

March, Christoph; Ziegelmeyer, Anthony

Working Paper

Behavioral social learning

Jena Economic Research Papers, No. 2009,105

Provided in Cooperation with:

Max Planck Institute of Economics

Suggested Citation: March, Christoph; Ziegelmeyer, Anthony (2009) : Behavioral social learning, Jena Economic Research Papers, No. 2009,105, Friedrich Schiller University Jena and Max Planck Institute of Economics, Jena

This Version is available at:

<https://hdl.handle.net/10419/32567>

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.



JENA ECONOMIC RESEARCH PAPERS



2009 – 105

Behavioral Social Learning

by

**Christoph March
Anthony Ziegelmeyer**

www.jenecon.de

ISSN 1864-7057

The JENA ECONOMIC RESEARCH PAPERS is a joint publication of the Friedrich Schiller University and the Max Planck Institute of Economics, Jena, Germany. For editorial correspondence please contact markus.pasche@uni-jena.de.

Impressum:

Friedrich Schiller University Jena
Carl-Zeiss-Str. 3
D-07743 Jena
www.uni-jena.de

Max Planck Institute of Economics
Kahlaische Str. 10
D-07745 Jena
www.econ.mpg.de

© by the author.

Behavioral Social Learning*

Christoph MARCH[†]

Anthony ZIEGELMEYER[‡]

Preliminary version. Comments are welcome.

This version: December 2009.

Abstract

We revisit the economic models of social learning by assuming that individuals update their beliefs in a non-Bayesian way. Individuals either overweigh or underweigh (in Bayesian terms) their private information relative to the public information revealed by the decisions of others and each individual's updating rule is private information. First, we consider a setting with perfectly rational individuals with a commonly known distribution of updating rules. We show that introducing heterogeneous updating rules in a simple social learning environment reconciles equilibrium predictions with laboratory evidence. Additionally, a model of social learning with bounded private beliefs and sufficiently rich updating rules corresponds to a model of social learning with unbounded private beliefs. A straightforward implication is that heterogeneity in updating rules is efficiency-enhancing in most social learning environments. Second, we investigate the implications of heterogeneous updating rules in social learning environments where individuals only understand the relation between the aggregate distribution of decisions and the state of the world. Unlike in rational social learning, heterogeneous updating rules do not lead to a substantial improvement of the societal welfare and there is always a non-negligible likelihood that individuals become extremely and wrongly confident about the state of the world.

KEYWORDS: Social learning; Non-Bayesian updating; Herding; Informational cascades.

JEL CLASSIFICATION: D82; D83.

*The authors benefited from discussions with the participants of the Schumpeter seminar (Humboldt University, Berlin).

[†]*Corresponding author:* Max Planck Institute of Economics, IMPRS "Uncertainty", Jena (Germany). Tel.: +49-3641-686674. Fax: +49-3641-686667. Email: MARCH@ECON.MPG.DE

[‡]Max Planck Institute of Economics, Strategic Interaction Group, Jena (Germany) and Technical University of Berlin, Faculty of Economics and Management, Berlin (Germany). Tel.: +49-3641-686630 and +49-30-314-23910. Fax: +49-3641-686667. Email: ZIEGELMEYER@ECON.MPG.DE

1 Introduction

Consider a large population of individuals, each makes only one decision and is endowed with a noisy private signal about an underlying state of the world. An individual's payoff depends on the underlying state of the world and her decision but is unaffected by the decisions of other individuals. Private signals are unbiased so that the aggregation of all private information resolves uncertainty completely. Individuals observe the decisions of those who decided earlier but not their private signals. In this situation, individuals care about the decisions of others because these decisions potentially reveal payoff-relevant information. Two natural questions, then, are: Can individuals learn enough to make the decision that yields the highest payoff by observing the decisions of others? And if so, how fast will individuals learn from others?

In the last two decades, economists have developed models of social learning to answer these two questions. Ultimately, the aim is to discuss the efficiency of the equilibrium in the presence of information externalities. An important class of social learning models assumes that each individual's decision reflects in a Bayes-rational fashion the content of her private signal and of the history of observed decisions. Using Bayes' rule, each individual forms her belief by combining her private belief, the probability estimate of the state of the world based solely on her private signal, with the public belief, the probability estimate of the state of the world based solely on the history of observed decisions. Accordingly, the process of social learning is the diffusion of the private beliefs to all individuals through the interactions of observed decisions, Bayesian updating and payoff-maximizing decisions.

In the seminal study on social learning by Bikhchandani, Hirshleifer, and Welch (1992) (henceforth BHW),¹ each individual in an infinite ordered sequence makes a one-shot binary decision which she may condition both on her private signal and on all her predecessors' decisions. BHW's main result is that the attempt to take advantage of the information of their predecessors prevents individuals from exploiting their private information in a socially optimal way. This likely consequence of social learning is what has been called *information cascades*. An information cascade occurs when the accumulated evidence from previous decisions is so conclusive that individuals rationally herd without regard to their private information. In an information cascade, decisions do not convey private information, the benefit of diversity of information is lost, social learning stops completely and the failure of information aggregation is spectacular.²

Smith and Sørensen (2000) provides the most comprehensive and exhaustive analysis of social learning in situations where players observe the full sequence of past decisions, and establishes that the failure of information aggregation is not a robust property. Information cascades depend on the fact that there is a maximum amount of information in any individual signal—private beliefs are bounded. Smith and Sørensen's main result is that when private beliefs are unbounded, information is asymptotically fully revealed (asymptotic learning). However, a consequence of the martingale convergence theorem is that herds take place for sure, eventually. In a herd all

¹See also Banerjee (1992).

²Since decisions are the words for the transmission of information between individuals, information cascades occur only if the information space is large relative to the decision space.

individuals make the same decision though there is the possibility that an individual has a private signal that induces her to make a different decision. Hence, there is some learning but the probability of breaking the herd must be extremely small for the herd to be realized which implies that social learning is very slow in a herd. As emphasized by Chamley (2004), the central message of models of Bayesian learning from others is the self-defeating property of social learning, or equivalently, the observation that informational externalities result in serious failures to achieve a desirable social outcome.

Numerous laboratory studies have checked the validity of the Bayesian rational view of herding (among others, Anderson and Holt, 1997; Kübler and Weizsäcker, 2004; Goeree, Palfrey, Rogers, and McKelvey, 2007; Ziegelmeyer, Koessler, Bracht, and Winter, 2009). Most of these economic experiments on social learning implement a simple environment which is based on BHW's specific model. Equilibrium predictions are only partially corroborated by the experimental evidence and the main observed regularities are: (i) Laboratory cascades emerge but they do so latter than predicted and, contrary to equilibrium predictions, a *short* laboratory cascade is often broken by participants with *low-accuracy contradictory* signals; (ii) Laboratory cascades are self-correcting meaning that after the break of an incorrect laboratory cascade the new laboratory cascade which emerges is often a correct one; (iii) *Long* laboratory cascades are stable and, contrary to equilibrium predictions, they are not broken by participants with *high-accuracy contradictory* signals; and (iv) Unlike in equilibrium, the more cascade choices participants observe the more they believe in the state favored by those choices. Several alternative theories of behavior to Bayesian rationality have been suggested in the experimental literature to account for these four stylized facts. None of the existing alternatives organizes well the bulk of the experimental evidence.

In this paper, we revisit the economic models of social learning by assuming that individuals update their beliefs in a non-Bayesian way. Individuals either overweigh or underweigh (in Bayesian terms) their private information relative to the public information revealed by the decisions of others. Each individual's updating rule is private information.

First, we consider a setting with perfectly rational individuals and where the distribution of updating rules is commonly known. We show that introducing heterogeneous updating rules in BHW's setting leads to equilibrium predictions which are more in line with the laboratory evidence on social learning than the original predictions. We also demonstrate that allowing for heterogeneous updating rules is equivalent to enlarging the support of private beliefs. In other words, a model of social learning with bounded private beliefs and sufficiently rich updating rules corresponds to a model of social learning with unbounded private beliefs. A straightforward implication is that heterogeneity in updating rules is efficiency-enhancing in most social learning environments. The reason is that society learns more if individuals are not Bayesian in their interpretation of others' behavior.

Second, we combine heterogeneous updating rules with the Analogy Based Expectation Equilibrium (ABEE) of Jehiel (2005). In such a setting, individuals only understand the relation between the aggregate distribution of decisions and the state of the world which leads them to update their beliefs according to a counting rule where the weight attached to each decision is determined by

the equilibrium frequencies of decisions. Unlike in rational social learning, heterogeneous updating rules do not lead to a substantial improvement of the societal welfare and there is always a non-negligible likelihood that players become extremely and wrongly confident about the state of the world.

Motivation and related literature

Psychologists and experimental economists have analyzed how subjects update probabilities in highly stylized situations to test whether respondents rely on Bayes' rule when provided with observations drawn from a sampling process, such as balls drawn from an urn (e.g., Tversky and Kahnemann, 1974; El-Gamal and Grether, 1995). Well-documented regularities show that, in addition to Bayes' rule, experimental participants employ certain heuristics to process probabilistic information. Compared to Bayesian individuals, some subjects are excessively conservative and do not adjust beliefs enough in light of new information (conservatism), others rely too heavily on recent information (base-rate neglect) or conduct some averaging between prior and conditional information. The most conclusive finding is that subjects exhibit considerable heterogeneity in the way they revise their expectations in light of the *same* information (for recent evidence on the heterogeneity in the updating process of beliefs see Delavande, 2008). Introducing heterogeneous updating rules in a model of social learning is clearly in line with this experimental evidence. Thus, a first and straightforward interpretation of the particular departure from Bayesian rationality that we consider is that individuals make mistakes in processing probabilistic information i.e. non-Bayesian updating rules reflect probability judgment biases. In this respect, our formal setting is related to the behavioral finance models which assume that overconfident investors overestimate the precision of their private information and predict that overconfidence leads to high trading volume (among others, Odean, 1998).³

Eyster and Rabin (2009) considers a social learning environment where individuals choose actions from a continuum and receive arbitrarily informative signals. Rational social learning predicts efficient information aggregation but the paper derives the possibility for an information cascade by assuming that individuals do not account for predecessors observing the same action history. In other words, individuals naively believe that each predecessor's action reflects solely that individual's private information. This form of inferential naivety is clearly related to the behavior of level-2 individuals who believe that others are level-1 in level- k thinking where level-0 individuals randomize (alternatively, inferential naivety is related to level-1's behavior with truthful level-0 play; see Crawford and Iriberri, 2007). Intuitively, heterogeneous updating rules provide an alternative way to capture the cognitive types of individuals who learn from observing others: The predictions of level- k thinking where Bayesian individuals assume mixtures of lower cognitive types (Strzalecki, 2009) are likely to be (almost) indistinguishable from the predictions of equilib-

³Bernardo and Welch (2001) and Kariv (2005) study a social learning model with a commonly know fraction of individuals who overweigh their private information relative to the public information revealed by the decisions of others. We generalize these theoretical models.

rium behavior where individuals update their beliefs in a non-Bayesian way. In the future, we hope to establish a precise link between the two formal frameworks.

2 A Basic Illustration

In this section, we illustrate the implications of heterogeneous updating rules for the process of social learning. We consider the same social learning environment as in BHW's specific model but we depart from the premise that all players update their private beliefs in a Bayesian way. We assume that only half of the players update their private beliefs in a Bayesian way while among the other half some overweigh their private information either weakly or strongly and the remaining players underweigh their private information. First, we provide a standard extension of BHW's specific model where the distribution of updating rules, the information structure, the payoffs and the perfect rationality of players are assumed to be commonly known. Second, along the lines of Guarino and Jehiel (2009), we consider the payoff-relevant reasoning extension of BHW's specific model where players need not be aware of the distribution of updating rules, the information structure, the payoffs and the rationality of other players. In this second extension, players only understand the relation between the aggregate distribution of decisions and the state of the world which leads them to update their beliefs according to a counting rule where the weight attached to each decision is determined by the equilibrium frequencies of decisions. Both extensions predict dynamics of beliefs and decisions which are better supported by the experimental evidence on social learning than BHW's predictions. Additionally, in both extensions, the presence of players who overweigh their private information improves the aggregation of information and is therefore efficiency-enhancing in large populations.

2.1 Rational Social Learning with Heterogeneous Updating Rules

Consider a setting where players face similar investment decisions under uncertainty and have private but imperfect information about the payoff of the investment. Players decide in sequence whether to invest and each player observes the decisions of all those ahead of her but not their private information. Payoffs from investing and rejecting are the same for all players. The investment payoff is denoted by the random variable $\tilde{\theta}$ with possible realizations 1 and (-1) which are equally likely, and θ , the realization of $\tilde{\theta}$, denotes the true value of the investment payoff. The payoff to rejecting is zero. Each player has private information in the form of a private signal which is the realization of a random variable whose distribution depends on true value of the investment payoff. Concretely, player i , $i = 1, 2, \dots$, observes a private signal (either high, H , or low, L) about θ which is the realization of the random variable \tilde{s}_i . If $\theta = 1$ then the probability that the signal is H is equal to $1 > q > 1/2$ and that the signal is L is $1 - q$. Similarly, if $\theta = (-1)$ then the signal realization is L with probability q (H with probability $1 - q$). Players' signals are independent conditional on the true value of the investment payoff and players aim at maximizing their expected payoffs. We assume that the information structure ($\tilde{\theta}$ and \tilde{s}_i) and the payoff structure are common knowledge. This social learning environment is isomorphic to the one considered by BHW in their specific model (see Bikhchandani, Hirshleifer, and Welch, 1992, section II.A).

Each player forms a belief, a probability estimate, about whether investing or rejecting is superior based on her private signal and on what she sees her predecessors do. She then makes her own investment decision. Concretely, player i observes her private signal s_i and the history h_i which consists of the investment decisions of all those ahead of her. $Pr(\tilde{\theta} = 1 | s_i, h_i)$ is player i 's belief that the true value of the investment payoff is one and $Pr(\tilde{\theta} = (-1) | s_i, h_i) = 1 - Pr(\tilde{\theta} = 1 | s_i, h_i)$. Since the expected payoff of investment is given by $E[\tilde{\theta} | s_i, h_i] = 2Pr(\tilde{\theta} = 1 | s_i, h_i) - 1$, player i invests if $Pr(\tilde{\theta} = 1 | s_i, h_i) > 1/2$. Said differently, player i invests if her likelihood ratio $\lambda(s_i, h_i) = Pr(\tilde{\theta} = 1 | s_i, h_i) / Pr(\tilde{\theta} = (-1) | s_i, h_i)$ is strictly greater than 1. Player i 's likelihood ratio is given by

$$\lambda(s_i, h_i) = \left(\frac{Pr(\tilde{s}_i = s_i | \tilde{\theta} = 1)}{Pr(\tilde{s}_i = s_i | \tilde{\theta} = (-1))} \right)^{\beta_i} * \frac{Pr(\tilde{h}_i = h_i | \tilde{\theta} = 1)}{Pr(\tilde{h}_i = h_i | \tilde{\theta} = (-1))},$$

where $\beta_i \in \mathbb{R}$ is the relative weight she assigns to her private information relative to the public information.

In BHW's specific model, players form their beliefs according to Bayes' rule i.e. $\beta_i = 1$ for all $i \in \mathbb{N}$ and this is commonly known. Denoting the difference between the number of investments and the number of rejections by I , players' optimal decision rule is characterized as follows: If $I = 0$ then the player invests if her private signal is H and she rejects if her private signal is L ; If $I = 1$ then the player invests if her private signal is H and she tosses a fair coin if her private signal is L ; If $I > 1$ then the player invests regardless of her private signal; The decisions for $I = -1$ and $I < -1$ are symmetric. Though the net number of investments evolves randomly, it will quickly reach either the amount of $+2$ and trigger an *investment (information) cascade* where all remaining players invest or the amount of -2 and trigger a *rejection (information) cascade* where all remaining players refrain from investing. Since decisions are uninformative once an information cascade has started, the informativeness of a cascade does not rise with the number of similar decisions. Thus, a small bulk of evidence causes the vast majority of players to either invest or reject, which might be the wrong decision. But the fallibility of information cascades causes them to be fragile: Assume, for example, than an investment cascade has started and that player i who decides late in the sequence observes two conditionally independent draws of the random variable \tilde{s}_i . If player i observes two L signals then she rejects.

Below, we extend BHW's specific model by assuming that half of the players are *Bayesians* while the remaining half is equally distributed among *conformists* ($\beta_i = 0$), *weak* ($\beta_i = 2$) and *strong overweighers* ($\beta_i = 3$) (i.e. conformists, weak and strong overweighers constitute $1/6$ of the population each). Relative weights are assumed to be private information while their distribution is commonly known. Finally, we make the assumption that players of type t , $t \in \{0, 1, 2, 3\}$, form their likelihood ratio according to $\beta_i = t + \epsilon$ for some small $\epsilon > 0$ in order to avoid ties.⁴

Dynamics of Investment Decisions

Anna, the first player, observes only a private signal. Whatever her type, Anna follows her private signal: If she observes H then she invests, if she observes L then she rejects.

⁴This assumption is equivalent to assuming that in case of a tie players act upon their private information.

Bob, the second player, as well as all other players can figure out Anna's private signal from observing her investment decision. Bob's information set consists of two private signals, his own private signal and the one he can infer from Anna's investment decision. From an objective point of view, both signals have the same informational value since we assume that all signals have the same precision (q). However, from Bob's subjective point of view, his own signal is private information whereas Anna's inferred signal is public information. So, Bob assigns different informational values to the two signals except in the case where Bob is Bayesian.

Assume that Bob's private signal confirms Anna's decision. In other words, either Bob's signal is H and Anna invested or Bob's signal is L and Anna rejected. Then, whatever his type, Bob follows his private signal or equivalently imitates Anna's investment decision.

Assume that Bob's signal contradicts Anna's decision i.e. either Bob's signal is L and Anna invested or Bob's signal is H and Anna rejected. Then Bob imitates Anna's decision if he is a conformist ($\beta_2 = 0$) and he follows his private signal otherwise.

In summary, if Bob is a conformist, which occurs with probability $1/6$, then he imitates Anna's investment decision regardless of his private signal. With probability $5/6$, Bob is not a conformist and therefore he follows his private signal.

Claire, the third player, faces one of three scenarios.

Assume that both Anna and Bob invested. Since Bob herds on Anna's decision whenever he is a conformist, Claire as well as all other players learn less from Bob's decision than from Anna's decision. Still, Bob follows his private signal with probability $5/6$ which implies that $\lambda(\emptyset, (1, 1)) > q/(1 - q)$. Therefore, if Claire is a conformist or a Bayesian then she invests regardless of her private signal. If Claire is an overweigher then she follows her private signal. In summary, Claire invests (i.e. she imitates Anna and Bob) regardless of her private signal with probability $2/3$ and she follows her private signal with probability $1/3$.

Assume that both Anna and Bob rejected. This scenario is symmetric to the first one. If Claire is a conformist or a Bayesian then she rejects (i.e. she imitates Anna and Bob) regardless of her private signal (which happens in $2/3$ of the cases). Otherwise, Claire follows her private signal (in $1/3$ of the cases).

Assume that Anna and Bob made opposite investment decisions. Consequently, Bob is not a conformist and Claire as well as all other players can identify the private signals of the first two players. These private signals contradict each other and the public likelihood ratio equals the prior. So, Claire is in the same position as Anna and she follows her private signal.

We end up by discussing the investment strategy of David, the fourth player, and we sketch the dynamics of investment decisions for the remaining players in the sequence.

Assume that Anna and Bob made opposite investment decisions. Since Claire is in the same position as Anna, David is in the same position as Bob, Emma is in the same position as Claire, and so forth.

Assume that Anna, Bob and Claire made the same investment decision. Less information can be inferred from Claire's decision than from Bob's decision since Claire herds not only if she is a conformist (like Bob) but also if she is a Bayesian (Bob imitates in $1/6$ of the cases whereas Claire

imitates in 2/3 of the cases). In fact, Claire's decision reveals so little information that David imitates his predecessors if and only if he is a conformist or a Bayesian (his investment strategy after having observed three identical investment decisions is the same as Claire's investment strategy after she observed two identical investment decisions). If David is an overweigher then he follows his private signal. The first weak overweigher who imitates after a sequence of identical investment decisions is Emma, the fifth player (assuming that q is not too large).⁵ Clearly, a strong overweigher will only imitate a long sequence of identical investment decisions since little information can be inferred from any herding decision.

Assume that Anna and Bob made the same investment decision but that Claire made a different one. Claire's decision reveals that she followed her private signal and hence is not a conformist or a Bayesian. Consequently, David infers Claire's private signal from her investment decision. This inferred private signal cancels out with the private signal inferred from Anna's decision and David is left with the information he inferred from Bob's decision. As already mentioned, Bob's decision reveals less information than a private signal which implies that David follows his private signal provided he is not a conformist. If David's decision differs from Anna and Bob's decisions then his decision conveys more information than Bob's decision and the rest of the players believe more in the state which is line with Claire and David's investment decisions.

The Emergence of Information Cascades

The further away from 1 the public likelihood ratio $\lambda(\emptyset, h_i)$ the more likely a player ignores her private signal and decides in accordance with this public information. Once the public likelihood ratio exceeds the threshold $q^3/(1-q)^3$ (falls below $(1-q)^3/q^3$), all subsequent players, whatever their type, invest (reject) regardless of their private signal.⁶ An investment cascade (rejection cascade) starts and lasts forever but, unlike in BHW's specific model, the difference between the number of investments and rejections does not suffice to characterize the information which can be inferred from a given history. The amount of information inferred from a history of investment decisions is strongly path-dependent.

Still, like in BHW's specific model, an information cascade arises in finite time with probability one. Indeed, conditional on the realized state being (-1) , the public likelihood ratio constitutes a Markov martingale which by the Martingale Convergence Theorem (MCT) converges to a limiting random variable. Any value in the support of this random variable must be invariant to updating following investment decisions. This however is possible only if any further investment decision does not convey additional information i.e. if all types decide regardless of their private signal. Thus, in the limit, an information cascade must arise almost surely and the information cascade arises in finite time as an infinite number of deviations from a herd would prohibit the convergence.

⁵If q is large then a private signal reveals so much information that it outweighs the informative value in David's decision which like for Claire is small as the probability that he ignores his private signal is rather large.

⁶Conformists invest (reject) regardless of their private signal once the public likelihood ratio lies strictly above (below) 1. Bayesians invest (reject) regardless of their private signal once the public likelihood ratio lies strictly above $q/(1-q)$ (strictly below $(1-q)/q$). Weak (strong) overweighers invest regardless of their private signal once the public likelihood ratio lies strictly above $q^2/(1-q)^2$ ($q^3/(1-q)^3$) and they reject regardless of their private signal once the public likelihood ratio lies strictly below $(1-q)^2/q^2$ ($(1-q)^3/q^3$).

Finally, with strictly positive probability, an information cascade on the less profitable investment decision arises (see proof 2.1.1 in the Appendix).

The Efficiency of the Social Learning Process

Introducing overweighers in BHW's specific model leads to more extreme public beliefs and to better information aggregation since longer sequences of identical investment decisions are needed for the emergence of information cascades. Consequently, the probability that a cascade starts on the less profitable investment decision (wrong information cascade) is smaller than in BHW's specific model (see proof 2.1.2 in the Appendix). Figure 1 shows the lower (upper) bound of the probability that a correct (incorrect) cascade arises in comparison to the probabilities in BHW's specific model. Heterogeneous updating rules are therefore efficiency-enhancing in large populations (the requirement of a large population is due to the suboptimal investment decisions of overweighers who follow their private signal when Bayes' rule predicts to imitate).




Figure 1: Lower (upper) bound upon the probability that a correct (incorrect) cascade arises compared to probabilities in BHW.

The Fragility of Information Cascades

Information cascades are less fragile compared to BHW's specific model where a player needs at least 2 contradictory private signals to decide against the herd. Indeed, in order to break an information cascade, conformists need infinitely precise information, Bayesians need at least 4 contradictory private signals, and overweighers need at least 2 contradictory private signals.

Summary

We have analyzed a straightforward extension of BHW's specific model where players rely on heterogeneous rules to update their beliefs, some overweighing their private information while

others underweigh it, and where the distribution of updating rules is assumed to be commonly known. The dynamics of investment decisions are more complex than in the original model. We believe that this loss in tractability is largely compensated by the gain in the accuracy of the predictions. Indeed, the predictions of BHW's specific model with heterogeneous updating rules nicely organize the experimental evidence on social learning. Compared to the original model, information cascades start later and they are less fragile since beliefs become more extreme. These predictions correspond exactly to the regularities observed in the laboratory studies on social learning. Finally, we have established that the presence of players who overweigh their private information improves the aggregation of information and is therefore efficiency-enhancing in large populations.

2.2 Social Learning with Coarse Inference and Heterogeneous Updating Rules

Though the standard equilibrium approach makes predictions well in line with the experimental evidence, there are numerous social learning interactions in the field where players are unlikely to know the distribution of updating rules as well as the information structure and payoffs of other players. Additionally, rational social learning requires an extremely high degree of cognitive sophistication on players' part which is commonly known, an assumption even less likely to be satisfied if some players update their private beliefs in a non-Bayesian way.

In this second extension of BHW's specific model, we combine heterogeneous updating rules with the analogy based expectation equilibrium. In line with Guarino and Jehiel (2009) (henceforth GJ), we assume that there are two analogy classes, one for each state of the world. In such a generalized social learning model, players understand only how the state of the world affects the aggregate distribution of decisions but not how it affects the sequence of decisions as a function of the history of observations. Players need not be aware of the distribution of updating rules, the information structure, the payoffs and the rationality of other players. GJ shows that, in BHW's specific model, a unique ABEE exists such that the first player follows her private signal while the remaining players in the sequence imitate her decision. Therefore with probability $(1 - q)$ all players take the less profitable decision and beliefs become completely wrong.

We now illustrate the properties of the ABEE with heterogeneous updating rules in BHW's specific model. The existence and unicity of the ABEE in the social learning environment considered by BHW's specific model with heterogeneous updating rules has been established with the help of simulations.⁷ Figure 2 shows the equilibrium frequencies of correct choices as a function of the signal's accuracy q for sequence lengths of 6 and 10 players. Equilibrium frequencies of correct choices are only slightly larger than the signal's accuracy q . Consequently, in equilibrium, players believe that each observed decision reflects that player's private information where the underlying

⁷In a simulation we assume an initial set of action distributions (one for each possible state of the world) which players use to derive the weight attached to each choice. Assuming that players update beliefs according to a counting rule using these weights and combine them with private beliefs according to the given distribution of private information weights, we then calculate analytically for all possible choice sequences of a fixed length their likelihood in either state. This yields a new set of action distributions that we use as an initial condition for the next run. We iterate this step until the assumed set of action distributions coincides with the resulting set. The procedure converged quickly and independently of the set of initial distributions chosen.

(fictitious) private signal is of a slightly larger accuracy than q .




Figure 2: Frequencies of correct choices in ABEE for sequence lengths of 6 (blue line) and 10 players (red line).

Dynamics of investment decisions are straightforward to characterize: If the observed number of investments equals the observed number of rejections then the player follows her private signal i.e. she invests if her private signal is H and she rejects if her private signal is L . If the difference between the number of investments and the number of rejections equals 1 (respectively -1) then conformists and Bayesians invest (respectively reject) regardless of their private signal whereas overweighters follow their private signal. If the difference between the number of investments and the number of rejections equals 2 (respectively -2) then conformists, Bayesians and weak overweighters invest (respectively reject) regardless of their private signal whereas strong overweighters follow their private signal. Finally, if the difference between the number of investments and the number of rejections equals 3 or more (respectively -3 or less) then all players invest (respectively reject) regardless of their private signal.

Like in BHW's specific model and its first extension, information cascades clearly emerge in this second extension. Compared to rational social learning, a much smaller sequence of identical investment decisions is needed to trigger an information cascade (compared to BHW's specific model, a slightly longer sequence is needed). The reason is that players apply a simple counting rule where the informational content of each observed decision is constant and independent of the decisions which precede it. Consequently, information cascades become infinitely robust i.e. no player, whatever his type and the accuracy of her private signal, will break a long herd. Figure 3 illustrates the evolution of the public belief with the depth of a herd (length of identical investment decisions).

As in rational social learning, the presence of players who overweigh their private information improves the aggregation of information and is therefore efficiency-enhancing in large populations (remember that all players imitate the first player's decision in GJ). However, unlike in rational




Figure 3: Public belief as a function of the depth of a herd.

social learning, heterogeneous updating rules do not lead to a substantial improvement of the societal welfare and there is always a non-negligible likelihood that players become extremely and wrongly confident about the state of the world.

3 A General Social Learning Environment

In this section, we consider a general social learning environment where players have bounded private beliefs and the set of updating rules is dense. First, we show that rational social learning with unbounded private weights is equivalent to rational social learning in an environment where Bayesian players have unbounded private beliefs. This equivalence result enables us to easily characterize the predicted dynamics of beliefs and decisions and to establish the genericity of the societal welfare improvement due to the presence of overweighers. Second, we study the payoff-relevant reasoning model with a dense set of updating rules and prove the existence of an analogy based expectation equilibrium. Again, the positive influence of overweighers on societal welfare in this second setting is much weaker.

3.1 Rational Social Learning with Rich Updating Rules

We first generalize the social learning environment introduced in the previous section, then we characterize the equilibrium behavior of fully rational players with rich updating rules, and finally we prove our equivalence result.

The state of the world is given by the random variable $\tilde{\theta}$ on $\Theta = \{(-1), 1\}$. It is distributed according to the flat prior $Pr(\tilde{\theta} = 1) = \frac{1}{2}$. Players $i = 1, 2, \dots$ sequentially choose from the set of actions $A = \{0, 1\}$. Payoffs are given by the mapping $u : A \times \Theta \rightarrow \mathbb{R}$ where $u(1, \theta) = \theta$ and $u(0, \theta) = 0$ for each $\theta \in \Theta$. We often refer to action $a = 1$ as “invest” and to action $a = 0$ as “reject”. Each player $i \in \mathbb{N}$ holds a private belief which is the realization of the random variable $b(\tilde{s}_i, \emptyset)$ on $[0, 1]$. Given the realization θ of $\tilde{\theta}$, $(b(\tilde{s}_i, \emptyset))_{i \in \mathbb{N}}$ is an i.i.d. stochastic process on $[0, 1]$ distributed according to the cumulative distribution function G_θ . $G_{(-1)}$ and G_1 satisfy the usual assumption i.e. they are absolutely continuous to one another and their Radon–Nikodym derivative satisfies $\frac{dG_{(-1)}}{dG_1}(b) = \frac{1-b}{b}$.

We assume that the convex hull of their common support is given by $\text{supp}(G_1) = \text{supp}(G_{(-1)}) = [\underline{b}, \bar{b}]$ where $\underline{b} > 0$ and $\bar{b} < 1$ meaning that private beliefs are bounded. Finally, each player i observes the history $h_i = (a_1, \dots, a_{i-1})$ of actions of all preceding players where $h_i \in H_i = A^{i-1}$.

Players aim at maximizing their expected utility $U(a) = E[u(a, \tilde{\theta})]$. Let $b(s_i, h_i)$ denote player i 's posterior belief that the true state is 1 given her private belief and the observed history of her predecessors' actions: $b(s_i, h_i) = \Pr(\tilde{\theta} = 1 | b(\tilde{s}_i, \emptyset) = b(s_i, \emptyset), \tilde{h}_i = h_i)$. The maximization of her expected utility leads player i to choose $a = 1$ if $b(s_i, h_i) > \frac{1}{2}$, $a = 0$ if $b(s_i, h_i) < \frac{1}{2}$ and to flip a fair coin if $b(s_i, h_i) = \frac{1}{2}$. In terms of the likelihood ratio $\lambda(s_i, h_i) = \frac{b(s_i, h_i)}{1-b(s_i, h_i)}$, the relevant posterior threshold equals one.

Each player $i \in \mathbb{N}$ updates her belief according to

$$\lambda(s_i, h_i) = (\lambda(s_i, \emptyset))^{\beta_i} \lambda(\emptyset, h_i) \quad (1)$$

where β_i is a player specific parameter. If $\beta_i = 1$ then this updating rule is equivalent to Bayesian updating. If $\beta_i < 1$ then player i puts too much weight on the public information relative to Bayes' rule. If $\beta_i > 1$ then player i overweights her private information. We assume that each player's weighting factor β_i is private information and thus unknown to other players. Formally, each player's weighting factor is the realization of a random variable $\tilde{\beta}_i$ where the $(\tilde{\beta}_i)_{i \in \mathbb{N}}$ form an i.i.d. stochastic process on $[0, \infty)$ distributed according to the cumulative distribution function W . Finally, the sequentially ordered set of players \mathbb{N} , the state space Θ together with the flat prior, the action set A , the utility function u , the private belief distributions $(G_{(-1)}, G_1)$ and the cumulative distribution function W are commonly known among the players.

An equilibrium of the *social learning game with rich updating rules* $\langle \mathbb{N}, \Theta, A, u, (G_{(-1)}, G_1), W \rangle$ is defined as follows.

Definition 3.1

An equilibrium of $\langle \mathbb{N}, \Theta, A, u, (G_{(-1)}, G_1), W \rangle$ is given by a behavioral strategy profile $(\sigma_i)_{i \in \mathbb{N}}$ where $\sigma_i : [0, 1] \times H_i \rightarrow \Delta A$ with $\sigma_i(a | b, h_i) = \Pr(\tilde{a}_i = a | b(\tilde{s}_i, \emptyset) = b, \tilde{h}_i = h_i)$, and a system of beliefs $(b(s_i, h_i))_{i \in \mathbb{N}}$ such that

(i) Beliefs are updated according to (1) where $\lambda(s_i, \emptyset) = \frac{b(s_i, \emptyset)}{1-b(s_i, \emptyset)}$ and

$$\lambda(\emptyset, h_i) = \prod_{j < i} \frac{\Pr(\tilde{a}_j = h_i(j) | \tilde{h}_j = (h_i(1), \dots, h_i(j-1)), \tilde{\theta} = 1)}{\Pr(\tilde{a}_j = h_i(j) | \tilde{h}_j = (h_i(1), \dots, h_i(j-1)), \tilde{\theta} = (-1))}$$

(ii) $\sigma_i(a = 1 | b(s_i, \emptyset), h_i) = \begin{cases} 1 & \text{if } b(s_i, h_i) > \frac{1}{2} \\ 0 & \text{if } b(s_i, h_i) < \frac{1}{2} \end{cases}$ (sequential rationality).

3.1.1 The Individual Decision Process and the Process of Social Learning

We first analyze the decision process of a single player and we discuss what other players might learn from her decision. Fix player $i \in \mathbb{N}$ with decision weight β_i . Using the updating rule (1) and the posterior LR threshold, the following lemma describes the player's strategy in terms of her realized private belief.

Lemma 3.1

For player $i \in \mathbb{N}$ let $\lambda_i = \lambda(\emptyset, h_i)$ be the realized public likelihood ratio and let β_i be the privately known weight he puts on the private belief. Define

$$t(\lambda_i, \beta_i) = \frac{1}{1 + (\lambda_i)^{\frac{1}{\beta_i}}}.$$

The player's strategy σ_i is given by

$$\sigma_i(a = 1 | s_i, h_i) = \begin{cases} 1 & \text{if } b(s_i, \emptyset) > t(\lambda(\emptyset, h_i), \beta_i) \\ \frac{1}{2} & \text{if } b(s_i, \emptyset) = t(\lambda(\emptyset, h_i), \beta_i) \\ 0 & \text{if } b(s_i, \emptyset) < t(\lambda(\emptyset, h_i), \beta_i) \end{cases}$$

Further $t : [0, \infty) \times [0, \infty) \rightarrow [0, 1]$ satisfies

- (i) $t(\lambda, 1) = 1/(1 + \lambda)$, $\lim_{\beta \rightarrow \infty} t(\lambda, \beta) = \frac{1}{2}$, $\lim_{\beta \searrow 0} t(\lambda, \beta) = \begin{cases} 1 & \text{if } \lambda < 1 \\ 0 & \text{if } \lambda > 1 \end{cases}$;
- (ii) $\lim_{\lambda \rightarrow 0} t(\lambda, \beta) = 1$, $\lim_{\lambda \rightarrow \infty} t(\lambda, \beta) = 0$, $t(1, \beta) = \frac{1}{2}$;
- (iii) $\frac{\partial t(\lambda, \beta)}{\partial \beta} = \begin{cases} > 0 & \text{if } \lambda > 1 \\ < 0 & \text{if } \lambda < 1 \end{cases}$;
- (iv) $\frac{\partial t(\lambda, \beta)}{\partial \lambda} < 0$.

Proof: See the Appendix. ■

Figure 4 shows the graph of the private belief thresholds as a function of the realized public belief for different values of the decision weight β_i . For $\beta_i > 1$, private belief thresholds regress towards $\frac{1}{2}$, the “informationally optimal private belief threshold” (Smith and Sørensen, 2008a, p.14). Intuitively, the player relies more on her private belief thus conveying more information to others. On the other hand for $0 < \beta_i < 1$ the player relies so much on the public information that only rather extreme private beliefs may overturn this. The smaller β_i , the more pronounced this behavior becomes. As $\beta_i = 0$, players rely only on the public information except when it is indecisive ($b(\emptyset, h_i) = \frac{1}{2}$).

We now turn to the process of social learning. While the history and thereby the associated public likelihood ratio $\lambda_i = \lambda(\emptyset, h_i)$ are publicly known, the realized private belief is not. It can however be inferred partially from the player's decision. More precisely, $a_i = 0$ conveys the message that $b(s_i, \emptyset) < t(\lambda_i, \beta_i)$ if the player is of type $\beta_i > 0$. This event has different likelihoods under different realizations of $\tilde{\theta}$ as

$$\begin{aligned} \Pr(\tilde{a}_i = 0 | \tilde{h}_i = h_i, \tilde{\beta}_i = \beta_i, \tilde{\theta} = \theta) &= \Pr(b(\tilde{s}_i, \emptyset) < t(\lambda_i, \beta_i) | \tilde{\theta} = \theta) \\ &= G_\theta(t(\lambda_i, \beta_i)). \end{aligned}$$

If weights were public information then players could perfectly infer these likelihoods by exactly computing the threshold function and updating the public belief accordingly. Given that β_i is private information, players have to form an average of these probabilities over the possible




Figure 4: Private belief thresholds as a function of the realized public belief $b(\emptyset, h_i)$ for weights $\beta_i = 0$ (solid grey line), $\beta_i = \frac{1}{2}$ (dashed grey line), $\beta_i = 1$ (solid black line), $\beta_i = 2$ (dashed black line) and $\beta_i \rightarrow \infty$ (dotted black line).

types using the cumulative distribution function W . Formally

$$Pr(\tilde{a}_i = a | \tilde{h}_i = h_i, \tilde{\theta} = \theta) = \int_0^\infty Pr(\tilde{a}_i = a | \tilde{h}_i = h_i, \tilde{\theta} = \theta, \tilde{\beta}_i = \beta) dW(\beta).$$

Consequently, social learning, i.e. the updating of the public belief, takes place according to

$$\lambda_{i+1} = \lambda_i * \begin{cases} \frac{\int_0^\infty G_1(t(\lambda_i, \beta)) dW(\beta)}{\int_0^\infty G_{(-1)}(t(\lambda_i, \beta)) dW(\beta)} & \text{if } a_i = 0, \\ \frac{1 - \int_0^\infty G_1(t(\lambda_i, \beta)) dW(\beta)}{1 - \int_0^\infty G_{(-1)}(t(\lambda_i, \beta)) dW(\beta)} & \text{if } a_i = 1. \end{cases} \quad (2)$$

3.1.2 An Equivalence Result

Let us go back to the social learning environment considered in BHW's specific model. Compare the decision process of player i with $\beta_i = 2$ to the decision process of player j with $\beta_j = 1$ but whose private signal has twice the precision of player i 's private signal. Assume that both players face the same realized public likelihood ratio $\lambda_i = \lambda_j$ and that both players are endowed with a high private signal $s_i = s_j = H$. Each player's posterior belief is given by $\lambda(s_i, h_i) = \frac{q^2}{(1-q)^2} * \lambda_i$. In other words, the posterior belief of a player who overweights her private information is identical to the posterior belief of a Bayesian player who is endowed with more accurate private information. Obviously, both players make the same decision. According to the following lemma, this equivalence property holds in general.

Lemma 3.2

Let the private belief distributions be given by $G_{(-1)}$ and G_1 with support $[b, \bar{b}] \subseteq [0, 1]$. The equilibrium of $\langle \mathbb{N}, \Theta, A, u, (G_{(-1)}, G_1), W \rangle$ is identical to the equilibrium of $\langle \mathbb{N}, \Theta, A, u, (G'_{(-1)}, G'_1), W' \rangle$ where

$$G'_\theta(b) = \int_0^\infty G_\theta \left(\frac{b^{\frac{1}{\beta}}}{b^{\frac{1}{\beta}} + (1-b)^{\frac{1}{\beta}}} \right) dW(\beta),$$

and W' is the Dirac measure concentrated at one.

Proof: See the Appendix. ■

Lemma 3.2 establishes that rational social learning with given private belief distributions and where players rely on heterogeneous updating rules is equivalent to rational social learning with modified private belief distributions and Bayesian players. Unlike in the standard social learning games studied in the literature (Smith and Sørensen, 2000), the modified private belief distributions are not self-fulfilling. A player that uses the different likelihoods of her realized private belief under different states to update her prior will in general not go back to her private belief. Formally, $b = 1/(1 + dG'_{(-1)}(b)/dG'_1(b))$ meaning that the well-known “no introspection condition” does not hold for arbitrary $G_{(-1)}$, G_1 and W . However, the condition does not constitute a necessity for the players to learn from observing others. From (2), $a_i = 0$ induces a change of the public belief by

$$\lambda_{i+1} = \lambda_i * \frac{G'_1\left(\frac{1}{1+\lambda_i}\right)}{G'_{(-1)}\left(\frac{1}{1+\lambda_i}\right)}.$$

Consequently, social learning takes place if $\frac{G'_1(b)}{G'_{(-1)}(b)} \neq 1$ for some beliefs b . The next lemma partially characterizes the private belief distribution with regard to social learning.

Lemma 3.3

Let $[\underline{\beta}, \bar{\beta}] \subseteq [0, \infty)$ denote the convex hull of the support of W . The convex hull of the (common) support of $G'_{(-1)}$ and G'_1 is given by $[\underline{b}', \bar{b}']$ where

$$\underline{b}' = \frac{\underline{b}^{\bar{\beta}}}{\underline{b}^{\bar{\beta}} + (1-\underline{b})^{\bar{\beta}}} \quad \text{and} \quad \bar{b}' = \frac{\bar{b}^{\underline{\beta}}}{\bar{b}^{\underline{\beta}} + (1-\bar{b})^{\underline{\beta}}}.$$

Furthermore it holds $G'_1(b) < G'_{(-1)}(b)$ for each $b \in (\underline{b}', \bar{b}')$.

Proof: See the Appendix. ■

Given initial private belief distributions $G_{(-1)}$ and G_1 , any modified distributions are attainable through a dense set of updating rules W . We say that private information weights are **bounded** if $\bar{\beta} < \infty$. Otherwise, private information weights are **unbounded**. As stated in the following corollary, if private information weights are **unbounded** then the modified distributions $G'_{(-1)}$ and G'_1 are unbounded.

Corollary 3.4

In the social learning game $\langle \mathbb{N}, \Theta, A, u, (G'_{(-1)}, G'_1), W' \rangle$, $G'_{(-1)}$ and G'_1 are unbounded if and only if W is unbounded in the social learning game with rich updating rules $\langle \mathbb{N}, \Theta, A, u, (G_{(-1)}, G_1), W \rangle$.

3.1.3 Dynamics of Beliefs and Decisions, Learning and Societal Welfare

The following results are provided for the sake of completeness since they derive directly from Lemma 3.2.

Learning

Corollary 3.5

The learning process converges eventually. It is complete, if and only if private information weights are unbounded.

Lemma 3.6

With bounded private information weights, the larger $\bar{\beta}$, the more extreme (i.e. the farther from $\frac{1}{2}$) beliefs are in the limit.

Proof: From the equivalence result (Lemma 3.2), the larger $\bar{\beta}$, the larger the support of the modified private beliefs. The boundaries of the cascade regions satisfy $\underline{\lambda} = \frac{1-\bar{b}}{\bar{b}} < 1$ and $\bar{\lambda} = \frac{1-b}{\underline{b}} > 1$. As $\bar{\beta}$ increases the RHS of the former decreases while the RHS of the latter increases. ■

With bounded decision weights the learning process may converge to a limit far away from the truth. The larger $\bar{\beta}$ the farther away this limit. As mentioned in the previous section, information cascades still emerge in BHW's specific model with heterogeneous updating rules. However, if decision weights are unbounded then there is complete learning even with bounded private beliefs. In the standard social learning games with unbounded beliefs (e.g. Chamley, 2003) the truth is attained very slowly. Compared to the case where players can observe private information, it is in fact exponentially slower. In the model we discuss here a degenerate distribution putting all the probability mass on $\beta = 1$ resembles the standard model while as the distribution puts more and more probability mass on very large β , we attain the regime where private information is fully observable. Therefore one can expect that the more probability mass the distribution W puts on larger β s, the faster the truth is learned. Indeed, we can show that in the special case where W puts all its mass on a single decision weight $\hat{\beta}$ and with unbounded private beliefs, beliefs converge slower to the truth than $t^{-\hat{\beta}}$ to zero (the proof is available from the authors upon request). We are currently working on a formal characterization of the speed of learning for the general case in terms of the properties of the distribution W .

Uniform Behavior

Corollary 3.7

Uniform behavior eventually arises in finite time.

Lemma 3.8

Uniform behavior is error-prone and idiosyncratic if and only if private information weights are bounded. Moreover the chance of uniform behavior on the less profitable action vanishes as $\bar{\beta} \rightarrow \infty$.

Proof: Given Lemma 3.2, the result is a direct consequence of Theorems 1 and 3 of Smith and Sørensen (2000). With bounded decision weights, the resulting modified private belief distribution is bounded. Consequently uniform behavior on a less profitable action arises with strictly positive probability. Hence, it is error-prone. Furthermore, which limit is achieved will clearly depend on the first few decisions and thus the first few private belief realizations, it is path-dependent and therefore idiosyncratic. Finally, the probability of uniform behavior on the less profitable action satisfies $\pi < \frac{1}{\lambda}$ and the RHS is strictly decreasing in \bar{b}' which in turn increases in $\bar{\beta}$. ■

Efficiency

The appropriate concept is given by Radner's (1962) team equilibrium, with variable discount factors as in Smith and Sørensen (2008a). Formally, we study the value of the discounted sum of ex-ante expected utilities $E \left[(1 - \delta) \sum_{i=1}^{\infty} \delta^{i-1} u(a_i, \tilde{\theta}) \right]$ in equilibrium.

From the above results, we can deduce that the more weight W puts on large β , the more a player relies on her private belief. The benefit of such a behavior is that the player conveys more information to her successors. On the other hand, the player has the disadvantage of relying too heavily on less precise information once enough has been accumulated in the public belief. However, once the public belief is close to the truth, even players with a very large but finite β_i choose accordingly. We have shown above that as $\bar{\beta} \rightarrow \infty$, the chance of a cascade on a less profitable action vanishes. On the other hand Smith and Sørensen (2008a) show that in the team equilibrium players lean more against the public belief than in the standard equilibrium when deciding. From the properties of the threshold function (Lemma 3.1), exactly the same happens for $\beta > 1$.

In conclusion, if the distribution W puts sufficient probability mass on large decision weights $\beta > 1$ then the welfare is higher than in the equilibrium of the social learning game where W is a Dirac measure concentrated at one. Furthermore, the team equilibrium is attainable by a sequence $(\check{\beta}_i)_{i \in \mathbb{N}}$ of mappings $\check{\beta}_i : [0, 1]^2 \rightarrow [0, \infty)$ such that $\check{\beta}_i(b(s_i, \emptyset), b(\emptyset, h_i))$ denotes the weight player $i \in \mathbb{N}$ puts on her private information.

3.2 Social Learning with Coarse Inference and Rich Updating Rules

We now combine rich updating rules with the Analogy Based Expectation Equilibrium in the general social learning environment. As in the previous section, we follow GJ in assuming that players group decision nodes of others into analogy classes according to the payoff-relevant analogy partition i.e. conditional on the underlying state of the world. Accordingly (see Definition 1 in Guarino and Jehiel, 2009), players update their (public) beliefs after observing a predecessor's action according to $\lambda(\emptyset, (h_i, a_i)) = \lambda(\emptyset, h_i) * \frac{\bar{\sigma}(a_i | \tilde{\theta}=1)}{\bar{\sigma}(a_i | \tilde{\theta}=(-1))}$ given aggregate action frequencies $\bar{\sigma}(a | \theta)$ for each $(\theta, a) \in \Theta \times A$. Hence, the information value $\bar{\sigma}(a | 1)/\bar{\sigma}(a | (-1))$ of an action is fixed and

independent of previous decisions. The updating of the public belief consequently takes the form of a counting rule.

First, we establish the existence of an Analogy Based Expectation Equilibrium in our general social learning environment where players rely on rich updating rules.

Lemma 3.9

For any finite sequence of players, there exists an analogy-based expectations equilibrium $(\bar{\sigma}^(1 \mid (-1)), \bar{\sigma}^*(1 \mid 1)) \in (0, 1)^2$ satisfying $\bar{\sigma}^*(1 \mid 1) > \bar{\sigma}^*(1 \mid (-1))$.*

Proof: See the Appendix. ■

We now discuss the predicted dynamics of beliefs and decisions. In equilibrium, given the realized private belief b_i and the realized history h_i , player i 's belief is given by

$$\lambda(b_i, h_i) = \left(\frac{b_i}{1 - b_i} \right)^{\beta_i} \left(\frac{\pi_1}{\pi_{(-1)}} \right)^{\Sigma_1(h_i)} \left(\frac{1 - \pi_1}{1 - \pi_{(-1)}} \right)^{\Sigma_0(h_i)}$$

where $\Sigma_a(h_i)$ denotes the number of times action $a \in \{0, 1\}$ occurred within history h_i , and π_θ denotes the average frequency of correct choices in state of the world $\theta \in \{(-1), 1\}$. For a given history, there are two opposing forces which influence player i 's belief. On the one hand, the larger $\pi_1/\pi_{(-1)} > 1$ the more information is attached to each action and the less identical actions are needed to herd on the history. On the other hand, the larger β_i the more influential private information and the more identical actions are needed to herd on the history. Having identified these two forces enables us to discuss how societal welfare evolves when W puts more probability mass on strong overweighters. Ceteris paribus, an increase in the fraction of strong overweighters will have two effects: (i) Early players are more likely to rely on their private information which leads to more information being aggregated and later players being more likely to choose correctly. Therefore, we expect the average frequency of correct choices i.e. the informational value of each action to increase; (ii) If the informational value of each action increases then beliefs rise faster in herds and wrong herds are more likely to persist. The first effect improves the societal welfare whereas the second effect is detrimental for the societal welfare. Overall, an increase in the fraction of strong overweighters may lead to choices being wrong with a higher probability. In fact, even in infinite sequences of players and for most cumulative distribution functions W , the average frequency of correct choices in equilibrium is bounded away from 1 as stated in the following lemma.

Lemma 3.10

$\pi_1 < 1$ and $\pi_{(-1)} > 0$ if

- (i) Private information weights are bounded ($\bar{\beta} < \infty$);
- (ii) Private information weights are unbounded ($\bar{\beta} = \infty$) and (a) have finite mean and variance, or (b) satisfy $1 - W(x) \approx x^{-(1+\alpha)}$ for some $0 < \alpha < \infty$.

With strictly positive probability a wrong cascade arises.

Proof: See the Appendix. ■

4 Conclusion

In this paper, we have revisited the economic models of social learning by assuming that individuals update their beliefs in a non-Bayesian way. We show that the introduction of heterogeneous updating rules in social learning improves drastically the predictive power of equilibrium predictions. Additionally, we provide a more satisfactory interpretation of unbounded beliefs by establishing that a model of social learning with bounded private beliefs and sufficiently rich updating rules corresponds to a model of social learning with unbounded private beliefs. This link also demonstrates that heterogeneity in updating rules is efficiency-enhancing in most social learning environments.

Future work will consider heterogeneous updating rules in social learning settings with continuous actions, flexible prices, or endogenous sequencing.

References

- ANDERSON, L., AND C. HOLT (1997): "Information cascades in the laboratory," *American Economic Review*, 87, 847–862.
- BANERJEE, A. V. (1992): "A simple model of herd behavior," *The Quarterly Journal of Economics*, 107, 797–817.
- BERNARDO, A., AND I. WELCH (2001): "On the Evolution of Overconfidence and Entrepreneurs," *Journal of Economics & Management Strategy*, 10, 301–330.
- 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.
- CHAMLEY, C. P. (2003): "Cascades and Slow Social Learning," Working paper, Boston University.
- (2004): *Rational herds: Economic models of social learning*. Cambridge, Cambridge University Press.
- CRAWFORD, V. P., AND N. IRIBERRI (2007): "Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," *Econometrica*, 75, 1721–1770.
- DELAVANDE, A. (2008): "Measuring Revisions to Subjective Expectations," *Journal of Risk and Uncertainty*, 36, 43–82.
- EL-GAMAL, M., AND D. GRETHER (1995): "Are People Bayesian? Uncovering Behavioral Strategies," *Journal of the American Statistical Association*, 90, 1137–1145.
- EYSTER, E., AND M. RABIN (2009): "Rational and Naive Herding," Cepr discussion paper 7351. london, centre for economic policy research.

- GOEREE, J., T. PALFREY, B. ROGERS, AND R. MCKELVEY (2007): "Self-correcting information cascades," *Review of Economic Studies*, 74, 733–762.
- GUARINO, A., AND P. JEHIEL (2009): "Social Learning with Coarse Inference," Working paper.
- JEHIEL, P. (2005): "Analogy-Based Expectation Equilibrium," *Journal of Economic Theory*, 123, 81–104.
- KARIV, S. (2005): "Overconfidence and Informational Cascades," mimeo, UC Berkeley.
- KÜBLER, D., AND G. WEIZSÄCKER (2004): "Limited depth of reasoning and failure of cascade formation in the laboratory," *Review of Economic Studies*, 71, 425–441.
- ODEAN, T. (1998): "Volume, Volatility, Price, and Profit When All Traders Are Above Average," *Journal of Finance*, 53, 1887–1934.
- RADNER, R. (1962): "Team Decision Problems," *Annals of Mathematical Statistics*, 33, 857–881.
- SMITH, L., AND P. SØRENSEN (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398.
- (2008a): "Informational Herding and Optimal Experimentation," Working paper.
- STRZALECKI, T. (2009): "Depth of Reasoning and Higher Order Beliefs," *mimeo*.
- TVERSKY, A., AND D. KAHNEMANN (1974): "Judgement under Uncertainty: Heuristics and Biases," *Science*, 185, 1124–1131.
- ZIEGELMEYER, A., F. KOESSLER, J. BRACHT, AND E. WINTER (2009, forthcoming): "Fragility of information cascades: An experimental study using elicited beliefs," *Experimental Economics*.

Appendix: Omitted Proofs

Proof 2.1.1: Assume the realized state is (-1) . We have shown that a cascade eventually arises, i.e. the process of public likelihood ratios converges to a limiting random variable $\tilde{\lambda}_\infty$ with support in the sets $[0, (1-q)^3/q^3)$ and $(q^3/(1-q)^3, \infty)$. Take the latter one. Clearly the last decision before a cascade starts must be an investment. Furthermore, the public likelihood ratio right before this investment must satisfy $q^2/(1-q)^2 < \lambda_i < q^3/(1-q)^3$ and all types except the strong overweighers already decide regardless of their private signal at this point of time. The final investment moves the public likelihood ratio by the factor $\frac{5+q}{5+(1-q)}$ which implies that, in the information cascade, the public likelihood ratio cannot exceed the value $\frac{q^3}{(1-q)^3} * \frac{5+q}{5+1-q}$. By the dominated convergence theorem, we have that

$$1 = \lambda_1 = E\left[\tilde{\lambda}_\infty \mid \tilde{\theta} = (-1)\right] = \pi E\left[\tilde{\lambda}_\infty \mid \tilde{\theta} = (-1), \tilde{\lambda}_\infty \in \left[0, \frac{(1-q)^3}{q^3}\right)\right] \\ + (1-\pi) E\left[\tilde{\lambda}_\infty \mid \tilde{\theta} = -1, \tilde{\lambda}_\infty \in \left(\frac{q^3}{(1-q)^3}, \infty\right)\right]$$

where $\pi = Pr\left(\tilde{\lambda}_\infty \in \left(-\infty, \frac{(1-q)^3}{q^3}\right) \mid \tilde{\theta} = (-1)\right)$ which leads us to conclude that $\pi < 1$ since $E\left[\tilde{\lambda}_\infty \mid \tilde{\theta} = -1, \tilde{\lambda}_\infty \in \left[0, \frac{(1-q)^3}{q^3}\right)\right] \leq \frac{(1-q)^3}{q^3} < 1$. ■

Proof 2.1.2: Consider the equation

$$Pr\left(\tilde{\lambda}_\infty \in \left[0, (1-q)^3/q^3\right) \mid \tilde{\theta} = (-1)\right) = \pi(\bar{\lambda}_U, \bar{\lambda}_L) = \frac{\bar{\lambda}_U - 1}{\bar{\lambda}_U - \bar{\lambda}_L}$$

where $\bar{\lambda}_U = E\left[\tilde{\lambda}_\infty \mid \tilde{\theta} = -1, \tilde{\lambda}_\infty \in \left(\frac{q^3}{(1-q)^3}, \infty\right)\right]$ and $\bar{\lambda}_L = E\left[\tilde{\lambda}_\infty \mid \tilde{\theta} = (-1), \tilde{\lambda}_\infty \in \left[0, \frac{(1-q)^3}{q^3}\right)\right]$. One can show that π is strictly increasing in both of its arguments. Furthermore $\bar{\lambda}_U > \frac{q^3}{(1-q)^3}$ and $\bar{\lambda}_L > \frac{(1-q)^3}{q^3} * \frac{6-q}{5+q}$. Therefore $\pi \geq \frac{q^3(5+q)(q^3-(1-q)^3)}{q^6(5+q)-(1-q)^6(6-q)}$. The RHS of the latter exceeds $\frac{q(1+q)}{2(1-q+q^2)}$, the probability in BHW's specific model, provided $q > \underline{q}$ where $\underline{q} < 0.505$. ■

Proof 2.1.3: Assume that an investment cascade has started and let the public likelihood ratio be given by λ_{IC} . In BHW's specific model, $\lambda_{IC} = \frac{q(1+q)}{2(1-q+q^2)}$ while in our model $\lambda_{IC} \in \left(\frac{q^3}{(1-q)^3}, \frac{q^3}{(1-q)^3} * \frac{5+q}{6-q}\right]$. Let i , the next player to decide, be better informed in the sense that given $\theta = 1$ ($\theta = (-1)$) her signal is H (L) with probability $q' > q$. We analyze, separately for each type, how large q' has to be in order to induce a player of this type to follow her private information. Clearly, the crucial case is the situation where the player receive the signal L . In this case a player of type β_i , $\beta_i \in \{0, 1, 2, 3\}$, holds the likelihood ratio $\lambda(L, h_i) = \left(\frac{1-q'}{q'}\right)^{\beta_i} * \lambda_{IC}$ and follows her signal provided $\lambda(s_i, h_i) < 1$. Hence, she invests if $q' > q'_{min}(\beta_i) = \frac{\lambda_{IC}^{1/\beta_i}}{1+\lambda_{IC}^{1/\beta_i}}$. We may reinterpret $q'_{min}(\beta_i)$ by solving the equation $q'_{min}(\beta_i) = \frac{q^z}{q^z + (1-q)^z}$ for z . The solution which we denote by $z'_{min}(\beta_i)$ is the number of private signals of initial quality the player needs at least to potentially break the cascade. In BHW's specific model, $\beta_i = 1$ for all players and $1 < z'_{min}(1) < 4/3$ i.e. less than two signals are required to break a cascade. In our first extension, conformists need infinitely precise information to break a cascade. Bayesians have $z'_{min}(1) > 3$ as $\lambda_{IC} > q^3/(1-q)^3$, hence they need at least 3 private signals to break the cascade (compare this to $z'_{min} < 4/3$ in BHW's analysis). For weak overweighers we obtain less fragility compared to the standard model since $z'_{min}(2) > 2$. For strong overweighers the cascade may be more fragile since $1 < z'_{min}(3) < 2$.

Alternatively, we could assume that a public agency emits new information in the form of a private signal of precision q'' and determine the minimal amount of q'' that causes a player to break an existing cascade. Both approaches are equivalent. ■

Proof of Lemma 3.1: Let $b_i = b(s_i, \emptyset)$ and let $\lambda_i = \lambda(\emptyset, h_i)$. The threshold function follows immediately from solving $\left(\frac{b_i}{1-b_i}\right)^{\beta_i} \lambda_i = \lambda(s_i, h_i) > 1$ for $b(s_i, \emptyset)$.

We therefore turn to the properties. (i) and (ii) are both easily calculated. For (iii) and (iv) we obtain

$$\begin{aligned} \frac{\partial t(\lambda, \beta)}{\partial \beta} &= \frac{\lambda^{1/\beta} \log(\lambda)}{\beta^2 [1 + \lambda^{1/\beta}]^2} \\ \text{and } \frac{\partial t(\lambda, \beta)}{\partial \lambda} &= -\frac{\lambda^{\frac{1}{\beta}-1}}{\beta [1 + \lambda^{1/\beta}]^2} \end{aligned}$$

from which the properties follow immediately. ■

Proof of Lemma 3.2: Let $b_i = b(s_i, \emptyset)$ denote a player's realized private belief and let $\lambda_i = \lambda(\emptyset, h_i)$ denote his realized public likelihood ratio. In the equilibrium of $\langle \mathbb{N}, \Theta, A, u, (G'_{(-1)}, G'_1), W' \rangle$ each player's decision is characterized by two assumptions. First, each player updates his belief about the state of the world $\tilde{\theta}$ given his (realized) private belief and the (realized) history of predecessors' decisions using Bayes' rule. Second, given the so formed posterior, the player makes his choice by maximizing his expected utility. On the other hand the equilibrium of the game $\langle \mathbb{N}, \Theta, A, u, (G_{(-1)}, G_1), W \rangle$ while satisfying the latter differs with regard to the first assumption in the sense that players update beliefs according to the LR updating rule $\lambda(s_i, h_i) = \left(\frac{b_i}{1-b_i}\right)^{\tilde{\beta}} * \lambda_i$ where $\tilde{\beta}$ is a r.v. distributed according to the CDF W .

Fix $i \in \mathbb{N}$ and let $h : [0, 1] \times [0, \infty) \rightarrow [0, \infty)$ be given by $h(b, \beta) = \frac{b^\beta}{b^\beta + (1-b)^\beta}$. Define further $\tilde{\mu}_i := h(\tilde{b}_i, \tilde{\beta})$ where $\tilde{b}_i = b(\tilde{s}_i, \emptyset)$. Clearly, $\tilde{\mu}$ is a random variable on $[0, 1]$. We compute its distribution conditional on θ being the realization of $\tilde{\theta}$. We have

$$\begin{aligned} G'_\theta(x) &= \Pr(\tilde{\mu} < x \mid \tilde{\theta} = \theta) = \Pr(h(\tilde{b}_i, \tilde{\beta}_i) < x \mid \tilde{\theta} = \theta) \\ &= \Pr\left(\frac{1}{1 + [(1 - \tilde{b}_i)/\tilde{b}_i]^\beta} < x \mid \tilde{\theta} = \theta\right) \\ &= \int_0^\infty \Pr\left(\frac{1 - \tilde{b}_i}{\tilde{b}_i} > \left(\frac{1-x}{x}\right)^{1/\beta} \mid \tilde{\theta} = \theta\right) dW(\beta) \\ &= \int_0^\infty \Pr\left(\tilde{b} < \frac{x^{1/\beta}}{x^{1/\beta} + (1-x)^{1/\beta}} \mid \tilde{\theta} = \theta\right) dW(\beta) \\ &= \int_0^\infty G_\theta\left(\frac{x^{1/\beta}}{x^{1/\beta} + (1-x)^{1/\beta}}\right) dW(\beta). \end{aligned}$$

We check first whether G'_θ constitutes a proper cumulative distribution function. First we notice that $\frac{x^{1/\beta}}{x^{1/\beta} + (1-x)^{1/\beta}} = 1 - t\left(\frac{x}{1-x}, \beta\right)$. Then from Lemma 3.1, $G'_\theta(0) = \int_0^\infty G_\theta(0) dW(\beta) = 0$. Similarly $G'_\theta(1) = 1$. Furthermore by Lemma 3.1 (iv) $t(\lambda, \beta)$ is decreasing in λ . As $x/(1-x)$ is increasing in x and G_θ is weakly increasing, the integrand is weakly increasing in x and hence G'_θ is. Finally, continuity from the right follows straightforwardly. Thus G'_θ is a proper cumulative distribution function. Moreover

given that $G_{(-1)}$ and G_1 are absolutely continuous to one another, so are G'_{-1} and G'_1 . Therefore the new distribution satisfies that no signal can perfectly reveal the state of the world.

In conclusion if private beliefs are distributed according to $G'_\theta(b)$ provided $\tilde{\theta} = \theta$, players in the equilibrium of $\langle \mathbb{N}, \Theta, A, u, (G'_{(-1)}, G'_1), W \rangle$ update beliefs the same way as players in the equilibrium of $\langle \mathbb{N}, \Theta, A, u, (G_{(-1)}, G_1), W \rangle$. Since the second requirement of sequential rationality is satisfied in either equilibrium this proves the Lemma. ■

Proof of Lemma 3.3: We first characterize the support of the new private belief distribution. Let $[\underline{\beta}, \bar{\beta}]$ be the support of W , the distribution of decision weights. First we have $G'_\theta(b) = 0$ if for each $\beta \in [\underline{\beta}, \bar{\beta}]$, $\frac{b^{1/\beta}}{b^{1/\beta} + (1-b)^{1/\beta}} < \underline{b}$. This holds exactly when $b \leq \frac{b^\beta}{b^\beta + (1-b)^\beta}$ for each $\beta \in [\underline{\beta}, \bar{\beta}]$. Given that $\underline{b} < \frac{1}{2}$ the RHS is strictly decreasing in β and thus attains its minimum at $\bar{\beta}$. Hence, the condition holds for all $\underline{\beta} \leq \beta \leq \bar{\beta}$ provided it holds for $\bar{\beta}$. Hence $G'_\theta(b) = 0$ if and only if $b \leq \frac{b^{\bar{\beta}}}{b^{\bar{\beta}} + (1-b)^{\bar{\beta}}}$.

On the other hand $G'_\theta(b) = 1$ if for each $\beta \in [\underline{\beta}, \bar{\beta}]$ it holds $\frac{b^{1/\beta}}{b^{1/\beta} + (1-b)^{1/\beta}} > \bar{b}$ which is equivalent to $b > \frac{\bar{b}^\beta}{\bar{b}^\beta + (1-\bar{b})^\beta}$ holding for each $\beta \in [\underline{\beta}, \bar{\beta}]$. As the RHS of it is strictly increasing in β in this case, it attains its maximum at $\bar{\beta}$ and thus the condition is satisfied exactly when $b > \frac{\bar{b}^{\bar{\beta}}}{\bar{b}^{\bar{\beta}} + (1-\bar{b})^{\bar{\beta}}}$.

Finally, if $\underline{b}' < b < \bar{b}'$, $\underline{b} < x < \bar{b}$ where $x = \frac{b^{1/\beta}}{b^{1/\beta} + (1-b)^{1/\beta}}$ for at least some $\beta \in [\underline{\beta}, \bar{\beta}]$. However, for such x , $G_1(x) < G_{(-1)}(x)$ and therefore the same inequality holds after integrating. ■

Proof of Lemma 3.9: We fix $T \in \mathbb{N}$ and for $\theta \in \{(-1), 1\}$ let $\pi_\theta = \bar{\sigma}(1|\theta)$ denote investment frequencies under state θ . The proof follows along the same lines as the one by Guarino and Jehiel (2009) (Proposition 7). Define the function $\Phi : [0, 1]^2 \rightarrow [0, 1]^2$ that for a given vector $(\pi_{(-1)}, \pi_1)$ where $\pi_\theta = \bar{\sigma}(1 | \theta)$ gives the aggregate investment frequencies in the social learning game of length T where players update beliefs according to $\lambda(s_i, h_i) = (\lambda(s_i, \emptyset))^{\beta_i} * \left(\frac{\pi_1}{\pi_{(-1)}}\right)^{\Sigma_1(h_i)} * \left(\frac{1-\pi_1}{1-\pi_{(-1)}}\right)^{\Sigma_0(h_i)}$ with $\Sigma_a(h_i)$ the number of times action $a \in A$ occurred in h_i and β_i distributed according to W . We show that for continuous distribution W , Φ is continuous and furthermore Φ maps the upper triangle of the unit interval $\Delta = \{(x, y) : 0 \leq x \leq 1 \wedge x \leq y \leq 1\}$ into itself. Then by Brouwer's fixed point theorem there exists $(\pi_{(-1)}^*, \pi_1^*) = \Phi(\pi_{(-1)}^*, \pi_1^*)$.

We denote by $\Phi_\theta(\pi_{(-1)}, \pi_1)$ the investment frequency Φ determines for state θ , i.e. if $\Phi(\pi_{(-1)}, \pi_1) = (\hat{\pi}_0, \hat{\pi}_1)$ then $\Phi_\theta(\pi_{(-1)}, \pi_1) = \hat{\pi}_\theta$. Then

$$\Phi_\theta(\pi_{(-1)}, \pi_1) = \frac{1}{T} \sum_{i=1}^T \sum_{h_i \in H_i} \Pr(\tilde{a}_i = 1 \mid \tilde{h}_i = h_i, \tilde{\theta} = \theta) \Pr(\tilde{h}_i = h_i \mid \tilde{\theta} = \theta)$$

and we have that

$$\begin{aligned} \Pr(\tilde{a}_i = 1 \mid \tilde{h}_i = h_i, \tilde{\theta} = \theta) &= \int_{\underline{\beta}}^{\bar{\beta}} \Pr\left(\left(\frac{\tilde{b}_i}{1-\tilde{b}_i}\right)^\beta > \left(\frac{\pi_{(-1)}}{\pi_1}\right)^{\Sigma_1(h_i)} \left(\frac{1-\pi_{(-1)}}{1-\pi_1}\right)^{\Sigma_0(h_i)} \mid \tilde{\theta} = \theta\right) dW(\beta) \\ &= 1 - \int_{\underline{\beta}}^{\bar{\beta}} G_\theta\left(\frac{\pi_{(-1)}^{\Sigma_1(h_i)/\beta} * (1-\pi_{(-1)})^{\Sigma_0(h_i)/\beta}}{\pi_{(-1)}^{\Sigma_1(h_i)/\beta} * (1-\pi_{(-1)})^{\Sigma_0(h_i)/\beta} + \pi_1^{\Sigma_1(h_i)/\beta} * (1-\pi_1)^{\Sigma_0(h_i)/\beta}}\right) dW(\beta) \\ &= 1 - G'_\theta\left(\frac{\pi_{(-1)}^{\Sigma_1(h_i)} (1-\pi_{(-1)})^{\Sigma_0(h_i)}}{\pi_{(-1)}^{\Sigma_1(h_i)} (1-\pi_{(-1)})^{\Sigma_0(h_i)} + \pi_1^{\Sigma_1(h_i)} (1-\pi_1)^{\Sigma_0(h_i)}}\right) \end{aligned}$$

where $G'_\theta(x) = \int_{\underline{\beta}}^{\bar{\beta}} G_\theta\left(\frac{x^{1/\beta}}{x^{1/\beta} + (1-x)^{1/\beta}}\right) dW(\beta)$.

First, continuity of G'_θ implies continuity of $Pr(\tilde{a}_i = 1 \mid \tilde{h}_i = h_i, \tilde{\theta} = \theta)$ and also of $Pr(\tilde{h}_i = h_i \mid \tilde{\theta} = \theta)$ ⁸ with respect to the vector $(\pi_{(-1)}, \pi_1)$. Hence, Φ_θ is continuous if G'_θ is. To show the latter notice first that trivially if G_θ is continuous, so is G'_θ . On the other hand G_θ can only possess countably many jumps. For each $x \in (0, 1)$ there thus exist countably many $\beta > 0$ such that G_θ has a jump at $\frac{x^{1/\beta}}{x^{1/\beta} + (1-x)^{1/\beta}}$. However, for a continuous distribution a countable subset is a null set. Thus G'_θ and thus Φ_θ is continuous for each $\theta \in \{(-1), 1\}$.

Second, we show that $\Phi : \Delta \rightarrow \Delta$. As $0 \leq \Phi_\theta(\pi_{(-1)}, \pi_1) \leq 1$ by definition the only thing we need to show is that $\pi_1 > \pi_{(-1)}$ implies $\Phi_1(\pi_{(-1)}, \pi_1) > \Phi_{-1}(\pi_{(-1)}, \pi_1)$. To proof the latter showing $Pr(a_i = 1 \mid \tilde{\theta} = 1) > Pr(a_i = 1 \mid \tilde{\theta} = 0)$ for any $i = 1, 2, \dots, T$ is sufficient. For each $i = 1, \dots, T$ we turn to the distribution of the differences $\Delta\Sigma(\tilde{h}_i) = \Sigma_1(\tilde{h}_i) - \Sigma_0(\tilde{h}_i)$ conditional on the state of the world. We first collect the following properties:

(i) Investment probabilities are constant across histories

$$h_i \in \{h_i \in H_i : \Delta\Sigma(h_i) = j \in \{-(i-1), -(i-1)+2, \dots, i-1\}\}$$

and given by

$$Pr(\tilde{a}_i = 1 \mid \Delta\Sigma(\tilde{h}_i) = j, \tilde{\theta} = \theta) = 1 - G'_\theta\left(\frac{\pi_{(-1)}^{(i-1+j)/2} (1 - \pi_{(-1)})^{(i-1-j)/2}}{\pi_{(-1)}^{(i-1+j)/2} (1 - \pi_{(-1)})^{(i-1-j)/2} + \pi_1^{(i-1+j)/2} (1 - \pi_1)^{(i-1-j)/2}}\right).$$

(ii) Investment probabilities are strictly increasing in $\Delta\Sigma(h_i)$, as $\left(\frac{\pi_1}{\pi_{(-1)}}\right)^{(i-1+j)/2}$ and $\left(\frac{1-\pi_1}{1-\pi_{(-1)}}\right)^{(i-1-j)/2}$ both are strictly increasing in j .

(iii) $Pr(\tilde{a}_i = 1 \mid \Delta\Sigma(\tilde{h}_i) = j, \tilde{\theta} = 1) > Pr(\tilde{a}_i = 1 \mid \Delta\Sigma(\tilde{h}_i) = j, \tilde{\theta} = (-1))$ for each $j = -(i-1), -(i-1)+2, \dots, i-1$ follows from $G'_{(-1)}(x) > G'_1(x)$ for each $x \in (\underline{b'}, \bar{b'})$.

Given these properties it suffices to show that the distribution of differences $\Delta\Sigma(\tilde{h}_i)$ conditional on $\tilde{\theta} = 1$, first-order stochastically dominates the associated distribution conditional on $\tilde{\theta} = (-1)$. Let $K_{i,\theta}$ denote the cumulative distribution function of differences for player $i \in \mathbb{N}$ conditional on $\tilde{\theta} = \theta$, i.e. $K_{i,\theta}(x) = Pr(\Delta\Sigma(\tilde{h}_i) < x \mid \tilde{\theta} = \theta)$. We show that $K_{i,(-1)}(x) > K_{i,1}(x)$ for any $x \in \{-(i-1), i-3\}$ and any $i = 1, 2, \dots, T$ via induction. For player 1 $\Delta\Sigma(\tilde{h}_1) \equiv 0$. For player 2 the difference takes value (-1) with probability $G_\theta(1/2)$ and value 1 with opposite probability $1 - G_\theta(1/2)$. Thus clearly for $-1 \leq x < 1$, $K_{2,(-1)}(x) > K_{2,1}(x)$. Now assume that for any $j < i$, $K_{j,(-1)}(x) > K_{j,1}(x)$. For player i and $z \in \{-(i-1), -(i-1)+2, \dots, i-3\}$ it holds

$$K_{i,\theta}(z) = K_{i-1,\theta}(z-1) + [K_{i-1,\theta}(z+1) - K_{i-1,\theta}(z-1)] * Pr(\tilde{a}_{i-1} = 0 \mid \Delta\Sigma(\tilde{h}_{i-1}) = z+1, \tilde{\theta} = \theta).$$

Then

$$\begin{aligned} K_{i,(-1)}(z) - K_{i,1}(z) &= [K_{i-1,(-1)}(z-1) - K_{i-1,1}(z-1)] \\ &\quad + [K_{i-1,(-1)}(z+1) - K_{i-1,1}(z-1)] Pr(\tilde{a}_{i-1} = 0 \mid \Delta\Sigma(\tilde{h}_{i-1}) = z+1, \tilde{\theta} = (-1)) \\ &\quad - [K_{i-1,1}(z+1) - K_{i-1,1}(z-1)] Pr(\tilde{a}_{i-1} = 0 \mid \Delta\Sigma(\tilde{h}_{i-1}) = z+1, \tilde{\theta} = 1) \end{aligned}$$

⁸This follows as $Pr(\tilde{h}_i = h_i \mid \tilde{\theta} = \theta)$ is a product of expressions $G'_\theta(x)$ and $1 - G'_\theta(x)$.

and rearranging terms yields

$$\begin{aligned}
 K_{i,(-1)}(z) - K_{i,1}(z) &= \left[K_{i-1,(-1)}(z-1) - K_{i-1,1}(z-1) \right] * Pr(\tilde{a}_{i-1} = 1 \mid \Delta\Sigma(\tilde{h}_{i-1}) = z+1, \tilde{\theta} = (-1)) \\
 &+ \left[K_{i-1,(-1)}(z+1) - K_{i-1,1}(z+1) \right] * Pr(\tilde{a}_{i-1} = 0 \mid \Delta\Sigma(\tilde{h}_{i-1}) = z+1, \tilde{\theta} = (-1)) \\
 &+ \left[K_{i-1,1}(z+1) - K_{i-1,1}(z-1) \right] * \Delta P_{rej}(z+1) \\
 &> 0
 \end{aligned}$$

where

$$\Delta P_{rej}(z+1) = \left[Pr(\tilde{a}_{i-1} = 0 \mid \Delta\Sigma(\tilde{h}_{i-1}) = z+1, \tilde{\theta} = (-1)) - Pr(\tilde{a}_{i-1} = 0 \mid \Delta\Sigma(\tilde{h}_{i-1}) = z+1, \tilde{\theta} = 1) \right].$$

The inequality follows from the induction assumption for the first and second term and from the properties of a c.d.f. and property (iii) of the investment probabilities discussed above for the third term. This finishes the proof. ■

Proof of Lemma 3.10: If private information weights are bounded, there exists $n \in \mathbb{N}$ such that if the first n choices are similar, all subsequent players follow suit with probability one. However, a finite number of choices can be wrong with strictly positive probability as a finite number of private signals indicates the wrong state with strictly positive probability.

We thus turn to the case of unbounded private information weights. Assume $\pi_1 = 1 - \pi_{(-1)} = 1 - \epsilon$ for some $\epsilon > 0$ small. Fix the length of the sequence T . As in the proof of Guarino and Jehiel (2009) (Proposition 8) we are going to show that all players are wrong with strictly positive probability. W.l.o.g. assume $\tilde{\theta} = (-1)$. Then the first player invests with probability $1 - G_{(-1)}(1/2)$ upon which the public LR rises to $\pi_1/\pi_{(-1)} > 1$. Given this choice the second player invests provided $\left(\frac{b_2}{1-b_2}\right)^{\beta_2} * \frac{\pi_1}{\pi_{(-1)}} > 1$. After the second player's investment the public LR increases to $\pi_1^2/\pi_{(-1)}^2$. In general after the first $(i-1)$ players invested, player i ($k = 1, 2, \dots, T$) faces a public LR of $\pi_1^{i-1}/\pi_{(-1)}^{i-1}$ and consequently invests provided $\left(\frac{b_i}{1-b_i}\right)^{\beta_i} * \left(\frac{\pi_1}{\pi_{(-1)}}\right)^{i-1} > 1$. As $\pi_1 > \pi_{(-1)}$ and $\beta_i > 0$, the latter is satisfied if either $b_i > 1/2$ or else if $b_i < 1/2$ and $\beta_i < (i-1) * \frac{\log(\pi_1/\pi_{(-1)})}{\log((1-b_i)/b_i)}$. Thus we may write the probability of player i investing after anyone has invested before by

$$\begin{aligned}
 Pr(\tilde{a}_i = 1 \mid \tilde{\theta} = (-1), \tilde{h}_i = (1, 1, \dots, 1)) &= 1 - G_{(-1)}(1/2) + \int_{\underline{b}}^{1/2} W\left((i-1) \frac{\log(\pi_1/\pi_{(-1)})}{\log((1-b)/b)}\right) dG_{(-1)}(b) \\
 &= 1 - \int_{\underline{b}}^{1/2} \left[1 - W\left((i-1) \frac{\log(\pi_1/\pi_{(-1)})}{\log((1-b)/b)}\right)\right] dG_{(-1)}(b).
 \end{aligned}$$

First assume (a) $\mu_\beta = E[\tilde{\beta}_i] < \infty$ and $\sigma_\beta^2 = Var[\tilde{\beta}_i] < \infty$. Then by Chebyshev's inequality $1 - W(x) \leq \frac{1}{1+(x-\mu_\beta)^2/\sigma_\beta^2} < \frac{\sigma_\beta^2}{(x-\mu_\beta)^2}$. Then as $\log(x) < x$

$$\begin{aligned}
 Pr(\tilde{a}_i = 1 \mid \tilde{\theta} = (-1), \tilde{h}_i = (1, 1, \dots, 1)) &\geq 1 - \int_{\underline{b}}^{1/2} \frac{\sigma_\beta^2 (\log((1-b)/b))^2}{\left[(i-1) \log(\pi_1/\pi_{(-1)}) - \mu_\beta \log((1-b)/b)\right]^2} dG_{(-1)}(b) \\
 &\geq 1 - \frac{\sigma_\beta^2 \log((1-\underline{b})/\underline{b}) G_1(1/2)}{\left[(i-1) \log(\pi_1/\pi_{(-1)}) - \mu_\beta \log((1-\underline{b})/\underline{b})\right]^2}.
 \end{aligned}$$

Taking the logarithm of this probability and using as in Guarino and Jehiel (2009) that by Taylor's series expansion $\log(1-x) \geq -ax$ for some $a > 1$ and $0 < x < 1$ we obtain for the probability that all

players invest conditional on $\tilde{\theta} = 0$

$$\begin{aligned} & Pr(\tilde{a}_i = 1 \forall i = 1, \dots, T \mid \tilde{\theta} = (-1)) \\ & \geq \left[1 - G_{(-1)}(1/2)\right] \exp \left(-a \sigma_\beta^2 \log((1 - \underline{b})/\underline{b}) G_1(1/2) \sum_{i=1}^{T-1} \frac{1}{\left[i \log(\pi_1/\pi_{(-1)}) - \mu_\beta \log((1 - \underline{b})/\underline{b})\right]^2} \right). \end{aligned}$$

As $\sum_{i=1}^{\infty} \frac{1}{(ai+b)^2} < \infty$ provided $b/a > -1$,⁹ μ_β and σ_β^2 are finite and as $\frac{\pi_1}{\pi_{(-1)}} = \frac{1-\epsilon}{\epsilon} \rightarrow \infty$ as $\epsilon \rightarrow 0$ the argument of the exponential function converges to zero as $T \rightarrow \infty$ and $\epsilon \rightarrow 0$. Hence, the probability that everyone invests in state 0 is bounded below by $1 - G_{(-1)}(1/2)$ which is strictly positive.

In the case of (b) $1 - W(x) \approx x^{-(1+\alpha)}$, we have that

$$\begin{aligned} Pr(\tilde{a}_i = 1 \mid \tilde{\theta} = (-1), \tilde{h}_i = (1, 1, \dots, 1)) &= 1 - \int_{\underline{b}}^{1/2} \left[1 - W\left((i-1) \frac{\log(\pi_1/\pi_{(-1)})}{\log((1-\underline{b})/\underline{b})}\right)\right] dG_{(-1)}(b) \\ &\approx 1 - \int_{\underline{b}}^{1/2} (i-1)^{-(1+\alpha)} \left(\frac{\log((1-\underline{b})/\underline{b})}{\log(\pi_1/\pi_{(-1)})}\right)^{1+\alpha} dG_{(-1)}(b) \\ &\geq 1 - G_{(-1)}(1/2) \left(\frac{\log((1-\underline{b})/\underline{b})}{\log(\pi_1/\pi_{(-1)})}\right)^{1+\alpha} (i-1)^{-(1+\alpha)}. \end{aligned}$$

Therefore by the same calculations as above

$$Pr(\tilde{a}_i = 1 \forall i = 1, \dots, T \mid \tilde{\theta} = (-1)) \geq \left[1 - G_{(-1)}(1/2)\right] \exp \left(-a G_{(-1)}(1/2) \left(\frac{\log((1-\underline{b})/\underline{b})}{\log(\pi_1/\pi_{(-1)})}\right)^{1+\alpha} \sum_{i=1}^{T-1} i^{-(1+\alpha)} \right).$$

Now, as $T \rightarrow \infty$ the sum converges to $\zeta(1+\alpha)$, the value of the Riemann Zeta function at the point $1+\alpha$. This is finite for any $\alpha > 0$. Thus as above the RHS converges to $1 - G_{(-1)}(1/2) > 0$ as $T \rightarrow \infty$ and $\epsilon \rightarrow 0$. With strictly positive probability, all players invest in state 0. ■

⁹The value of the sum equals the value $\Psi'(1 + b/a)$ of the first derivative of the digamma function at the point $1 + \frac{b}{a}$. This derivative is finite for any argument $x > 0$.