

March, Christoph; Ziegelmeyer, Anthony

Working Paper

Excessive Herding in the Laboratory: The Role of Intuitive Judgments

CESifo Working Paper, No. 6855

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: March, Christoph; Ziegelmeyer, Anthony (2018) : Excessive Herding in the Laboratory: The Role of Intuitive Judgments, CESifo Working Paper, No. 6855, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/174978>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Excessive Herding in the Laboratory: The Role of Intuitive Judgments

Christoph March, Anthony Ziegelmeyer

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Excessive Herding in the Laboratory: The Role of Intuitive Judgments

Abstract

We designed four observational learning experiments to identify the key channels that, along with Bayes-rational inferences, drive herd behavior. In Experiment 1, unobserved, whose actions remain private, learn from the public actions made in turn by subjects endowed with private signals of medium quality. We find that when unobserved face a handful of identical actions that contradict their high quality signals they herd more extensively than predicted by Bayes-rational herding. Deviations from the normative solution result in severe expected losses and unobserved would be better off without the chance to learn from others. When unobserved are endowed with medium quality signals they learn rather successfully from public actions, but they overweight their low quality signals relative to public information. Experiments 2-4 reveal that non-Bayesian updating and informational misinferences are the two channels that drive excessive herding, while the strong (resp. mild) overemphasis on low (resp. medium) quality signals is caused by wrong expectations about others' strategy. A model of intuitive observational learning accounts for the phenomenon of excessive herding, it captures well herd behavior with medium quality signals, but it fails to predict that the reluctance to contradict private signals is stronger for low than for medium quality.

JEL-Codes: C920, D820, D830, D840.

Keywords: observational learning, herd behavior, intuitive judgments, experiments.

Christoph March
TUM School of Management
Technical University of Munich
Arcisstraße 21
Germany – 80333 Munich
christoph.march@tum.de

*Anthony Ziegelmeyer**
Queen's Management School
Queen's University Belfast
185 Stranmillis Road
United Kingdom – Belfast BT9 5EE
email@anthonyziegelmeyer.com

*corresponding author

January 8, 2018

We are grateful to Jörg Cyriax, Anh Do, Fabian Frank, Florian Inderst, Maria Mircea, Leonard Przybilla, Alexander Schlimm, Martha Splitthoff, and Christina Strobel for excellent research assistance. We thank Thomas Daske, Dirk Engelmänn, Bernd Irlenbusch, Andreas Ostermaier, Stefan Penczynski, Georg Weizsäcker, and seminar participants at University of Bamberg, Humboldt University Berlin, Technical University of Munich, University of Munich, the Thurgau Experimental Economics Meeting in Kreuzlingen, the EEA Annual Congress in Mannheim, the ESA European Meeting in Heidelberg, the European Winter Meeting of the Econometric Society in Milan, and the European Meeting of the Econometric Society in Geneva for valuable comments.

1 Introduction

Economists have long recognized that people often adjust their behavior to conform with the choices of others and that social costs become substantial when herding spreads through large segments of the population. Early informal analyses insisted upon the psychological and sociological aspects of herding (Keynes, 1930; Kindleberger, 1978). In contrast, a more recent theoretical literature, beginning with Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (BHW 1992), shows that herd behavior can be a rational response to the information contained in others' actions. When a sequence of Bayes-rational agents each in turn take one of several actions, and learn from their predecessors' actions, an information cascade quickly occurs in which they abandon their private information and take imitative actions.¹ This striking implication of Bayes-rational observational learning has been tested in a series of economic experiments. The bulk of this laboratory evidence is summarized in Weizsäcker (2010) and Ziegelmeyer, March, and Kruegel (2013), two meta-studies that measure the success of observational learning by controlling for the empirical profitability of actions. The main finding is that laboratory cascades eventually emerge, though subjects follow others only when the monetary incentives to do so are strong enough.

The experimental evidence on simple cascade games informs only partially our understanding of herd behavior since these experiments merely test whether subjects imitate their predecessors or not. However, as repeatedly argued in BHW (1992), Bayes-rational herding is unique in that the conformity of behavior caused by an information cascade does not become more robust as the number of imitative actions increases. This property allows some of the benefit of information diversity to be recaptured as Bayes-rational herds are easily overturned by better informed agents.² In this paper, we report the results from four experiments that separate Bayes-rational herding from less rational forms of herding by investigating whether subjects imitate others excessively.

In Section 2, we show that when subjects learn from a handful of identical actions that contradict their precise signal they herd more extensively than predicted by Bayes-rational herding, and we quantify the inefficiencies associated with such a stronger and more robust form of herding. Our experimental setting builds on a rich information structure which, for a given string of previous actions, distinguishes between situations where Bayes-rational herding occurs and situations where it does not. Concretely, we implement an observational learning scenario with two parallel sequences of players who submit binary guesses (either B or O) about an unknown payoff-relevant state (either $\mathcal{B}$ or $\mathcal{O}$ with $\Pr(\mathcal{B}) = 0.55$). Seven *observed* players receive binary private signals (either b or o) of medium quality, and their guesses are made public meaning that they learn from previous *observed* guesses as in simple cascade games. Eight *unobserved* players also learn from the public guesses made by the *observed*, but their own guesses remain private and they receive binary private signals whose quality is either low, medium or high (we have that $\Pr(o | \mathcal{O}) = \Pr(b | \mathcal{B}) = 1 - \Pr(b | \mathcal{O}) = 1 - \Pr(o | \mathcal{B}) = 12/21, 14/21$ and $18/21$ when the quality is low, medium and high respectively). If the player's guess correctly predicts the state—her guess is B when the state is $\mathcal{B}$ or her guess is O when the state is $\mathcal{O}$ —the player gets a payoff of 1, otherwise she gets nothing. Bayesian rationality predicts that, for most strings of public guesses, *unobserved* endowed with low or medium quality signals imitate the most recent guess when the latter conflicts with their private information. Bayes-rational herding often occurs with low and medium quality signals. Per contra, Bayes-rational *unobserved* who receive signals of high

¹If agents' actions are always sufficient statistics for their information, then learning from others' actions is efficient (Lee, 1993). In settings where economic outcomes are inefficient, information cascades need not arise as observational learning is asymptotically complete with unbounded private signals (Smith and Sørensen, 2000).

²Said differently, Bayes-rational herding is a weak form of herding as it denies the possibility of extreme confidence in wrong beliefs. Eyster and Rabin (2014) even show that Bayes-rational herding is a limited form of herding since, in the case of observation structures more general than the canonical single-file setting, the logic of social inference requires that Bayes-rational agents greatly limit the scope of their imitation.

quality never imitate the guesses they observe, but systematically follow their private information. Bayes-rational herding never occurs with high quality signals.

To measure how successful subjects are in learning from others, we first estimate the incentives to contradict private information and then we assess the proportion of guesses that contradict private information in situations where it is empirically optimal to do so and in situations where following private information is empirically optimal. A reliable measure of the empirical success of observational learning requires a large dataset. Hence, we employ a strategy method-like procedure according to which, in the same repetition of the scenario, subjects learn from various strings of public guesses while endowed with the same private signal. Estimated incentives reveal that herds of public guesses aggregate at most two private signals which implies that following private information is always empirically optimal for *unobserved* with high quality signals.

The main results of our first experiment are as follows. First and foremost, *unobserved* with high quality signals herd excessively: when she observes an excess of at least four public guesses that contradict her high quality signal, the average *unobserved* follows the public majority. Deviations from Bayes-rational herding result in severe expected losses and *unobserved* with high quality signals would be better off without the chance to learn from others. Second, though they fall short of better responding to the value of their available information, *unobserved* with medium quality signals are quite successful in learning from others. In comparison, the average *observed* is clearly reluctant to contradict her private information when the expected monetary costs of making an informative guess are low to moderate. Third, given identical incentives to follow others, *unobserved* are more reluctant to contradict their private information with low than with medium quality signals. In particular, the average *unobserved* follows a single contrary guess slightly more often than her low quality signal, though her incentives to contradict private information are 1.75 times stronger than her incentives to follow private information. Fourth, individual herd behavior is substantially heterogeneous. *Unobserved* are almost equally divided into successful observational learners, conformists who tend to herd excessively, and dissenters who respond too strongly to their private information. This being said, the tendencies to herd excessively with high quality signals and to respond too strongly to low quality signals are common among *unobserved*.

In view of this evidence, we conclude that *unobserved* mislearn from public guesses conflicting with their private signals, especially when the latter are of high quality, because they wrongly assess the value of their available information in those challenging situations. Misvaluations of the available information may be caused by each of three potential deviations from Bayes-rational herding. First, *unobserved* may have an incorrect model of how others make guesses (Kübler and Weizsäcker, 2004; Bohren, 2016). Second, even if their expectations about others' strategy are correct, *unobserved* may use improper inference rules leading them to infer signals that differ systematically from the ones actually received by others (Eyster and Rabin, 2005, 2010). Third, *unobserved* may not adhere to Bayes' rule when incorporating into their beliefs the signals inferred from public guesses (Huck and Oechssler, 2000; Goeree, Palfrey, Rogers, and McKelvey, 2007). In Section 3, we report three experiments that uncover the nature of these deviations in our observational learning scenario, and we evaluate the impact that each deviation has on herd behavior. As in Experiment 1, *unobserved* are endowed with private signals of low, medium or high quality, and they learn from strings of public guesses. However, the three experiments rely more extensively on the strategy method-like procedure to offer subjects ample opportunities to learn from strings of contrary guesses. Experiments 2-4's evidence aims at clarifying how herd behavior is shaped by the interplay of the three components of belief formation, namely belief updating, informational inferences, and expectations about others.

In Experiment 2, *unobserved* learn from public signals. Thus, the components informational inferences and expectations about others are turned off. The observable-signals scenario unveils the

nature of the belief updating rules used by *unobserved* and the extent to which they undermine their observational learning success. In Experiment 3, *unobserved* learn from public guesses known to be made by computer-*observed* that adopt the Bayes-rational strategy. Informational misinferences and non-Bayesian updating potentially impact herd behavior in the second scenario, but the component expectations about others remains turned off. The comparison of the observational learning success in Experiments 2 and 3 reveals the nature of informational misinferences and the extent to which they cause *unobserved* to misvalue their available information over and above non-Bayesian updating. Finally, in Experiment 4, *unobserved* learn from guesses made by other subjects which implies that the expectations component is turned on. By comparing their observational learning success in Experiments 3 and 4, we shed light on the model that *unobserved* have about how others make guesses and we assess the extent to which uncertainty about others’ strategy impacts herd behavior.

Experiments 2-4’s results show that non-Bayesian updating, informational misinferences, and incorrect expectations about others all undermine the success of observational learning, though in different ways. First, non-Bayesian updating and informational misinferences are the two channels that drive excessive herding with high quality signals. In Experiment 2, the average *unobserved* acts against her precise signal when the incentives to follow private information are 1.40 times stronger than the incentives to contradict private information. Excessive herding is even more pronounced in Experiment 3 since the average *unobserved* contradicts her high quality signal when the incentives to contradict private information are further reduced by 60%. On the other hand, *unobserved* with high quality signals learn as successfully from guesses submitted by other subjects as from public signals. Uncertainty about others’ strategy mitigates excessive herding. Second, the overemphasis on low and medium quality signals is mainly caused by incorrect expectations about others’ strategy. In Experiment 3, the average *unobserved* with a low or medium quality signal learns rather successfully from short strings of contrary guesses indicating that she correctly appreciates the connection between the signals of Bayes-rational *observed* and their informative guesses. In Experiment 4, however, the average *unobserved* wrongly believes that *observed* respond weakly to their private information which reduces her success when learning from a few contrary guesses. As a final observation, and in line with herd behavior at the aggregate level, we find the largest proportion of conformists in Experiment 3 (44% versus 14% and 34% in Experiment 2 and 4 respectively) and the largest proportion of dissenters in Experiment 4 (49% versus 32% and 16% in Experiment 2 and 3 respectively).

In Section 4, we present a structural model of intuitive observational learning which incorporates non-Bayesian updating, informational misinferences, and wrong expectations about others’ strategy. In line with the central idea of the “heuristics and biases” program (Kahneman, Slovic, and Tversky, 1982), we posit that when people learn from others’ actions they form probability judgments that are reflexive and effortless and that these intuitive judgments are often supplemented and sometimes overridden by more deliberate and taxing judgments. Still, deliberate probability judgments are likely to remain anchored on their intuitive counterparts (Kahneman, 2003). We should emphasize that the cognitive processes underlying intuitive observational learning are not merely simpler than the ones underlying Bayes-rational observational learning, they are different in nature since they rest on general-purpose heuristics such as the representativeness heuristic. Because of their non-normative nature, intuitive valuations of the available information are systematically biased.

We consider a logit quantal response version of intuitive observational learning where each belief component is formulated as a one-parameter extension of its Bayes-rational counterpart. So, our alternative model embeds the normative one as a constellation of parameter values, which are individual-specific to accommodate the rich behavioral heterogeneity found in our experiments. As to the expectations about how public guesses are made, we postulate that intuitive observational learners—henceforth *Intuitive*—don’t entertain the possibility that others’ beliefs are distorted, but they simply

expect others to play (homogeneous) logit quantal-response equilibrium strategies. *Intuitive* attribute to others a payoff-responsiveness that is either strictly smaller or larger than their own meaning that expectations are heterogeneous. Additionally, we assume that when *Intuitive* draw an inference from a public guess, the signal more frequently associated with that guess comes foremost to their mind. This assumption captures the logic of the representativeness heuristic. Indeed, a public guess that is more frequently associated with a given signal is more representative of that signal than of the other signal. Given their simple view of how others make guesses, *Intuitive* regard guess B as more representative of signal b than of signal o . Therefore, *Intuitive* make an informational inference that is biased towards signal b when they observe guess B (and they make a biased inference towards signal o after guess O). Our approach follows Gennaioli and Shleifer (2010) who examine how representativeness distorts the assessed probabilities of alternative hypotheses. Borrowing from their terminology, *Intuitive* are “local thinkers” who make informational inferences in light of what comes to mind, but not of what does not. The extent to which an *Intuitive* misinfers from public guesses depends on her degree of local thinking. Finally, we posit that *Intuitive* weight the signals inferred from public guesses differently than advocated by Bayes’ rule. Depending on the value of their public information weight, *Intuitive* treat inferred signals as either less or more informative than would a Bayesian updater. Though the weighted updating rule is an overly simplified formalization of non-Bayesian updating, it captures typical probability judgment biases such as conservatism (Edwards, 1968) or the “belief in the law of small numbers” (Rabin, 2002), the latter being a possible consequence of the reliance on the representativeness heuristic (Tversky and Kahneman, 1971).

We use maximum-likelihood techniques to estimate, for each *unobserved* in Experiment 4, her parameters of the intuitive observational learning model. Our structural estimation results reveal a rich diversity in the weighting of public information and in the degree of local thinking. Slightly more than a quarter of the *unobserved* perceive the informativeness of inferred signals as (at most) four times lower than would a Bayesian updater, whereas almost a sixth of them treat inferred signals as (at least) fifty percent more informative than would a Bayesian updater. And though more than half of the *unobserved* have a high degree of local thinking, almost a fifth of them make proper informational inferences. On the other hand, expectations about others’ strategy are quite homogeneous since more than two-thirds of the *unobserved* attribute to others the same payoff-responsiveness as their own. Lastly, we measure the accuracy of our estimated model’s predictions relative to the guesses made by *unobserved* in Experiment 4. We find that intuitive observational learning has a much higher predictive power than the theoretical benchmark which assumes that players probabilistically best respond to the correct value of their available information. Most notably, the estimated model captures well the phenomenon of excessive herding with high quality signals. Its strong predictive power in those guessing situations is driven by the high degrees of local thinking and the non-negligible overweighting of public information. The behavior of *unobserved* with medium quality signals is also well captured by intuitive observational learning. In contrast, the model fails to predict that, when the incentives to follow others are identical, observational learners are more reluctant to contradict their low than their medium quality signals. By comparing the predictive power of alternative specifications of the expectations component, we also find that the nature of expectations about others’ strategy hardly affects the ability of intuitive observational learning to predict accurately.

We conclude in Section 5 by discussing the robustness of our prediction results with respect to the modelling assumptions of intuitive observational learning. The supplementary material contains a series of appendices with complementary data analyses and proofs.³

³The supplementary material is available at https://www.vwl.wi.tum.de/fileadmin/w00bpu/www/team/march/ExcessiveHerding_SM.pdf.

Related literature

Since Weizsäcker (2010) summarized the early literature on laboratory cascades, several studies, complementary to ours, attempted at disentangling rational from less rational forms of herding.⁴

Brunner and Goeree (BG 2011) test the model of Callander and Hörner (2009) which predicts that, when they are differentially informed and they observe the number of times each option has been chosen (rather than the sequence of previous choices as in BHW), rational observational learners may discard their private information and follow the minority instead of the majority. Contrary to this prediction, their subjects tend to follow the majority, i.e., they don't believe that majorities are more likely to be wrong than right. Likewise, our subjects doubt the fallibility of crowds when they contradict their high quality signals and follow large contrary majorities. Yet, BG (2011) report that deviations from theoretical predictions are approximate best responses to the empirical distribution of play whereas excessive herding with high quality signals results in severe expected losses.

Eyster, Rabin, and Weizsäcker (ERW 2015) investigate experimentally whether people appreciate the redundancy of information conveyed by others' actions. A novel observational learning scenario is introduced where subjects in the current period observe the set of actions chosen by a group of subjects in the previous period. Though rational observational learning requires anti-imitation, there is strong evidence for redundancy neglect, which creates excessive imitation, and the average subject would earn more by ignoring others' actions. Our finding that *unobserved* with high quality signals herd excessively and would be better off without the chance to learn from others clearly corroborates their results. In their setting, however, rational players aim at learning the signals of their predecessors rather than an underlying payoff-relevant state and observational learning does not require the use of Bayes' rule. Also, ERW (2015) interpret the deviations from optimality as errors in higher-order reasoning whereas we emphasize the role of intuitive judgments in observational learning.

In an observational learning scenario with private signals of identical quality, Angrisani, Guarino, Jehiel, and Kitagawa (AGJK 2017) elicit a subject's beliefs both before and after she receives her private information. Their experimental results are at odds with the interpretation that early decisions in the sequence have an undue influence on later decisions (Eyster and Rabin, 2010), but they support a model where players believe they have a higher ability to understand the private signal than their predecessors. Our finding that the overemphasis on low and medium quality signals is mainly caused by incorrect expectations about others' strategy is clearly in line with AGJK (2017)'s results.

Duffy, Hopkins, Kornienko, and Ma (2017) report an observational learning experiment where subjects must choose between receiving a private signal or seeing their predecessors' guesses. The authors find that aggregate behavior approximates the equilibrium predictions fairly closely, though some subjects consistently choose social information and others consistently choose private information. Heterogeneity in individual herd behavior is also a recurrent finding in our experiments.

2 Laboratory Evidence on Excessive Herding

Laboratory experiments on simple cascade games have confirmed that people herd for informational reasons, at least when the actions of their predecessors strongly conflict with their private signal. Yet, Bayes-rational herding is pervasive in these observational learning scenarios, implying that alternative

⁴Kübler and Weizsäcker (2005) report that in many experiments that follow the basic model by BHW (1992) the strength of laboratory cascades is positively correlated with their length. They also document that this correlation leads subjects to imitate their predecessors though the normative solution prescribes to follow one's own information. However, past evidence on excessive herding is not entirely convincing. Indeed, to the best of our knowledge, in all former experimental settings where excessive herding has been observed the empirical incentives to follow private information were (at best) slightly larger than the empirical incentives to follow others.

forms of herding were mostly unexplored. Per contra, our first experiment has been expressly designed to explore the extent to which forces other than Bayes-rational inferences drive herd behavior, in a setting where the expected costs of excessive imitation are substantial.

2.1 Theory and Experimental Design

Experiment 1 extends the classic “ball and urn” setup, initially proposed by Anderson and Holt (1997), in two major ways. First, (almost) each subject submits multiple guesses about the unknown event in different periods, though only one of her guesses is payoff-relevant. This strategy method-like procedure, introduced in Cipriani and Guarino (2009), generates a dataset large enough to assess the empirical success of observational learning. Second, similarly to Ziegelmeyer, Koessler, Bracht, and Winter (2010), there are two groups of subjects who submit their guesses in two parallel sequences. *Observed* receive private signals of medium quality and they learn from previous *observed* guesses as in BHW’s cascade game. *Unobserved* also learn from *observed* guesses but their own guesses remain private and they receive private signals whose quality is either low, medium or high.

After introducing our laboratory game, we detail the progress of a session and the experimental procedures.

2.1.1 The Laboratory Game and Its Predictions

We first describe the “2 Sequences and 3 Qualities” observational learning game (henceforth *2S3Q* game) played by subjects during non-practice rounds and then we derive its Bayes-rational predictions.

There are two possible payoff-relevant states of Nature—state *BLUE* and state *ORANGE*—which we denote by states $\mathcal{B}$ and $\mathcal{O}$. There are two groups of players: A group of seven *observed* players and a group of eight *unobserved* players. Each player receives a private signal which is a ball drawn from an urn whose composition depends on the state and on the player’s group. In state $\mathcal{B}$ each *observed* draws a ball from an urn which contains n_{Obs} blue balls and $(21 - n_{Obs})$ orange balls whereas each *unobserved* draws a ball from an urn which contains n_{Unobs} blue balls and $(21 - n_{Unobs})$ orange balls with $21 > n_{Obs}, n_{Unobs} \geq 12$. In state $\mathcal{O}$ each *observed* draws a ball from an urn which contains $(21 - n_{Obs})$ blue balls and n_{Obs} orange balls whereas each *unobserved* draws a ball from an urn which contains $(21 - n_{Unobs})$ blue balls and n_{Unobs} orange balls. If a player draws a blue ball then we denote her signal by b and if a player draws an orange ball then we denote her signal by o . The ratio $n_{Obs}/21$ corresponds to the signal quality of *observed* and the ratio $n_{Unobs}/21$ corresponds to the signal quality of *unobserved*. Different parametrizations of the game rely on different values of the *observed* and *unobserved* signal qualities where, for a given parametrization, both are public knowledge. As detailed below, in each session of Experiment 1 subjects play three parametrized versions of the game with $n_{Obs} = 14$ and $n_{Unobs} \in \{12, 14, 18\}$.

At the start of the game Nature randomly selects one of the two states with $\Pr(\mathcal{B}) = 0.55$ and each player receives a private signal. The randomly selected state remains unknown to the players. Then players submit guesses about the state randomly selected by Nature over eight periods. There are two possible guesses—“guess state *BLUE*” and “guess state *ORANGE*”—which we denote by guesses B and O . In period 1 all fifteen players simultaneously submit a guess. At the beginning of period 2, one of the seven *observed* guesses submitted in period 1 is publicly revealed and the *observed* whose guess becomes public stops guessing. Then all fourteen remaining players simultaneously submit a guess. At the beginning of period 3, one of the six *observed* guesses submitted in period 2 is publicly revealed and the *observed* whose guess becomes public stops guessing. Then all thirteen remaining players simultaneously submit a guess. And so on, until period 8 where all *unobserved* simultaneously submit a guess. Note that the guesses of *unobserved* remain private and that players who act in period

$t \in \{1, \dots, 8\}$ face history $h_t = (g_1, \dots, g_{t-1}) \in H_t = \{B, O\}^{t-1}$, with $h_1 = \emptyset$, where g_τ is the *observed* guess submitted in period $\tau \in \{1, \dots, t-1\}$ and made public at the beginning of the next period. Every *observed* guess submitted in period $t \in \{1, \dots, 7\}$ is equally likely to be publicly revealed at the beginning of the next period meaning that the selection probability equals $1/(8-t)$.

Once all guesses have been submitted, payoffs are realized. For each player only one of her guesses is *payoff-relevant*. The payoff-relevant guess of each *observed* is the last guess she submitted, i.e., the guess which has been publicly revealed. On the other hand, each *unobserved* is randomly assigned to one of the eight periods, a different one for each *unobserved*, and her payoff-relevant guess is the one she submitted in that period (which she learns about only after all guesses have been made). We assume that players are expected utility maximizers and that each payoff-relevant guess B has vNM payoffs $u(B, B) = 1$ and $u(B, O) = 0$ whereas each payoff-relevant guess O has vNM payoffs $u(O, B) = 0$ and $u(O, O) = 1$. Thus, our laboratory game is strategically equivalent to an observational learning game where, in each group, players submit guesses in an exogenously given order.

Predictions. Assume that players update their beliefs using Bayes' rule and that they maximize their expected payoffs conditional on those beliefs. If we assume further that this fact and the game structure are commonly known, then the 2S3Q game has a unique rationalizable outcome. From now on, we refer to this joint assumption simply as Bayesian rationality. In period 1, *observed* guess in accordance with their private information, signal o leads to guess O whereas signal b leads to guess B . In period 2, *observed* follow their private information if the history is (O) and they submit guess B if the history is (B) . Hence, if guess B is publicly revealed at the beginning of period 2, an information cascade starts with all subsequent *observed* submitting guess B . In period 3, an information cascade starts at history (O, O) with all subsequent *observed* submitting guess O whereas *observed* follow their private information at history (O, B) . In later periods, a B -cascade (resp. O -cascade) starts after any history with one B (resp. two O s) not canceled out by previous guesses. The only history which does not lead to an information cascade is the history $(O, B, \dots, O, B)$, and the probability of no cascade decreases exponentially.⁵

Unobserved with signal quality 12/21 imitate the most recent public guess except in the first period or after history $(O, B \dots O, B)$ where they follow their private information. *Unobserved* with signal quality 14/21 act like *observed*. Finally, *unobserved* with signal quality 18/21 follow their private information at all histories.

2.1.2 The Progress of an Experimental Session

Each session starts with three paid practice rounds during which the 15 subjects familiarize themselves with the observational learning task. In a given practice round, each subject submits a guess about one of two possible states after having observed a sequence of previous guesses (the public information) and a ball drawn from a physical urn whose composition depends on the unknown state (her private information). After practice, subjects play six repetitions of the 2S3Q game in each of three parts which differ according to the quality of private signals that *unobserved* receive. The strategy method-like procedure allows subjects to learn from various sequences of previous guesses for a given private signal. In particular, *unobserved* gain extensive experience with the combination of private and public information as they learn from various sequences of *observed* guesses with three different signal qualities. Note that to reduce confusion and errors, there is no mentioning of the three parametrizations of the 2S3Q game at the start of the session meaning that subjects are informed

⁵Having an asymmetric prior in the 2S3Q game has two advantages. First, there is no need to consider a commonly known tie-breaking rule as Bayes-rational players are never indifferent. Second, information cascades are more likely to emerge than in a game with a flat prior where indifferent players follow their private information.

about the prevailing parametrization only at the beginning of the part.

Practice rounds: In each of the three practice rounds *unobserved* receive private signals of quality 14/21 and each subject submits only one guess. Concretely, in each of the first seven decision periods one *observed* and one *unobserved* submit a guess, and in the last decision period only the remaining *unobserved* submits a guess. Assignments to decision periods are random. Participants receive 1 Euro for a correct guess and 0 Euro otherwise. At the end of the practice round, each subject is reminded of her private signal, her guessing period, the guess she made and the sequence of *observed* guesses, and she is informed about the realized state and her earnings. Additionally, feedback screens of *observed* display the composition of the urn used in the *observed* sequence whereas feedback screens of *unobserved* display the urn composition in each sequence.

Part 1: Medium quality signals for *unobserved*. In part 1 subjects play six repetitions of the 2S3Q game where *unobserved* receive private signals of quality 14/21. Concretely, all 15 subjects submit a guess in the first decision period. The guess of one *observed* is randomly selected to be made public at the beginning of the next period and this subject stops from submitting guesses. In the second decision period, all remaining 14 subjects submit a guess. The guess of one *observed* is randomly selected to be made public at the beginning of the next period and this subject stops from submitting guesses. And so on, until the last decision period where all *unobserved* submit a guess. For each subject, only one randomly selected guess is paid in each repetition. If the guess is correct the subject receives 1 Euro, otherwise she receives nothing. Feedback screens are identical to those in the practice rounds except that each subject is only reminded of her payoff-relevant guess and of the *observed* guesses which were made public.

Part 2: High quality signals for *unobserved*. The second part is identical to the first one except that *unobserved* receive private signals of quality 18/21.

Part 3: Low quality signals for *unobserved*. The third part is identical to the first one except that *unobserved* receive private signals of quality 12/21.

We designed Experiment 1 to measure the observational learning success in rich-information scenarios with a particular interest for the scenario where *unobserved* receive high quality signals. Experimental sessions therefore start with the simple scenario where *observed* and *unobserved* receive medium quality signals to promote successful observational learning. And they end with *unobserved* receiving low quality signals to avoid biasing their behavior towards excessive herding in earlier parts.

2.1.3 Experimental Procedures

The experimental sessions took place at the laboratory for experimental economics of the Technische Universität München (experimenTUM) in April 2014, February and April 2015. Students from the Technische Universität München and the Ludwig-Maximilians-Universität München were invited using the ORSEE recruitment system (Greiner, 2015). We conducted nine sessions with 16 subjects in each session. One subject was randomly selected to serve as the laboratory assistant and the other fifteen were randomly assigned to computer terminals in isolated booths. Experiment 1 was programmed in zTree (Fischbacher, 2007).

Each session started with short demonstrations of the state-selection procedure. An experimenter shuffled a deck of 20 cards—11 cards with a blue front and 9 cards with an orange front—and laid the

cards face down on a table. The assistant then picked 1 card out of the 20 cards, and the front color of the picked card determined the state.⁶ After the demonstrations, paper instructions for the practice rounds were distributed and subjects were given time to read them at their own pace. Instructions were then read aloud and finally subjects learned about their role, *observed* or *unobserved*, which they kept during the entire experimental session.

The procedures of the three practice rounds closely follow those used by Anderson and Holt (1997) in their baseline experiment except for the two parallel sequences of subjects and the fact that guesses were collected and transmitted through computer terminals. Concretely, after the assistant randomly selected a state, experimenters went to each *unobserved* with the physical “UNOBSERVED” urn containing 14 correct balls and 7 incorrect balls. The participant was asked to draw one ball, return it to the urn and confirm its color in an input box on her computer screen. At the same time each *observed* was approached by another experimenter with the physical “OBSERVED” urn, containing 14 correct balls and 7 incorrect balls, to learn the color of one ball. Guesses were then made at the computer.

Once the practice rounds were over, paper instructions for part 1 were distributed and subjects were given time to read them at their own pace. A summary of the instructions was then read aloud. A short on-screen-demonstration of the draws from the virtual urns followed (to save time, subjects drew private signals from virtual urns displayed on their computer screens during non-practice parts). Again, one of the experimenters summarized aloud the main points of the demonstration. After that, the six repetitions of part 1 were run. The second part of the experiment was conducted in a similar way as the first one except that only short paper instructions were distributed. Part 2 was followed by a short break. Subjects were offered soft drinks and water, and a paper questionnaire was distributed asking for gender, month and year of birth, academic major, mother tongue, and citizenship. Short paper instructions for part 3 were then distributed and the six repetitions were conducted. Finally, subjects privately retrieved their earnings.

In each session we collected 45 guesses from the three practice rounds and 1,656 guesses from the 18 repetitions of the *2S3Q* game for a total of 4,725 *observed* and 10,584 *unobserved* guesses. On average, *observed* and *unobserved* earned 17.33 Euro and 17.96 Euro respectively, including a show-up fee of 3 Euro. A session lasted for about 105 minutes. During the entire session, subjects interacted only through the computers and no other communication was permitted. Appendix A contains a translated version of the instructions.

2.2 Results

The data analysis reported in the main text excludes the few guesses submitted during the practice rounds and groups together the *observed* guesses over the three non-practice parts. Our analysis therefore abstracts from the potential changes in the observational learning behavior of *observed* over the course of the experiment and it compares the behavior of *unobserved* with medium quality signals to the *observed* behavior averaged across the three parts. March and Ziegelmeyer (2016) examine thoroughly the dynamics of the *observed* and *unobserved* behavior in a setting similar to Experiment 1 except that the two roles always receive medium quality signals and interactions take place over two non-practice parts. They show that *observed* contradict their private information significantly less often than *unobserved* in situations where the monetary incentives to follow others are moderately weak and that once the incentives to follow others are strong enough both roles contradict their private information to the same extent. Moreover, behavioral differences between *observed* and *unobserved*

⁶The laboratory assistant randomly selected the state in each practice and non-practice round. The assistant also helped with the drawing of signals from the physical urns during the practice rounds and she monitored the progress of the session on her own computer terminal.

significantly increase as the session progresses. They conclude that the exaggerate response to private information mainly originates from subjects recognizing the future benefits of informative guesses and behaving altruistically. Complementary data analyses are provided in Appendix B and they confirm that *observed* guesses become more informative over time.

Note that the strategy method-like procedure implies that subjects face random payments with known probabilities in each repetition of the 2S3Q game. We assume that subjects evaluate objective lotteries according to expected utility, a sufficient condition to ensure the incentive compatibility of our payment protocol.

2.2.1 Descriptive Analysis

First, we examine the nature of the histories of public guesses in the different decision periods, i.e., the *observed* guesses that have been publicly revealed up to the (beginning of the) relevant period. Second, we assess the influence of public guesses on subjects' propensity to contradict their private information and on their ability to make (ex post) correct guesses.

Table 1 shows the distributions of histories of public guesses in each period derived from the 162 repetitions of the 2S3Q game. For the sake of space, we shorten the notation of histories—for example histories *BBBB* and *OOOB* are shortened to *4B* and *3OB*—and from period 5 on we only report histories which occur at least 3 times.

Period 2	<i>B</i> 53%				<i>O</i> 47%					
Period 3	<i>2B</i> 32%		<i>BO</i> 21%			<i>OB</i> 24%		<i>2O</i> 23%		
Period 4	<i>3B</i> 29%	<i>2BO</i> 03%	<i>BOB</i> 12%	<i>B2O</i> 09%	<i>O2B</i> 09%	<i>OBO</i> 15%	<i>2OB</i> 02%	<i>3O</i> 21%		
Period 5	<i>4B</i> 28%	<i>BO2B</i> 05%		<i>BOBO</i> 07%	<i>B3O</i> 08%	<i>O3B</i> 08%	<i>OBOB</i> 04%	<i>OB2O</i> 11%	<i>3OB</i> 03%	<i>4O</i> 18%
Period 6	<i>5B</i> 25%	<i>4BO</i> 03%	<i>BO3B</i> 04%	<i>BOBOB</i> 02%	<i>BOB2O</i> 04%	<i>B4O</i> 07%	<i>O4B</i> 07%	<i>OB3O</i> 11%	<i>3OBO</i> 02%	<i>5O</i> 17%
Period 7	<i>6B</i> 23%	<i>4BOB</i> 02%	<i>BO4B</i> 04%	<i>BOBO2B</i> 02%	<i>BOB3O</i> 03%	<i>B5O</i> 07%	<i>O5B</i> 07%	<i>OB4O</i> 11%	<i>6O</i> 17%	
Period 8	<i>7B</i> 22%	<i>4BO2B</i> 02%	<i>BO5B</i> 04%	<i>BOBO3B</i> 02%	<i>BOB4O</i> 02%	<i>B6O</i> 07%	<i>O6B</i> 06%	<i>OB5O</i> 10%	<i>7O</i> 15%	

Table 1: Distributions of Public Histories

We observe that in many empirical histories guess *O* follows guess *B* (a fifth of the histories that occur in period 3 or later) and that in some empirical histories either guess *O* follows guesses *BB* or guess *B* follows guesses *OO* or even guesses *OOO* (together, almost one tenth of the histories that occur in period 4 or later). These non-Bayes rational guesses imply that empirical histories are more diverse than predicted. In particular, Bayesian rationality predicts that 75% of the final histories are full cascades, i.e. *7B* or *7O*, and that about 20% of final histories are either *O6B* or *OB5O*. The predicted distribution of final histories differs significantly from the empirical one (Chi-square test; p -value < 0.01) where only 37% of the final histories are full laboratory cascades. Our results are perfectly in line with those of past cascade experiments and they indicate that *observed* guesses rely too much on private information, as confirmed by the next analysis.

In a given period, the information set of a subject corresponds to the couple (private signal, history of public guesses). As a convention we denote the size of the majority of public guesses by $\Delta = \#blue - \#orange$ where $\#blue$ and $\#orange$ is the number of blue and orange guesses in the

public history with $\Delta \in \{-(t-1), \dots, t-1\}$ at the beginning of period $t \in \{1, \dots, 8\}$. Moreover, we refer to the majority of public guesses as a contrary majority (resp. favoring majority) in cases where the subject's private signal and the majority of public guesses are conflicting (resp. concordant) pieces of information. Thus, the subject faces a contrary (resp. favoring) majority either when she is endowed with a blue signal and $\Delta < 0$ (resp. $\Delta > 0$) or when she is endowed with an orange signal and $\Delta > 0$ (resp. $\Delta < 0$). If $\Delta = 0$, there is no majority in the history of public guesses. Table 2 reports the percentage of guesses that contradict private information by the signal of each role and for the different majorities of public guesses. Note that *observed* (resp. *unobserved*) face majorities of size at most 6 (resp. 7). We don't differentiate between favoring majorities since the percentages of guesses contradicting private information hardly change with the size of the favoring majority or between large contrary majorities as fewer data are available for contrary majorities of size 5 or more. We say that a subject herds if she contradicts her private signal when she faces a contrary majority of public guesses, thereby excluding imitative guesses which accord with private information.

History of public guesses		<i>Observed</i>		<i>Unobserved</i>					
		Medium quality		Low quality		Medium quality		High quality	
		<i>b</i>	<i>o</i>	<i>b</i>	<i>o</i>	<i>b</i>	<i>o</i>	<i>b</i>	<i>o</i>
Favoring majority		02% (889)	02% (745)	02% (769)	01% (741)	03% (825)	03% (751)	01% (968)	02% (791)
No majority		03% (796)	06% (781)	06% (372)	10% (340)	05% (350)	11% (330)	01% (335)	03% (393)
Contrary majority of size	1	14% (292)	20% (296)	53% (204)	53% (163)	27% (155)	33% (144)	02% (123)	06% (184)
	2	63% (136)	63% (141)	79% (130)	85% (098)	77% (094)	67% (104)	26% (072)	12% (083)
	3	81% (106)	84% (107)	86% (111)	91% (098)	83% (103)	82% (106)	47% (072)	25% (075)
	4	89% (062)	87% (076)	88% (072)	94% (082)	91% (080)	88% (097)	52% (064)	47% (062)
	≥ 5	91% (043)	92% (066)	93% (134)	94% (142)	94% (137)	88% (180)	57% (102)	50% (132)

Note: In each cell, the first row reports the percentage of guesses that contradict private information and the second row reports the number of guesses.

Table 2: Percentages of Guesses that Contradict Private Information

Several observations can be made from Table 2. First, whenever the private signal and the history of public guesses do not constitute conflicting pieces of information, i.e., at favoring and no majorities, almost all guesses follow private information. Second, subjects' propensity to contradict their private information systematically increases with the size of the contrary majority. This observation entails that once the public evidence is conclusive enough herd behavior is qualitatively consistent with Bayes-rational herding for low and medium quality signals and that *unobserved* with high quality signals herd excessively. Indeed, the majority of *unobserved* guesses contradict high quality private information at large contrary majorities. Third, at small and medium contrary majorities, the higher their signal quality the more *unobserved* follow their private information. Reassuringly, *unobserved* guesses take into account the quality of the signal. Fourth, subjects contradict their private information more frequently with orange than with blue signals at contrary majorities of size 1 but the difference vanishes at larger contrary majorities. This suggests that guesses account to some extent for the asymmetric prior. Finally, in line with the main finding of March and Ziegelmeyer (2016), *unobserved*

with medium quality signals contradict their private information more often than *observed* at short contrary majorities but the difference vanishes at larger contrary majorities.

In Appendix B, we report the fraction of (ex-post) correct guesses for *observed* and for each signal quality in the case of *unobserved* (fractions are averaged across signals). We find that, whatever the size of the contrary majority, the more profitable guess consists in following private information for *unobserved* with high quality signals. Also, it is more profitable for *observed* (resp. *unobserved* with low or medium quality signals) to contradict private information when the contrary majority is of size ≥ 2 (resp. ≥ 1).

2.2.2 Measuring the Success of Observational Learning by Controlling for Incentives

As rightly argued by Weizsäcker (2010), existing theoretical models of observational learning have a rather limited descriptive power which implies that in many cases they falsely predict the incentives of laboratory subjects. It is therefore problematic to measure the success of observational learning by assuming that subjects learn from players who obey a descriptively inaccurate model solution. The alternative approach we follow here first estimates subjects' incentives to contradict their private information and then measures the success of observational learning by controlling for these incentives.

In each of the three non-practice parts, we estimate separately for *observed* and *unobserved* the (expected monetary) value of a guess that contradicts private information across all observations with the same history and private signal. This empirical value of contradicting private information approaches the true value of contradicting private information as the number of occurrences of the *guessing situation* increases in the dataset. Formally, a guessing situation is characterized by part $p \in \{1, 2, 3\}$, role $r \in \{\text{observed}, \text{unobserved}\}$, signal $s \in \{b, o\}$ and history of public guesses $h_t = (g_1, \dots, g_{t-1}) \in \{B, O\}^{t-1}$, and the empirical value of contradicting private information at guessing situation (p, r, s, h_t) is defined by

$$\text{value_contra_PI}(p, r, s, h_t) = \begin{cases} \left[1 + \frac{11q}{9(1-q)} \prod_{\tau < t} \frac{2\hat{P}r(g_\tau | b, h_\tau, \mathcal{B}) + \hat{P}r(g_\tau | o, h_\tau, \mathcal{B})}{\hat{P}r(g_\tau | b, h_\tau, \mathcal{O}) + 2\hat{P}r(g_\tau | o, h_\tau, \mathcal{O})} \right]^{-1} & \text{if } s = b \\ \left[1 + \frac{9q}{11(1-q)} \prod_{\tau < t} \frac{\hat{P}r(g_\tau | b, h_\tau, \mathcal{O}) + 2\hat{P}r(g_\tau | o, h_\tau, \mathcal{O})}{2\hat{P}r(g_\tau | b, h_\tau, \mathcal{B}) + \hat{P}r(g_\tau | o, h_\tau, \mathcal{B})} \right]^{-1} & \text{if } s = o, \end{cases}$$

where $q \in \{12/21, 14/21, 18/21\}$ is the quality of signal s , and $\hat{P}r(g_\tau | s', h_\tau, \omega)$ is the fraction of g_τ guesses, $g_\tau \in \{B, O\}$, among all *observed* guesses with signal $s' \in \{b, o\}$ of quality 14/21, at history $h_\tau \subset h_t$ and state $\omega \in \{\mathcal{B}, \mathcal{O}\}$ (products over $\tau < t$ are assumed equal to one in the first period).

Note that the computation of $\text{value_contra_PI}(p, r, s, h_t)$ requires that the fraction of g_τ guesses exists for each couple (s', ω) at each sub-history $h_\tau \subset h_t$, a more stringent requirement in later periods. We are unable to compute value_contra_PI at 133 of the 327 guessing situations encountered by *observed* and at 303 of the 506 guessing situations encountered by *unobserved*. Still, the guessing situations for which value_contra_PI can be computed cover about 94% and 74% of the *observed* and *unobserved* guesses respectively. The omitted guessing situations occur rather infrequently as they are mostly encountered in the last period of each sequence.⁷ More importantly, $\text{value_contra_PI}(p, r, s, h_t)$ is an imperfect measure of the true underlying monetary incentives whose precision relates to the number of occurrences of the guessing situation in the dataset. We denote by $\text{sitcount}(p, r, s, h_t)$ the number of occurrences of guessing situation (p, r, s, h_t) in the dataset. The

⁷For example, half of the left out *unobserved* guesses are made in period 8 where the empirical value of contradicting private information cannot be computed for (almost) any guessing situation.

data analysis reported in the main text has been performed on the sample of guessing situations with $sitcount \geq 10$ as $value_contra_PI$ is likely to be far away from the true expected value of contradicting private information for rarely occurring situations. Robustness checks based on different samples of guessing situations deliver the same qualitative results (see Appendix B).

Note also that subjects cannot derive the empirical value of contradicting private information after the first period since $value_contra_PI(p, r, s, h_t)$ is computed using the occurrences of the guessing situation in all sessions and using all *observed* guesses, publicly revealed or not. However, our aim is not to examine how well subjects uncover the true behavior of others as their experiences accumulate, but rather to compare their success in learning from others with the success of players who, on average, correctly react to the true value of contradicting information. Concretely, our theoretical benchmark assumes that players know the informational value of public guesses, whether these guesses are compatible with Bayesian rationality or not, that they update their beliefs using Bayes' rule, and that they make probabilistic money-maximizing guesses conditional on their beliefs (where the choice probability of a guess increases with its expected monetary payoff). If the empirical value estimates perfectly the true value of contradicting private information, the theoretical correspondence between $value_contra_PI$ and the probability of contradicting private information is an *S*-shaped curve through (0.5, 0.5). By controlling for the underlying incentives, our analysis measures how successful subjects are in learning from others compared to players in the theoretical benchmark.

Empirical Values of Contradicting Private Information

Before analyzing the subjects' responses to the empirical value of contradicting private information, we examine the level of the underlying incentives in various guessing situations. Table 3 reports the empirical values of contradicting private information by majorities of public guesses and by role, distinguishing between signal qualities for *unobserved*. In each cell, we display the mean of $value_contra_PI$, the first and ninth deciles of $value_contra_PI$, and the total number of individual observations for all guessing situations in the first, second, and third row respectively.

Incentive levels in the top row of Table 3 clearly imply that, for each role and each signal quality, the empirically optimal guess at favoring majorities consists in following private information. The average incentives to act in accordance with private information are at least four times stronger than the average incentives to contradict private information. Further down the table incentives to contradict private information firmly increase till the contrary majority reaches size 2 but then incentive levels hardly vary with additional contrary guesses. According to the estimated values of contradicting private information, contrary majorities of size ≥ 2 aggregate about two private signals.⁸ In the last two columns of Table 3, the incentives to follow private information are at least 1.5 times stronger than the incentives to contradict private information. *Unobserved* should therefore always follow their private information when endowed with high quality signals. On the other hand, *unobserved* with low quality signals should herd at contrary majorities of any size and they should follow their private information otherwise (columns 4-5). Still, incentives to make the empirically optimal guess are weak in the case of no majority and an orange signal. Given the average incentive levels in columns 2-3 and 6-7 of Table 3, subjects with medium quality signals should follow (resp. contradict) their private information at favoring and no majorities (resp. at contrary majorities of size ≥ 2). At contrary majorities of size 1 they should follow (resp. contradict) their private information with a blue signal (resp. an orange signal) though, as expected, incentives to make the empirically

⁸The value of contradicting private information for an individual with a blue signal of low (resp. medium and high) quality who infers (a net of) two orange signals of medium quality from the public guesses equals 0.71 (resp. 0.62 and 0.35). The value of contradicting private information for an individual with an orange signal of low (resp. medium and high) quality who infers two blue signals of medium quality from the public guesses equals 0.79 (resp. 0.71 and 0.45).

History of public guesses	<i>Observed</i>		<i>Unobserved</i>					
	Medium quality		Low quality		Medium quality		High quality	
	<i>b</i>	<i>o</i>	<i>b</i>	<i>o</i>	<i>b</i>	<i>o</i>	<i>b</i>	<i>o</i>
Favoring majority	.16	.20	.18	.22	.17	.18	.05	.09
	.11 – .19 (704)	.13 – .25 (596)	.14 – .25 (517)	.15 – .32 (538)	.16 – .19 (524)	.12 – .25 (540)	.04 – .07 (655)	.05 – .11 (464)
No majority	.29	.38	.38	.48	.28	.39	.12	.17
	.28 – .29 (751)	.38 – .39 (748)	.37 – .40 (331)	.46 – .49 (301)	.27 – .29 (321)	.38 – .40 (311)	.12 – .12 (319)	.17 – .17 (359)
1	.45	.53	.57	.64	.43	.52	.20	.27
Contrary	.43 – .48 (250)	.51 – .53 (266)	.55 – .64 (150)	.60 – .74 (147)	.37 – .46 (135)	.50 – .61 (120)	.20 – .23 (115)	.26 – .28 (156)
2	.59	.63	.71	.77	.58	.58	.29	.38
majority	.56 – .63 (076)	.58 – .67 (110)	.60 – .77 (084)	.75 – .81 (069)	.49 – .63 (080)	.58 – .58 (064)	.19 – .38 (064)	.38 – .39 (060)
3	.63	.64	.73	.76	.59	.56	.27	.41
of	.59 – .69 (064)	.56 – .68 (088)	.56 – .80 (076)	.76 – .76 (045)	.49 – .64 (078)	.56 – .56 (061)	.18 – .33 (046)	.41 – .41 (039)
4	.61	.64	.75	.78	.56	.57	.26	.39
size	.59 – .66 (042)	.57 – .70 (068)	.75 – .75 (032)	.78 – .78 (042)	.49 – .59 (049)	.57 – .57 (060)	.18 – .33 (039)	.39 – .39 (039)
≥ 5	.61	.63	.73	.77	.63	.55	—	.38
	.61 – .61 (012)	.56 – .70 (044)	.73 – .73 (096)	.77 – .78 (068)	.61 – .65 (068)	.54 – .56 (099)	— — — (000)	.38 – .38 (037)

Note: In each cell, from top to bottom: mean of *value_contra_PI*, 1st – 9th deciles of *value_contra_PI*, and number of individual observations.


Table 3: Empirical Values of Contradicting Private Information

optimal guess are rather weak.

Responses to the Empirical Value of Contradicting Private Information

Figure 1 plots the empirical value of contradicting private information against the proportion of contradictions collected in identical guessing situations. The abscissae and sizes of bubbles are given by levels of *value_contra_PI* and *sitcount*. The ordinates of red, black and green bubbles are given by the proportions of contradictions for *unobserved* with high, medium and low quality signals respectively (for the sake of readability, *observed* bubbles are omitted). Figure 1 also superimposes fitted curves from a weighted linear regression that includes a cubic polynomial in *value_contra_PI* fully interacted with indicator variables for *unobserved* in each part. To correct for the fact that *value_contra_PI* imperfectly measures the true expected value of contradicting private information, we follow the split-sample instrumental variable (IV) method described in Weizsäcker (2010) which obtains an instrument by partitioning the dataset in two subsamples. The grey, red, black, and green curve is the fitted curve for *observed*, *unobserved* with high, medium and low quality signals respectively (the *observed* response is averaged over the three parts). The sample of guessing situations with *sitcount* ≥ 10 consists of 239 distinct guessing situations for a total of 10,039 individual observations.

First, we compare the responses of *unobserved* with medium quality signals to those of *observed*. In situations where their private information happens to support the empirically optimal guess both roles largely follow their signal though *unobserved* do so less often than *observed*. Averaging across observations where *value_contra_PI* ≤ 0.5, the proportion of guesses that are optimal is 0.92 and 0.96



- Notes: i) ●, ●, ●: *Unobserved* guesses made with high, medium and low quality signals respectively;
 ●: *Observed* guesses.
- ii) Empirical values of contradicting private information fall into the range $[0.10, 0.70]$, $[0.11, 0.81]$, $[0.12, 0.65]$, and $[0.04, 0.41]$ for *observed* across all three parts, *unobserved* with low, medium and high quality signals respectively.

Figure 1: Responses to the Empirical Value of Contradicting Private Information

for *unobserved* and *observed*, respectively. By contrast, in situations where they should contradict their private information *unobserved* guess optimally and follow others far more often than *observed*. Averaging across observations where $value_contra_PI > 0.5$, the proportion of guesses that are optimal is 0.76 and 0.55 for *unobserved* and *observed*, respectively. Still, the black curve is not an *S*-shaped curve through $(0.5, 0.5)$, and we reject the hypothesis that *unobserved* with medium quality signals probabilistically best respond to the value of their available information as the vertical distance between the black curve and $(0.5, 0.5)$ is strongly significant (two-tailed p -value < 0.01). Back to comparing the observational learning behavior of the two roles, we note that the reluctance of *observed* to contradict their private information is especially pronounced in situations where the expected monetary costs of making an informative guess are moderate at most. Averaging across observations where $0.5 < value_contra_PI \leq 0.6$, the proportion of *observed* contradictions is only 0.45 whereas the proportion of *unobserved* contradictions is 0.73. However, once the incentives to follow others are strong enough both roles largely contradict their private information: Averaging across observations where $value_contra_PI > 0.6$, the proportion of contradictions is 0.87 and 0.73 for *unobserved* and *observed*, respectively. We test whether *observed* make more informative guesses than *unobserved* by comparing the predicted frequencies to contradict private information at $value_contra_PI = 0.55$ which corresponds to the level of monetary incentives necessary for *unobserved* to follow others with more than probability one-half. We find that *observed* act more informatively than *unobserved* since the vertical distance between the grey curve and $(0.55, 0.5)$ is strongly significant (two-tailed p -value < 0.01). Importantly enough, in 431 out of the 503 observations where $0.5 < value_contra_PI \leq 0.6$

observed face a contrary majority of size ≤ 3 and they act in period 4 or earlier. Our findings therefore support the conclusion of March and Ziegelmeyer (2016) that subjects overweight their private information largely because they recognize the future benefits of informative guesses and behave altruistically.

We now examine the responses of *unobserved* with low quality signals. In the left half of Figure 1, *unobserved* almost always make the empirically optimal guess: Averaging across observations where $value_contra_PI \leq 0.5$, the proportion of guesses that are optimal is 0.97. Thus, in situations where their private information supports the empirically optimal guess *unobserved* respond more strongly to the underlying incentives with low than with medium quality signals. On the contrary, in situations where they should contradict their private information *unobserved* guess less optimally with low than with medium quality signals. Averaging across observations where $0.5 < value_contra_PI \leq 0.65$, the proportion of *unobserved* contradictions with low quality signals equals 0.56 which is only three-quarter of the proportion of *unobserved* contradictions with medium quality signals. Most notably, the proportion of contradictions equals 0.55 for the 150 observations with $0.6 < value_contra_PI \leq 0.65$ which correspond to guessing situations where the size of the contrary majority equals 1 (in 123 out of the 150 cases the signal is orange). We also find that the vertical distance between the green curve and (0.5, 0.5) is strongly significant (two-tailed p-value < 0.01) leading to the rejection of the hypothesis that *unobserved* with low quality signals probabilistically best respond to the value of their available information. Still, once the incentives to follow others are strong enough, *unobserved* with low quality signals largely contradict their private information as the average proportion of contradictions is 0.88 across observations where $value_contra_PI > 0.65$.

Finally, we examine the responses of *unobserved* with high quality signals. The red curve firmly increases with the empirical value of contradicting private information and it reaches the level of 0.5 at $value_contra_PI = 0.39$. Thus, in guessing situations where the incentives to follow private information are 1.5 times larger than the incentives to contradict private information, the average *unobserved* acts against her signal. This result occurs because *unobserved* with high quality signals fail to respond correctly to the value of their available information when they face large contrary majorities. Indeed, at contrary majorities of size 5, $value_contra_PI$ equals 0.38 and the proportion of contradictions reaches 0.54 meaning that most guesses are incorrect responses to the underlying incentives. Even at contrary majorities of size 4 the proportion of contradictions is larger than one-half (0.51) though the mean of $value_contra_PI$ only equals 0.32 when averaged across signals. On the other hand, at contrary majorities of size 2, the mean of $value_contra_PI$ equals 0.33 and *unobserved* often make the empirically optimal guess as the proportion of contradictions is 0.18. In our descriptive analysis we already highlighted that the herding of *unobserved* with high quality signals is particularly pronounced when they face large contrary majorities. By controlling for the underlying incentives our analysis shows that *unobserved* with high quality signals wrongly assess the informational value of public guesses in large contrary majorities and that their excessive herding is severely harmful. In fact, *unobserved* with high quality signals would be better off by *not* learning from others as the proportion of optimal guesses in period 1 is greater than the average proportion of optimal guesses across later periods (0.98 versus 0.92), even though *unobserved* mostly face favoring majorities when endowed with high quality signals.

2.2.3 Herd Behavior at the Individual Level

To examine herd behavior at the individual level, we assign *unobserved* into six decision rules based on their profile of guesses across ten groups of guessing situations (see Appendix D for details). In group 1 (resp. 2 and 3), the history of public guesses induces a favoring or no majority and the signals'

quality is low (resp. medium and high). In groups 4 to 10, the history of public guesses always induces a contrary majority. In group 4 (resp. 5), the signals' quality is low and the majority size is less than 2 (resp. more than 3). In both groups the empirically optimal guess is herding and the average incentives to do are weaker in group 4 than in group 5 ($vcPI = 0.65$ vs. $vcPI = 0.75$). In group 6, the signal is blue medium and the majority size is 1 which implies that *unobserved* should follow their signal. In group 7, the empirically optimal guess is herding since either the signal is orange medium and the majority size is 1 or the signals' quality is medium and the majority size is 2. The incentives to guess optimally are rather weak in groups 6-7. In group 8, the signals' quality is medium and the majority size is larger than 3 which implies that *unobserved* should herd ($vcPI = 0.58$). Finally, in group 9 (resp. 10), the signals' quality is high and the majority size is less than 2 (resp. more than 3). In both groups following private information is empirically optimal (on average, $vcPI = 0.30$).

Our classification proceeds in two steps. First, an *unobserved* is classified as noisy if less than half of her guesses are empirically optimal both across groups 1-3 and across groups 4-10. Second, each of the remaining *unobserved* is assigned to one of five non-noisy decision rules based on her profile of guesses across groups 4-10. The first decision rule is successful observational learning (SOL) which guesses optimally in every group of guessing situations. The next two rules herd excessively compared to SOL. The weak conformism rule (WC) guesses like SOL except that it herds in groups 6 and 10, and the strong conformism rule (SC) guesses like WC except that it also herds in group 9. On the other hand, the last two rules follow private information excessively compared to SOL. The weak following-private-information rule (WFPI) guesses like SOL except that it follows private information in groups 4 and 7, and the strong following-private-information rule (SFPI) guesses like WFPI except that it also follows private information in groups 5 and 8. For each *unobserved* we compute 5 scores where each score reflects the adequacy between her guesses and the guesses made by one of the five non-noisy decision rules.⁹ The highest score determines the rule to which the *unobserved* is assigned.

We find that 32% of the *unobserved* are assigned to the SOL rule, 33% (resp. 03%) of them are assigned to the WC (resp. SC) rule, and 25% (resp. 06%) of them are assigned to the WFPI (resp. SFPI) rule. Only one *unobserved* is classified as noisy. Thus, *unobserved* are almost equally divided into observational learners who guess optimally, observational learners who herd too often, and observational learners who respond too strongly to their private information. Unsurprisingly, excessive herding is most pronounced among conformists with an average proportion of contradictions of private information equal to 0.83 in group 10. We note, however, that even SOL and WFPI exhibit a non-negligible tendency to herd with high quality signals at large contrary majorities (the average proportion of contradictions is 0.17 for SOL and 0.43 for WFPI). In a similar vein, though the overemphasis on low or medium quality signals is most pronounced among FPIs, SOL and WC also tend to respond too strongly to private information at small contrary majorities (when averaged across groups 4 and 7, the proportion of guesses that follow private information is 0.20 for SOL and 0.28 for WC). In sum, herd behavior is substantially heterogeneous at the individual level but the tendencies to herd excessively with high quality signals and to respond too strongly to low or medium quality signals are rather widespread among *unobserved*.

2.3 Discussion

There are four main insights that come out of Experiment 1. First and foremost, *unobserved* with high quality signals wrongly respond to the informational value of public guesses in large contrary majorities which leads them to herd excessively. Deviations from Bayes-rational herding result in

⁹Concretely, if in a given situation the *unobserved* guess matches the guess of the decision rule then we add one to the score, otherwise the score remains unchanged.

severe expected losses and *unobserved* with high quality signals would be better off without the chance to learn from others. Second, though they fall short of better responding to the value of their available information, *unobserved* with medium quality signals are quite successful in learning from others. In particular, their overemphasis on private information is less pronounced than for *observed* whose behavior is partly driven by efficiency concerns. Third, *unobserved* are more reluctant to contradict their private information with low than with medium quality signals when facing identical incentives to follow others. Still, the proportion of *unobserved* contradictions with low quality signals is larger than one-half whenever the empirical value of contradicting private information exceeds 0.55. Fourth, there is substantial heterogeneity in the individual herd behavior. *Unobserved* are almost equally divided into observational learners who guess optimally, observational learners who herd too often, and observational learners who respond too strongly to their private information.

As we just emphasized, the observational learning success of *unobserved* is rather modest in the challenging situations where public guesses conflict with their private signals, especially when signals are of high quality. We interpret this modest success as a failure by *unobserved* to properly assess the value of their available information in the face of contradictory public information. Secondary insights from Experiment 1 support our interpretation, rather than the alternative interpretation according to which *unobserved* fail to better respond to properly valued information endowments. Indeed, in many of the situations where *unobserved* learn from contrary majorities their guesses respond to the quality of the private signals or to the asymmetric prior as qualitatively advocated by our theoretical benchmark. We found that i) the higher their signal quality the more often *unobserved* follow their private information at small and medium contrary majorities; and that ii) *unobserved* contradict their private information more frequently with orange than with blue signals at contrary majorities of size 1 though the difference vanishes at larger contrary majorities. We also note that *unobserved* predominantly make the empirically optimal guess when they face favoring or no majorities. Overall, the evidence from Experiment 1 indicates that *unobserved* guesses largely respond to the information structure of the $2S3Q$ game and that systematic deviations from the theoretical benchmark play are restricted to situations where *unobserved* face contrary majorities.

3 Laboratory Evidence on the Driving Forces of Herd Behavior

This section reports on three laboratory experiments designed to identify the elements of human behavior that best account for the descriptive failure of Bayes-rational herding. We focus on the behavior of the *unobserved* and therefore neglect efficiency concerns which, as confirmed by Experiment 1, partly drive the behavior of the *observed*. To guide our identification strategy, we endorse a belief-based view of herd behavior. Concretely, we assume that *unobserved* form probabilistic beliefs about the payoff-relevant states and that they better respond to their beliefs.

Economists have offered three complementary explanations for the descriptive failure of Bayes-rational herding. First, people may wrongly incorporate the signals inferred from public guesses into their beliefs because they suffer from biases in statistical reasoning (Huck and Oechssler, 2000; Goeree, Palfrey, Rogers, and McKelvey, 2007). Second, observational learners may infer signals from public guesses that differ in systematic ways from the signals actually received by others because they rely on “informationally naïve” inference rules (Eyster and Rabin, 2005, 2010). Third, the model of how others make guesses may be incorrect. Wrong expectations about others’ strategy can induce people to either overappreciate or underappreciate the informational value of public guesses (Kübler and Weizsäcker, 2004; Bohren, 2016).

Each of the aforementioned explanations has the potential to account by itself for our two main findings in Experiment 1, namely excessive herding with high quality signals and overemphasis on low

and medium quality signals. To isolate the impact that each deviation from Bayesian rationality has on herd behavior, we conducted three new experiments. The new evidence should ultimately further our understanding of how herd behavior is shaped by the interplay of the three components of belief formation: belief updating, informational inferences, and expectations about others. To achieve this goal, we compare the behavior of *unobserved* in three scenarios which differ in the set of deviations that can impact the observational learning success. In Experiment 2 *unobserved* learn from public signals, implying that the two components informational inferences and expectations about others are turned off. Thus, the observable-signals scenario informs us about the nature of non-Bayesian updating and the extent to which it undermines the observational learning success. In Experiment 3 *unobserved* learn from public guesses known to be made by computer-*observed* that adopt the Bayesian strategy. In the second scenario both informational misinferences and non-Bayesian updating potentially impact herd behavior, but the component expectations about others remains turned off. The comparison of the observational learning success in Experiments 2 and 3 reveals the nature of informational misinferences and the extent to which they cause *unobserved* to improperly assess the value of their available information over and above non-Bayesian updating. Finally, in Experiment 4 *unobserved* learn from public guesses made by subjects in the *observed* sequence meaning that the component expectations about others is turned on. By comparing their observational learning success in Experiments 3 and 4, we shed light on the model that *unobserved* have about how others make guesses and we assess the extent to which uncertainty about others' strategy impacts herd behavior.

3.1 Experimental Designs and Procedures

As in Experiment 1, *unobserved* play repeatedly the 2S3Q game in Experiments 2-4. However, compared to Experiment 1, the three experiments rely more extensively on the strategy method-like procedure to offer ample learning opportunities from contrary majorities. The main differences between the non-practice parts of a laboratory session in the three experiments and in Experiment 1 are as follows. The first non-practice part is similar to the non-practice parts in Experiment 1 except that *unobserved* receive private signals of different qualities in subsequent rounds. In the second non-practice part *unobserved* remain uninformed of their private signal till the end of the round. They must therefore submit two guesses in each decision period, one guess with signal b and one guess with signal o . Thus, in a given decision period each *unobserved* faces either twice a no majority with different signals or once a contrary majority and once a favoring majority of the same size (Cipriani and Guarino, 2009, first considered this extension). Finally, in the third non-practice part *unobserved* remain uninformed of both the signal and its quality till the end of the round. They must therefore submit a total of six guesses in each decision period, which ensures that each *unobserved* regularly faces a contrary majority even when endowed with high quality signals.

Below we first describe each experimental setting, then we outline the progress of a laboratory session, and finally we detail our experimental procedures.

Experiment 2: Learning from signals

In Experiment 2 *unobserved* learn from guesses made by computer-*observed* and it is public knowledge that the latter guess B if and only if their private signal is b . Since public guesses perfectly reveal private signals, *unobserved* are in effect learning from strings of public signals.

A different parametrization of the 2S3Q game is played in Experiment 2 than in Experiment 1 with $n_{Obs} = 12$ and $n_{Unobs} \in \{12, 14, 18\}$. In other words, computer-*observed* receive low quality signals in the second experiment. Indeed, since public information keeps accumulating in an observable-signals scenario, the evidence from strings of medium quality signals would swamp the private information

of *unobserved* rather quickly. By lowering the quality of computer-*observed* signals, we ensure that following private information remains the Bayes-rational strategy for *unobserved* with high quality signals in (almost) all of Experiment 2’s situations. Specifically, Bayesian rationality prescribes to guess in accordance with private information except in the following situations: With high (resp. medium and low) quality signals, guess B if $\Delta \geq 6$ (resp. $\Delta \geq 2$ and $\Delta \geq 1$) and guess O if $\Delta = -7$ (resp. $\Delta \leq -4$ and $\Delta \leq -2$) where Δ denotes the difference between the number of blue and orange public guesses.

The two components informational inferences and expectations about the strategy that generates public guesses are turned off in the observable-signals scenario. Deviations from better responses to the true value of information merely reflect failures to use Bayes’ rule when aggregating multiple signals. For example, if they believe in the “law of small numbers” (Tversky and Kahneman, 1971; Rabin, 2002) then *unobserved* overestimate the informational value of relatively small contrary majorities which leads them to herd excessively when endowed with medium or high quality signals. On the other hand, if they are prone to the conservatism bias (Edwards, 1968) then *unobserved* underestimate the informational value of large contrary majorities and in turn they insufficiently contradict their low or medium private signals. By measuring the observational learning success in Experiment 2 we can identify the errors in statistical reasoning that prevent *unobserved* from forming Bayesian beliefs when they learn from strings of public signals.

Experiment 3: Learning from Bayes-rational guesses

In Experiment 3 *unobserved* learn from guesses made by computer-*observed* and it is public knowledge that the latter adopt the Bayes-rational strategy described in subsection 2.1.1. The parametrization of the 2S3Q game is the same as in Experiment 1 with *unobserved* receiving either low, medium or high quality signals and computer-*observed* receiving medium quality signals. Accordingly, Bayes-rational players endowed with low quality signals always imitate the most recent public guess except in the first period or after history $(O, B, \dots, O, B)$ where they follow private information. When they receive medium quality signals Bayes-rational players act like the computer-*observed*, and they follow their private information at all histories of public guesses when endowed with high quality signals.

Though *unobserved* know the strategy that generates public guesses in both experiments, the task of evaluating their available information is cognitively more demanding in the third than in the second experiment. Indeed, to properly assess their available information in Experiment 3 *unobserved* must not only update their beliefs in a Bayesian manner but they must also correctly infer signals from public guesses. Two opposite departures from Bayesian rationality can cause *unobserved* to infer signals from public guesses that differ from the signals actually received by the computer-*observed*. First, *unobserved* may fail to realize that, except when they are part of an information cascade, Bayes-rational guesses reflect the signals of computer-*observed*, i.e., they are informative. If they underappreciate the connection between the signals of computer-*observed* and their non-cascade guesses, then *unobserved* will be reluctant to contradict their low or medium quality signals. Second, *unobserved* may fail to understand that in an information cascade Bayes-rational guesses become uninformative. *Unobserved* who don’t understand that cascade guesses have no informational value will herd excessively when endowed with high quality signals. By comparing the deviations from better responses to the true value of information in Experiments 2 and 3 we can identify the nature of informational misinferences and their impact on the observational learning success.

There is a final aspect of the experimental setting that we would like to mention since it bears on the theoretical rationalization of informational misinferences in Experiment 3. To infer the correct signal from the most recent public guess, *unobserved* don’t have to reason through how computer-

observed make informational inferences from previous public guesses. Indeed, subjects get a sheet of paper with a table that reports the guess made by the computer depending on its signal and the history of previous public guesses (see our experimental procedures for details). To determine whether the most recent public guess is informative or not, *unobserved* simply have to check whether, given the history of previous public guesses, the computer makes a different guess with each signal or the same guess with both signals. This design feature casts doubt on the appropriateness of assuming that *unobserved* have limited depth of iterated reasoning to rationalize informational misinferences in Experiment 3.

Experiment 4: Learning from human guesses

In Experiment 4 *unobserved* learn from public guesses made by other subjects in the *observed* sequence. Thus, *unobserved* are uninformed of the strategy that generates public guesses. They have to form expectations about how *observed* play in order to seize the informational benefits of public guesses. Note that uncovering the strategy adopted by *observed* is particularly challenging in the 2S3Q game since the signals they receive in a given repetition of the game are never made public even at the end of the repetition.

The parametrization of the 2S3Q game is identical in Experiments 1 and 4 which implies that Bayesian rationality makes the same predictions in both experiments (see subsection 2.1.1). Behavioral differences in Experiments 1 and 4 inform us about the extent to which observational learning is affected by a more extensive use of the strategy method-like procedure.

Even if they subscribe to the Bayes-rational view of herding in settings where they know how public guesses are generated, *unobserved* may fail to better respond to their available information in Experiment 4 because of wrong expectations about others' strategy. Indeed, *unobserved* may be reluctant to contradict their low or medium quality signals if they wrongly believe that *observed* always respond weakly to their private information. Alternatively, *unobserved* may herd excessively with high quality signals if they wrongly believe that *observed* always respond strongly to their private information. By comparing the observational learning success in Experiments 3 and 4, we can isolate the impact that expectations about others have on herd behavior.

3.1.1 The Progress of an Experimental Session

As in Experiment 1, each session starts with three paid practice rounds. After practice, subjects play twelve repetitions of their respective 2S3Q game in each experiment. There are six (resp. three) repetitions of the game in the first non-practice part (resp. second and third non-practice parts) and the number of hypothetical guesses increases from one non-practice part to the next.

Practice rounds: Practice rounds proceed as in Experiment 1 except that the quality of the private signals received by *unobserved* differs in each round. In the first (resp. second and third) practice round *unobserved* receive medium (resp. high and low) quality signals.

Part 1: Strategy method with respect to the decision period. In part 1 the six repetitions of the 2S3Q game proceed as in Experiment 1 except that *unobserved* receive private signals of different qualities in subsequent rounds. *Unobserved* receive low quality signals in rounds 1 and 4, medium quality signals in rounds 2 and 5, and high quality signals in rounds 3 and 6.

Part 2: Strategy method with respect to the decision period and the signal. The second part is identical to the first one except that the 2S3Q game is repeated only three times and that

unobserved make two guesses in each decision period, one guess with signal b and one guess with signal o . *Unobserved* receive high quality (resp. low quality and medium quality) signals in round 1 (resp. round 2 and round 3). In Experiment 4 an *observed* whose guess has not been publicly revealed yet also makes one guess with signal b and one guess with signal o in the current decision period. Subjects are informed of their payoff-relevant signal at the end of the round.

Part 3: Strategy method with respect to the decision period, the signal, and the quality.

The third part is identical to the second one except that *unobserved* make a total of six guesses in each decision period and that subjects receive 3 Euro for a correct guess and 0 Euro otherwise. *Unobserved* are informed of the payoff-relevant quality of their signal at the end of the round.

We implemented the strategy method-like procedure so as to limit its potential impact on behavior. First, the strategy method-like procedure is gradually extended over the course of a session. Subjects start by familiarizing themselves with the observational learning task through playing the 2S3Q game with the direct-response method in the three practice rounds. After having hopefully acquired a good understanding of the scenario, subjects make multiple guesses in each round with the number of hypothetical guesses increasing from part 1 to part 2 and from part 2 to part 3 (except for the *observed* in Experiment 4). Second, we refrain from eliciting complete strategies for the 2S3Q game. In each decision period subjects observe the actual history of public guesses which largely restricts the information sets they have to consider. Third, we tripled the monetary stakes in the third non-practice part where *unobserved* make the highest number of hypothetical guesses.

3.1.2 Experimental Procedures

The experimental sessions took place between November 2013 and September 2016 at the same laboratory as for Experiment 1 (experimenTUM), students from the same recruitment sample as Experiment 1 were invited using ORSEE, and all three experiments were programmed in zTree. We conducted six sessions in each experiment with 9 (resp. 16) subjects per session in Experiments 2 and 3 (resp. 4). One subject was randomly selected to serve as the laboratory assistant and the other subjects were randomly assigned to computer terminals in isolated booths.

Sessions proceeded in a similar way as in Experiment 1. In particular, they started with short demonstrations of the state-selection procedure, physical urns were used in practice rounds and virtual urns were used in non-practice rounds, the second-practice part was followed by a short break during which subjects answered demographic questions, and at the end of each session subjects privately retrieved their earnings. Still, several procedural aspects of Experiment 1 had to be tailored to the specifics of Experiments 2-4. In Experiment 1 *unobserved* were informed about the two possible compositions of the “UNOBSERVED” urn only in the instructions distributed at the beginning of each part since their private signals were of the same quality in every decision period of the same part. In Experiments 2-4, however, the signal quality changes from round to round up to the second non-practice part, and in the third non-practice part *unobserved* make a guess for each possible quality in the same decision period. To minimize confusion and errors, every time they had to make a guess for a given signal quality in Experiments 2-4, *unobserved* were reminded of the two possible urns on their computer screen. Additionally, each subject in Experiments 2-3 received at the beginning of the session a sheet of paper with the decision rule adopted by computer-*observed*. In Experiment 2 we made clear that computer-*observed* who receive a blue signal always guess B whereas those who receive an orange signal always guess O . In Experiment 3, we refrained from describing the Bayes-rational strategy algorithmically to alleviate the need for subjects to engage in higher-order thinking.

Instead, each *unobserved* was given a sheet of paper with a table showing the guess made by the computer-*observed* depending on its signal and the history of previous public guesses. Concretely, the first two columns of the table reported for each period $t \in \{1, \dots, 7\}$ the $2t$ couples (signal, history of public guesses) and the last column reported the guess made by the computer-*observed* in that period. Thus, for each history of previous public guesses, one row of the table reported the guess made by the computer-*observed* when endowed with signal b and another row reported the guess made by the computer-*observed* when endowed with signal o (see Appendix A). Finally, the signals of computer-*observed* were drawn by the assistant in Experiments 2-3. An experimenter visited the assistant with the physical “OBSERVED” urn and asked her to draw one ball 7 times (with replacement) in the practice rounds. The assistant drew from a virtual urn on her computer screen with re-shuffling after each draw in the non-practice rounds.

In Experiment 2 and 3 we collected in each session 24 *unobserved* guesses from the three practice rounds and 1,920 *unobserved* guesses from the 12 repetitions of the 2S3Q game for a total of 11,664 *unobserved* guesses. In Experiment 4 we collected in each session 21 *observed* guesses from the three practice rounds and 504 *observed* guesses from the 12 repetitions of the 2S3Q game for a total of 3,150 *observed* guesses. We only collected a total of 11,472 *unobserved* guesses in Experiment 4 as one subject could not submit all her guesses due to a technical error in the first session. On average, subjects earned 16.79 Euro (resp. 18.19 Euro and 18.35 Euro) in Experiment 2 (resp. 3 and 4), including a show-up fee of 3 Euro, and a session lasted for about 110 minutes. During the entire session, subjects interacted only through the computers and no other communication was permitted.

3.2 Results

We summarize here the results of our data analysis. We start by outlining the nature of public histories, then we report on the aggregate success of observational learning, and lastly we investigate herd behavior at the individual level. Additional figures and tables as well as details of the data analysis are provided in Appendix C.

3.2.1 Histories of Public Guesses

Public guesses in Experiment 2 are the signals received by *observed* in the different periods and they correspond to independent draws from the state-dependent Bernoulli distribution with parameter value $12/21$. As expected, public histories are particularly diverse—we identify 56 different final histories in the 72 repetitions of the 2S3Q game—and *unobserved* usually face short majorities of public guesses with about 85% of the final majorities being of size at most 3.

On the other hand, the 72 repetitions of the 2S3Q game generate only 6 different final histories in Experiment 3, and *unobserved* usually face large majorities of public guesses. From period 3 on, 80% of public histories are strings of identical guesses. This is, of course, due to the fact that public guesses are Bayes-rational in Experiment 3 and, as described in subsection 2.1.1, almost all histories lead to the emergence of information cascades.

We find that histories of public guesses in Experiment 4 are of similar nature as histories of public guesses in Experiment 1. Most notably, they are more diverse than predicted by Bayesian rationality with only 38% of the final histories being full laboratory cascades. This observation suggests that, as in Experiment 1, the guesses made by subjects who acted as *observed* relied too much on private information in Experiment 4.


3.2.2 The Success of Observational Learning

To measure the success of observational learning in Experiments 2-3, we rely on the entire sample of guessing situations since the values of contradicting private information are derived theoretically in these two experiments (for the sake of brevity, we often refer to the true value of contradicting private information as *tvcPI*). On the other hand, the expected monetary value of a guess that contradicts private information is estimated in Experiment 4. This empirical value of contradicting private information, often referred to as *vcPI*, is derived differently in Experiment 4 than in Experiment 1 to accommodate the procedural variations in the two experiments. Indeed, a guessing situation in Experiment 4 is characterized by the quadruple (quality, role, signal, history of public guesses) rather than by the quadruple (part, role, signal, history of public guesses). Each value of contradicting private information is therefore estimated over the three non-practice parts of Experiment 4. We discuss here the analysis of the observational learning success that has been performed on the sample of guessing situations with *sitcount* ≥ 10 . Robustness checks on different samples of guessing situations are provided in Appendix C.

Figure 2 depicts the responses to the value of contradicting private information in the three experiments. Each subfigure plots *tvcPI* or *vcPI* against the proportion of contradictions collected in identical guessing situations, and it superimposes fitted curves from a weighted linear regression that includes a cubic polynomial in the value of contradicting private information fully interacted with indicator variables for the signal quality of *unobserved* (and the role in Experiment 4). There are 658 (resp. 186 and 217) guessing situations in Experiment 2 (3 and 4) for a total of 11,520 (11,520 and 11,650) individual observations.

Before summing up the key insights from Figure 2, we note that several of the regularities found in Experiment 1 are also present in Experiments 2-4. The following observations emerge from examining the percentage of (human) guesses that contradict private information in each experiment by the signal of each role and for different majorities of public guesses (these percentages are reported in a table which, for the sake of space, has been relegated to Appendix C). First, *unobserved*, as well as *observed* in Experiment 4, often guess in accordance with their private information at favoring and no majorities. When averaged across the two signals, the proportion of contradictions ranges between 1%—in Experiment 2 with high quality signals at large favoring majorities—and 16%—in Experiment 3 with low quality signals at no majority. Second, the higher their signal quality the more often *unobserved* follow their private information. For example, 87% (resp. 94% and 97%) of their guesses accord with private information in Experiment 2 when *unobserved* face a no majority with low (resp. medium and high) quality signals (in Experiment 3 the relative frequencies are 84%, 93% and 94% respectively, and in Experiment 4 the relative frequencies are 90%, 94% and 98% respectively). Third, guesses account to some extent for the asymmetric prior since subjects contradict their private information more frequently with orange than with blue signals at contrary majorities of size 1 but the difference vanishes at larger contrary majorities. Fourth, subjects' propensity to contradict private information increases with the size of the contrary majority (except for low quality signals in Experiments 2-3 when moving from size 3 to size 4 or more). This final observation implies that, as in Experiment 1, once the public evidence is conclusive enough herd behavior is qualitatively consistent with Bayes-rational herding for low and medium quality signals whereas *unobserved* with high quality signals herd excessively.

To uncover the mechanisms underlying herd behavior, we now compare how successfully *unobserved* learn from contrary majorities in the three experiments. In Experiment 2, we find that *unobserved* herd excessively with high quality signals. The average *unobserved* acts against her high quality signal in guessing situations where the incentives to follow private information are 1.40 times



Note: ●, ●, ●: *Unobserved* guesses made with high, medium and low quality signals respectively; ●: *Observed* guesses.

Figure 2: Responses to the Value of Contradicting Private Information

stronger than the incentives to contradict private information (in Figure 2 the red curve reaches the level of 0.5 at $tv\text{c}PI = 0.41$). Moreover, the proportion of contradictions grows at an almost constant rate with the number of contrary public signals. When averaged across the two signals, the proportion

of contradictions equals 0.09, 0.17, 0.31, and 0.52 at contrary majorities of size 1, 2, 3 and 4 or more, respectively. Together with the fact that *tvcPI* rises slowly with the size of contrary majorities, this regularity implies that *unobserved* with high quality signals are quite unsuccessful even when the incentives to contradict private information are relatively small (e.g., the average *tvcPI* is only 0.30 at contrary majorities of size 3). By contrast, the majority of guesses contradict (resp. accord with) the low or medium quality signal when *tvcPI* > 0.5 (resp. *tvcPI* < 0.5) at every size of contrary majorities (with the minor exception of the blue medium signal at contrary majorities of size 3 where 54% of the guesses are contradictions and *tvcPI* = 0.49). Though our regression results indicate that the vertical distance between the black or green curve and (0.5, 0.5) is statistically significant (two-tailed p-values < 0.05), the average *unobserved* learns well from contrary majorities when endowed with a low or medium quality signal (in Figure 2 the black and green curve reaches the level of 0.5 at *tvcPI* equal to 0.53 and 0.55 respectively).¹⁰ In sum, our findings in Experiment 2 reveal that non-Bayesian updating induces the average *unobserved* to overestimate the informational value of medium-sized and large contrary majorities.

In Experiment 3, the average *unobserved* acts against her high quality signal in guessing situations where the incentives to follow private information are 2.25 times stronger than the incentives to contradict private information (in Figure 2 the red curve reaches the level of 0.5 at *tvcPI* = 0.31). Thus, *unobserved* herd excessively with high quality signals in Experiment 3 even when their relative incentives to contradict private information are 60% weaker than in Experiment 2. *Unobserved* are markedly unsuccessful when learning from contrary majorities in Experiment 3 because as contrary majorities grow larger their propensity to contradict private information increases continuously while *tvcPI* barely rises. When averaging across the two signals, the proportion of contradictions increases from 0.23 at contrary majorities of size 1 to 0.53 at contrary majorities of size 4 whereas *tvcPI* only rises from 0.27 at contrary majorities of size 1 to 0.30 at contrary majorities of size 2 or more. On the other hand, the average *unobserved* learns well from contrary majorities when endowed with a low or medium quality signal in Experiment 3. Indeed, the majority of guesses contradict (resp. accord with) the low or medium quality signal when *tvcPI* > 0.5 (resp. *tvcPI* < 0.5) at every size of contrary majorities, and in Figure 2 the black (resp. green) curve reaches the level of 0.5 at *tvcPI* = 0.48 (resp. *tvcPI* = 0.55). Still, we note a clear increase in the proportion of contradictions with medium quality signals when the size of contrary majorities grows from 2 to 4 even though the incentives to contradict private information remain constant (when averaging across the two signals, the proportion of contradictions increases from 0.75 to 0.83 at *tvcPI* = 0.56). We conclude that *unobserved* tend to overinfer from a handful of contrary Bayes-rational guesses and that these informational misinferences severely undermine their observational learning success when they are endowed with high quality signals.

In Experiment 4, we find that, at a given size of the contrary majority, *unobserved* contradict their high quality signals to a similar extent as in Experiment 2. At a contrary majority of size 1, 2, 3 and 4, the proportion of contradictions equals 0.07, 0.20, 0.34, and 0.45 when averaged across the two signals. Also in line with our findings in Experiment 2, the average *unobserved* acts against her high quality signal in guessing situations where the incentives to follow private information are 1.50 times larger than the incentives to contradict private information (in Figure 2 the red curve reaches the level of 0.5 at *vcPI* = 0.40).¹¹ Hence, compared to Experiment 3, the relative incentives to

¹⁰Part of the reason why the green curve does not go through (0.5, 0.5) is that *unobserved* with an orange signal of low quality make the optimal guess *too often* when facing no majority. Indeed, in these guessing situations we find that the proportion of contradictions is only 0.19 though *tvcPI* = 0.48.

¹¹Still, *unobserved* learn more successfully from public guesses in Experiment 4 than in Experiment 2 when incentives to follow high quality signals are strong, but the reverse is true when incentives are weaker (in Figure 2 the red curve in Experiment 2 and 4 has a linear and convex shape, respectively). The reason is that, compared to those in Experiment

contradict private information have to be increased by a third for *unobserved* to herd excessively with high quality signals in Experiment 4. Moreover, *unobserved* in Experiment 4 fail to better respond to the value of their available information in situations where they should contradict their low or medium quality signals and the incentives to do so are weak. For example, at contrary majorities of size 1, only 43% of *unobserved* guesses contradict the blue signal of low quality ($vcPI = 0.54$) and only 29% of *unobserved* guesses contradict the orange signal of medium quality ($vcPI = 0.52$). Averaging across observations where $0.5 < vcPI \leq 0.55$, the proportion of contradictions equals 0.27 and 0.30 with low and medium quality signals respectively. However, once the incentives to follow others become stronger the majority of guesses contradict private information though *unobserved* remain less successful in learning from others with low than with medium quality signals (averaging across observations where $0.55 < vcPI \leq 0.65$, the proportion of contradictions equals 0.53 with low quality signals and 0.66 with medium quality signals). In short, Experiment 4’s results indicate that the average *unobserved* wrongly believes that *observed* respond weakly to their private information. Compared to a scenario where they learn from Bayes-rational guesses, these wrong expectations about others enhance the observational learning success of *unobserved* with high quality signals, but they reduce their success in learning from short contrary majorities with low or medium quality signals.

Finally, we outline the main differences and similarities between Experiments 1 and 4. To test for the statistical significance of differences in the proportion of contradictions between the two experiments, we report the results of two-sided permutation tests that use session averages as the unit of observation. First, *observed* herd slightly more in Experiment 4 than in Experiment 1 when they should, both at small and large contrary majorities, but none of the differences is statistically significant (p-values > 0.10). As a consequence, for each quality of their private signals, *unobserved* face almost identical incentives to contradict private information in the two experiments. Second, the observational learning success of *unobserved* with high quality signals is almost undistinguishable in the two experiments (see how similar the red curve in Figure 1 and the red curve of Experiment 4 in Figure 2 are). Third, in guessing situations where the contrary majority is small and $vcPI > 0.5$, *unobserved* contradict their low quality signals slightly less often in Experiment 4 than in Experiment 1 but the difference is not statistically significant (p-value > 0.10), and they contradict their medium quality signals significantly less often in Experiment 4 than in Experiment 1 (p-value = 0.08). Fourth, in guessing situations where the contrary majority is large and $vcPI > 0.5$, *unobserved* contradict their low and medium quality signals significantly less often in Experiment 4 than in Experiment 1 (p-values ≤ 0.01). In line with these observations, we find that, contrary to Experiment 1, *unobserved* with medium quality signals act as informatively as *observed* in Experiment 4. To summarize, a more extensive use of the strategy method-like procedure mainly affects the herd behavior of *unobserved* with medium quality signals who become less successful in learning from contrary majorities.

3.2.3 Herd Behavior at the Individual Level

To examine herd behavior at the individual level in Experiments 2-4, we follow the same rule classification as in Experiment 1 (groups of guessing situations in Experiment 2 differ from those in the other experiments to account for the low quality of public signals; see Appendix D for details).

In Experiment 2, 54% of the *unobserved* are successful observational learners (SOL) who guess optimally, 19% and 13% are weak and strong followers of private information (WFPI and SFPI) who, compared to SOL, respond too strongly to their private information, and 10% and 04% are weak and strong conformists (WC and SC) who, compared to SOL, herd too often. There are two apparent

2, the incentives to contradict high quality signals in Experiment 4 are stronger at contrary majorities of size less than 3 but weaker at contrary majorities of size 4 or more.

discrepancies between herd behavior at the aggregate and individual level. First, our aggregate results substantiate that *unobserved* with low or medium quality signals learn well from contrary majorities, but almost a third of them tend to follow their private information in situations where herding is the optimal guess. This discrepancy is easily accounted for since, in situations where they should herd, the average proportion of optimal guesses made by WFPI is 0.48 and only SFPI guess almost always in accordance with their private information (their average proportion of optimal guesses is 0.14). Second, though *unobserved* herd excessively on average, less than a sixth of them tend to contradict their private information in situations where it is suboptimal to do so. We note, however, that the average proportion of guesses that contradict high quality signals is 0.82 for conformists who face large contrary majorities and that in the same situations SOL also exhibit a substantial tendency to herd (their average proportion of contradictions is 0.38).

In Experiment 3, 38% of the *unobserved* are SOL, 06% and 10% are WFPI and SFPI, and 19% and 25% are WC and SC (subject 3411 guesses suboptimally more often than not and is classified as noisy). In line with our aggregate results, the largest share of *unobserved* herd too often and only a sixth of them respond too strongly to their private information. Moreover, when they are endowed with high quality signals and they face large contrary majorities, conformists almost always contradict their private information and SOL exhibit a non-negligible tendency to herd (the average proportion of contradictions is 0.92 and 0.23, respectively). Lastly, as in Experiment 2, in situations where private-information-followers should herd only SFPI guess almost always in accordance with their private information (their average proportion of optimal guesses is 0.23).

The distribution of decision rules in Experiment 4 resembles the one in Experiment 1 except for a lower share of SOL (15% vs. 32%) and a larger share of SFPI (28% vs. 06%). This difference is clearly in line with the fact that, on average, *unobserved* with medium quality signals respond more strongly to their private information in Experiment 4 than in Experiment 1. In the next section, we discuss in more details the proportion of optimal guesses by decision rule in Experiment 4.

3.3 Discussion

The results of Experiments 2-4 confirm that the observational learning success is undermined by non-Bayesian updating, informational misinferences, and also incorrect expectations about others' strategy. More importantly, our new evidence outlines the nature of these deviations from Bayesian rationality, and it shows that they impact herd behavior differently.

First, non-Bayesian updating and informational misinferences are the two channels that drive excessive herding with high quality signals. In Experiment 2, the proportion of contradictions rises too strongly with the (net) number of contrary signals which leads the average *unobserved* to herd excessively once she faces a sufficiently large number of contrary signals. On average, *unobserved* overinfer from medium-sized and large strings of contrary signals. In Experiment 3, *unobserved* herd excessively with high quality signals even when the relative incentives to contradict private information are 60% weaker than in Experiment 2. The manner in which she draws informational inferences from Bayes-rational guesses prevents the average *unobserved* from appreciating that cascade guesses have no informational value. Compared to Experiment 3, the relative incentives to contradict private information have to be increased by a third in Experiment 4 for the average *unobserved* to act against her quality signal. When endowed with high quality signals, *unobserved* learn more successfully from guesses submitted by other subjects than from Bayes-rational guesses.

Second, the overemphasis on low and medium quality signals is mainly caused by incorrect expectations about others' strategy. In Experiments 2 and 3, *unobserved* endowed with low or medium

quality signals essentially better respond to the true value of their available information.¹² Accordingly, underinferences from short strings of contrary signals are rare and *unobserved* appreciate the connection between the signals of Bayes-rational *observed* and their informative guesses. In Experiment 4, however, the average *unobserved* wrongly believes that human *observed* respond weakly to their private information which reduces her success when learning from short contrary majorities.

4 Intuitive Observational Learning

This section presents a structural model of intuitive observational learning which builds on the three belief distortions that drive herd behavior: non-Bayesian updating, informational misinferences, and wrong expectations about others’ strategy. In line with the central idea of the “heuristics and biases” program (Kahneman, Slovic, and Tversky, 1982; Gilovich, Griffin, and Kahneman, 2002),¹³ we posit that belief distortions result from observational learners making use of their intuitions when they extract signals from public guesses and incorporate those signals into their beliefs. Intuitive valuations of the available information are systematically biased because the cognitive processes underlying intuitive observational learning rest on general-purpose heuristics.

Each belief distortion is formulated as a one-parameter extension of its Bayes-rational counterpart so that our alternative model embeds the normative one as a constellation of parameter values. And to accommodate the rich behavioral heterogeneity found in our experiments, we allow for individual-specific parameter values. By formulating a simple parametric extension of Bayesian rationality we ensure its “*portability*” meaning that, once parameter values are fixed, new predictions can be made in various observational learning scenarios (for a detailed argumentation supportive of portable extensions of normative models, see Rabin, 2013).

Below, we first expose the formal details of our model, then we estimate the model econometrically, and finally we measure the gain in predictive power that results from belief distortions.

4.1 A Structural Model of Intuitive Observational Learning

As in our theoretical benchmark, we assume that intuitive observational learners—henceforth simply *Intuitive*—use their available information to form beliefs about the payoff-relevant states and that they make probabilistic money-maximizing guesses conditional on their beliefs. To facilitate exposition, we first describe formally the nature of the belief distortion when *Intuitive* learn from informative signals (as in Experiment 2). Next, we expand our formal description of intuitive beliefs by including informational misinferences when learning from Bayes-rational guesses (as in Experiment 3). Lastly, we complete the model by formalizing the wrong expectations that *Intuitive* have about the guessing strategy of others. In the interest of clarity, we restrict the exposition of intuitive observational learning to the informational and interactive structure of our experimental settings though we abstract from their specific parametrizations.

¹²Angrisan, Guarino, Jehiel, and Kitagawa (2017) report a similar finding in their “individual decision making” treatment as in our second experiment with low quality signals.

¹³This program explores the heuristics that people use and the biases to which they are prone in various tasks of judgment under uncertainty. It offers a cognitive alternative to the normative model of statistical reasoning according to which subjective judgments of probability are intuitive and often rest on a limited number of simplifying judgmental heuristics. Contrary to their deliberative counterparts, intuitive assessments of probability come to mind quickly and they directly reflect impressions of the characteristics of the available information. For example, some people expect the key characteristics of a random sample to be similar to those of its parent population, i.e., they view random samples as extremely representative. The tendency to rely on the representativeness heuristic leads to the “belief in the law of small numbers,” i.e., people expect small random samples to be more similar to their parent population than predicted by sampling theory (Tversky and Kahneman, 1971).

Before detailing each belief distortion, we introduce the elements common to every observational learning scenario. In period $t \in \{1, \dots, T\}$, the information available to *Intuitive* $i \in \{1, \dots, I\}$ consists of the (non-doctrinaire) prior $\mathbf{p} = (\Pr(\mathcal{B}), \Pr(\mathcal{O}))$ with $\Pr(\mathcal{B}) = 1 - \Pr(\mathcal{O}) \in (1/2, 1)$, her private signal $s_i \in \{b, o\}$ of quality $q_i \in (\Pr(\mathcal{B}), 1)$, and the history of public guesses $h_t = (g_1, \dots, g_{t-1}) \in H_t = \{\mathcal{B}, \mathcal{O}\}^{t-1}$ with $h_1 = \emptyset$. We assume that *Intuitive* understand that public guesses are submitted sequentially, and that they know the prior and the qualities of signals.

The beliefs of *Intuitive* i are shaped by her belief distortion type Ψ_i which specifies in a parametric form the relevant distortion(s) in the considered scenario. Formally, *Intuitive* i 's subjective probability that $\mathcal{B}$ is the payoff-relevant state in period t , conditional on her available information $(\mathbf{p}, s_i, h_t)$ and her type Ψ_i , is given by

$$\Pr_i(\mathcal{B} \mid \mathbf{p}, s_i, h_t; \Psi_i) = \left[1 + \frac{\Pr_i(\mathcal{O}, s_i, h_t; \Psi_i)}{\Pr_i(\mathcal{B}, s_i, h_t; \Psi_i)} \right]^{-1} \in [0, 1]$$

which we denote by $\mu_i(\mathbf{p}, s_i, h_t; \Psi_i)$ with $\Pr_i(\mathcal{O} \mid \mathbf{p}, s_i, h_t; \Psi_i) = 1 - \mu_i(\mathbf{p}, s_i, h_t; \Psi_i)$.

For simplicity we assume that private beliefs are Bayesian.¹⁴ This simplification entails that $\Pr_i(\theta, s_i, h_t; \Psi_i) = \Pr(\theta) \Pr_i(s_i \mid \theta) \Pr_i(h_t \mid \theta; \Psi_i)$ for each $\theta \in \{\mathcal{B}, \mathcal{O}\}$ where $\Pr_i(o \mid \mathcal{O}) = \Pr_i(b \mid \mathcal{B}) = q_i = 1 - \Pr_i(b \mid \mathcal{O}) = 1 - \Pr_i(o \mid \mathcal{B})$ and $\Pr_i(h_t \mid \theta; \Psi_i)$ denotes the subjective probability assigned by type Ψ_i to the occurrence of history h_t conditional on state θ . In period t , *Intuitive* i 's belief takes the form

$$\mu_i(\mathbf{p}, s_i, h_t; \Psi_i) = \left[1 + \frac{\Pr(\mathcal{O})}{\Pr(\mathcal{B})} \frac{\Pr_i(s_i \mid \mathcal{O})}{\Pr_i(s_i \mid \mathcal{B})} \frac{\Pr_i(h_t \mid \mathcal{O}; \Psi_i)}{\Pr_i(h_t \mid \mathcal{B}; \Psi_i)} \right]^{-1}, \quad (1)$$

and belief distortions are entirely reflected in the discrepancy between the subjective and the objective public likelihood ratio $\Pr(h_t \mid \mathcal{O}) / \Pr(h_t \mid \mathcal{B})$. As detailed below, this discrepancy originates from the failure to properly infer the informational content of public guesses or from the failure to incorporate inferred signals into beliefs as advocated by Bayes' rule.

We further assume that, given her belief $\mu_i(\mathbf{p}, s_i, h_t; \Psi_i)$ and her *payoff-responsiveness* $\lambda_i \geq 0$, *Intuitive* i submits guess B in period t with probability

$$\sigma_i(B \mid \mu_i(\mathbf{p}, s_i, h_t; \Psi_i); \lambda_i) = \frac{1}{1 + \exp(\lambda_i (1 - 2\mu_i(\mathbf{p}, s_i, h_t; \Psi_i)))},$$

and that she submits guess O with the complementary to one probability. Given λ_i , the stronger the belief of *Intuitive* i the larger her probability to submit guess B . Moreover, guesses approach best responses to beliefs as λ_i goes to infinity and they approach uniform randomness as it goes to zero. Though our main interest lies in the impact of belief distortions on observational learning behavior, there are two reasons for considering a logit quantal response version of intuitive observational learning. First, quantal responses provide an error structure for the structural estimation of belief distortion types. Second, logit QRE predictions match the behavioral patterns observed in cascade experiments significantly better than standard predictions (e.g., Goeree, Palfrey, Rogers, and McKelvey, 2007). Thus, a model of intuitive beliefs augmented with logistic decision errors enables us to measure the predictive power that is gained by introducing belief distortions into a logit quantal

¹⁴Though a less restrictive approach could easily be considered, the assumption that private beliefs are Bayesian seems appropriate to rationalize observational learning behavior in our experimental settings. Indeed, *unobserved* largely follow their private information in the first period of Experiments 2 to 4. While this evidence only provides equivocal support for our restriction, we conjecture that most subjects were able to properly combine the (almost flat) prior with their private signal as the latter corresponds to a single draw from an urn. More importantly, our experimental settings have not been designed to shed light on the judgment biases that might distort private beliefs.

response model of observational learning. The vector of parameters (Ψ_i, λ_i) summarizes *Intuitive i*'s behavioral type.

Intuitive Learning from Signals. Consider the simplest observational learning scenario where *Intuitive* know that public guesses perfectly reveal private signals whose quality is denoted by $q_{\text{PUB}} \in (\Pr(\mathcal{B}), 1)$ ($\Pr(\mathcal{B}) = 11/20$ and $q_{\text{PUB}} = 12/21$ in Experiment 2). In this observable-signals scenario, the objective public likelihood ratio in period t is simply given by $((1 - q_{\text{PUB}})/q_{\text{PUB}})^{\Delta(h_t)}$ where $\Delta(h_t)$ is the difference between the number of blue and orange public guesses in history h_t .

If their intuitions about random sampling satisfy the law of small numbers then the herd behavior of *Intuitive* matches qualitatively the herd behavior found at the aggregate level in Experiment 2. Indeed, overinferences from medium-sized strings of contrary public signals lead to excessive herding with high quality signals, but they don't undermine the observational learning success with medium and low quality signals. However, our analysis of herd behavior at the individual level has revealed that about a third of the *unobserved* tend to underinfer from very short strings of contrary public signals when they are endowed with low or medium quality signals. These subjects are conservative in the sense that they do not extract enough information from short contrary majorities. Former balls-and-urns experiments conducted by economists also find a substantial degree of individual heterogeneity in updating behavior (for recent evidence, see Holt and Smith, 2009). In particular, El-Gamal and Grether (1995) introduce a classification procedure to investigate the extent to which updating behavior is shaped by various judgmental heuristics. They conclude that the most prominent updating rules used by subjects (in order of prominence) are Bayes' rule, representativeness, and, to a lower extent, conservatism.

To accommodate the rich heterogeneity in individual behavior, we propose that *Intuitive* have different perceptions about the informativeness of public signals. These perceptions are entirely reflected in the weight assigned to public information relative to private information. Depending on the value of her public information weight, an *Intuitive* either makes Bayesian inferences, or she treats public signals as less informative than would a Bayesian, or she treats public signals as more informative than would a Bayesian. For simplicity, we assume that *Intuitive* understand that public signals are i.i.d. conditional on the payoff-relevant state, and that they assign the same weight to each public signal independently of their private signal. Though only the public information is weighted, our interpretation is not that *Intuitive* view private signals as different in nature from public signals. The weighted updating rule merely allows for the possibility that multiple informative signals are combined in a non-Bayesian manner.¹⁵

Formally, *Intuitive i*'s behavioral type is given by (w_i, λ_i) where $w_i \in [0, \infty)$ is the weight she assigns to each public signal, and Equation (1) simplifies to

$$\mu_i(\mathbf{p}, s_i, h_t; w_i) = \left[1 + \frac{\Pr(\mathcal{O})}{\Pr(\mathcal{B})} \frac{\Pr_i(s_i | \mathcal{O})}{\Pr_i(s_i | \mathcal{B})} \left(\frac{1 - q_{\text{PUB}}}{q_{\text{PUB}}} \right)^{w_i \cdot \Delta(h_t)} \right]^{-1}. \quad (2)$$

Weighted updating leads to heterogeneity in how *Intuitive* update their private beliefs in response to identical strings of public signals. If $w_i = 1$ then *Intuitive i* is a Bayesian observational learner. And if $w_i \in (1, \infty)$ then *Intuitive i* overweights public information, possibly because her intuitions about random sampling satisfy the law of small numbers. On the other hand, if $w_i \in [0, 1)$ then *Intuitive i* underweights public information as she updates her private beliefs too conservatively upon observing strings of public signals. Note that weights are log-symmetric since the distortion of beliefs for types

¹⁵Adaptive and information-theoretic foundations for the weighted updating rule are provided in March (2016) and Zinn (2016) respectively.

$w_i > 1$ is of the same magnitude as for types $1/w_i$. For example, when they are endowed with the same private signal, type $w_i = 2$ who faces a majority of one blue public signal holds the same beliefs as type $w_j = 1/2$ who faces a majority of four blue public signals, and those beliefs match the beliefs of a Bayesian observational learner who faces a majority of two blue public signals.

For the sake of parsimony, we have assumed that private beliefs are Bayesian meaning that only the public information is weighted.¹⁶ More importantly, the weighted updating rule reflects the view that normative assessments of probability provide good first approximations to subjective ones. Though this view sits well with conservatism, it is inconsistent with the representativeness heuristic (as already acknowledged by Grether, 1980). Indeed, Kahneman and Tversky have repeatedly argued that the most salient property of a binomial sample is the sample proportion and that subjective assessments of probability are quite insensitive to the sample size. Representativeness therefore implies that beliefs will depend primarily on the similarity between the proportion of blue public signals and q_{PUB} , rather than on the difference between the number of blue and orange public signals. In fact, the marked variability of the responses to a given *tvcPI* in Experiment 2 indicates that *unobserved* do not rely solely on the difference between the number of blue and orange signals to assess the value of their available information (see top panel in Figure 2).¹⁷ Despite its shortcomings, we decided to employ a restricted version of the weighted updating rule because it provides a tractable rationalization of both under- and over-inferences from public signals.

Intuitive Learning from Bayes-rational Guesses. Consider now the observational learning scenario where *Intuitive* know that public guesses are Bayes-rational. Let $\sigma^*(g_t | s_t, h_t) \in \{0, 1\}$ denote the probability that the Bayes-rational guess submitted in period t is g_t given private signal s_t of quality $q_{\text{PUB}} \in (\Pr(\mathcal{B}), 1)$ and public history h_t . Once guess g_t is publicly revealed, the objective public likelihood ratio $\Pr(h_t | \mathcal{O}) / \Pr(h_t | \mathcal{B})$ is updated with the multiplier

$$\begin{aligned} \frac{\Pr(g_t | h_t, \mathcal{O})}{\Pr(g_t | h_t, \mathcal{B})} &= \frac{\Pr(s_t = b | \mathcal{O}) \sigma^*(g_t | b, h_t) + \Pr(s_t = o | \mathcal{O}) \sigma^*(g_t | o, h_t)}{\Pr(s_t = b | \mathcal{B}) \sigma^*(g_t | b, h_t) + \Pr(s_t = o | \mathcal{B}) \sigma^*(g_t | o, h_t)} \\ &= \frac{(1 - q_{\text{PUB}}) \sigma^*(g_t | b, h_t) + q_{\text{PUB}} \sigma^*(g_t | o, h_t)}{q_{\text{PUB}} \sigma^*(g_t | b, h_t) + (1 - q_{\text{PUB}}) \sigma^*(g_t | o, h_t)} \end{aligned}$$

which, in case the guess is informative, simplifies to $(1 - q_{\text{PUB}})/q_{\text{PUB}}$ if $g_t = B$ and to $q_{\text{PUB}}/(1 - q_{\text{PUB}})$ if $g_t = O$ (in Experiment 3, $q_{\text{PUB}} = 2/3$ so that the multiplier equals 1/2 and 2 respectively). In contrast, the objective public likelihood ratio remains unchanged after an uninformative guess ($\sigma^*(g_t | b, h_t) = \sigma^*(g_t | o, h_t) = 1$).

Excessive herding is more pronounced in the third than in the second experiment. Though they draw proper informational inferences from public guesses that are informative, *unobserved* tend to overinfer from public guesses that are part of an informational cascade and are therefore uninformative. Before exposing our own approach to informational misinferences, we discuss two alternatives that are complementary in certain observational learning settings but are less relevant in the Bayes-rational-guesses scenario.

First, wrong expectations about others' strategy don't constitute a convincing explanation for why

¹⁶Previous investigations of the nature of subjective assessments of probability have employed a more flexible version of the weighted updating rule. For example, Grether (1980) introduced a two-weights updating rule to test whether subjects, who observe a sample of informative signals drawn from an unknown urn, rely on the representativeness heuristic when forming beliefs. The author concluded that, though priors are not ignored, representativeness captures reasonably well belief updating behavior in the aggregate since the estimated weight of the likelihood ratio is significantly larger than the estimated weight of the prior odds.

¹⁷Also, we estimated in Experiment 2 the public information weight for each subject as well as an extended version with a dummy variable that singles out the observations such that the proportion of blue public signals equals q_{PUB} . For one-third of the subjects, the coefficient of the dummy variable is positive, which is supportive of representativeness.

unobserved misinfer from cascade guesses since we provide an exhaustive description of how public guesses are generated. Second, our description of the generating process of public guesses nullifies, or at least severely limits, the need for *unobserved* to reason through how computer-*observed* make inferences from previous public guesses (see subsection 3.1.2 for details). Therefore, we don't view bounds on higher-order reasoning as the most convincing explanation for why *unobserved* do not fully account for the fact that cascade guesses are redundant (Eyster and Rabin, 2010, 2014).

Following Gennaioli and Shleifer (2010), we assume that the extent to which public guesses are representative of the underlying signals shapes the informational inferences made by *Intuitive*. Once guess g_t is publicly revealed, *Intuitive* make an inference that, compared to the Bayes-rational one, leans towards the signal of which g_t is most representative. In line with Kahneman and Tversky's representativeness relation between causal systems and their outcomes (Tversky and Kahneman, 1982), we say that guess g_t is more representative of signal s_t than of signal $\bar{s}_t$ if it is more frequently associated with signal s_t than with signal $\bar{s}_t$. Formally, the representativeness of public guess $g_t \in \{B, O\}$ for signal $s_t \in \{b, o\}$ is defined as $R(s_t, g_t) \equiv \Pr(g_t | s_t) \in [0, 1]$, $t \in \{1, \dots, T\}$. In Appendix E we show that, $\forall t \in \{1, \dots, T\}$, $R(s_t = b, g_t = B) > R(s_t = o, g_t = B)$ meaning that Bayes-rational guess B is more representative of signal b than of signal o . Likewise, Bayes-rational guess O is more representative of signal o than of signal b .

When they draw an inference from public guess g_t , *Intuitive* know that g_t is Bayes-rational. Still, their inference is potentially biased because the signal of which g_t is most representative first comes to mind and *Intuitive* anchor their belief to this signal. Using Gennaioli and Shleifer's terminology, *Intuitive* are "local thinkers" who make informational inferences in light of what comes to mind (the signal more frequently associated with the public guess), but not of what does not (the signal less frequently associated with the public guess). Given h_t and her degree of local thinking $\ell_i \geq 0$, *Intuitive* i updates her subjective likelihood ratio upon observing g_t with the multiplier

$$\frac{\Pr(g_t | h_t, \mathcal{O}; \ell_i)}{\Pr(g_t | h_t, \mathcal{B}; \ell_i)} = \frac{(1 - q_{\text{PUB}}) R(b, g_t)^{\ell_i} \sigma^*(g_t | b, h_t) + q_{\text{PUB}} R(o, g_t)^{\ell_i} \sigma^*(g_t | o, h_t)}{q_{\text{PUB}} R(b, g_t)^{\ell_i} \sigma^*(g_t | b, h_t) + (1 - q_{\text{PUB}}) R(o, g_t)^{\ell_i} \sigma^*(g_t | o, h_t)}.$$

Local thinking affects differently the updating of the subjective public likelihood ratio depending on the informativeness of the public guess. Indeed, if the public guess is informative then, for any degree of local thinking, the objective and subjective likelihood ratios are updated identically. Following an informative public guess, local thinking does not distort beliefs. On the other hand, if the public guess is uninformative then the higher the degree of local thinking the more distorted beliefs are. First, following uninformative guess B , the updating factor of the subjective public likelihood ratio belongs to the interval $(\frac{1-q_{\text{PUB}}}{q_{\text{PUB}}}, 1]$ and it approaches the lower bound of the interval as the degree of local thinking goes to infinity. Extreme local thinkers infer signal b from uninformative guess B . Second, following uninformative guess O , the updating factor of the subjective public likelihood ratio belongs to the interval $[1, \frac{q_{\text{PUB}}}{1-q_{\text{PUB}}})$ and it approaches the upper bound of the interval as the degree of local thinking goes to infinity. Extreme local thinkers infer signal o from uninformative guess O .

Finally, given $\Psi_i = (w_i, \ell_i)$, *Intuitive i*'s belief $\mu_i(\mathbf{p}, s_i, h_t; \Psi_i)$ in period $t \geq 2$ takes the form

$$= \left[1 + \left(\frac{1 - \mu_i(\mathbf{p}, s_i, h_{t-1}; \Psi_i)}{\mu_i(\mathbf{p}, s_i, h_{t-1}; \Psi_i)} \right) \left(\frac{\Pr(g_{t-1} | h_{t-1}, \mathcal{O}; \ell_i)}{\Pr(g_{t-1} | h_{t-1}, \mathcal{B}; \ell_i)} \right)^{w_i} \right]^{-1} \\ = \left[1 + \left(\frac{1 - \mu_i(\mathbf{p}, s_i, h_{t-1}; \Psi_i)}{\mu_i(\mathbf{p}, s_i, h_{t-1}; \Psi_i)} \right) \left(\frac{\sum_{s_{t-1}} \Pr(s_{t-1} | \mathcal{O}) R(s_{t-1}, g_{t-1})^{\ell_i} \sigma^*(g_{t-1} | s_{t-1}, h_{t-1})}{\sum_{s_{t-1}} \Pr(s_{t-1} | \mathcal{B}) R(s_{t-1}, g_{t-1})^{\ell_i} \sigma^*(g_{t-1} | s_{t-1}, h_{t-1})} \right)^{w_i} \right]^{-1} \quad (3)$$

with $\mu_i(\mathbf{p}, s_i, h_1 = \emptyset) = 1 / \left(1 + \frac{\Pr(\mathcal{O}) \Pr_i(s_i | \mathcal{O})}{\Pr(\mathcal{B}) \Pr_i(s_i | \mathcal{B})} \right)$, and

$$\begin{aligned} \sigma^*(B | b, h_{t-1}) &= 1 \text{ if } \Delta(h_{t-1}) \in \{-1, 0\}, \\ \sigma^*(B | o, h_{t-1}) &= 0 \text{ if } \Delta(h_{t-1}) \in \{-1, 0\}, \\ \sigma^*(B | s_{t-1}, h_{t-1}) &= 1 \quad \forall s_{t-1} \in \{b, o\} \text{ if } \Delta(h_{t-1}) \geq 1, \\ \sigma^*(B | s_{t-1}, h_{t-1}) &= 0 \quad \forall s_{t-1} \in \{b, o\} \text{ if } \Delta(h_{t-1}) \leq -2 \end{aligned}$$

where $\Delta(h_{t-1})$ is the difference between the number of blue and orange guesses in history h_{t-1} . In Appendix E, we show that the difference $R(b, B) - R(o, B)$ decreases over periods. Thus, the intuitive belief increases at a decreasing rate along a B -cascade: Each B guess that is part of an informational cascade is perceived as informative but less so than the previous one. Similarly, *Intuitive* perceive each O cascade guess as informative but less so than the previous one.

Intuitive Learning from Human Guesses. Lastly, consider the observational learning scenario where *Intuitive* learn from public guesses made by human subjects. In this human-*observed* scenario, *Intuitive* are uninformed of the decision process that generates public guesses. Therefore, to make inferences from public guesses, they need to form expectations about the strategies played by *observed*.

In view of the evidence collected in Experiments 2 and 3, we argued that non-Bayesian updating and informational misinferences reflect intuitive judgments that build on other attributes than normative judgments—for example, the more accessible attribute of representativeness is often substituted for the required target attribute of probability. It is our contention that intuitive reasoning also affects the formation of expectations about others' strategy. This leads us to postulate that *Intuitive* have a simple model of how public guesses are made which abstracts from the aforementioned belief distortions. Still, intuitive expectations should not starkly contradict *observed* play since we attempt at rationalizing the behavior of experienced observational learners who repeatedly learn from public guesses submitted by the same group of people. To accommodate both requirements, we assume that *Intuitive* expect *observed* to play (homogeneous) logit quantal-response equilibrium strategies.

Let $\lambda_{\text{PUB}}^{E_i} \geq 0$ denote the payoff-responsiveness that *Intuitive i* attributes to *observed* and that she believes to be commonly known among them. When $\lambda_{\text{PUB}}^{E_i} < \lambda_i$ (resp. $\lambda_{\text{PUB}}^{E_i} \geq \lambda_i$) *Intuitive i* expects *observed* to be less (resp. weakly more) responsive to their expected payoffs than herself. And as $\lambda_{\text{PUB}}^{E_i} \rightarrow 0$ *Intuitive i* tends to believe that public guesses are uninformative, while she tends to believe that *observed* play Bayes-rational equilibrium strategies as $\lambda_{\text{PUB}}^{E_i} \rightarrow \infty$. Conditional on $\lambda_{\text{PUB}}^{E_i}$, *Intuitive i* expects the *observed* who acts in period $\tau < t$ to submit guess B with probability

$$\sigma_{\text{PUB}}^{E_i}(B | s_\tau, h_\tau; \lambda_{\text{PUB}}^{E_i}) = \frac{1}{1 + \exp(\lambda_{\text{PUB}}^{E_i} (1 - 2\mu_{\text{PUB}}^{E_i}(\mathbf{p}, s_\tau, h_\tau; \lambda_{\text{PUB}}^{E_i})))}$$

and that guess O is submitted with the complementary probability $1 - \sigma_{\text{PUB}}^{E_i}(B | s_\tau, h_\tau; \lambda_{\text{PUB}}^{E_i})$ where

$$\mu_{\text{PUB}}^{E_i}(\mathbf{p}, s_\tau, h_\tau; \lambda_{\text{PUB}}^{E_i}) = \left[1 + \frac{\Pr(\mathcal{O})}{\Pr(\mathcal{B})} \frac{\Pr(s_\tau | \mathcal{O})}{\Pr(s_\tau | \mathcal{B})} \prod_{\rho < \tau} \frac{\sum_{s_\rho \in \{b, o\}} \Pr(s_\rho | \mathcal{O}) \sigma_{\text{PUB}}^{E_i}(g_\rho | s_\rho, h_\rho; \lambda_{\text{PUB}}^{E_i})}{\sum_{s_\rho \in \{b, o\}} \Pr(s_\rho | \mathcal{B}) \sigma_{\text{PUB}}^{E_i}(g_\rho | s_\rho, h_\rho; \lambda_{\text{PUB}}^{E_i})} \right]^{-1}$$

is the belief that *Intuitive* i expects the *observed* to form when the latter is endowed with private signal $s_\tau \in \{b, o\}$ with $\Pr(o | \mathcal{O}) = \Pr(b | \mathcal{B}) = q_{\text{PUB}} = 1 - \Pr(b | \mathcal{O}) = 1 - \Pr(o | \mathcal{B})$ and the history of public guesses is h_τ (the product over $\rho < \tau$ is assumed equal to one when $\tau = 1$). Therefore, *Intuitive* i 's belief $\mu_i(\mathbf{p}, s_i, h_t; (w_i, \ell_i, \lambda_{\text{PUB}}^{E_i}))$ derives from Equation (3) where $\sigma^*(g_{t-1} | s_{t-1}, h_{t-1})$ is replaced with $\sigma_{\text{PUB}}^{E_i}(g_{t-1} | s_{t-1}, h_{t-1}; \lambda_{\text{PUB}}^{E_i})$ and the representativeness of logit quantal-response equilibrium guesses for signals is a function of $\lambda_{\text{PUB}}^{E_i}$ whose properties are specified in Appendix E.

Illustrative Predictions of Intuitive Observational Learning. We predicted the responses to *vcPI* for 18 different behavioral types across the guessing situations of Experiment 4. Behavioral type $i = 1, \dots, 18$ is characterized by $\lambda_i = 10$, $w_i \in \{0.25, 1, 2.5\}$, $\lambda_{\text{PUB}}^{E_i} \in \{2.5, 10, 40\}$, and either a null degree of local thinking ($\ell_i = 0$) or an extreme degree of local thinking ($\ell_i \rightarrow \infty$). For the sake of space, we provide here only a summary of our prediction results (see Appendix E for details).

First, we find that behavioral types with $w_i = 0.25$ are always reluctant to contradict their low and medium quality signals. These behavioral types are less reluctant to contradict their low and medium quality signals when they are extreme local thinkers, but the level of noise they attribute to public guesses hardly affects their behavior. Behavioral type ($w_i = 1, \ell_i = 0, \lambda_{\text{PUB}}^{E_i} = 2.5$) is the only type with $w_i > 0.25$ that also overweights its low and medium quality signals, though mildly so.

Second, for any $\lambda_{\text{PUB}}^{E_i} \in \{2.5, 10, 40\}$, excessive herding with high quality signals results either from $w_i \geq 1$ and extreme local thinking or from $w_i = 2.5$ and the absence of local thinking. The difference is that behavioral types with $w_i = 2.5$ and $\ell_i = 0$ exhibit a stronger tendency to herd excessively with low and medium quality signals than extreme local thinkers who properly weight public information. Moreover, large public information weights lead to excessive herding with high quality signals even at short contrary majorities, a prediction that does not hold for large degrees of local thinking.

In sum, our illustrative predictions suggest that a small public information weight best captures the reluctance to contradict low and medium quality signals, while both pronounced local thinking and large public information weights lead to excessive herding with high quality signals. On the other hand, the level of noise assigned to public guesses seems to have little impact on the predictions made by intuitive observational learning.

4.2 Estimating Intuitive Observational Learning

We first report the estimation results for our structural model of intuitive observational learning, and then we examine the relative influence of belief distortions on herd behavior.

The behavioral type of each *unobserved* in Experiment 4 is estimated using maximum-likelihood techniques. Estimations are performed at the individual level to avoid making restrictive assumptions about the joint distribution of the parameters, and they rely on guesses made in Experiment 4 since each *unobserved* faced many short and large contrary majorities with all three signal qualities. We employ a step-wise estimation procedure to mitigate identification problems. For each *unobserved*, we repeatedly estimate her public information weight, her degree of local thinking, and her payoff-responsiveness while holding the ratio $\lambda_{\text{PUB}}^E/\lambda$ fixed at each value in the grid $\{0.1, 0.2, \dots, 0.9, 1\} \cup \{1/0.9, 1/0.8, \dots, 5, 10\}$. We then select the ratio that maximizes the log-likelihood across all 19

estimation runs. Finally, parameter estimates with $\lambda_{\text{PUB}}^E \neq \lambda$ are kept only if they significantly improve the fit over the parameter estimates with $\lambda_{\text{PUB}}^E = \lambda$ (Appendix F details our estimation approach).

Table 4 reports our estimation results.¹⁸ Note that, to facilitate the comparison between subjects, we report estimates of a rescaled version of the degree of local thinking, namely $\ell/(1+\ell)$. Henceforth, we simply refer to $\ell/(1+\ell)$ as the degree of local thinking. For each *unobserved*, columns 2-3 (resp. 4-5 and 6-7) report the selected estimate of parameter w (resp. $\ell/(1+\ell)$ and λ) and its bootstrapped standard error, column 8 reports the selected estimate of the ratio $\lambda_{\text{PUB}}^E/\lambda$, column 9 reports the log-likelihood of the model with the selected ratio, and column 10 reports the p -value of the likelihood ratio test.

There is a rich diversity in the weighting of public information and in the degree of local thinking. Slightly more than a quarter of the *unobserved* decidedly underweight public information ($\hat{w} \leq 1/4$) whereas almost a sixth of them quite strongly overweight public information ($\hat{w} \geq 3/2$). And though more than half of the *unobserved* have a pronounced degree of local thinking ($\hat{\ell}/(1+\hat{\ell}) \geq 1/2$), about one sixth of them make proper informational inferences ($\hat{\ell}/(1+\hat{\ell}) \leq 1/20$). On the other hand, for more than two-thirds of the *unobserved* $\hat{\lambda}_{\text{PUB}}^E \neq \hat{\lambda}$ does not significantly improve the fit of the model which indicates that *unobserved* mostly believe that *observed* have the same payoff-responsiveness as them (the 1st quartile, median, and 3rd quartile of the payoff-responsiveness distribution is 5.94, 7.60, and 13.03 respectively).

Belief distortions and herd behavior

To determine the relative influence of belief distortions on herd behavior, we explore the relationship between the decision rules *unobserved* have been assigned to and their estimated belief distortion types. Estimates of payoff-responsiveness are rarely mentioned since, except for subject 4410 who has been classified as noisy, $\hat{\lambda}$ is quite high for most *unobserved* (the median value is 6.30, 8.23, 14.55, 7.31 and 8.70 for the SFPI, WFPI, SOL, WC, and SC rule respectively).

For each *unobserved* in Experiment 4 (with the exception of subject 4410), Figure 3 shows her estimated belief distortion type in relation to the decision rule she has been assigned to. The figure contains a marker for each *unobserved* with x -value $\hat{w}$, y -value $\hat{\ell}/(1+\hat{\ell})$, and label $\hat{\lambda}_{\text{PUB}}^E/\hat{\lambda}$ where dark (resp. light) brown dots correspond to the SFPI (resp. WFPI) rule, white diamonds correspond to the SOL rule, and dark (resp. light) purple squares correspond to the SC (resp. WC) rule.

For almost all *unobserved* assigned to the SOL rule, neither the weighting of public information nor the degree of local thinking is too extreme ($0.62 < \hat{w} < 1.43$ and $0.29 < \hat{\ell}/(1+\hat{\ell}) < 0.71$), and the two estimates tend to be negatively correlated. This is clearly in line with the fact that the average proportion of optimal guesses is large with high quality signals (0.98 and 0.74 at contrary majorities of size less than 2 and more than 3 respectively) and with low or medium quality signals (0.84 and 0.98, when averaged across the two qualities, at contrary majorities of size less than 2 and more than 3 respectively). We also note that five out of the seven successful observational learners believe that *observed* have the same payoff-responsiveness as them.

For *unobserved* who have been assigned to the WC rule, $\hat{w}$ and $\hat{\ell}/(1+\hat{\ell})$ are also negatively correlated (of course, the median values of $\hat{w}$ and $\hat{\ell}/(1+\hat{\ell})$ are larger for weak conformists, 1.39 and 0.50, than for successful observational learners, 0.84 and 0.45). Thus, weak conformism often results either from the combination of a high degree of local thinking and mild overweighting (or even mild underweighting) of public information or from the combination of strong overweighting of public information and a rather low degree of local thinking. Still, in order for them to herd with

¹⁸Due to a network error in the first session of Experiment 4, subject 4109 was prevented from submitting guesses in non-practice parts 2 and 3 and is therefore excluded from the estimation.

<i>Unobserved</i>	$\hat{w}$		$\hat{\ell}/(1+\hat{\ell})$		$\hat{\lambda}$		$\hat{\lambda}_{\text{PUB}}^E/\hat{\lambda}$	<i>LL</i>	p-value
	Est.	SE	Est.	SE	Est.	SE			
4108	1.279	(0.159)	0.465	(0.070)	6.163	(0.552)	0.3	-60.1	0.067
4110	2.278	(0.204)	0.500	(0.048)	7.584	(0.577)	0.2	-41.2	0.011
4111	1.089	(0.104)	0.044	(0.049)	15.796	(1.942)	0.2	-23.6	0.001
4112	1.228	(0.173)	0.323	(0.099)	8.009	(0.856)	0.2	-53.7	0.070
4113	0.177	(0.089)	0.000	(0.175)	5.777	(1.024)	1	-88.1	0.154
4114	0.321	(0.064)	0.787	(0.165)	7.289	(1.156)	1	-59.6	0.352
4115	0.016	(0.028)	0.606	(0.267)	7.382	(0.855)	1	-72.0	0.644
4208	0.372	(0.109)	0.000	(0.095)	5.732	(0.751)	1	-82.3	0.358
4209	2.020	(0.282)	0.317	(0.082)	4.308	(0.414)	1	-66.3	0.127
4210	1.426	(0.098)	0.289	(0.053)	28.911	(8.863)	0.1	-10.7	0.008
4211	0.089	(0.032)	0.584	(0.215)	9.531	(2.289)	1	-53.7	0.815
4212	0.052	(0.036)	0.993	(0.460)	4.427	(0.456)	1	-101.2	1.000
4213	2.646	(0.238)	0.122	(0.070)	6.538	(0.973)	1	-37.2	0.382
4214	0.041	(0.008)	0.766	(0.107)	71.592	(49.977)	1	-15.3	0.923
4215	0.767	(0.057)	0.539	(0.053)	9.774	(1.745)	1	-37.8	0.586
4308	0.346	(0.027)	0.994	(0.103)	14.810	(5.457)	1	-27.0	0.204
4309	1.349	(0.191)	0.433	(0.107)	7.514	(1.689)	1	-28.0	0.285
4310	0.067	(0.053)	0.275	(0.210)	6.102	(3.178)	1	-81.0	0.170
4311	0.045	(0.016)	0.992	(0.280)	14.527	(8.866)	1	-38.7	0.849
4312	0.242	(0.022)	0.968	(0.062)	21.097	(6.104)	1	-23.3	0.993
4313	1.104	(0.079)	0.994	(0.103)	46.877	(119.44)	1	-5.6	1.000
4314	0.544	(0.073)	0.899	(0.121)	4.861	(0.565)	1	-67.6	0.692
4315	1.000	(0.077)	0.303	(0.048)	11.717	(11.510)	1	-24.3	0.379
4408	0.196	(0.068)	0.691	(0.208)	8.438	(1.544)	1	-54.3	0.873
4409	0.621	(0.031)	0.706	(0.017)	22.172	(7.017)	10	-17.8	0.094
4410	0.000	(7.303)	0.000	(0.482)	0.170	(0.212)	1	-166.2	0.884
4411	0.562	(0.200)	0.000	(0.118)	4.811	(0.543)	1	-91.8	1.000
4412	1.115	(0.101)	0.687	(0.048)	7.596	(1.204)	10	-45.5	0.001
4413	1.413	(0.107)	0.592	(0.034)	7.435	(0.522)	0.2	-49.8	0.095
4414	0.853	(0.079)	0.799	(0.037)	5.089	(0.611)	10	-68.8	0.079
4415	0.742	(0.077)	0.815	(0.037)	6.310	(0.718)	10	-59.7	0.017
4508	1.019	(0.062)	0.447	(0.043)	14.903	(2.210)	1	-23.1	0.328
4509	0.368	(0.041)	0.945	(0.097)	8.355	(1.311)	1	-49.1	0.391
4510	0.000	(0.035)	0.000	(0.270)	6.833	(0.941)	1	-73.0	0.999
4511	0.838	(0.073)	0.307	(0.059)	8.004	(1.058)	1	-51.7	0.505
4512	0.686	(0.072)	0.266	(0.057)	13.287	(1.304)	0.2	-34.4	0.003
4513	0.000	(0.000)	0.000	(0.000)	2.524	(0.342)	1	-134.9	1.000
4514	0.136	(0.024)	0.700	(0.085)	17.074	(3.232)	1	-31.1	0.284
4515	1.807	(0.166)	0.365	(0.056)	7.231	(0.593)	0.3	-40.3	0.003
4608	1.857	(0.162)	0.068	(0.062)	6.114	(0.633)	1	-43.8	0.366
4609	1.394	(0.169)	0.988	(0.122)	9.957	(26.639)	1	-20.7	0.379
4610	1.270	(0.065)	0.338	(0.035)	14.682	(1.172)	0.2	-22.8	0.004
4611	2.181	(3.897)	0.990	(0.053)	8.203	(5.033)	1	-20.4	1.000
4612	0.651	(0.254)	0.000	(0.148)	2.280	(0.336)	1	-137.3	0.968
4613	0.286	(0.051)	0.988	(0.196)	3.896	(0.505)	1	-100.7	1.000
4614	0.376	(0.091)	0.681	(0.121)	5.184	(0.587)	10	-82.8	0.044
4615	3.146	(0.257)	0.407	(0.053)	12.769	(0.933)	0.1	-21.7	0.040

Note: Bootstrapped standard errors in parentheses.

Table 4: Parameter Estimates for Intuitive Observational Learning

high quality signals at large contrary majorities, weak conformists with $\hat{w} < 1.11$ have to believe that *observed* best respond (among those weak conformists, the median value of $\hat{\lambda}_{\text{PUB}}^E$ is 50.89). On the other hand, weak conformists who tend to strongly overweight public information also tend to perceive *observed* guesses as noisy (among those weak conformists, the median value of $\hat{\lambda}_{\text{PUB}}^E$ is 3.62).

private information is well rationalized by a substantial underweighting of public information.

4.3 Assessing the Predictive Power of Intuitive Observational Learning

To assess the predictive power of our model of intuitive observational learning (IOL), we measure the accuracy of its predictions relative to the guesses made by *unobserved* in Experiment 4. We also investigate whether the nature of expectations affects IOL’s ability to predict accurately by comparing the predictive power of alternative models that differ in the richness of the expectations that *Intuitive* have about others’ strategy (details about our prediction framework and results are provided in Appendix F).


We consider 4 structural models of intuitive observational learning. In each model, the behavioral type of *Intuitive* i is summarized by five parameters $(w_i, \ell_i, \lambda_{\text{PUB}}^{E_i}, \lambda_{\text{PUB}}^{E_i^2}, \lambda_i)$. The two parameters $\lambda_{\text{PUB}}^{E_i} \geq 0$ and $\lambda_{\text{PUB}}^{E_i^2} \geq 0$ characterize how expectations are formed. Concretely, *Intuitive* expect *observed* to play logit quantal-response strategies with payoff-responsiveness $\lambda_{\text{PUB}}^{E_i}$ and they believe that *observed* expect others to play logit quantal-response equilibrium strategies with payoff-responsiveness $\lambda_{\text{PUB}}^{E_i^2}$. In IOL, we have that $\lambda_{\text{PUB}}^{E_i} = \lambda_{\text{PUB}}^{E_i^2}$ since *Intuitive* expect *observed* to play equilibrium strategies. And $\lambda_{\text{PUB}}^{E_i}$ is deemed different from λ_i only if the fit of the model is significantly improved. We consider a first alternative to IOL, referred to as 2 λ s-QRE, where $\lambda_{\text{PUB}}^{E_i}$ always equals its estimated value. The second alternative, referred to as 3 λ s-QR, is a more drastic variation of IOL where we allow all three payoff-responsivenesses to differ $(\lambda_i \neq \lambda_{\text{PUB}}^{E_i}, \lambda_i \neq \lambda_{\text{PUB}}^{E_i^2}, \text{ and } \lambda_{\text{PUB}}^{E_i} \neq \lambda_{\text{PUB}}^{E_i^2})$ meaning that *Intuitive* i might expect that *observed* play non-equilibrium strategies. In particular, *observed* are expected to (almost) always guess in accordance with their private information if $\lambda_{\text{PUB}}^{E_i^2}$ is negligible but $\lambda_{\text{PUB}}^{E_i}$ isn’t. Finally, we consider a third alternative, referred to as 1 λ -QRE, where *Intuitive* i expects *observed* to play equilibrium strategies with the same payoff-responsiveness as her own meaning that $\lambda_i = \lambda_{\text{PUB}}^{E_i} = \lambda_{\text{PUB}}^{E_i^2}$.

To measure the predictive power of a model, we first estimate the behavioral type of each of the 47 *unobserved* (subject 4109 is excluded). Second, we predict for each behavioral type its probability to contradict private information in each guessing situation. Third, we take the average of these 47 predicted probabilities in each guessing situation, and we compute the weighted sum of squared differences (*SSD*) between the average predicted probabilities and the empirical propensities to contradict private information where the weight of a guessing situation is given by the number of its occurrences. Finally, we compute the model’s predictive power as $1 - \text{SSD}/\text{SSD}_B$ where SSD_B is the weighted sum of squared differences for our theoretical benchmark. Remember that benchmark players know the informational value of public guesses, form Bayesian beliefs, and make probabilistic money-maximizing guesses conditional on their beliefs. To compute SSD_B , we assume more specifically that benchmark guesses are logit quantal-responses to *vcPI* and we estimate the payoff-responsiveness of each *unobserved* which enables us to predict the average benchmark probabilities to contradict private information.¹⁹ Given our definition of the predictive power of a model, if *unobserved* guesses are perfectly matched by the predictions of a model then its predictive power equals 100 percent. And since the predictive power of our theoretical benchmark is null, a model that predicts worse than our theoretical benchmark has a negative predictive power.

Table 5 reports the predictive powers of the 4 models of intuitive observational learning by the quality of private signals and averaged across qualities. Figure 4 shows the responses to *vcPI* predicted by IOL where only guessing situations with *sitcount* ≥ 10 are considered and fitted lines of weighted IV regressions are superimposed.

¹⁹To derive SSD_B from the same sample of guesses as each model’s *SSD*, we replace *vcPI* by a value based on the state-contingent relative frequencies of histories in every guessing situation where it cannot be calculated. See Appendix F.3 for details.

Signal Quality	IOL	1 λ -QRE	2 λ s-QRE	3 λ -QR
Low	60.8%	60.2%	60.7%	61.4%
Medium	86.3%	85.6%	86.5%	86.7%
High	83.1%	82.9%	83.6%	84.3%
All	76.8%	76.4%	77.0%	77.6%



Note: ●, ●, ●: Predicted responses with high, medium and low quality signals respectively.

Figure 4: Predicted Responses to the Empirical Value of Contradicting Private Information

Two main observations can be made from Table 5. First, IOL’s predictive power is substantially stronger than the predictive power of our theoretical benchmark, though the increase in predictive power varies with the quality of private signals. When *unobserved* are endowed with medium or high quality signals, we find that IOL achieves an impressive predictive power of about 85% meaning that its SSD is only one-sixth of SSD_B . With low quality signals, however, IOL’s predictive power drops to 60.8% which implies that its SSD is almost two-fifth of SSD_B . Second, the nature of expectations hardly affects the ability of intuitive observational learning to predict accurately. Indeed, for every signal quality, the difference between the strongest predictive power of 3 λ s-QR and the weakest predictive power of 1 λ -QRE is less than 2 percentage points. We performed a simulation exercise to check whether the small differences in predictive power are statistically significant. For each model, we randomly drew guesses from a binomial distribution where the probability of a contradictory guess is the model’s average predicted probability to contradict private information. We then computed in each guessing situation the relative frequency of simulated guesses that contradict private information and we derived the model’s predictive power from these relative frequencies. We repeated the simulation process 1,000 times to construct the 90%-confidence intervals of the model’s predictive power. We find that the simulated confidence intervals largely overlap and we cannot reject the null hypothesis that the 4 models of intuitive observational learning have identical predictive power.²⁰

²⁰Unsurprisingly, our estimation results show that allowing for rich expectations about others’ strategy can partly substitute for non-Bayesian updating and local thinking in capturing the behavior of *unobserved*. Indeed, as expectations become more flexible we find higher estimates of public information weights and lower estimates of local thinking degrees.

We now examine in more details the extent to which IOL is able to predict excessive herding with high quality signals and overweighting of low and medium quality private signals. First, in guessing situations where *unobserved* face contrary majorities with high quality signals and $vcPI < 0.4$, IOL achieves an average predictive power of about 57%. It correctly predicts that the proportion of contradictions increases sharply with the size of the majority—on average, the empirical (resp. predicted) proportion is 0.11 (resp. 0.08) at small contrary majorities and 0.42 (0.38) at large contrary majorities—whereas our theoretical benchmark predicts only a mild increase (from 0.20 to 0.26). Thus, as illustrated in Figure 4, IOL captures well the phenomenon of excessive herding with high quality signals, and its strong predictive power in those guessing situations results from both the overweighting of public information and local thinking.

Second, in guessing situations where *unobserved* face small contrary majorities with medium quality signals and $0.6 > vcPI \geq 0.5$, IOL achieves an average predictive power of 69%. In contrast to our theoretical benchmark, IOL correctly predicts that a minority of guesses contradict private information. IOL also correctly predicts that only about half of the guesses are contradictions when $vcPI \geq 0.6$ —on average, the empirical (resp. predicted) proportion is 0.54 (resp. 0.53) whereas our theoretical benchmark predicts that most guesses contradict private information (the benchmark proportion equals 0.78), and IOL’s predictive power reaches 83%. Thus, IOL also captures well the phenomenon of overweighting of private information when signals are of medium quality. We note that in those guessing situations the average predictive power of IOL is comparable to the average predictive power of both 1λ -QRE and 3λ s-QR which indicates that the strong predictive power of intuitive observational learning results from the underweighting of private information.

Third, in guessing situations where *unobserved* face small contrary majorities with low quality signals and $vcPI \geq 0.5$, IOL achieves an average predictive power of only 25%. Though IOL’s predictive power is positive, it is rather low partly because our theoretical benchmark also predicts well. Both IOL and our theoretical benchmark exaggerate the proportion of contradictions when $vcPI < 0.6$ —the empirical, IOL predicted, and benchmark proportion is 0.50, 0.56, and 0.61 respectively—and even more so when $vcPI \geq 0.6$ —the empirical, IOL predicted, and benchmark proportion is 0.61, 0.74, and 0.77 respectively.

At last, there are guessing situations where IOL’s predictions are quite at odds with *unobserved* guesses. Mostly, IOL overpredicts the propensity to contradict low quality signals in the absence of a majority when $0.4 \leq vcPI < 0.5$, though our theoretical benchmark predicts even worse—the empirical, IOL predicted, and benchmark proportion of contradictions is 0.12, 0.38, and 0.45 respectively. Note that most of the corresponding guesses (241 out of 312) are made in period 1 where none of the belief distortions affects the predicted probability to contradict private information.²¹

5 Concluding Discussion

We designed four experiments to explore whether forces other than Bayes-rational inferences drive herd behavior and, if so, to delineate the nature of these forces. In Experiment 1, *unobserved* learn from public guesses made by other subjects who are endowed with medium quality signals. We find

See Appendix F.2 for details.

²¹Regarding the many guessing situations where *unobserved* face a favoring or no majority and $vcPI < 0.4$, IOL achieves an average predictive power of about 90% with medium or high quality signals and almost 80% with low quality signals. Indeed, IOL predicts proportions of contradictions that are almost as tiny as the empirical ones whereas our theoretical benchmark predicts proportions that are between 3 and 5 times larger than the empirical ones. The two reasons are: i) local thinking causes the predicted probability to contradict private information to sharply decrease with the size of the favoring majority, a property that is shared by the data; and ii) estimated payoff-responsivenesses are considerably larger for IOL than for our theoretical benchmark (about three-quarters of the IOL payoff-responsivenesses are larger than all of the benchmark payoff-responsivenesses).

that *unobserved* herd excessively with high quality signals, they learn rather successfully from public guesses with medium quality signals, and they overweight their low quality signals relative to public information. Experiments 2-4 reveal that non-Bayesian updating and informational misinferences are the two channels that drive excessive herding, while the strong (resp. mild) overemphasis on low (resp. medium) quality signals is mainly caused by incorrect expectations about others' strategy. A model of intuitive observational learning (IOL) accounts for the phenomenon of excessive herding and it also captures well observational learning with medium quality signals. However, IOL fails to predict that *unobserved* are more reluctant to contradict their low than their medium quality signals.

We conclude by summarizing the content of i) two appendices that investigate the robustness of our prediction results with respect to the modelling assumptions of IOL; and ii) one appendix that evaluates the increase in IOL's predictive power due to the inclusion of efficiency concerns.

Appendix G evaluates the predictive value of heterogeneity in the belief distortions that compose IOL. We find that allowing for individual-specific weights in Experiment 2 is of considerable predictive value. Indeed, non-Bayesian updating with a single weight has the same predictive power as Bayesian updating which is approximately 20% lower than the predictive power of heterogeneous non-Bayesian updating. In contrast, local thinking with a common degree, whose estimate is 0.572, predicts as well as heterogeneous local thinking in Experiment 3, despite considerable diversity in the individual degrees of local thinking. Restricting IOL to a single degree of local thinking still permits a good description of herd behavior while removing many degrees of freedom.

In Appendix H, we measure the predictive power of two restricted versions of 1λ -QRE and IOL where i) belief updating is Bayesian and ii) local thinking is absent. We find that the predictive power of 1λ -QRE markedly decreases when either non-Bayesian updating or local thinking are dispensed with. Conversely, the full and restricted versions of IOL predict equally well for each signal quality. These findings confirm that, once they are sufficiently rich, expectations about others' strategy can substitute for non-Bayesian updating or local thinking in describing herd behavior. These findings also help understand why deviations from Bayesian rationality in cascade experiments have often been interpreted as errors in higher-order reasoning. Still, the fact that expectation models of how others draw informational inferences are flexible enough to be descriptively accurate does not entail that they pinpoint the main principles behind herd behavior. Though we acknowledge that people are unlikely to have a correct model of how others learn from public guesses, our new evidence indicates that non-Bayesian updating and local thinking are vital principles governing herd behavior.

Appendix I measures the out-of-sample predictive power of (extensions of) IOL by calibrating the model from *unobserved* guesses in Experiment 4 and predicting *observed* guesses in the same experiment as well as *observed* and *unobserved* guesses in Experiment 1. As expected, we find that the predictive power of IOL in Experiment 4 is significantly lower for *observed* guesses than for *unobserved* guesses with medium quality signals. To capture the behavioral differences between *observed* and *unobserved*, we propose a simple extension of IOL that incorporates efficiency concerns. The extension increases IOL's predictive power for *observed* guesses by almost 12% in Experiment 1 and by only 5% in Experiment 4. These prediction results confirm that the observational learning behavior of *observed* is partially driven by efficiency concerns and that their impact is stronger in Experiment 1 than in Experiment 4.

References

ANDERSON, L., AND C. HOLT (1997): "Information Cascades in the Laboratory," *American Economic Review*, 87, 847–62.

- ANGRISANI, M., A. GUARINO, P. JEHIEL, AND T. KITAGAWA (2017): “Information Redundancy Neglect versus Overconfidence: A Social Learning Experiment,” CEMMAP Working Paper 32/17.
- BANERJEE, A. V. (1992): “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 107, 797–817.
- 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.
- BOHREN, J. A. (2016): “Informational Herding with Model Misspecification,” *Journal of Economic Theory*, 163, 222–47.
- BRUNNER, C., AND J. K. GOEREE (2011): “Wise Crowds or Wise Minorities?,” Unpublished.
- CALLANDER, S., AND J. HÖRNER (2009): “The Wisdom of the Minority,” *Journal of Economic Theory*, 144, 1421–39.
- CIPRIANI, M., AND A. GUARINO (2009): “Herd Behavior in Financial Markets: An Experiment with Financial Market Professionals,” *Journal of the European Economic Association*, 7, 206–33.
- DUFFY, J., E. HOPKINS, T. KORNIENKO, AND M. MA (2017): “More is Better? Information Choice in a Social Learning Experiment,” Unpublished.
- EDWARDS, W. (1968): “Conservatism in Human Information Processing,” in *Formal Representation of Human Judgment*, ed. by B. Kleinmuntz. John Wiley and Sons: New York.
- EL-GAMAL, M. A., AND D. M. GREETHER (1995): “Are People Bayesian? Uncovering Behavioral Strategies,” *Journal of the American Statistical Association*, 90, 1137–45.
- EYSTER, E., AND M. RABIN (2005): “Cursed Equilibrium,” *Econometrica*, 73, 1623–72.
- (2010): “Naïve Herding in Rich-Information Settings,” *American Economic Journal: Microeconomics*, 2, 221–43.
- (2014): “Extensive Imitation is Irrational and Harmful,” *Quarterly Journal of Economics*, 129, 1861–98.
- EYSTER, E., M. RABIN, AND G. WEIZSÄCKER (2015): “An Experiment on Social Mislearning,” Unpublished.
- FISCHBACHER, U. (2007): “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, 10, 172–178.
- GENNAIOLI, N., AND A. SHLEIFER (2010): “What Comes to Mind,” *Quarterly Journal of Economics*, 125, 1399–1433.
- GILOVICH, T., D. GRIFFIN, AND D. KAHNEMAN (2002): *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge University Press.
- GOEREE, J. K., T. PALFREY, B. ROGERS, AND R. MCKELVEY (2007): “Self-Correcting Information Cascades,” *Review of Economic Studies*, 74, 733–762.
- GREINER, B. (2015): “Subject pool recruitment procedures: Organizing experiments with ORSEE,” *Journal of the Economic Science Association*, 1, 114–25.

- GRETHER, D. M. (1980): “Bayes Rule as a Descriptive Model: The Representativeness Heuristic,” *Quarterly Journal of Economics*, 95, 537–57.
- HOLT, C. A., AND A. M. SMITH (2009): “An Update on Bayesian Updating,” *Journal of Economic Behavior & Organization*, 69, 125–34.
- HUCK, S., AND J. OECHSSLER (2000): “Informational Cascades in the Laboratory: Do They Occur for the Right Reasons?,” *Journal of Economic Psychology*, 21, 661–71.
- KAHNEMAN, D. (2003): “Maps of Bounded Rationality: Psychology for Behavioral Economics,” *The American Economic Review*, 93, 1449–75.
- KAHNEMAN, D., P. SLOVIC, AND A. TVERSKY (1982): *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press.
- KEYNES, J. M. (1930): *A Treatise on Money*. London: Macmillan.
- KINDLEBERGER, C. P. (1978): *Manias, Panics and Crashes: A History of Financial Crises*. London: Macmillan.
- KÜBLER, D., AND G. WEIZSÄCKER (2004): “Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory,” *Review of Economic Studies*, 71, 425–41.
- (2005): “Are Longer Cascades More Stable?,” *Journal of the European Economic Association*, 3, 330–39.
- LEE, I. H. (1993): “On the Convergence of Informational Cascades,” *Journal of Economic Theory*, 61, 395–411.
- MARCH, C. (2016): “Adaptive Social Learning,” CESifo Working Paper No. 5783.
- RABIN, M. (2002): “Inference by Believers in the Law of Small Numbers,” *Quarterly Journal of Economics*, 117, 775–816.
- (2013): “An Approach to Incorporating Psychology into Economics,” *American Economic Review*, 103, 617–622.
- SMITH, L., AND P. SØRENSEN (2000): “Pathological Outcomes of Observational Learning,” *Econometrica*, 68, 371–98.
- TVERSKY, A., AND D. KAHNEMAN (1971): “Belief in the Law of Small Numbers,” *Psychological Bulletin*, 76, 105–10.
- (1982): “Judgments Of and By Representativeness,” in *Judgment under Uncertainty: Heuristics and Biases*, ed. by D. Kahneman, P. Slovic, and A. Tversky, pp. 84–98. Cambridge University Press.
- 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.
- ZIEGELMEYER, A., C. MARCH, AND S. KRUEGEL (2013): ““Do We Follow Others when We Should? A Simple Test of Rational Expectations”: Comment,” *American Economic Review*, 103, 2633–42.

ZINN, J. A. (2016): “An Information-Theoretic Foundation for the Weighted Updating Model,”
Unpublished.