 1. Note that the urn in place in periods 1 and 2 is the same urn for all $n = 9$ members of a matching group.[6] Then 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 made by the other 8 subjects in his/her matching group for period 1 (and *not* the 8 signals the other 8 players received).

[6]For example if we had two matching groups, the urn in period 1 would be randomly chosen to be black(red) for one group an then red (black) for the other group and persistence in period 2 would follow the same switching pattern, e.g., if in period 2 there was a switch from the black to the red urn in the first matching group, the second matching group would have a switch in period 2 from the red to the black urn.

The Period 1 decision screen contained the following text. "For this period, the actual urn is either black or red, with a 50-50 chance. From this actual urn, one of the 10 balls has been randomly drawn and has the color: [Black/Red]. Which urn do you think is the actual urn in period 1? [Black/Red]".

This was followed by the Period 2 decision screen: "There is a p percent chance that the actual urn in period 2 will stay the same and a $1 - p$ percent chance it will instead be the other urn. What would you like to do? Draw a ball from the actual urn from period 2 or See which urn(s) the other 8 participants in your group chose in period 1."

After making this information choice ("info") we move to a third decision screen. If the subject chose to draw a new ball (i.e., to go it alone), 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. In the latter case, the third decision screen is similar to the first decision screen. On the other hand, if the subject chose to look at the urn choices made by the other 8 subjects in period 1, then on the third decision screen the subject is shown the number of the other subjects who chose the Black urn in period 1 and the number of subjects who chose the Red urn in period 1. Below this information, the subject is reminded of her own choice for period 1 and asked to make an urn choice for period 2.

After all period 2 choices were submitted, the round was over and subjects received feedback on the outcome 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 were now informed about the actual color of the urn in period 1. They were further reminded of their information choice prior to period 2 (new ball draw or group information from period 1), the content of their chosen information, their guess of the urn color for period 2 and they now learned the actual color of the urn in period 2. In addition, we provided subjects with the content of the *other* piece of information they could have chosen prior to period 2 but did not chose, either the group information or a random ball draw from the period 2 urn. We provided this information so that subjects had an opportunity to assess whether their choice of information prior to period 2 was optimal or not.[7] Finally, subjects were also informed of 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 their guesses for the urn colors in periods 1 and 2. Following the first round of play of the main task, a complete history of outcomes from all prior rounds of play of this main task was found at the bottom of subjects' decision screens. In particular, subjects saw a history of 1) the color of the ball drawn in period 1, 2) the subject's guess of which urn was selected in each period, 3) the information the subject chose to view (New Draw /Group), 4) the other piece of information the subject did not choose

[7]In particular, by providing this feedback, subjects could accurately assess whether private or social information was more accurate in predicting the color of the urn in period 2. In this manner, any possible biases to social information could be quickly detected.

to view, 5) the actual group urn that was selected for periods 1 and 2 and 6) the subject's points earned for the round.

3.2 Signal Experience (Between Design Only)

In our between-subject design, the experiment started with a signal experience part, followed by a persistence experience part, and, finally, i the main task part. In the signal experience part (part 1), subjects participated in 24 rounds of an individual-choice experiment in which their task was to correctly guess the color of the urn from which a ball was drawn in each round. Note that this set-up is equivalent to the period 1 choice in the main task, except that subjects were not interacting as part of a group in this signal experience treatment. Specifically, subjects were told that there were two urns, black or red, and that in each of the 24 round one of these two urns would be randomly chosen (a 50 percent chance each round). While they would not know which urn was chosen, they were told that the black urn contained 7 black balls and 3 red balls and the red urn contained 3 black balls and 7 red balls, again consistent with our choice of $q = 0.7$. A ball from the chosen urn was randomly chosen and shown to each subject. Their task was to correctly guess the color of the urn from which that ball was chosen. In the first session, the actual sequence of random draws of the two urns was determined in real-time (i.e., "on the fly") according to the stated random process, as was the ball that was drawn from the chosen urn for all subjects. Thereafter, these choices were hard-coded so that there were no differences in these random realizations across sessions/subjects. After all subjects had guessed the color of the urn from which the ball was drawn, the round was over. Subjects were then informed of the outcome of the round. Specifically they were reminded of the ball color drawn, their guess of the color of the urn from which the ball was drawn and they then learned the true color of urn from which the ball was drawn. If they guess the correct urn they got 1 point, and 0 points otherwise. Following the first round, a scrollable history of outcomes from prior rounds was shown to subjects on their first decision screen. The purpose of this individual-choice experiment was to give subjects considerable experience (24 rounds) with our signal precision choice of $q = 0.7$ that was again used in the main task.

3.3 Persistence Experience (Between Design Only)

In our between-subjects design, the signal experience part was followed by a persistence experience part (part 2). In this persistence experience part, subjects participated in 24 rounds of an individual-choice experiment in which their task was to correctly guess whether the color of an urn (red or black) had changed from period 1 to period 2, given knowledge of the persistence parameter, p. Specifically, in each of 24 rounds, subjects were instructed that in period 1, one of the two urns, black or red would be randomly chosen, with each urn having an equal chance of being chosen. The color of the urn chosen in period 1 was then revealed to subjects. Subjects' task in this second part was simply to guess the color of the urn in period 2. For this guess, no ball (or signal) was drawn from the urn chosen for period 2. Instead, subjects were informed that from period 1 to period 2, the chance that the urn remained the same color was p, the persistence probability and the chance that it changed color was $1-p$. The value of p was fixed over all 24 rounds of this second part of the session. At the start of the next 2-period round, the color of the urn in period 1 was again chosen randomly (50% red, 50% black) and with probability p it remained the same color in period 2. As in the signal persistence part, in the first session, the actual sequence of random draws of the urns in periods 1 and 2 was determined in real-time (i.e., "on the fly") according to the stated random process, for the urn in period 1 and the persistence parameter p for determining the urn color in period 2. Thereafter, these choices were hard-coded so that there were no differences in these random realizations across sessions/subjects. After all subjects had guessed the color of the urn in period 2, the round was over. Subjects were then informed of the outcome of that round. Specifically they were reminded of the color of the urn drawn for period 1 and their guess of the color of the urn for period 2 and they then learned the true color of urn for period 2. If their guess was correct they got 1 point, and 0 otherwise. Following the first round, a scrollable history of outcomes from prior rounds was shown to subjects on their first decision screen. The purpose of this individual-choice experiment was to give subjects considerable experience (24 rounds) with our persistence parameter choice, p.

In the third part of the between-subjects experiment involving the main task (described above) we used the same value for p that we used in this, second signal experience part, so that subjects were quite experienced with the treatment specific persistence of the urn color between the two periods for the main task prior to entering that task. We only used one value for p in this third part (main task) of our between-subjects design which involved 48 rounds of play.

3.4 Questionnaire and Payment (Within and Between Designs)

At the end of the experiment, one round was randomly chosen from each of part of the experiment (the two parts of the within-subject experiment and the three parts in the between-subjects experiment). For the within subject design, one round was drawn from each main task part. Subjects could earn up to 2 points for each round, and since two rounds were chosen–one from each part–they could earn up to a maximum of 4 points. For the between-subjects design, one round was randomly drawn from part 1, one round was randomly drawn from part 2, and one round was randomly drawn from part 3, the main task. Subjects could earn 1 point (for a correct guess) in each of parts 1 and 2 and up to 2 points for the main task in part 3. Thus, as in the within-subject design, subjects could earn up to a maximum of 4 points in the between subjects design. Points were converted into money payments at the fixed and known rate of 1 point=\$6. In addition, all subjects earn a fixed show-up payment of $6 that required completion of an ex-post experimental survey. Thus, maximum total earnings were $30 for an experiment lasting approximately 90 minutes.

The ex-post survey for the within subject design varied across sessions, with only data on major collected for all sessions, as well as a subset of questions designed to elicit individual differences. In addition, some sessions contained questions on gender, age, 3 CRT questions, 3 statistics questions, and additional personality questions. The survey for the between subjects design consisted of questions on gender, age, major, 3 CRT questions, 3 statistics questions and 17 personality questions.

4 Results from the Within Subject Experiment

In the within-subject sessions, subjects performed the main task in two different persistence environments. At the beginning of the experiment, they were informed that there would be two parts of the experiment, but were not told about the content of the second part. Specifically, when they were performing the first part of the experiment, they were not aware that the second part would involve the same main task, but with a different persistence rate. Different pairs of environments were run in different sessions and the order of each pair was varied. Specifically, there were four within subject treatments: 96, PERSISTENT then ERRATIC; 69, ERRATIC then PERSISTENT; 16, NON-CONFORMIST then ERRATIC; 61, ERRATIC then NON-CONFORMIST.

	Part 1				Part 2				Total			
p	Nobs	Mean	SD	%	Nobs	Mean	SD	%	Nobs	Mean	SD	%
p=0.9	36	47.61	0.99	99.19	36	46.86	4.79	97.63	72	47.24	3.45	98.41
p=0.6	72	46.58	3.87	97.05	72	47.00	2.95	97.92	144	46.79	3.44	97.48
p=0.1	36	46.25	3.43	96.35	36	47.61	1.29	99.19	72	46.93	2.66	97.77
Total	144	46.76	3.29	97.41	144	47.12	3.23	98.16	288	46.94	3.26	97.79

Table 3: Within subjects: Summary statistics for Period 1 urn choices which followed the random signal.

The main task takes place over two time periods. In the first period, subjects receive a private signal and must simply decide whether to follow it. This is important as the optimality of choosing social information in the second period of the PERSISTENT and NON-CONFORMIST environments depends on subjects correctly following their signal in the first. Overall, about 97.8 percent of decisions were to follow the signal. As noted earlier (see Table 1) this frequency is high enough for the optimal policy in the second period to remain choosing social information in the PERSISTENT and NON-CONFORMIST

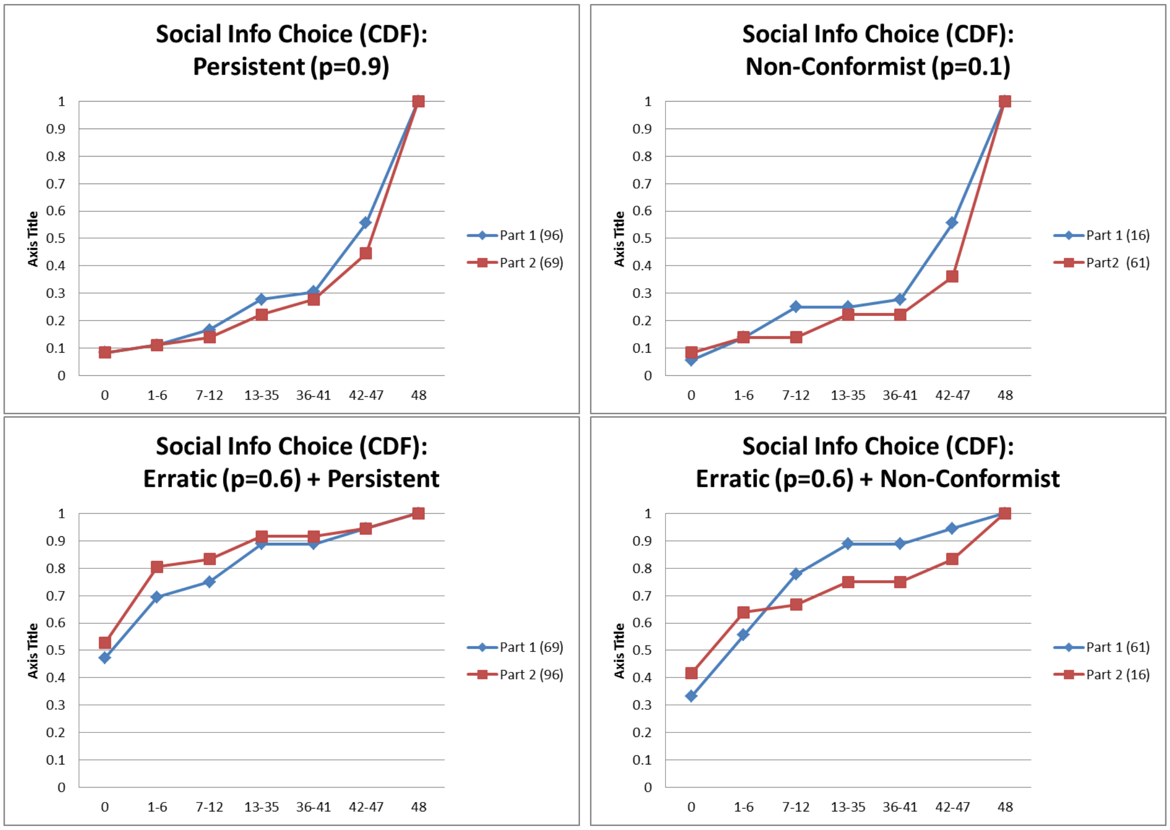

Figure 1: Within subjects: Cumulative frequencies of choice of social information for each of the three environments, for each part.

environments.

We now turn to the main task. A first basic question is whether there was an order effect in the main task. As Figure 1, demonstrates, the biggest difference between subjects' choice of social information in a particular persistence treatment in part 1 versus part 2 happens in those sessions where at least one of the parts involved the NON-CONFORMIST environment (p=0.1). Indeed, for disaggregated cumulative frequencies, the absolute value of the biggest difference is 0.1944. However, using a Kolmogorov-Smirnov two-tailed test for two independent samples involving 36 subjects each, this difference is insignificant. Thus, we can ignore order effects and pool together the observations for a particular environment from all relevant sessions.

In Figure 2 (left panel), we graph the frequencies with which subjects choose social information in the second period. Since they each made 48 such choices, this frequency runs from 0 to 48. One can see that for the ERRATIC ($p = 0.6$) environment, there is a big spike at 0 (no social information) the optimal choice in this context. In contrast, for the PERSISTENT ($p = 0.9$) and NON-CONFORMIST ($p = 0.1$) environments, the modes are at

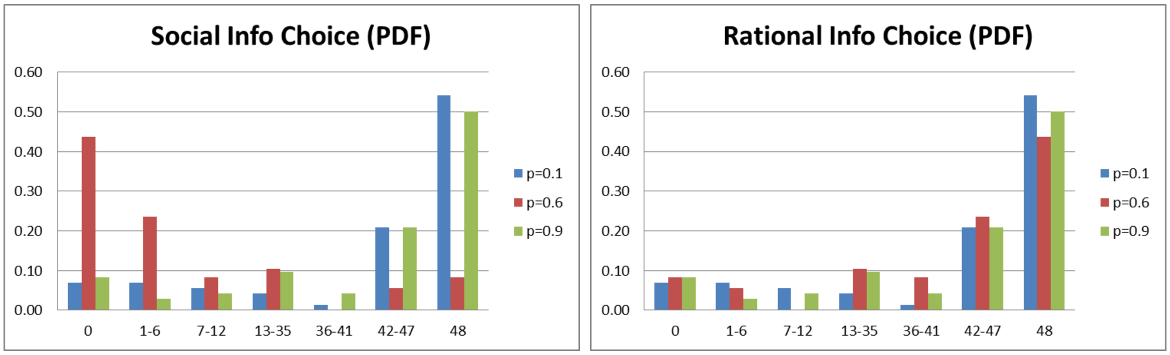

Figure 2: Within subjects: Frequencies of choice of social information (left panel) and optimal information (right panel) in each of the three environments.

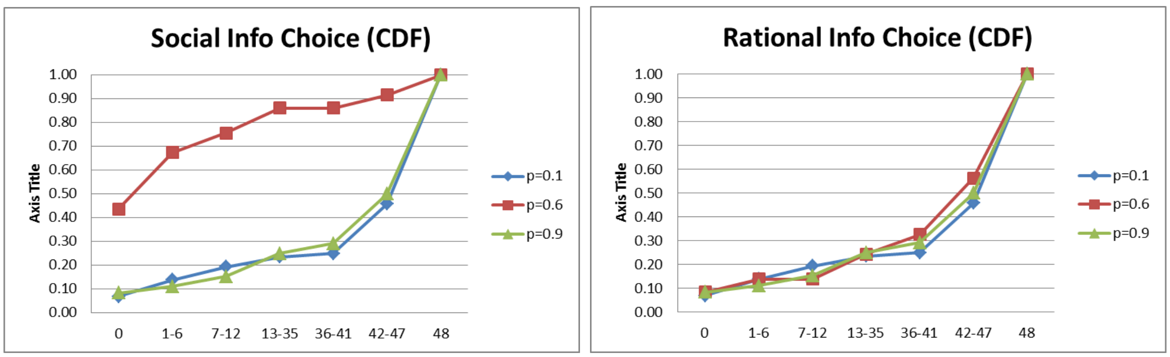

Figure 3: Within subjects: Cumulative frequencies of choice of social information (left panel) and optimal information (right panel) in each of the three environments, aggregated.

48 (always choose social information). That is, in all three environments, the modal choice is to always choose the information that conforms with the optimal policy. This is particularly clear if we consider the frequencies of rational choice of information (i.e., private information in the erratic environment and social information in the other two environments). Indeed, as Figure 2 (right panel) shows, the modal choice is at 48. That is, the modal outcome involves 100% rational choice of information. This tendency for rational choice of information can be further observed when one explores the cumulative frequency distributions as in Figure 3.

Table 4 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. The strategies are ordered in terms of their expected payoff, with type 1 being the optimal policy in that environment and type 4 being the worst. One can see that the frequencies of strategies chosen are largely ordered in terms of their relative payoffs, with the optimal strategy being the far most frequent, with an overall average of 74.96%. However, note that the frequency of choosing the correct information but then not using it correctly is non-zero in all environments.

		Part 1				Part 2				Total			
p=0.9	Strategy	N	Mean	SD	%	N	Mean	SD	%	N	Mean	SD	%
Type 1	SF	36	35.58	16.78	74.13	36	37.50	17.20	78.13	72	36.54	16.90	76.13
Type 2	PF	36	8.64	13.69	18.00	36	7.89	14.71	16.44	72	8.26	14.12	17.22
Type 3	PN	36	2.22	5.11	4.63	36	1.50	4.18	3.13	72	1.86	4.65	3.88
Type 4	SN	36	1.56	2.43	3.24	36	1.11	2.93	2.31	72	1.33	2.68	2.78
p=0.6	Strategy	N	Mean	SD	%	N	Mean	SD	%	N	Mean	SD	%
Type 1	PF	72	35.60	15.76	74.16	72	35.53	17.46	74.02	144	35.56	16.57	74.09
Type 2	SF	72	9.18	14.53	19.13	72	9.92	17.02	20.66	144	9.55	15.77	19.89
Type 3	SN	72	0.92	1.41	1.91	72	0.54	1.19	1.13	144	0.73	1.31	1.52
Type 4	PN	72	2.31	4.04	4.80	72	2.01	4.19	4.20	144	2.16	4.11	4.50
p=0.1	Strategy	N	Mean	SD	%	N	Mean	SD	%	N	Mean	SD	%
Type 1	SN	36	34.50	18.64	71.88	36	38.00	16.83	79.17	72	36.25	17.72	75.52
Type 2	PF	36	9.61	15.66	20.02	36	8.28	16.70	17.25	72	8.94	16.09	18.63
Type 3	PN	36	2.28	4.53	4.75	36	0.42	1.40	0.87	72	1.35	3.46	2.81
Type 4	SF	36	1.61	1.69	3.36	36	1.31	1.74	2.72	72	1.46	1.71	3.04
Total		N	Mean	SD	%	N	Mean	SD	%	N	Mean	SD	%
Type 1		144	35.32	16.66	73.58	144	36.64	17.16	76.33	288	35.98	16.89	74.96
Type 2		144	9.15	14.52	19.07	144	9.00	16.31	18.75	288	9.08	15.41	18.91
Type 3		144	1.58	3.59	3.30	144	0.75	2.38	1.56	288	1.17	3.06	2.43
Type 4		144	1.94	3.23	4.05	144	1.61	3.43	3.36	288	1.78	3.33	3.70

Table 4: Table of summary statistics for Period 2 strategy types (within treatment).

Indeed, this “between subjects” analysis provides a remarkable support for rationality among experimental subjects. However, we can utilize our within subject design to explore the level of individual rationality. Specifically, we are interested in whether subjects exhibit any bias towards private information. To investigate the level of bias among subjects we introduce a “lone wolf index”. This is defined for each subject as the total number of choices of private information across the two environments she or he faces less 48, which is the number of decisions faced in each part (96 total). Thus, the possible values of this lone wolf index run from -48 to 48. If a subject chooses optimally, for example, in the 16 treatment, then in the first part, social information should be chosen 48 times when the environment is NON-CONFORMIST and in the second part, private information should be chosen 48 times when the environment switches to ERRATIC so that her index value will be 0, that is, unbiased. If, however, she always chooses social information, she will have a score of -48, biased toward social. If she always chooses private information, then her score would be 48, a fully antisocial lone wolf.

If we graph our constructed data on the lone wolf index as in Figure 4, we see that the distribution is broadly unimodal and symmetric around zero. Many subjects are unbiased,

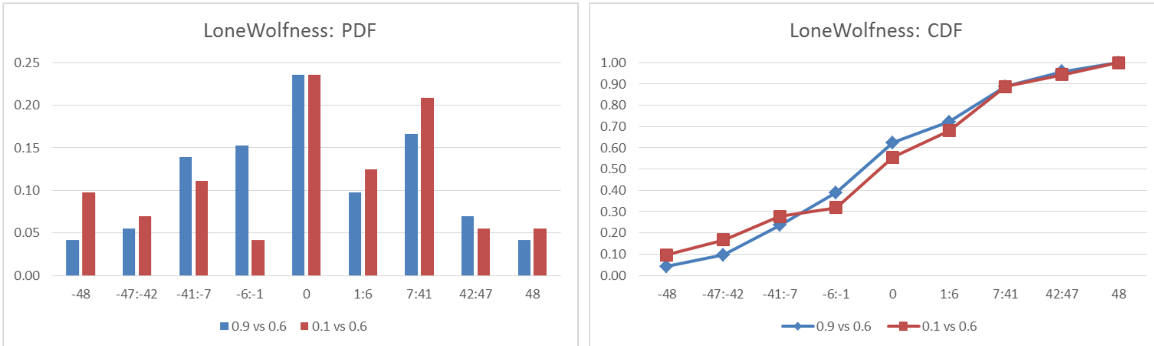

Figure 4: Within subject: Distribution of the Lone Wolf Index.

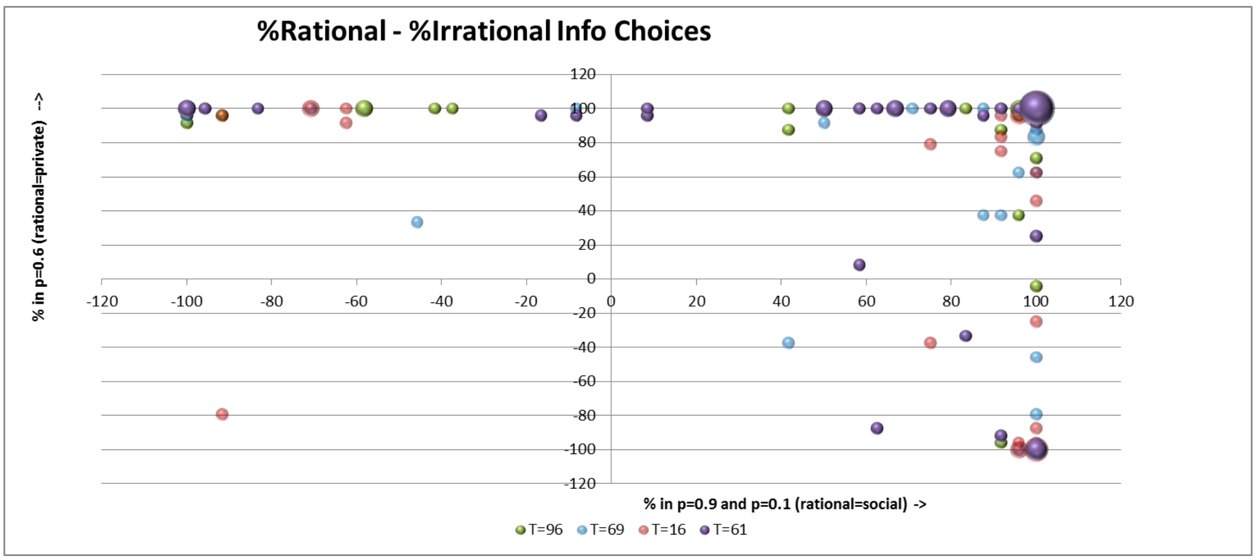

Figure 5: Within subject: Distribution of the two-dimensional Rationality Index.

with more than 20% exactly at zero. In the aggregate, the population of subjects is effectively unbiased. But again this is because the numbers of lone wolves and herd animals are approximately equal. The average (st. dev.) and median of this distribution are -0.694 (26.795) and 0 respectively. Thus, there is no significant bias in either direction.

We also construct a two-dimensional index of rationality for each subject, by subtracting the percentage of incorrect choices from the percentage of correct choices in each environment he faces. Thus, if a subject chooses social information when it is optimal and private information when that is optimal, his index would be (100,100). If he is always wrong, his index would be (-100, -100). Lone wolves have the index score of (-100,100): all choices are private info (thus not optimal when $p = 0.9$ and $p = 0.1$ and optimal when $p = 0.6$). A herd animal has score (100,-100): all choices are for social information (optimal when $p = 0.9$ and $p = 0.1$ but not optimal when $p = 0.6$).

We can see how the population of subjects is distributed in Figure 5. The biggest concentration is at (100, 100), optimal behavior. But there are also significant numbers of subjects at the other “corners” (-100, 100) and (100, -100). That is, there are both lone wolves and herd animals. Again the distribution is broadly symmetric, with both types being roughly equally frequent. Specifically most, more than 60% of subjects are rational, with about 25% completely rational with scores of (100,100). The number of lone wolves at (-100,100) is about 15% and the number of herd animals at (100,-100) is about 12%.

Overall, there is considerable evidence for both rationality and heterogeneity. Most subjects, about two thirds, choose information to maximize their monetary payoff. Those who fail to do so, about 30%, choose the type of information compatible with their apparent preferences over the level of social interaction, with roughly equal numbers being social and anti-social.

4.1 Level-k and QRE Predictions

In this section we consider some well-known behavioral models and ask whether they predict deviations from the optimal policy in the task that we consider. In a Level-k model of differing levels of strategic sophistication, Level 1 agents believe that Level 0 types act at random. Applying this belief to our setting, Level 1 subjects would reason that the first period actions of others are uninformative. Consequently, Level 1 subjects should always choose the private signal over social information in period 2, regardless of the persistence parameter. By contrast, Level 2 (and higher) subjects would act optimally because they see other subjects as Level 1, and those types follow their own signal in period 1. Thus, overall, Level-k produces a bias towards lone-wolf behavior. Given that most studies that have fitted the Level-k model to experimental data have found a large proportion (20-30% in many studies) of subjects to be Level 1, there should be that proportion being lone-wolf, with the remaining more sophisticated subjects behaving optimally (any Level 0 subjects would randomise uniformly over the two forms of information and thus would not provide an overall bias).

As noted in the discussion of the rationality index plotted in Figure 5, about 15% of subjects always choose private information. It is tempting to identify these as Level 1 players. However, this number is rather low as compared to other studies. More importantly, there are an almost equal number (12%) who always choose social information. Such behavior is not consistent with the Level-k model.

Another behavioral theory used in social learning is the “social confirmation bias” as proposed by Eyster and Rabin (2010) does not offer distinct predictions in the current environment. It assumes that individuals believe others choose entirely according to the private signals they receive and do not learn from third parties. In our game, where there is not much sequential play, this simply implies that subjects should believe that other subjects follow their signal in the first period, which is in fact the optimal policy. Therefore, this theory's predictions correspond to play of the optimal policy by all subjects.

In a quantal response equilibrium (QRE) model (see for example, Goeree et al. (2007) for further details) all subjects play noisy best responses to the play of others. Since QRE involves a positive chance of error in any direction, it potentially can explain both the mistaken choice of social information as well as the excessive use of private information. We will quickly see how its predictions compare with the data.

Since QRE involves a positive error rate, it predicts that subjects with a positive probability will not follow their signal in the first period, reducing the accuracy of social information. Thus it might seem that again there would be a bias towards lone-wolf behavior. However, this is not the case because reducing the payoff to choosing social information and using it correctly, also raises the payoff to the (very stupid) strategy of choosing social information and not following it. Thus, in fact, a logit QRE predicts an almost symmetric error rate, so that, for example, the predicted frequency with which the optimal strategy is chosen in the PERSISTENT environment is roughly the same as in the ERRATIC environment. This is broadly in line with our experimental data in terms of average behavior.[8] However, as we will see, this prediction fails to explain the heterogeneity we observe in individual subject behavior.

We first calculate the logit QRE for all possible values of the logit precision parameter β. In Figure 6, there is a plot of the predicted frequencies of choosing the correct information in the different environments. The axes are effectively the same as in Figure 5 - the probability of choosing social in PERSISTENT and NON-CONFORMIST environments on the horizontal axis, the probability of private in ERRATIC on the vertical. The set of QRE is the line running from (0.5, 0.5) to (1, 1). When β is zero, in QRE agents choose at random, giving choice probabilities equal to 0.5. As β becomes very large, individuals choose the correct information with a probability close to one. These calculations take into account noisy behavior in the first round. The kink near (0.5, 0.5) is due to the first round behavior being so noisy for low values of β that the expected return to social information is lower than private in the PERSISTENT environment. However, one can also see that for the most part the predicted equilibrium frequencies are very close to being symmetric. The basic form of QRE does not predict any particular bias toward or against social info.

Second, we fit the logit QRE to overall average second round behavior, as given in Table 4, jointly across all environments. We calculate payoffs to choosing social information using the empirical average first round accuracies as given in Table 3. The β that maximizes the log-likelihood is 7.955. This gives equilibrium frequencies in the PERSISTENT environment of $\{SF, PF, PN, SN\} = (0.699, 0.284, 0.012, 0.005)$ where the strategies are labeled S for

[8] Goeree et al. (2007) find that a modified QRE better fits the data in their sequential choice social learning experiments. The modification, base rate neglect, allows QRE to match the bias towards the private information found in their data. Again here there is no aggregate bias in that direction.

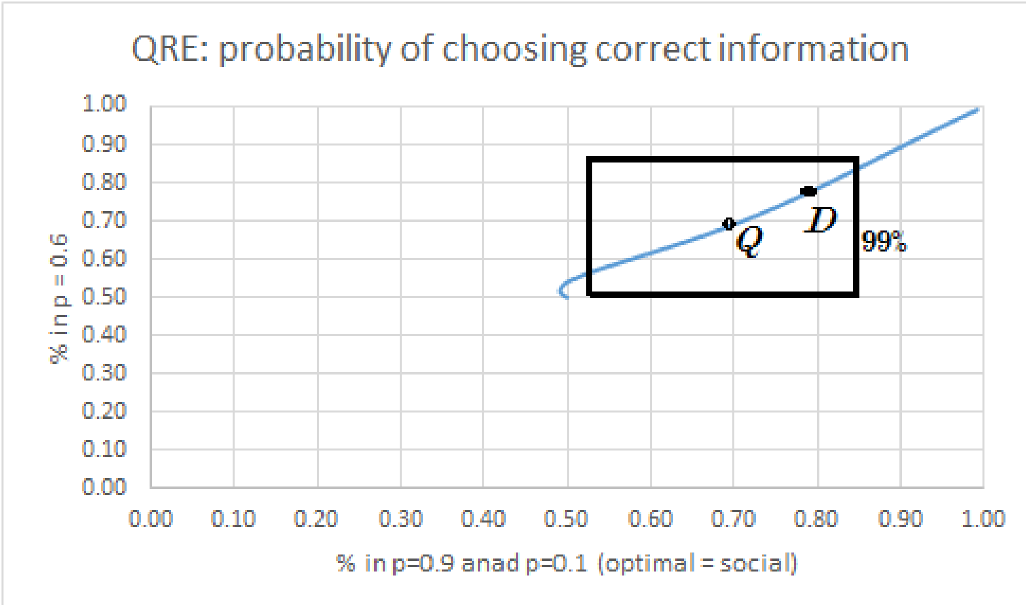

Figure 6: The set of logit QRE for all values of precision. The point Q is the maximum likelihood estimate. The point D is the average of the data across all subjects and sessions. The box gives a 99% interval for accuracy given 48 repetitions and the frequencies in Q.

social, P for private, F for follow the signal received and N for not following. Overall, the frequency of choosing social information is about 70%. For the ERRATIC environment, the predicted frequencies are $\{SF, PF, PN, SN\} = (0.247, 0.654, 0.027, 0.072)$, which gives the frequency of private information at about 68%. This estimated QRE is represented by point Q in Figure 6. Further, these frequencies are a good approximation to those observed - see Table 4. For example, in the PERSISTENT environment, the overall percentage of choices of social information is 78.9, in the ERRATIC private information was chosen 78.6% of the time. The average across all sessions is plotted as point D in the figure.[9]

However, a good fit at the aggregate level hides a 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 Figures 5 and 6. If the data had been generated by a symmetric QRE then there should be a unique mode near around the estimated QRE, point Q in Figure 6. Indeed, on Figure 6, we have drawn a box such that if an individual was choosing information according to the estimated QRE frequencies across 48 repetitions, the frequency with which she chose the correct information would be within the box with 99% probability.

[9]One reason why the QRE does not match the empirical frequencies of information choice more closely is that it also has to try to fit the frequencies of the suboptimal strategies, such as choosing private information and not following it.

However, the experimental data in Figure 5 looks very different. The mode is not around (0.7, 0.7) but at (1,1). There are many subjects that are far more accurate than is predicted by QRE. There are further modes in two other corners representing individuals apparently biased for or against social information. 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 information when it is optimal, implying a very high precision level, and also always choose social when it is not optimal, implying a low or even negative precision level. In conclusion, a symmetric QRE cannot explain the distribution of behavior across subjects, in that it cannot explain the simultaneous existence of lone wolves and herd animals.

4.2 Rational Inattention

One alternative model that might help to explain what we see is one based on the theory of rational inattention developed in Matejka and McKay (2015). This assumes a two-stage cognitive process. First, an individual decides how much information to receive depending on the cost of information. Second, she receives information according to the chosen accuracy and takes an action that is optimal given the information and her prior.

We can apply a model of this kind to our experiments if we assume the first stage is about how much mental effort to expend to determine the relative accuracy of social and private information. Specifically, assume that subjects know that the expected accuracy of private information is $q = 0.7$. Further, despite being given a full description of the environment, assume they do not know for sure the accuracy of social information, because even approximating the calculations we make in Section 2 requires effort. Suppose that an individual i thinks that with probability g_i the accuracy of social information is $0.7 + k$ with $k > 0$, and with probability $1 - g_i$ that $k < 0$. Then in effect they choose how much mental effort to exert. Different subjects face different costs of information λ, where it takes values two possible values $\lambda_H > \lambda_L$. They choose an optimal effort given these costs, update their beliefs from the results of this introspection, and then choose between private and social information.

The broad predictions of the model are first that those subjects with high information costs never gain precise information about whether social information is more accurate than private. Instead, they are likely to go with their priors and choose private information for sure if g_i is low and social information if g_i is high. If g_i takes an intermediate value, then they randomize. Second, subjects with low information costs will correctly determine the optimal information and choose it. Thus, a model can rationalize the four broad types of subject behavior observed. Lone wolf behavior is predicted by high information cost λ_H and low prior g_i; herd animal by high information cost and high prior; optimal choosers by low information cost; learners and experimenters, i.e. those who are often but not always choose optimally, by high information cost and intermediate prior.

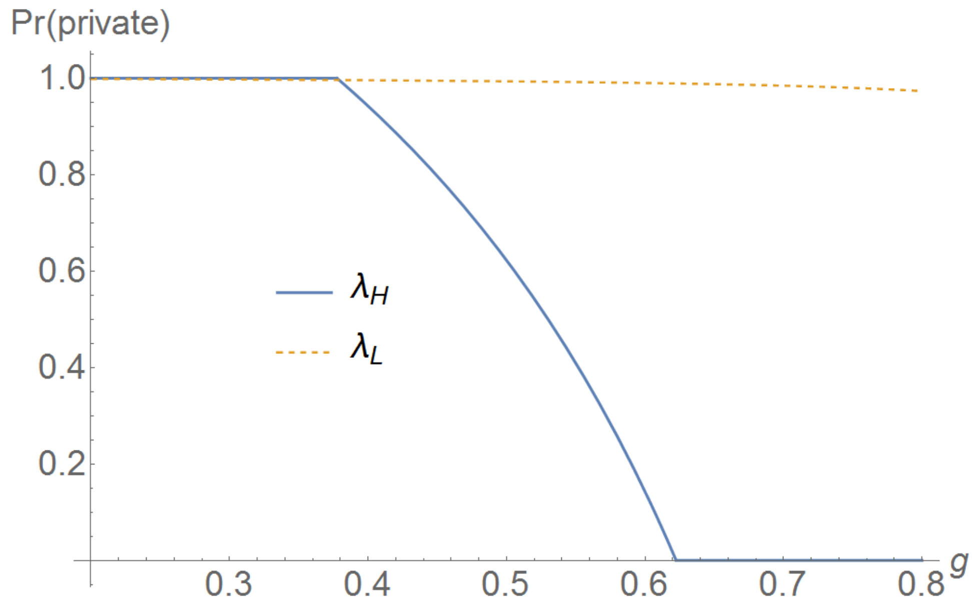

Figure 7: Probability of choosing private information in the erratic environment as a function of Prior g for Different Information Costs

This is illustrated in Figure 7.[10] This assumes the ERRATIC environment so that private information is optimal. With low information costs (dotted line), an individual ends up choosing the optimal action with probability close to one, largely irrespective of her prior. However, with high information cost (solid line), choices depend heavily on one's prior so that if one's g_i is only a bit above one half, one chooses social information with probability one, even though in reality it is not optimal. Note that even these high information cost individuals do learn something from introspection. Notice that given a prior of $g_i = 0.5$, no initial bias, the probability of choosing private information is somewhat above 0.6 - choosing correctly is more common than not.

We do not attempt to fit this model to the data, but just observe that it can generate predictions that are qualitatively similar. Note that to do so we have in effect three parameters to play with, λ_H, λ_L and a distribution of priors g_i. This is what enables us to match the wide heterogeneity in subject behaviour. This is in contrast with the QRE that has only one parameter and is symmetric across subjects.

[10]This is derived from Matejka and McKay's (2015) Problem S1 in their Appendix F. There are two actions, one of which (here choosing private information) pays a fixed payoff R, and the other high or low, depending on the state of the world. In their notation, the Figure 7 plots $\mathcal{P}_1(1, R)$ as a function of g_0 for $\lambda_H = 1$ and $\lambda_L = 0.1$.

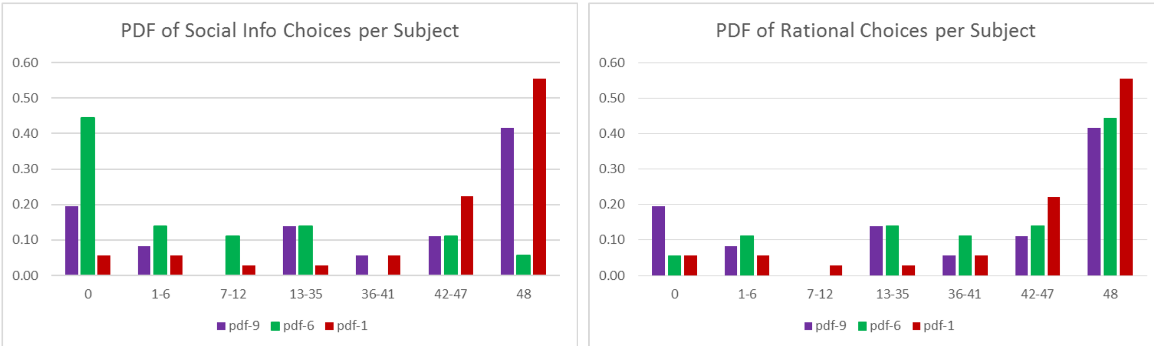

Figure 8: Between subjects: Frequencies of choice of social information (left panel) and optimal information (right panel) in each of the three environments.

5 Results from the Between Subject Treatments

We now will turn to the results of our between subjects design. The between subject sessions started with 24 rounds of signal training and 24 rounds of persistence training. The main task takes place over two time periods. In the first, subjects receive a private signal and must simply decide whether to follow it or not. This is important, as the optimality of choosing social information in the second period of the PERSISTENT and NON-CONFORMIST environments depends on subjects correctly following their signal in the first period. Overall, we find that about 98 percent of decisions were to follow the signal, which is about the same frequency we found in the within-subjects experiments. This is certainly high enough for the optimal policy in the second period to be to choose social information in the PERSISTENT and NON-CONFORMIST environments. [Add something about persistence training!]

We now turn to the main task. In Figure 8, we graph the frequencies with which subjects choose social information in the second period. Since they each made 48 such choices, the frequency runs from 0 to 48. One can see that for the ERRATIC environment, labeled 6 for $p = 0.6$, there is a big spike at 0, the optimal choice in this context. In contrast, for the PERSISTENT and NON-CONFORMIST environments, the modes are at 48. That is, in all three environments, the modal choice is again to always choose the information that conforms with the optimal policy. Further support for this finding is provided in Figure 9.

The choice of social information over time is graphed in Figure 10. The three different panels represent the three different environments, PERSISTENT ($p = 0.9$), NON-CONFORMIST ($p = 0.1$) and ERRATIC ($p = 0.6$). Time in terms of the number of periods is represented on the horizontal axis, the number of subjects choosing social information (from zero to nine) is on the vertical axis. Each group of nine subjects is graphed separately. One can see that there are no strong trends in choice over time. However, some of those who do not choose optimally initially do change their strategy with experience. This may not be obvious to the eye but it is picked up in regression analysis (see below).

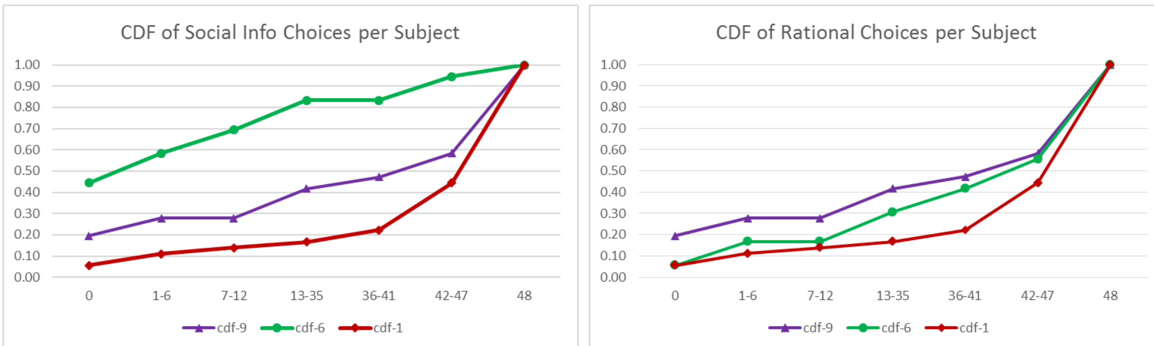

Figure 9: Between subjects: Cumulative frequencies of choice of social information (left panel) and optimal information (right panel) in each of the three environments.

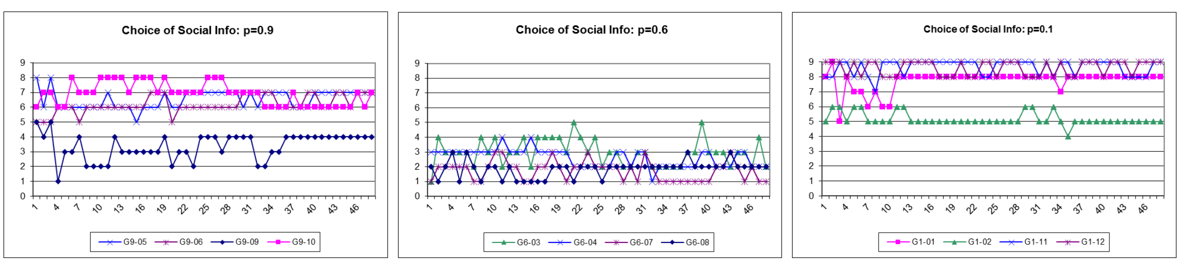

Figure 10: Between subjects: Dynamics of choices of social information by subjects groups in each of the three environments.

We now formulate our main experimental hypotheses and test them. The first hypothesis is that overall subjects behave like herd animals and always choose social information, or they behave like lone wolves and always choose private information, irrespective of the environment.

Hypothesis 1: Choice of information is independent of environment.

As we have seen, Figures 8, 9 and 10 cast much doubt on this hypothesis, as clearly the frequency with which subjects chose social information is higher in the PERSISTENT and NON-CONFORMIST environments than in the ERRATIC. More formally, consider the regressions reported in the third column of Table 5. There, the frequency of choice of social information, the variable Dinfosocial, is regressed on dummies for the environments, DS1 for NON-CONFORMIST and DS9 for PERSISTENT (ERRATIC is the baseline). The coefficients are positive and highly significant, indicating that social information was chosen more frequently in the NON-CONFORMIST and PERSISTENT environments, when it was optimal to do so, than in the ERRATIC treatment where it was not optimal. Thus, the hypothesis that behavior is entirely driven by fixed norms or rules of thumb of social behavior is, not surprisingly, rejected.

	Dinfocorrect		OptStrategy		Dinfosocial	
	Coef.	R.S.E.	Coef.	R.S.E.	Coef.	R.S.E.
DS1	1.131	1.307	0.484	0.372	4.545***	1.057
DS9	-0.488	0.902	-0.34	0.353	3.621***	0.989
Dfemale	-1.423***	0.538	-0.656*	0.365	-0.19	0.354
Age	0.326**	0.155	0.218**	0.093	-0.082	0.115
Dtech	-0.029	0.509	-0.217	0.313	1.653***	0.48
Dsoft	-0.098	0.546	-0.157	0.358	1.53**	0.632
CRT	0.324*	0.186	0.099	0.121	0.162	0.162
Stat	-0.156	0.306	-0.067	0.259	0.08	0.209
FHerd	0.004	0.202	0.038	0.182	0.616*	0.324
FRival	0.206	0.327	-0.079	0.166	0.50***	0.187
Round	0.008**	0.004	0.004*	0.003	0.002	0.003
LogL		-968.9		-2796.4		-973.7
Prob > chi2		0.0002***		0.000***		0.000***

Table 5: Regressions: the dependent variables are Dinfocorrect, the number of optimal choices of information made; Dinfosocial, the number of times social info is chosen. For the discrete independent variables, the baseline is a male social science major in the ERRATIC ($p = 0.6$) treatment. Significance at *: 0.10; **: 0.05; ***: 0.01

Clearly, subjects' aggregate choices are responsive to material incentives. But is this consistent across environments? The most important reason why this might not be the case is the possibility that there is a bias toward lone wolf or herd animal behavior. For example, as noted in Weizsäcker (2010), in previous social learning experiments subjects overweight private signals. So, if errors were biased in a similar direction here, then one would expect the choice of the optimal policy to be higher in the ERRATIC environment, where choosing private information is optimal, than in the PERSISTENT environment where it is not. Another possibility is differing complexity. For example, the optimal policy in the NON-CONFORMIST environment seems ex ante more difficult to calculate. So, one might think that the choice of optimal information might be lower here. Alternatively, the choices of entirely rational subjects would always be optimal and would not be affected by bias or complexity. Thus, the frequency of optimal choices should not depend on the environment. This is the way the hypothesis is formulated.

Hypothesis 2: The frequency of choice of optimal information (i.e., social information in the PERSISTENT and NON-CONFORMIST treatments and private information in ERRATIC treatment) is independent of the environment.

The first column of Table 5 reports regressions with the dependent variable Dinfocorrect being the frequency of optimal choice of information. One can see that the environment dummies DS1 for NON-CONFORMIST and DS9 for PERSISTENT are not significant. Thus, we cannot reject the hypothesis. Looking at Figures 8, 9 and 10, one can see that more subjects choose optimally in the NON-CONFORMIST environment than in the other two. However, this difference is not significant with the additional demographic controls included in the regressions. For example, the subjects in the NON-CONFORMIST environment seem to have been older, a factor associated with a greater likelihood of optimal choice. Note that this result in itself does not imply that subjects are rational. It is also consistent with the polar opposite where no-one ever chooses optimally. Rather, it just shows that the frequency of optimal choice does not vary significantly across environments.

We also remark that the time trend, Round, in Table 5, is small, positive and significant. So, optimal choices are increasing with time.

Our third main hypothesis is whether the deviations from optimality are systematic or completely random. Our strategy is to identify the difference based on several measures. First, we have demographic measures: age, gender and subject major. Second, we conducted some cognitive tests: a cognitive reflection test (CRT) and a test with problems from statistics, designed to test probabilistic reasoning. Third, we conducted a personality survey. We then ask whether these measures can predict either optimal choice of information or choice of social information. Note that if all subjects chose optimally, then none of these factors would be significant in the first regression in Table 5. Further, none of these factors should be significant in the second regression, choice of social information, only the environment dummies should be.

Hypothesis 3: Choice of social information depends on individual characteristics.

	Ave Age	M	F	Tech	Soc Sci	Soft	Ave CRT	Ave Stat
PERSISTENT	21.3	18	18	7	14	15	1.03	1
ERRATIC	20.8	20	16	15	10	11	1.11	1.11
NON-CONFORMIST	20.7	22	14	12	15	9	1.14	0.91

Table 6: Demographic Summary Statistics

We cannot reject this hypothesis because some demographic, cognitive and non-cognitive factors were clearly significant in subject choices. We start with demographics. Descriptive statistics for subjects by environment are given in Table 6: average age in years, then numbers of subjects by gender and subject major, and average scores on the two cognitive tests. Their subject majors have been aggregated into three broad classes: "Tech" includes science, mathematics and engineering; "Soc Sci" includes social sciences, particularly economics and business; "Soft" includes the humanities and arts.

We had no particular ex ante hypotheses about the demographic measures. However, we found that older subjects are apparently wiser, even though there is not much variation in age among our subjects who were undergraduates. Women performed worse than men.[11] But gender is not significant in the second regression on type of information chosen. Despite stereotypes, in this task women are not more socially inclined than men. Subject major is significant here, with science and humanities majors more likely to choose social information than social science majors. Thus, one stereotype, that business and economics majors are more individualistic, is in accordance with the data.

The two tests were a cognitive reflection test and a test on statistics and probability.Statistics scores, Stat in Table 5, are never significant, perhaps surprisingly. But the CRT scores do predict choice of optimal information. Thus, those subjects choosing suboptimally may be doing so because of lack of reflection and impulsive choice.

Lastly we consider the personality measures. We conducted a personality survey after the session. From the answers subjects gave, we constructed two factors, one relating to conformity FHerd, one to social rivalry, FRival. From Table 5, one can see that both factors are positively associated with choosing social information. The link with conformity, i.e., wanting to do the same as others, is very intuitive. The link with social rivalry can be understood by realizing that those that want to be ahead of others must have some interest in where others are. Thus, subjects who score relatively highly on one or both measures were more likely to choose social information, controlling for what is optimal. Thus, at least some of the deviations from what was monetarily optimal can be explained by preferences over social interaction.

[11] These results are unexplained but seem broadly consistent with those in common-value auctions, where also Bayesian-type inference may be important in outcomes. See for example Casari et al. (2007).

6 Conclusions

We have conducted experiments on social learning using a novel experimental design. Subjects have to choose whether to observe a private signal or the past choices of fellow subjects. This presents the two types of information symmetrically and thus avoid possible demand effects created by providing social information but giving each subject a private signal. By altering the persistence of the state, we alter the optimal choice of information to use. Mistakes can therefore run both ways: subjects can choose private information when social information is optimal and vice versa. This allows for a clearer identification of biases in subject behavior.

We find that there is no overall bias towards private information as has been reported in previous social learning experiments. This finding suggests that a demand effect may have played a role in these previous results. However, this interpretation should be treated with caution as there are several further differences between our experimental design and the previous social learning experiments which were all based on sequential choice. Further experiments, allowing a choice of information in sequential choice problem and modifying our current design to give subjects both forms of information, would be useful in disentangling these differences.

Most importantly, we find that there is considerable subject heterogeneity, with what we call lone wolves and herd animals being present alongside rational individuals. We argue that these deviations from optimal behavior are driven at least in part by taste, because of the relationships found between the choice of information and personality measures and subject majors. Lone-wolf behavior is more common among social science majors, many of whom are studying economics and business. This in accordance with the stereotype of economists being more individualist than, for example, humanities majors. Finally, our personality measures of both conformism and rivalry are associated with herd animal behavior. The conformism is easy to explain but note that in order to surpass others one needs to know what they are doing. Thus both the desire to be the same and the desire to outdo can drive the pursuit of social information.

We realize that some readers may view the observed heterogeneity in our experiment as reflecting "rules of thumb" rather than preferences with respect to social interaction. In contexts like this, it is difficult to separate these two. For example, if when trying to find one's way in an unknown city, some people always try to find their way with a map, while others ask passers-by for directions, does this express a preference over social interaction or over different rules of thumb? Equally here, having a preference for social over private information could be considered as a preference over different ways of approaching decision problem. We are open to either way of describing the observed behavior.

In conclusion, we find evidence in support of the notion that social influence is an important aspect of human behavior, but at the same time we find that its reverse, an aversion to social influence, also exists. Further, many subjects appear to have no intrinsic interest in or an aversion to learning from others' behavior, but simply use it when it useful. This parallels the literature on social preferences where inequality averse and competitive agents have been identified to exist alongside the entirely self-interested. We hope that the menagerie of social types introduced here will also find widespread applications.

References

Anderson, Lisa R., and Charles A. Holt (1997) "Information Cascades in the Laboratory", *American Economic Review*, 87(5), 847-62.

Banerjee, A.V. (1992) "A simple model of herd behavior", *Quarterly Journal of Economics*, 107 (3), 797-817.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992) "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades", *Journal of Political Economy*, 100(5), 992-1026.

Casari, Marco, John C. Ham and John H. Kagel (2007) "Selection Bias, Demographic Effects, and Ability Effects in Common Value Auction Experiments", *American Economic Review*, 97(4), 1278-1304.

Casari, M., Ham, J. C., and Kagel, J. H. (2007) "Selection bias, demographic effects, and ability effects in common value auction experiments", *American Economic Review*, 97(4), 1278-1304.

Celen, Bogachan, and Kyle Hyndman (2012) "Social learning through endogenous information acquisition: an experiment", *Management Science*, 58(8), 1525-1548.

Celen, Bogachan, and Shachar Kariv (2004) "Distinguishing Informational Cascades from Herd Behavior in the Laboratory", *American Economic Review*, 94(3), 484-98.

Cooper, David and Mari Rege (2011) "Misery loves company: Social regret and social interaction effects in choices under risk and uncertainty", *Games and Economic Behavior*, 73 (1), 91-110.

Duffy, John and Tatiana Kornienko (2010) "Does Competition Affect Giving?", *Journal of Economic Behavior and Organization* 74, 82-103.

Eyster, E., and M. Rabin (2010) "Naïve Herding in Rich-Information Settings", *American Economic Journal: Microeconomics*, 2, 221-243.

Goeree, Jacob K., Thomas R. Palfrey, Brian W. Rogers, and Richard D. McKelvey (2007) "Self-Correcting Information Cascades", *Review of Economic Studies*, 74(3), 733-62.

Goeree, Jacob K. and Leeat Yariv (2015) "Conformity in the Lab", *Journal of the Economic Science Association*, 1(1), 15-28.

Hertwig, Ralph, Greg Barron, Elke U. Weber and Ido Erev (2004), "Decisions from Experience and the Effect of Rare Events in Risky Choice", *Psychological Science*, 15 (8): pp. 534-539.

Hopkins, Ed and Tatiana Kornienko (2004) "Running to Keep in the Same Place: Consumer Choice as a Game of Status", *American Economic Review*, 94 (4), 1085-1107.

Huck, Steffen, and Jörg Oechssler (2000) "Informational Cascades in the Laboratory: Do They Occur for the Right Reasons?", *Journal of Economic Psychology*, 21(6), 661-71.

Kübler, Dorothea, and Georg Weizsäcker (2004) "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory", *Review of Economic Studies*, 71(2), 425-41.

Lahno, Amrei M. and Marta Serra-Garcia (2015) "Peer Effects in Risk Taking", *Journal of Risk and Uncertainty*, 50(1), 73-95.

Linde, J. and J. Sonnemans (2012) "Social comparison and risky choices", *Journal of Risk and Uncertainty*, 44 (1), 45-72.

Matějka, Filip and Alisdair McKay (2015) "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model", *American Economic Review*, 105(1), 272-298.

Miller, Joshua B. and Zacharias Maniadis (2012) "The Weight of Personal Experience: an Experimental Measurement", working paper.

Samuelson, Larry (2004) "Information-Based Relative Consumption Effects", *Econometrica*, 72 (1), 93-118.

Weizsäcker, Georg (2010) "Do We Follow Others when We Should? A Simple Test of Rational Expectations", *American Economic Review*, 100, 2340-2360.

Ziegelmeyer, Anthony, Frédéric Koessler, Juergen Bracht and Eyal Winter (2010) "Fragility of information cascades: an experimental study using elicited beliefs", *Experimental Economics*, 13, 121-145.