Even Experts Follow Large Herds: Naïve Herding in the Laboratory*

Christoph March[†] and Anthony Ziegelmeyer[‡]

Motivation

With the help of a large meta-dataset covering 13 experiments on social learning games, Weizsäcker (2010) investigates whether participants follow others and contradict their private information in situations where it is *empirically* optimal to do so. Weizsäcker finds that the success of social learning is very modest. The average participant follows others only in situations where the evidence conveyed by their observable choices is so strong that the private information is wrong more than twice as often as it is correct. Economic experiments on social learning games have repeatedly concluded that Bayesian rationality organizes well most of participants' choices except for an inflated tendency to follow private information (among others, Nöth and Weber, 2003; Goeree, Palfrey, Rogers, and McKelvey, 2007). By estimating the lucrativeness of the available actions, the meta-study additionally shows that participants forgo substantial parts of earnings when falling prey to the 'overweighting-of-private-information' bias.

The bulk of Weizsäcker's meta-dataset consists of experimental treatments that implement the stripped-down model of information cascades developed by Bikhchandani, Hirshleifer, and Welch (1992, henceforth BHW). In this simple social learning environment, a sequence of participants each in turn choose one of two options with each participant observing all of her predecessors' choices. Induced preferences over the two equally likely options are common, and participants receive independent and equally strong private binary signals about the correct option. According to Bayesian rationality, once the pattern of signals leads to two identical choices not canceled out by previous ones, all subsequent participants should ignore their signals and follow the herd. Though of interest, the experimental evidence on social learning behavior provided by the existing literature is too restrictive. Of particular concern is the coarseness of the social learning environment which favors the emergence of the 'overweighting-of-private-information' bias.[1]

*We are grateful to Sebastian Krügel who helped us conduct the pilot sessions of this study.

[†]Corresponding author: TUM-Munich University of Technology, Department of Business and Economics, Arcisstraße 21, 80333 Munich, Germany. Email: christoph.march@wi.tum.de

[‡]Queen's University Management School, Queen's University Belfast, 185 Stranmillis Road, Belfast BT9 5EE, United Kingdom. Email: email@anthonyziegelmeyer.com

[1]In situations where predecessors' choices do not point in any direction or point in the same direction as private information, following private information seems the only reasonable choice. In these situations, the few experimental choices not in line with private information have been understood as resulting from confusion. Moreover, in situations where an option is favored by exactly one choice over the other option and private information points in the opposite direction, Bayesian rationality is silent about the optimal choice. Note that about one third of the data in Weizsäcker's meta-dataset stem from Nöth and Weber (2003) which considers two privately known signal precisions. However, the predictions in this variant of BHW's stripped-down model closely match the original ones.

In this paper, we investigate whether participants follow their private information and contradict the herd in situations where it is empirically optimal to do so. To address this complementary issue, our social learning game relies on a richer information structure than BHW's stripped-down model. Following Ziegelmeyer, Koessler, Bracht, and Winter (2010), we consider two sequences of decision makers, an *observed* and an *unobserved* sequence. Observed decision makers sequentially guess which of two options has been randomly chosen with the help of a medium quality binary signal (quality equals 14/21). At the end of each decision period, the choice of one observed is made public knowledge. In a matched pairs design, unobserved decision makers guess which of two options has been randomly chosen knowing previous public choices and with the help of a low, medium or high quality binary signal (quality equals 12/21, 14/21 or 18/21 respectively). Their choices remain private.

Design

Our laboratory experiment uses an expanding strategy method-like procedure which allows us to detect herding behavior directly, allows participants to gain experience with each information set {precision and realization of signal, history of observed choices}, and generates a large dataset (see also Cipriani and Guarino, 2009). In the first part of each of these sessions, the signal's quality for the unobserved is fixed at the beginning of each of the three rounds, each decision maker observes only one signal realization and makes only one choice. Each participant earns 1 (0) Euro for each correct (wrong) guess. The second part of a session is identical to the first part except that i) all unobserved make one choice in each decision period (8 choices in total); ii) all seven observed make one choice in decision period 1 and one choice is randomly selected to be made public, the remaining six observed make one choice in decision period 2 and one choice is randomly selected to be made public, and so on till decision period 7 where the remaining observed makes a last choice; and iii) for each participant, only one randomly selected choice is paid in each round. The third part of a session is identical to the second part except that each choice is made for *both* realizations of the private signal. Decision makers are informed of the payoff-relevant realization of their private signal after having made their last choice. Finally, the fourth part of each session is identical to the third part except that i) unobserved make their choices for *each* quality of the private signal and they are informed of the payoff-relevant quality of their private signal at the end of each round; and iii) for each participant, only one randomly selected choice is paid and each participant earns 3 (0) Euro for a correct (wrong) guess.

A second novelty of our design is the use of preprogrammed computers as observed decision makers in 9 out of the 15 experimental sessions. Unobserved decision makers, on the other hand, are always embodied by human participants. Since the computers' strategy is fully revealed to the unobserved participants, they face lower behavioral uncertainty in the computer-human treatment than in the human-human treatment. Computers act in line with the Bayesian equilibrium strategy in six sessions which implies that, no matter how big the contradicting herd is, it is always beneficial for the unobserved participants to follow their high quality signal. Given the experimental choices of the observed participants, the estimation of the empirical value of actions leads to the conclusion that following the high quality signal is also the empirically optimal action for the unobserved participants in the human-human treatment no matter how big the contradicting herd is. In the remaining three sessions computers follow their private signal at each history (observed decision makers observe signals

Treatment Name	Observed	# Sessions
Human-Human Strategy Method (*HHSM*)	Human Participants	6
Computer-Human Strategy Method (*CHSM*)	Bayes-Rational Computers	6
Computer-Human Strategy Method with Revealed Private Signals (*CHSM signals*)	Computers following private signal	3

Table 1: Overview of the treatments.

of low signal quality) which further reduces behavioral uncertainty and enables us to study the impact of biases in belief updating.

Overall we collected 32,118 decisions made by 162 participants in 15 sessions and 3 treatments. Table 1 gives an overview of the treatments.

The richness of our dataset enables us to measure the success of social learning both in situations where it is empirically optimal to follow others (and contradict private information) and in situations where it is empirically optimal to follow private information (and contradict the herd).[2]

Results

In the *HHSM* treatment we find that conditional on being endowed with a *low* or *medium* quality signal and observing a contradicting herd of size at least 2, participants make the empirically optimal choice in 72 percent of the cases. In contrast, conditional on being endowed with a *high* quality signal and observing a contradicting herd of size at least 2, unobserved participants choose optimally only in 57 percent of the cases. In the latter situations, the evidence conveyed by the observable choices is so weak that the private information is correct almost twice as often as it is wrong. Our new evidence therefore suggests that participants are prone to a 'social-confirmation' bias and it gives support to the argument that they naïvely believe that each observable choice reveals a substantial amount of that person's private information (Eyster and Rabin, 2010). Though both the 'overweighting-of-private-information' and the 'social-confirmation' bias coexist in our data, participants forgo much larger parts of earnings when herding naïvely than when relying too much on their private information. Concretely, when endowed with a *low* or *medium* quality signal participants earn 0.67 of the normalized pie size when they could have earned 0.73 had they always followed the empirically optimal action. Moreover conditional on observing a contradicting herd of size at least 2 participants earn 0.62 out of the 0.71 they could have earned by choosing optimally. On the other hand conditional on being endowed with a *high* quality signal participants earn 0.62 out of the 0.83 they could have earned by choosing optimally, and they earn 0.40 out of 0.68 conditional on observing a contradicting herd of size at least 2.

The results are confirmed in the *CHSM* treatment where the presence of computer players enables us to exactly determine the incentives participants face. Conditional on being endowed with a *low* or *medium* quality signal and observing a contradicting herd of size at least 2 participants make

[2]We rely on Weizsäcker's (2010)'s estimation technique to control for incentives when assessing the empirically optimal action. When few observations are available this estimate may be far from its expected value. Unless otherwise specified the results for the *HHSM* treatment rely on situations which occur in at least three distinct repetitions of the social learning game (out of 72 repetitions).

the empirically optimal choice in 83 percent of the cases, and they earn more than what they could have earned had they always picked the empirically optimal choice. In contrast, conditional on being endowed with a *high* quality signal and observing a contradicting herd of size at least 2, unobserved participants choose optimally only in 49 percent of the cases, and they earn only 64 percent of what they could have earned by choosing optimally.

Finally, the results for the treatment *CHSM signals*, where the preprogrammed computers always choose in line with their private signal, reveal that the 'social-confirmation' bias is only partly due to participants' lack of understanding of the herding decisions of the observed. In this treatment conditional on observing a contradicting herd of size at least 2, participants make the empirically optimal choice in 73 percent of the cases when endowed with a *low* or *medium* quality signal, and they choose optimally in 68 percent of the cases when endowed with *high* quality signal.

Overall, unobserved participants make the empirically optimal choice in 85 percent of the cases in the *HHSM* and the *CHSM* treatment, and in 88 percent of the cases in the *CHSM signals* treatment which suggests that social learning improves only slightly in the presence of lower behavioral uncertainty.

Thanks to the large amount of data collected at the individual level, we are also able to classify the social learning behavior of each of our unobserved participants. Though a substantial fraction of participants (almost) always follow their private information, we also find a substantial fraction of social-conformists who follow contradicting herds of any size. Clearly, some unobserved participants drew opposite conclusions from the same evidence. Further experimental work on social learning should dig deeper into this heterogeneity in informational inferences.

References

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992–1026.

Cipriani, M., and A. Guarino (2009): "Herd Behavior in Financial Markets: An Experiment with Financial Market Professionals," *Journal of the European Economic Association*, 7, 206–233.

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

Goeree, J., T. Palfrey, B. Rogers, and R. McKelvey (2007): "Self-Correcting Information Cascades," *Review of Economic Studies*, 74, 733–762.

Nöth, M., and M. Weber (2003): "Information Aggregation with Random Ordering: Cascades and Overconfidence," *Economic Journal*, 113, 166–189.

Weizsäcker, G. (2010): "Do We Follow Others when We Should? A Simple Test of Rational Expectations," *American Economic Review*, 100, 2340–60.

Ziegelmeyer, A., F. Koessler, J. Bracht, and E. Winter (2010): "Fragility of Information Cascades: An Experimental Study Using Elicited Beliefs," *Experimental Economics*, 13, 121–45.