# A New Perspective on Social Learning

vorgelegt von
Diplom-Wirtschaftsmathematiker
**Christoph March**
aus Erfurt

Von der Fakultät VII – Wirtschaft & Management
der Technischen Universität Berlin
zur Erlangung des akademischen Grades
Doktor der Wirtschaftswissenschaften
Dr. rer. oec.

genehmigte Dissertation

Promotionsausschuss:

| | |
|---|---|
| Vorsitzender: | Prof. Dr. Frank Heinemann |
| Gutachterin: | Prof. Dr. Dorothea Kübler |
| Gutachter: | Dr. Anthony Ziegelmeyer |
| Gutachter: | Prof. Christophe Chamley |

Tag der wissenschaftlichen Aussprache: 15.12.2010

Berlin 2011

D 83

## Abstract

Home prices in the United States experienced unprecedented growth beginning in the late 1990s, followed by a dramatic downturn in 2006 which caused a global economic recession. Though market fundamentals such as low interest rates and easy mortgage terms boosted housing prices, the full magnitude of the housing bubble can only be reconciled with a social epidemic of optimism for real estate. From a policy perspective, a solid understanding of the idea of social contagion is therefore desirable. Existing economic models of social learning have contributed to this understanding by showing that in many situations the decisions of rational individuals tend to converge quickly and the conformist outcome is often wrong. We argue in this dissertation, however, that some of the conclusions reached by the rational view of social learning are unsound which limits its capacity to explain social epidemics. A thorough investigation of the learning foundations of rational herding enables us to develop an alternative theoretical framework which holds greater promise for explaining social epidemics.

Chapter 1 re-examines the classic model of Bayesian rational social learning most commonly investigated in the literature. Previous studies have mainly focused on the long-run properties of the equilibrium outcome in relation with the support of the private belief distribution. From a welfare perspective, this focus on the long-run learning outcome is misplaced. To address this concern, we introduce new measures of uniformity and fragility which enable us to discuss the medium-run properties of the equilibrium outcome. We find that less uniform investment decisions lead to higher welfare levels of the equilibrium outcome but lower fragility of rational herds.

Chapter 2, the core of the thesis, investigates the learning foundations of economic models of social learning. We pursue the prevalent idea in economics that rational play is the outcome of a dynamic process of adaptation. Our learning approach offers us the possibility to clarify when and why equilibrium is likely to capture observed regularities in the field. Contrary to the eductive justification for equilibrium, a learning-theoretic model must address the issue of individual and interactive knowledge before the adaptive process starts. We argue that knowledge about the private belief distribution is unlikely to be shared in most social learning contexts. Absent this mutual knowledge, we show that the long-run outcome of the adaptive process favors non-Bayesian rational play.

Chapter 3 introduces a simple extension to the rational model of social learning by assuming that players update their beliefs in a non-Bayesian way. We examine the properties of this generalized social learning model and show that it is able to capture the experimental regularities. Unlike standard models of rational social learning, our model predicts that as a herd evolves players tend to quickly become extremely confident about the appropriateness of the chosen action which is why the fragility of the social learning process vanishes over time. Both regularities seem to prevail in social epidemics.

## Zusammenfassung

Ende der 1990er Jahre setzte ein zu diesem Zeitpunkt beispielloser Preisanstieg auf dem US-Eigenheimmarkt ein. Dessen dramatischer Einbruch im Jahre 2006 gilt als einer der Hauptgründe für die noch immer anhaltende Rezession der Weltwirtschaft. Obschon durch niedrige Zinssätze und Kreditvergabekonditionen begünstigt, liefert erst eine Betrachtung der epidemieartigen Ausbreitung von Optimismus über die zukünftige Marktentwicklung eine Erklärung für die Ausmaße dieser "Eigenheim-Blase". Das Verständnis solcher sozialen Herdenphänomene ist daher ökonomisch und sozialpolitisch von großer Bedeutung.

Bestehende ökonomische Modelle zum "Sozialen Lernen" nehmen in diesem Bestreben eine zentrale Rolle ein. Diese Modelle zeigen, dass die schnelle Entwicklung von konformem Verhalten oft rational, gleichzeitig jedoch häufig ineffizient oder gar hochgradig schädlich ist. Allerdings sind nicht alle Schlussfolgerungen, welche sich aus den Modellen ergeben, plausibel. Die vorliegende Dissertation stellt daher die Eignung der Modelle zur Erklärung der Phänomene in Frage. Insbesondere werden die lernorientierten Grundlagen des angenommenen strategischen Verhaltens kritisch beleuchtet. Die Ergebnisse münden in der Entwicklung eines Alternativmodells, welches eher in der Lage ist, soziale Herdenphänomene zu erklären.

Kapitel 1 ist den bestehenden klassischen Modellen sozialen Lernens gewidmet, welche vollkommene Rationalität und Bayes'sches Verhalten der Entscheider annehmen. Die Schwerpunktsetzung dieser Modelle auf Langfrist-Vorhersagen und den maximalen Informationsgehalt privater Informationen wird hinterfragt. Eine neue Perspektive mit Blick auf mittelfristige Vorhersagen und soziale Wohlfahrt und basierend auf neuen Definitionen u.a. von Konformität wird erarbeitet. Es wird gezeigt, dass höhere Wohlfahrt eng mit der langsameren Entwicklung konformen Verhaltens verknüpft ist.

Das zentrale Kapitel 2 untersucht die lerntheoretischen Grundlagen der Modelle. Es steht damit in der ökonomischen Tradition, welche rationales Verhalten als Ergebnis von Anpassungsprozessen versteht. Eine solche dynamische Perspektive ermöglicht daher eine Antwort auf die Frage, ob und warum ein Gleichgewichtsmodell in der Lage ist, tatsächlich beobachtbares Verhalten zu erfassen. Insbesondere kann untersucht werden, welche Kenntnis der Entscheider über die Struktur der Entscheidungssituation vorausgesetzt werden kann. Die Ergebnisse werfen in dieser Hinsicht große Zweifel über die Annahmen bestehender Modelle auf. Genauer gesagt zeigen die Ergebnisse, dass die Vorhersagen der Modelle entscheidend durch eine als zu groß angenommene Kenntnis der Entscheider beeinflusst sind. Dies macht eine Neubewertung sowohl der theoretischen als auch der experimentellen Resultate notwendig.

Kapitel 3 schließlich erweitert die bestehenden Modelle um die Präsenz von Individuen, welche Informationen nicht im Sinne von Bayes kombinieren. Es wird gezeigt, dass die Vorhersagen des derart erweiterten Modells besser mit dem in Experimenten beobachtbaren Verhalten übereinstimmen. Insbesondere liefert das erweiterte Modell im Gegensatz zu den bestehenden Modellen eine Erklärung für das wachsende Vertrauen von Individuen in die Richtigkeit des Herdenverhaltens und die daraus resultierende Stabilität desselbigen.

## Acknowledgements

Numerous people have supported me during the past 3 years within which this thesis was completed.

I am particularly grateful to my supervisor, Anthony Ziegelmeyer, who has also co-authored two of the chapters. This thesis could not have been completed without our innumerable discussions, his aptitude to always ask the right questions, the calmness with which he faced my concerns, and his guidance into the field of economics and through the process of writing this dissertation.

I am furthermore indebted to Dorothea Kübler and Christophe Chamley for placing their trust in me and my work by agreeing to take over the posts of first and third supervisor.

I gratefully acknowledge financial support from the graduate program IMPRS "Uncertainty". I thank Linan Diao, Christiane Ehses-Friedrich, Charlotte Klempt, Sebastian Krügel, Laxmi Natarjan, Lauri Sääksvuori, Matthias Uhl, Ori Weisel, Johannes Weisser, Fabian Winter, and all the other former or current students from the IMPRS-program as well as the faculty and our coordinators Susanne Büchner and Tobias Uske for providing a stimulating research atmosphere inside the office and especially during summer schools and workshops. I also want to thank participants at the 2009 Summer School on Limited Cognition, Strategic Thinking, and Learning in Games in Bonn, Germany as well as the speakers Vincent Crawford and Philippe Jehiel.

To my friends for being there and providing necessary distraction, to my mother for her endless unconditional support, and to Constanze for all of the above and much more, my deep gratitude.

# Contents

# Introduction

Imagine a tourist in a foreign city around dinner time. There are two restaurants facing each other on the main street. The tourist may consult his travel guide as to obtain some information about the quality of either of the two. Or she may take a look through the windows to find one of them well-filled, the other one sparsely occupied. Upon which information is the tourist going to act? Is it possible that a crowded place will cause her to disregard even a strong recommendation of the other still empty place? And may the better restaurant remain empty?

Such questions are at the heart of economists' research on *social learning* – learning by observing others. In the past two decades since the seminal papers of Bikhchandani, Hirshleifer, and Welch (1992) and Banerjee (1992) economists have developed, extended and tested models to answer such questions. The canonical economic model of social learning comprises a set of players sequentially choosing from a finite set of actions encumbered solely by uncertainty about the state of the world. Each player receives a private signal correlated with the state of the world. In addition players may observe the choices of those players which acted already but not their private information. In such settings it is rational for players to learn from one another.

### Sales, Fashion, and Cultural Change

But social learning is about more than just restaurant choices, it pervades the social life. People consult bestseller lists when deciding which book to read or movie to watch, firms observe other firms before making an investment. Various examples of learning through the observation of others have been reported in the literature.[1]

A very striking one is the case of oil drilling in the U.S.(see Hendricks and Kovenock, 1989). The leasing practices of the federal government allow firms to lease tracts in an area which may contain an oil deposit for a period of five years. Firms invest in private seismic surveys which provide some information about the existence of oil. But ultimately this can only be determined by drilling an exploratory well, which is very costly. Frequently two or more firms lease adjoining tracts. In this case any firm would want the other(s) to incur the expense of the well. During the period of 1954 to 1970 more than 300 leases off the coast of Texas and Louisiana expired without any drilling despite leasing costs of approximately $800, 000 (see Hendricks, Porter, and Boudreau, 1987).

Also in financial markets herd behavior is a well known force. Bubbles, crashes, and bank runs are a few examples. In fashion, cultural and political movements, conformity arises as regularly as it disappears.[2] Models of social learning are able to accommodate these phenomena better than

[1]See for instance the survey of Bikhchandani, Hirshleifer, and Welch (1998).

[2]See Lohmann (1994) for an interesting application to the fall of the Berlin wall.

many other explanations offered.

### Rational Social Learning

An important class of social learning models assumes that each individualŠs decision reflects in a Bayes-rational fashion the content of her private signal and of the history of observed decisions. Using BayesŠ 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.

Early studies demonstrate how under these assumptions the attempt to take advantage of information of other players prevents individuals from exploiting their private information in a socially optimal way. This likely consequence of rational social learning has been termed *informational 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.

### Brittle, Error-prone Conformity

An information cascade induces a *herd*; all players choose the same action. A cascade starts exactly when the accumulated evidence from previous choices is just larger than the strongest amount of (contradictory) private information. But once a cascade starts learning stops completely. Thus any player in a cascade chooses the same action based on the same small amount of information. This has two major implications: On one hand with positive probability players herd on a less profitable action. On the other hand conformity is brittle. If even a little new information arises, for example released by the government, behavior might shift radically. This property distinguishes cascades from other explanations of conformity. For instance assuming positive payoff externalities or having people exhibit a preference for conformity[3] tend to bring about a rigid conformity which is more robust the longer the bandwagon. Informational cascades hence provide a better explanation for various real world phenomena such as fads, trends, or cultural movements which emerge today to disappear tomorrow.

### The Robustness of Cascades

Many studies have extended the basic social learning paradigm. The most comprehensive and exhaustive analysis of social learning where players can observe the complete sequence of past decisions has been provided by Smith and Sørensen (2000). They show that the complete failure

[3]Bikhchandani, Hirshleifer, and Welch (1992, p.993) provide a more thorough discussion of this issue together with a list of references.

of information aggregation is not a robust property. Indeed, if players can obtain private signals of unbounded strength, the truth is asymptotically fully revealed. However, the convergence of players' actions in a herd still takes place eventually. In this case all players make the same choice despite a strictly positive probability that some player receives a private signal which induces her to make a different one. Since this probability must necessarily be small learning takes place very slowly in a herd. The main message of social learning models with finite action space is the *self-defeating property*, the serious failure to achieve a socially desirable outcome due to information externalities.

The outcome of social learning models is not robust on a second account. If the action space is continuous and the payoff function non-degenerate players are able to perfectly adjust their choices to posterior beliefs. Hence, any new information is fully revealed and consequently convergence to the truth obtains at the same speed as if all information was fully revealed. Indivisibilities of the action space are crucial for the self-defeating property of social learning. When individuals learn by observing others, actions are the filter through which information is communicated. If this filter is very coarse, a lot of information gets lost and the process of social learning is inefficient.

### Rationality and Inefficiency

The inefficiency of social learning with finite action space arises due to the *information externality* a player's action creates for the information of others and which is not taken into account by rational players. This has led to the question, how a socially more desirable outcome can be brought about. Studies of Sgroi (2002) and Smith and Sørensen (2008a) show that in the social optimum players should overweight their private information compared to what Bayesian rationality dictates. Hence, the process of social learning is inefficient and may fail completely not merely despite the rationality of individual players, but often because of that rationality. This hints at social benefits of nonrational behavior.

### Social Learning in the Laboratory

Social learning models have been extensively studied not only theoretically but also experimentally. Anderson and Holt (1997) were the first to investigate the emergence of informational cascades in the laboratory. They find evidence that informational cascades occur in the laboratory and interpret their data as a clear support for Bayesian rationality. Yet, their findings also show that participants most often deviate from Bayesian rationality when the latter predicts a choice inconsistent with the one based only on the private signal. More precisely participants are reluctant to rationally ignore their private information when the evidence from previous decisions is small or mixed. On the other hand they have no problem to follow an pattern established in a long, pure sequence of previous decisions. Hence, participants act as if they form their belief by assigning to much weight to their private belief relative to the public belief. We refer to this as the "overweight"-phenomenon. Several subsequent studies confirmed the findings of Anderson and Holt. More importantly by introducing relevant complications to the cascade process these studies were able to rule out several simple explanations for the "overweight"-phenomenon and question the success of others such as logit quantal response equilibrium (LQRE, McKelvey and Palfrey, 1995, 1998). In a recent meta-study, Weizsäcker (2010) aggregates the results of the most important experimental studies of informational cascades and produces a clear picture of the "overweight"-phenomenon. His findings illustrate the failure of Bayesian rationality and LQRE to predict behavior and strongly demonstrate the need for an alternative theory.

### Overview of the Thesis

**Chapter 1** (joint with Anthony Ziegelmeyer) surveys the bulk of the economic literature on social learning. Though this literature has greatly improved our understanding of rational herding, its regular emphasis on the long-run learning outcome has limited the development of accurate assessments of the equilibrium welfare. As a remedy, this chapter suggests an approach to discuss the medium-run properties of the equilibrium outcome by introducing new measures of uniformity and fragility.

**Chapter 2**, the heart of the thesis, investigates the learning foundations of economic models of social learning. An epistemic learning process favors non-Bayesian play when players face significant structural uncertainty. Since in most social learning environments significant structural uncertainty prevails, Bayesian rational social learning is unlikely to be the long-run outcome of adaptation.

**Chapter 3** (joint with Anthony Ziegelmeyer) extends the economic model of social learning by incorporating players' adaptive response to structural uncertainty. Players either overweigh or underweigh (in Bayesian terms) their private information relative to the public information revealed by the decisions of others. Introducing heterogeneous updating rules reconciles equilibrium predictions with evidence from the laboratory and the field.

# Chapter I.

# Bayesian (Rational) Social Learning

## I.1. Introduction

Worldwide real estate markets have recently experienced the most pronounced turbulence ever. There have been wild swings in prices, a wave of foreclosures, and countless failed investments which caused a global credit crunch. Despite its severity and its ample effects, the recent economic crisis is hardly exceptional. The history of industrialized economies is the history of recurrent economic crises which originate in speculative price movements (Reinhart and Rogoff, 2009). Sound theoretical models of investment booms and subsequent recessions are therefore needed to prevent economic crises and reduce their effects once they occur.

Extreme and rapid market fluctuations are often claimed to be inconsistent with rational investors and are attributed to psychological biases, behavioral rules and persistent mistakes made by boundedly rational investors. However, economic models of social learning have established that boundedly rational investors with behavioral biases are not needed to explain booms and crashes. Idiosyncratic and fragile herding behavior can be the result of imperfectly informed rational investors who try to learn from observing others' investment decisions. Socially inefficient market outcomes occur due to informational externalities not taken into account by rational investors. The economic literature on social learning suggests that market frenzies and crashes should not be attributed fully to the irrational behavior of investors.

With the help of a game-theoretic framework, this chapter discusses simple economic models of social learning. We restrict attention to situations in which Bayesian rational investors are endowed with private signals about a payoff-relevant state of Nature and choose irreversible investment options in an exogenous order after having observed their predecessors' investment decisions. The price of each investment option is fixed, i.e. payoff externalities are absent, and private signals are unbiased meaning that the pooled information of investors reveals the most profitable investment option. The objective of this chapter is twofold: First, to present the key results in the economic literature on social learning; and second, to suggest an alternative measure of the informational efficiency of the market which relies exclusively on observable variables. The main insights extend to situations in which investors choose endogenously the time of their investment (Chamley, 2004a) and, albeit with some qualification, to situations with flexible prices (Chamley, 2004b; Vives, 2008).

Section I.2 introduces the social learning game and defines its equilibrium. The social learning game is a dynamic game of incomplete information with a countable number of investors and two payoff-relevant states of Nature which are assumed equiprobable. Three features distinguish the various versions of the social learning game: (i) the distribution of private signals; (ii) the set of investment options; and (iii) the set of preference types. We essentially consider social learning games with homogeneous preferences and study the properties of the equilibrium outcome for different pairs of the first two features. Appendix A and B provide additional technical details.

Section I.3 discusses the properties of the equilibrium outcome in relation with the distribution of private signals and the cardinality of the set of investment options. Whatever the distribution of private signals, if investors choose from a continuum of investment options and are rewarded according to the proximity of their choice to the most profitable investment option then private signals can be perfectly inferred from choices and social learning is efficient (Lee, 1993).[1] In situations where the investment space is discrete, the equilibrium outcome is always inefficient and a herd almost surely arises. Other characteristics of the equilibrium outcome depend on the distribution of private signals. First, if the distribution of private signals is bounded and discrete then an *informational cascade* eventually occurs in which investors choose investment options which do not convey private information, i.e. social learning stops, and herd on a wrong option with positive probability (Banerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992).[2] Second, the equilibrium outcome is almost identical for a bounded but atomless distribution of private signals except that information cascades do not occur i.e. social learning is incomplete but does not stop (Chamley, 2004b, chap. 4). Third, if the distribution of private signals is unbounded then social learning is complete meaning that society "learns the truth" and the most profitable investment option is chosen asymptotically (Smith and Sørensen, 2000; Acemoglu, Munther, Lobel, and Ozdaglar, 2010).

The material in Section I.3 makes apparent that the economic literature on social learning has carefully studied how the long-run properties of the equilibrium outcome vary with the distribution of private signals (in the most relevant case of a discrete set of investment options). Though these theoretical analyzes have greatly improved our understanding of rational herding, they seem unsuited for the development of a positive model of social learning and this for two reasons. First, we agree with Gale (1996) that the emphasis on the long-run learning outcome is largely misplaced since, from a social welfare point of view, it matters little whether incorrect herds arise or whether information is fully revealed but extremely slowly. Both phenomena are manifestations of the *self-defeating property* and both phenomena result from the presence of *informational externalities* (Vives, 1996). Second, the distribution of private signals is of limited use for guiding the identification of correct herds since it is unlikely to be observable. As an attempt to address these concerns, Section I.4 introduces new measures of uniformity and fragility to discuss the medium-run properties of the equilibrium outcome. We find that less uniform investment decisions lead to higher welfare levels of the equilibrium outcome but lower fragility of rational herds.

[1] This result is not robust. Vives (1993) shows that social learning is inefficient when investors choose from a continuum of investment options but the observation of investment decisions is noisy.

[2] Though investors choose from a continuum of investment options in Banerjee (1992), the model shares the properties of a discrete choice model due to degenerate payoffs.

Section I.5 concludes and Appendix C contains omitted proofs.

## I.2. The Game-Theoretic Framework

In this section we adopt a game-theoretical approach to social learning since a similar approach has been pursued in the economic literature. For the sake of exposition, investors (resp. investment options) are referred to as players (resp. actions).

### I.2.1. Players, Payoff-Relevant States of Nature and Private Beliefs

Players $i = 1, 2, \dots, n$ sequentially make decisions in a preordained order encumbered solely by uncertainty about the payoff-relevant state of Nature (simply state of Nature).

We take as given a background probability space $(\Omega, \mathfrak{F}, P)$ underlying the definition of all random objects (the construction of such a probability space is discussed in Appendix I.A). Throughout the paper we denote random variables by letters endorsed with a tilde and their realizations without the tilde. Vectors are denoted by bold variables.

The *state of Nature* is given by $\tilde{\theta}$ distributed on the (finite) state space $\Theta$ according to the prior $\mathbf{p}$. In line with most of the economic literature, we assume that $\Theta = \{0, 1\}$ and $\mathbf{p} = (1/2, 1/2)$ (the extension to an arbitrary finite number of states comes at significant algebraic cost without delivering additional insights).

#### Private Information

Every player $i \in \{1, \dots, n\}$ is endowed with private information about the state of Nature. This may for instance take the form of a sequence of random i.i.d. *private signals* $(\tilde{s}_i)_{i=1}^n$ distributed on some signal space $S$ according to conditional probability measures $(\mu_\theta)_{\theta \in \Theta}$ with common support. Given her realized signal $s_i$, player $i$ updates her assessment of the distribution of $\tilde{\theta}$ via

$$b(s_i, \emptyset) = Pr\left(\tilde{\theta} = 1 \mid \tilde{}_i = s_i\right) = \frac{\mu_1(s_i)}{\mu_1(s_i) + \mu_0(s_i)}.$$

We refer to $b(s_i, \emptyset)$ as the player's (realized) *private belief* given signal $s_i$. Its random counterpart $b(\tilde{s}_i, \emptyset)$ on $[0, 1]$ is straightforwardly defined and permits the representation of any source of private information directly via (random) private beliefs. Accordingly, we represent the distribution of private beliefs conditional on realization $\theta$ of $\tilde{\theta}$ via the cumulative distribution function (CDF) $G_\theta$. $G_0$ and $G_1$ are mutually absolutely continuous to one another. Common knowledge of the distribution of private beliefs implies that

$$\frac{dG_0}{dG_1}(s) = \frac{1 - s}{s} \tag{I.1}$$

where $dG_0/dG_1$ denotes the positive finite Radon-Nikodym derivative of the distributions. This *no introspection condition* (Smith and Sørensen, 2000) asserts that a player which uses the differences of likelihoods of her private belief to update her prior in a Bayesian fashion learns nothing new. Useful consequences of the no introspection condition are derived in Appendix I.B.

We denote by $[\underline{b}, \overline{b}] \subseteq [0, 1]$ the convex hull of the common support of $G_0$ and $G_1$, $[\underline{b}, \overline{b}] = co(supp(G))$. The following distinctions are essential in economic models of social learning.

**Definition I.1.** *We call the private beliefs*

(i) ***(Un)Bounded*** *if* $\underline{b} > 0$ *(*$\underline{b} = 0$*) and* $\overline{b} < 1$ *(*$\overline{b} = 1$*); and*

(ii) ***Atomic (Atomless)*** *if* $b(\tilde{s}_1, \emptyset)$ *is a discrete (continuous) random variable.*[3]

In the case of atomless private beliefs we denote the density of $G_\theta$ by $g_\theta$.

### I.2.2. Actions, Preferences, and Payoffs

Each player makes one choice from a finite *set of actions* $A = \{1, 2, \ldots, m\}$ where $m \geq 2$ (assuming a single set of actions is not crucial to the analysis). We allow for heterogeneity in preferences in the following way. Let $T = \{t : A \times \Theta \to \mathbb{R}\}$. Player $i$'s *preference type* $t_i$ is the realization of a random variable on $T$ distributed according to the distribution $\pi \in \Delta(T)$. A player's preference type is private information and straightforwardly defines her *utility function* via $u(a, \theta, t) = t(a, \theta)$. Furthermore, preference types are realized independently from private beliefs and the state of Nature so that they do not convey information about the state. We add the following regularity condition on $\pi$: For each preference type $t \in T$, let $A(t)$ denote the set of actions which are optimal for type $t$ at some posterior belief, i.e.

$$A(t) = \left\{ a \in A : \begin{array}{c} \exists R(a \mid t) \subseteq [0,1] \text{ s.t. } (1-r)\, t(a,0) \;+\; r\, t(a,1) \;>\; (1-r)\, t(a',0) \;+\; r\, t(a',1) \\ \text{for each } r \in R(a \mid t) \text{ and } a' \in A \end{array} \right\}.$$

Call $t \in T$ a *rational type* if $|A(t)| \geq 2$, otherwise call $t$ a *noise type*. Finally, we require that $\pi$ assigns strictly positive probability to a set of rational types.[4]

For each $t \in T$ assume w.l.o.g. that $A(t) = \{a^1(t), \ldots, a^{m(t)}(t)\}$ where $m(t) = |A(t)|$ is sorted according to payoffs in state $\theta = 1$, that is $t(a^1(t), 1) < t(a^2(t), 1) < \ldots < t(a^{m(t)}(t), 1)$.[5] The following lemma further characterizes $A(t)$ for rational types $t$ (see Smith and Sørensen (2000, Lemma 1)).

**Lemma I.1.** *Fix rational type t and define for each* $1 \leq j < k \leq m(t)$

$$r(j, k \mid t) = \frac{u(a^j(t), 0, t) - u(a^k(t), 0, t)}{u(a^j(t), 0, t) - u(a^k(t), 0, t) \;+\; u(a^k(t), 1, t) - u(a^j(t), 1, t)}.$$

*It holds that* $r(k-1, k \mid t) \geq r(j, k \mid t)$ *and* $r(j, j+1 \mid t) \leq r(j, k \mid t)$ *for each* $1 \leq j < k \leq m(t)$ *and*

$$a^k(t) \in \arg\max_{a' \in A(t)} \left[ r\, u(a', 1, t) \;+\; (1-r)\, u(a', 0, t) \right] \;\Leftrightarrow\; r(k-1, k \mid t) < r < r(k, k+1 \mid t).$$

*Proof.* See appendix. □

Finally we use this result to define for each $a \in A$

$$\begin{aligned} \underline{r}(a \mid t) &= \begin{cases} r(j-1, j \mid t) & \text{if } a = a^j(t) \in A(t) \\ 1 & \text{else} \end{cases}, \\ \overline{r}(a \mid t) &= \begin{cases} r(j, j+1 \mid t) & \text{if } a = a^j(t) \in A(t) \\ 0 & \text{else} \end{cases}. \end{aligned} \tag{I.2}$$

[3] See for instance Athreya and Lahiri (2006, Chapters 2 and 4).

[4] Both the finite type space of Smith and Sørensen (2000) and the type space with private and common values of Goeree, Palfrey, and Rogers (2006) are special cases of our preference type space.

[5] We use superscripts to distinguish type $t$'s actions from action $a_i$ chosen in period $i$.

In the following, we often restrict ourselves to a single rational type $t$ satisfying $A(t) = A$ in which case we omit the type-dependence.

### I.2.3. Observations

We assume that every player observes the complete sequence of previously chosen actions (more general observation structures require different techniques, see Smith and Sørensen, 2008b; Acemoglu, Munther, Lobel, and Ozdaglar, 2010). Hence, player $i \in \{1, \ldots, n\}$ observes a (realized) *history* $h_i \in H_i = A^{i-1}$. Since randomly determined preference types and private beliefs have a crucial influence on a player's choice, $h_i$ is the realization of a random variable $\tilde{h}_i$ on $H_i$ whose distribution we denote by (conditional) probabilities $Pr(h_i \mid \tilde{\theta} = \theta)$. Obviously, $\tilde{h}_1 \equiv \emptyset$. Furthermore we let $H = \bigcup_{i=1}^{n} H_i$ denote the set of all histories.

### I.2.4. Strategies, Beliefs, and Equilibria

Let $\mathscr{G} = \langle n, H, \Theta, \mathbf{p}, (G_0, G_1), T, \pi, u \rangle$ denote the *social learning game with exogeneous timing and complete observations* (simply social learning game).

The sequence of moves in $\mathscr{G}$ is as follows. Nature moves first and determines (i) the realized state of Nature, (ii) the sequence of private beliefs, and (iii) the sequence of preference types. Second, players move according to the exogenously determined order, w.l.o.g. player 1 decides first, player 2 decides second and so forth. Player $i$'s information set consists of her realized preference type $t_i$, her realized private belief $b(s_i, \emptyset)$ and the realized history of previous actions $h_i$. We assume that the structure of the game and the composition of the players' information sets are commonly known among players.

The behavior of players in $\mathscr{G}$ is captured by *behavioral strategies* $\sigma_i : [\underline{b}, \bar{b}] \times H_i \times T \rightarrow \Delta(A)$ such that

$$\sigma_i(a \mid s_i, h_i, t_i) = Pr\left(\tilde{a}_i = a \mid b(\tilde{s}_i, \emptyset) = b(s_i, \emptyset), \tilde{h}_i = h_i, \tilde{t}_i = t_i\right).$$

Let $\Sigma_i$ denote the space of behavioral strategies available to player $i$. Furthermore we denote by $\boldsymbol{\sigma} = (\sigma_i)_{i=1}^{n} \in \Sigma$ a *profile of behavioral strategies* where $\Sigma = \bigtimes_{i=1}^{n} \Sigma_i$. The utility function $u$ can be straightforwardly extended to $\Sigma$ via

$$\begin{aligned} U(\sigma_i, \boldsymbol{\sigma}_{-i}) &= \int_{t_i \in T} \pi(dt_i) \sum_{\theta \in \Theta} \frac{1}{2} \int_{\underline{b}}^{\bar{b}} dG_\theta(s_i) \sum_{h_i \in H_i} Pr(h_i \mid \tilde{\theta} = \theta, \boldsymbol{\sigma}) \sum_{a \in A} \sigma(a \mid s_i, h_i, t_i)\, u(a, \theta, t_i) \\ &= \frac{1}{2} \int_{t \in T} \int_{\underline{b}}^{\bar{b}} \sum_{h_i \in H_i} \sum_{a \in A} \sigma(a \mid s_i, h_i, t_i)\, U(a \mid s_i, h_i, t_i, \boldsymbol{\sigma})\, \pi(dt) \qquad \text{(I.3)} \end{aligned}$$

where

$$U(a \mid s_i, h_i, t_i, \boldsymbol{\sigma}) = \left[Pr(h_i \mid \tilde{\theta} = 1, \boldsymbol{\sigma})\, dG_1(s_i)\, u(a, 1, t_i) \; + \; Pr(h_i \mid \tilde{\theta} = 0, \boldsymbol{\sigma})\, dG_0(s_i)\, u(a, 0, t_i)\right]. \qquad \text{(I.4)}$$

In this chapter we assume that players are endowed with instrumental and cognitive rationality in a strong sense. Representing knowledge by probabilities, cognitive rationality means that players form *beliefs* $b : [\underline{b}, \overline{b}] \times H \to [0,1]$ where

$$b(s,h) = Pr\left(\tilde{\theta} = 1 \mid b(\tilde{s}, \emptyset) = s, \tilde{h} = h\right)$$

about the realized state of Nature by using the laws of conditional probability. Instrumental rationality means that once a belief is formed a decision is made which maximizes expected payoff . We assume furthermore that players' degree of rationality is common knowledge. In a nutshell players are Bayesian rational and Bayesian rationality is commonly known. Proposition I.1 characterizes the set of rationalizables outcomes (Bernheim, 1984; Pearce, 1984) of the social learning game.

**Proposition I.1.** *Any rationalizable outcome of the game $\mathcal{G}$ is characterized by a profile of behavioral strategies $\boldsymbol{\sigma}^*$ and a system $\mathbf{b}^*$ of beliefs such that for each $i = 1, \ldots, n$, each rational $t_i \in supp(\pi)$, each $b(s_i, \emptyset) \in [\underline{b}, \overline{b}]$, and each $h_i \in H_i$ it holds*

***(i) Bayes' Rule:*** $b_i(s_i, h_i) = \frac{b(s_i,\emptyset)\ Pr\left(h_i \mid \tilde{\theta}=1, \sigma^*\right)}{b(s_i,\emptyset)\ Pr\left(h_i \mid \tilde{\theta}=1,\sigma^*\right) + (1-b(s_i,\emptyset))\ Pr\left(h_i \mid \tilde{\theta}=0,\sigma^*\right)}$ *provided* $Pr\left(h_i \mid \tilde{\theta} = \theta, \boldsymbol{\sigma}^*\right) > 0$ *for some* $\theta \in \Theta$ *where* $Pr\left(h_i \mid \tilde{\theta} = \theta, \boldsymbol{\sigma}^*\right) = \prod_{j<i} \int_{\underline{b}}^{\overline{b}} \sigma_j^*(a_j \mid s_j, h_j)\, dG_\theta(s_j)$

***(ii) Sequential Rationality:***

$$\sigma_i\left(a|s_i, h_i, t_i\right) = \begin{cases} 1 & if\ \underline{r}(a \mid t_i) < b(s_i, h_i) < \overline{r}(a, t_i) \\ 0 & if\ b(s_i, h_i) \notin \left[\underline{r}(a \mid t_i), \overline{r}(a \mid t_i)\right] \end{cases}.$$

*Furthermore each noise type $t_i \in supp(\pi)$ trivially chooses the single $a \in A(t_i)$ at each history $h_i$ and private belief $b(s_i, \emptyset)$. Generically the rationalizable outcome of the social learning game is unique.*

*Proof.* See appendix. □

Note that rationalizability does not determine a player's strategy whenever she is indifferent between two actions, i.e. whenever $b(s_i, h_i) = \underline{r}(a \mid t_i)$ $(\overline{r}(a \mid t_i))$ for some $a \in A$. In this case, a *tie-breaking rule* that unambiguously makes a (possibly randomized) choice at such posteriors has to be assumed. However, the occurrence of ties is a non-generic property of the social learning game.[6] In the few simplified versions of the social learning game where we cannot avoid specifying a tie-breaking rule we follow Bikhchandani, Hirshleifer, and Welch (1992) and assume that a player when indifferent will randomize uniformly.

Absent ties the rationalizable outcome is unique. However, if private beliefs are bounded then certain histories occur with probability zero. In such settings, there exist infinitely many rationalizable strategy profiles since expected payoff is zero for any action at such histories. With unbounded private beliefs only a unique rationalizable strategy profile exists.

Finally, since the rationalizable outcome of the social learning game is unique it coincides with the unique Bayesian equilibrium outcome. For convenience we will stick with the equilibrium notion henceforth.

[6]Following Smith and Sørensen (2000), we say that a property is generic if it holds for an open and dense subset of parameters of the game.

# I.3. Key Results on Bayes-Rational Social Learning

## I.3.1. Learning

Among others, models of rational herding have investigated whether the decentralized information held by the players is aggregated through the process of social learning. Following Chamley (2004b), the process of social learning can be defined as a process of observation, learning at the individual level (simply individual learning) and action choices revealing something about the individual learning outcome. Individual learning corresponds to *Bayesian learning* (see for instance Cyert and DeGroot (1970, 1974) and Kalai and Lehrer (1993)) defined in the following way.

**Definition I.2.** *Individual learning is* ***Bayesian*** *if (i) players have a prior on the set of states of Nature; (ii) they make an observation that is informative about the state of Nature (in the sense of having different conditional likelihoods); and (iii) they form a posterior by combining this observation with their prior using Bayes' rule.*

In the social learning game, a player's prior is given by her realized private belief while the history of previously chosen actions constitutes her observation. The history is informative about the state of Nature provided it occurs with different likelihoods under different states of Nature. A necessary condition is that the decision of at least one of the player's predecessors depends in a non-trivial way on private information. Formally, there must exist $j < i$ and rational preference type $t \in supp(\pi)$ s.t. either $b(\underline{s}, h_j) < \underline{r}(a_j \mid t)$ or $b(\overline{s}, h_j) > r(a_j \mid t)$ where

$$b(\underline{s}, h) = \frac{\underline{b}\, Pr(h \mid \tilde{\theta} = 1)}{\underline{b}\, Pr(h \mid \tilde{\theta} = 1) + \left(1 - \underline{b}\right) Pr(h \mid \tilde{\theta} = 0)}$$

and analogously for $b(\overline{s}, h)$. Equivalently

$$\begin{aligned} \underline{b} \quad &< \quad \underline{s}(a_j \mid t, h_j) = \frac{\underline{r}(a_j \mid t)\, Pr(h_j \mid \tilde{\theta} = 0)}{\underline{r}(a_j \mid t)\, Pr(h_j \mid \tilde{\theta} = 0) + \left[1 - \underline{r}(a_j \mid t)\right] Pr(h_j \mid \tilde{\theta} = 1)} \\ \text{or } \overline{b} \quad &> \quad \overline{s}(a_j \mid t, h_j) = \frac{\overline{r}(a_j \mid t)\, Pr(h_j \mid \tilde{\theta} = 0)}{\overline{r}(a_j \mid t)\, Pr(h_j \mid \tilde{\theta} = 0) + \left[1 - \overline{r}(a_j \mid t)\right] Pr(h_j \mid \tilde{\theta} = 1)}. \end{aligned}$$

We refer to the expressions $\underline{s}(a \mid t, h)$ and $\overline{s}(a \mid t, h)$ as type $t$'s *private belief thresholds* for action $a$ at history $h$. Clearly, rational type $t$ satisfies $\underline{s}(a^j(t) \mid t, h) > 0$ for any $1 < j \leq m(t)$ and any history $h$ such that $Pr\left(h \mid \tilde{\theta} = 0\right) > 0$. Equivalently $\overline{s}(a^j(t) \mid t, h) < 1$ for any $1 \leq j < m(t)$ and any $h$ such that $Pr\left(h \mid \tilde{\theta} = 1\right) > 0$. Accordingly, a necessary condition for choice $a$ after history $h$ to be informative is that for some subset $T \subseteq supp(\pi)$ of rational preference types it holds $[\underline{b}, \overline{b}] \not\subseteq [\underline{s}(a \mid t, h), \overline{s}(a \mid t, h)]$. Given complete observation of histories the definition of social learning is straightforward.

**Definition I.3.** *The* ***process of social learning*** *is the sequence of individual acts consisting of (i) observation of information from a common information pool, (ii) individual learning from this observation, and (iii) act of choice adding to the common information pool and causing a change in the probability distribution over the the set of states conditional on the information source.*

The common information source is the history of choices. Every player observes the entire history of choices, learns individually as defined above and makes a choice. Formally, the process of social learning is captured by the evolution of *public beliefs* $b(\emptyset, h)$ respectively *public likelhood ratios* $\lambda(\emptyset, h)$ given by

$$b(\emptyset, h) \quad = \quad \frac{Pr\left(h \mid \tilde{\theta} = 1\right)}{Pr\left(h \mid \tilde{\theta} = 1\right) + Pr\left(h \mid \tilde{\theta} = 0\right)}, \tag{I.5}$$

$$\lambda(\emptyset, h) \quad = \quad \frac{Pr(h \mid \tilde{\theta} = 1)}{Pr(h \mid \tilde{\theta} = 0)}. \tag{I.6}$$

Denote by $T_{rat} \subseteq T$ the set of rational preference types and by $T_a = \{t : A(t) = \{a\}\} \subset T$ the set of noise types for which $a$ is the single dominant action. Define

$$\psi(a \mid h, \theta) = \pi(T_a) + \int_{t \in T_{rat}} \left[G_\theta\left(\bar{s}(a \mid t, h)\right) - G_\theta\left(\underline{s}(a \mid t, h)\right)\right] \pi(dt) \tag{I.7}$$

and let $\tilde{\theta} = 0$ w.l.o.g. The process of social learning is the Markov process $\left\{\lambda(\emptyset, \tilde{h}_i)\right\}_{i=1}^{n}$ determined by transitions

$$\lambda(\emptyset, h, a) = \lambda(\emptyset, h) * \frac{\psi(a \mid h, 1)}{\psi(h \mid h, 0)} \tag{I.8}$$

with associated transition probabilities $\psi(a \mid h, 0)$. Since choices depend on $h$ only via $\lambda(\emptyset, h)$ we denote in an abuse of notation $\psi(a \mid \lambda(\emptyset, h), \theta) = \psi(a \mid h, \theta)$.

Various dynamics of the social learning process are possible.

**Definition I.4.** *The social learning process* ***converges*** *if there exists $\tilde{\lambda}_\infty$ s.t. $\lim_{i\to\infty} \lambda(\emptyset, \tilde{h}_i) = \tilde{\lambda}_\infty$ almost surely. Furthermore (i) social learning is* ***complete*** *if and only if $\tilde{\lambda}_\infty \equiv 0$; (ii) social learning* ***stops in finite time*** *if and only if the stopping time $N = \min\{i : \lambda(\emptyset, \tilde{h}_i) = \lambda(\emptyset, \tilde{h}_{i+1})\}$ is bounded.*[7]

According to this definition, learning is complete if, given the flat prior $\mu$ on $\Theta$, $(\Theta, \mu)$ is consistent in the sense of Diaconis and Freedman (1986) subject to the way private beliefs are observed in our case.[8] Furthermore the stopping of social learning is closely related to the concept of an *informational cascade* introduced by Bikhchandani, Hirshleifer, and Welch (1992). Indeed, if an informational cascade starts in some period $i$ then learning stops in the very same period.[9]

**Proposition I.2.** *The social learning process converges s.t. $supp(\tilde{\lambda}_\infty) = [0, \infty)$. Furthermore with a single rational type and no noise:*

***(i)*** *Social learning is complete, if and only if private beliefs are unbounded.*

***(ii)*** *Social learning stops in finite time if private beliefs are bounded and discrete.*

***(iii)*** *Social learning is incomplete but does not stop in finite time if private beliefs are bounded and the continuation function $\varphi(\lambda, a) = \lambda\, \psi(a \mid \lambda, 1)/\psi(a \mid \lambda, 0)$ is strictly increasing in $\lambda$ for any $a \in A$ and any $\lambda > 0$.*

[7]See for instance Athreya and Lahiri (2006), Definition 8.5.3.

[8]This distinguishes it from the notion of *learnability* (see e.g. Jackson, Kalai, and Smorodinsky, 1999) which states that given sufficiently large observations a player will forecast as if he knows the correct distribution.

[9]The convergence of the social learning process has first been stated in a general way by Smith and Sørensen (2000, Lemma 3). The proof of Bikhchandani, Hirshleifer, and Welch's (1992) Proposition 1 relies on similar arguments.

*Proof.* Convergence of the social learning process to a non-fully incorrect limit follows from the martingale property of the process (i.e. $E[\lambda(\emptyset, \tilde{h}_{i+1}) \mid \lambda(\emptyset, h_i)] = \lambda(\emptyset, h_i)$) by the Martingale Convergence Theorem (see e.g. Athreya and Lahiri, 2006, Theorem 13.3.2).

If private beliefs are unbounded it must hold that $supp(\tilde{\lambda}_\infty) = \{0\}$ since every public belief $0 < b(\emptyset, h) < 1$ can be swamped by some private belief – formally for any $a \in A$, $[\underline{s}(a \mid h), \bar{s}(a \mid h)] \subset [0,1] = [\underline{b}, \bar{b}]$ – and since $b(\emptyset, \tilde{h}_i) \to 1$ a.s. is not possible. Conversely if learning is complete, public beliefs eventually become arbitrarily close to 0. Therefore if private beliefs are bounded any private belief will eventually be swamped by the public belief such that $a = 1$ is the dominant choice for all private beliefs in which case learning stops – a contradiction. Formally, as $b(\emptyset, h_i) \to 0$, $\bar{s}(a \mid h) \to 1 > \bar{b}$.

From the argumentation above if private beliefs are bounded the process must eventually reach the absorbing set

$$C(a) = \left\{ \lambda : \underline{r}(a) < \frac{\underline{b}\,\lambda}{\underline{b}\,\lambda + 1 - \underline{b}} < \frac{\bar{b}\,\lambda}{\bar{b}\,\lambda + 1 - \bar{b}} < \bar{r}(a) \right\} = [\underline{\lambda}(a), \bar{\lambda}(a)]$$

of some action $a \in A$ where

$$\underline{\lambda}(a) = \frac{\underline{r}(a)}{1 - \underline{r}(a)} \frac{1 - \underline{b}}{\underline{b}} \qquad \text{and} \qquad \bar{\lambda}(a) = \frac{\bar{r}(a)}{1 - \bar{r}(a)} \frac{1 - \bar{b}}{\bar{b}}.$$

This has been termed the *cascade set* of action $a$. $\bar{b} < 1$ implies that $0 < \bar{\lambda}(a) < \infty$ for each $a < m$ while $\underline{b} > 0$ induces $0 < \underline{\lambda}(a) < \infty$ for each $a > 1$. If private beliefs are discrete the social learning process evolves in finite steps as long as $\lambda(\emptyset, \tilde{h}_i) \notin C(a)$ for some $a \in A$ and consequently it must reach some cascade set in a finite number of steps. On the other hand if the continuation function $\varphi(\lambda, a)$ is strictly increasing for any $\lambda > 0$, $\lambda < \underline{\lambda}(a)$ (resp. $\lambda > \bar{\lambda}(a)$) implies that $\varphi(\lambda, a) < \varphi(\underline{\lambda}(a), a) = \underline{\lambda}(a)$ (resp. $\varphi(\lambda, a) > \bar{\lambda}(a)$). Intuitively as the process gets arbitrarily close to a cascade set, its transitions become negligible. Therefore by the boundedness of beliefs learning cannot be complete, but it stops only in the limit. □

In characterizing the limits of the process we have restricted ourselves to a single (rational) preference type. The results would still hold in the presence of noise types or various rational preference types whose preferences are sufficiently similar. However, Smith and Sørensen (2000) show that the existence of multiple rational types may induce the process of social learning to reach a point where each action is equilikely in each state while no one action is chosen with probability 1. This has been termed a *confounding outcome*. While choices at a confounding outcome generically depend on private beliefs and thus reflect information uncertainty about preference types chokes off the possibility to infer this information. Smith and Sørensen (2000) provide a comprehensive characterization of confounding outcomes. On the other hand Goeree, Palfrey, and Rogers (2006) establish conditions under which heterogeneity in preferences allows for complete learning even if private beliefs are bounded.

The proposition gives three possible modes of social learning: complete learning, incomplete learning in finite time, and convergence to a possibly incorrect limit. While conditions for complete learning are well-established, a complete classification of models of incomplete learning is missing. For instance, it is unclear whether models can be constructed which allow for both: Incomplete learning in finite time and incomplete learning which converges to a possibly incorrect limit.

Though the distinction between complete and incomplete learning is theoretically compelling, it seems of minor relevance from a welfare perspective. Indeed, Chamley (2004b) has shown that even if private beliefs are unbounded and consequently the truth is revealed asymptotically, revelation takes place at a very slow rate due to the occurrence of herds (see next section). The central message with regard to the learning aspect of social learning is the *self-defeating* property. Learning by observing others is self-defeating for Bayesian rational individuals because the more information has been accumulated in predecessors' decisions the less weight is given to private information in the current decision which in turn leads to a lower increase in public information.

#### Social Learning with a Continuum of Actions

Lee (1993) establishes that if players choose from a continuum of (undominated) actions and are rewarded according to their closeness to the realization of the state of Nature, social learning is complete independent of the distribution of private beliefs. Intuitively a continuum of action with this payoff structure allows players to fully reveal their private information via their choice. Accordingly the social learning process is equivalent to the process of learning from an i.i.d. sequence of signals. We may capture this result by taking the limit as $m \to \infty$ s.t. $m(t) \to \infty$ for some $t$ with $\pi(\{t\}) > 0$. Take a single (rational) preference type. From Lemma I.1 as $m$ increases the partitition of the set of posteriors becomes increasingly fine. In particular $\bar{r}(1) \to 0$ and $\underline{r}(m) \to 1$. From the definition of cascade sets this implies that $C(1) \to \{0\}$ and $C(m) \to \{1\}$ while obviously all interior cascade sets vanish. Therefore the limit of the social learning process satisfies $supp(\tilde{\lambda}_\infty) = C(1) \cup C(m) \to \{0, 1\}$. Since fully incorrect learning a.s. is not possible, social learning must be complete. Notice however that for any finite $m$ cascades eventually arise. Yet, for any $\lambda(\emptyset, h_i)$ there exists $M \in \mathbb{N}$ such that for each $m > M$, $\lambda(\emptyset, h_i)$ is not part of any cascade set.

### I.3.2. Uniformity

A second crucial aspect of social learning is the uniformity of players' actions. Uniformity is straightforwardly defined in models with informational cascades but providing a definition of uniformity in more general models of social learning seems difficult. Here we collect the different notions which have been suggested in the literature to assess the uniformity of players' actions.[10] In Section I.4 we suggest an encompassing notion.

**Definition I.5.** ***(i)*** *An* ***informational cascade*** *arises in period i, if player i's action does not depend on her private information. That is for each $t_i \in supp(\pi)$ there exists $1 \leq j \leq m(t_i)$ such that $\lambda(\emptyset, h_i) \in C(a^j(t_i) \mid t)$ where $C(a \mid t)$ is the cascade set of action a for type t.*

***(ii)*** *A* ***limit cascade*** *takes place, if eventually the social learning process reaches a limit where no players choice depends on the private information. That is $\tilde{\lambda}_\infty \in \bigcap_{t \in T_{rat} \cap supp(\pi)} \bigcup_{a \in A(t)} C(a \mid t)$ almost surely.*

[10] The notion of an informational cascade has been introduced by Bikhchandani, Hirshleifer, and Welch (1992). Banerjee (1992) defines herd behavior. Smith and Sørensen (2000) introduce the concept of limit cascades (though a similar definition can already be found in Lee, 1993) and of action convergence. Finally, proportional herds are studied by Smith and Sørensen (2008b). A nice overview can be found in both Chamley (2004b) and Çelen and Kariv (2004).

***(iii)*** *A **herd** starts in period $i$ if after this period players of the same (rational) preference type act alike. That is for each $j, k \geq i$, $t_j = t_k$ implies $a_j = a_k$.*

***(iv)*** ***Action convergence*** *obtains, if the relative frequency with which an action is chosen is convergent. That is for each $a \in A$ the limit $\lim_{i\to\infty} |\{j \leq i : a_j = a\}|/i$ exists. A **proportionate herd** is given by $\lim_{i\to\infty}(|\{j \leq i : a_j = a\}|/i) = 1$ for some $a \in A$.*

**Proposition I.3.** *Assume a single (i.e. rational) preference type. A herd almost surely arises in finite time and a limit cascade takes place eventually. Furthermore*

***(i)*** *If private beliefs are unbounded, the herd takes place on the most profitable action a.s.*

***(ii)*** *If private beliefs are bounded, absent a cascade on $a = 1$ in period 1 with strictly positive probability the limit cascade is not on the most profitable action.*

***(iii)*** *If private beliefs are bounded and discrete, an informational cascade almost surely arises in finite time and with strictly positive probability it is not on the most profitable action. An informational cascade induces a herd.*

*Proof.* A herd almost surely arises because beliefs converge and Smith and Sørensen (2000, Proof of Theorem 3) show that the *overturning principle* applies: After player $i$ chooses action $a_i$, $\underline{r}(a_i) < b(\emptyset, h_i) < \overline{r}(a_i)$. Hence, any switch in behavior causes a discrete jump in the social learning process. This cannot happen infinitely often, as it would prevent convergence. A more direct proof is given by Chamley (2004b) who shows that the probability that a herd has not started in period $i$ tends to zero like $1/i$ (Theorem 4.3). A similar reasoning applies to the arisal of limit cascades: Smith and Sørensen (2000, Theorem 1(a)) show that if $\lambda(\emptyset, h)$ lies outside a cascade set any action arising with strictly positive probability is informative. Since for any $\lambda \in supp(\tilde{\lambda}_\infty)$ and any action $a \in A$ it must hold that either $\psi(a \mid \lambda, \theta) = 0$ for each $\theta \in \Theta$ or $\lambda = \lambda\, \psi(a \mid \lambda, 1)/\psi(a \mid \lambda, 0)$, it must be that $\psi(a \mid \lambda, \theta) = 1$ for each $\theta \in \Theta$ for some $a \in A$, i.e. a cascade on some action $a$ must arise in the limit.

The rest of the results follow straightforwardly from these results and Proposition I.2: With unbounded private beliefs only the extreme cascade sets are non-empty and satisfy $C(1) = \{0\}$ and $C(m) = \{1\}$. Since $\tilde{\lambda}_\infty < \infty$ a.s., the herd must arise on $a = 1$. With bounded private beliefs since by the Dominated Convergence Theorem the mean of $\lambda(\emptyset, \tilde{h}_i)$ is preserved in the limit, $\lambda(\emptyset, h_1) > \overline{\lambda}(1)$ and $supp(\tilde{\lambda}_\infty) \subset [0, \overline{\lambda}(1)]$ a.s. are not possible. Hence, the cascade arises not on the most profitable action with strictly positive probability. Finally, if private beliefs are bounded and discrete the social learning process evolves in discrete steps and jumps into a cascade set within a finite number of steps. This is Proposition 1 of Bikhchandani, Hirshleifer, and Welch (1992). Obviously, the signals of any finite sequence of players may favor an action other than the most profitable one and consequently induce a cascade on this less profitable action. Furthermore in an informational cascade every player's action does only depend on the public belief. Since the latter does not change, every player must choose the same action. □

A straightforward implication of this theorem is that uniform behavior is error-prone and idiosyncratic if, and only if private beliefs are bounded. For in this case a cascade may arise with strictly positive probability on a less profitable action and clearly which cascade set is reached depends on the course of actions.

The Theorem is restricted to the case of a single rational preference type and no noise types. Intuitively, heterogeneity in preferences and (strong notions of) uniform behavior are at odds. In particular with noise types or multiple rational types the overturning principle fails. If a sequence of similar choices is disrupted by a different action, this will be attributed to differences in preferences rather than a strongly opposed private information. Hence, convergence of the social learning process does not necessarily induce a herd.[11]

Yet, if only a single rational type exists limit cascades still arise eventually and results (i), (ii), and (iii) of the proposition hold. With multiple rational types however even limit cascades need not occur almost surely due to the possibility of confounding outcomes at which at least two (generically exactly two) different actions are chosen infinitely often. Exceptions are the cases of bounded and discrete private beliefs, and unbounded private beliefs with $m > 3$ actions for which confounding outcomes are nongeneric(Smith and Sørensen, 2000, Theorem 2(d),(e)). On the other hand for atomless private beliefs and $m \geq 2$ actions only action convergence is guaranteed to obtain almost surely (Smith and Sørensen, 2000, Corollary to Lemma 2).

A common property of all notions of uniformity studied in the literature is their focus on limit behavior. Hence, while their distinction is clear from a mathematical point of view, it is dubious from a practical perspective. Indeed whether a sequence of similar decisions reflects an informational cascade, a herd or some other temporary phenomenon may not be observable. Still any clustering of decisions is likely to have the same detrimental effects on the accumulation of information emphasized at the end of the previous section.

### I.3.3. Fragility

In their seminal paper, Bikhchandani, Hirshleifer, and Welch (1992, p. 993) note that "mass behavior is often fragile in the sense that small shocks can frequently lead to large shifts in behavior." Their model of informational cascades provides an explanation for the fragility of uniform behavior: Once an informational cascade occurs, social learning stops and since cascades arise fast, the process of social learning accumulates little information. Accordingly, only a small amount of information needs to be injected by an external source in order to radically shift behavior.

The concept of fragility has received considerably less attention in subsequent studies on social learning. Some authors went as far as identifying fragility with the overturning principle which asserts that a single individual may overturn any given sequence of similar choices provided her private information is sufficiently strong (see for instance Gale, 1996). We do not consider this an appropriate notion of fragility since the shift of behavior arises endogenously according to the overturning principle while it has to be induced *exogenously*.

In Section I.4, we argue that from a welfare perspective fragility is an important property of the equilibrium outcome and we suggest a new measure of fragility.

[11] Though to our knowledge the possibility that the process converges while more than a single action is chosen an infinite number of times has not been established formally.

### Fads

Bikhchandani, Hirshleifer, and Welch (1992) consider a second explanation for seemingly whimsical shifts in behavior, a possible change in the state of Nature. Indeed they show that behavior may change even if the state does not and that the probability with which a change of behavior occurs may be larger than the probability of the change of the state itself. Therefore this extension of their model provides an explanation of why fads might occur.

Assume that after player $i-1$'s decision the state of Nature changes with probability $\epsilon > 0$. Then the public belief in period $i$ is given by

$$b'(\emptyset, h_i) = (1-\epsilon)\, b(\emptyset, h_i) \; + \; \epsilon \; [1 - b(\emptyset, h_i)]$$

It can easily be shown that $b'(\emptyset, h_i) > (<) b(\emptyset, h_i)$ if $b(\emptyset, h_i) < (>) 1/2$. Hence, the effect is similar to the introduction of contrarian information. Our measure of fragility introduced in Section I.4 corresponds to the average contrarian information necessary to introduce a change in behavior. It therefore encompasses the idea of fads.

The idea that the state of Nature may change has been carefully studied by Moscarini, Ottaviani, and Smith (1998). Changes of the state are possible after each player's decision. They show that in this case any cascade must be temporary. Moreover depending on the probabilities of the changes new phenomena might arise. In particular if the probabilies are very high behavior keeps constantly changing between the two most profitable actions $a = 1$ and $a = m$ (in their case $m = 2$). On the other hand if changes occur with intermediate probabilities no social learning is possible since any information from previous periods is outdated after the state has possibly changed. Therefore while this model might be interesting in its own right we do not think that environments which exhibit a constant change of fundamentals are the ones where social learning is most relevant.

## I.3.4. Social Welfare

In the social learning game players do not take into account the informational effects of their decisions on successors. In other words, learning by observing others involves informational externalities which leads to a slow accumulation of information and the occurrence of detrimental phenomena such as herds or confounding outcomes. Consequently, the outcome of Bayes-rational social learning is inefficient. Alternatively, the aggregation of all the decentralized information would enable players to become sufficiently convinced about the true realized state of Nature and to choose the most profitable action. To discuss the inefficiencies of rational herding, we introduce a social welfare function.

**Definition I.6.** *The social welfare function is given by*

$$W_n = E\left[\sum_{i=1}^{n} \delta^{i-1}\, u\left(\tilde{a}_i, \tilde{\theta}, \tilde{t}_i\right)\right] \tag{I.9}$$

*where $0 < \delta < 1$ is a discount factor and the expectation is taken over the possible states of Nature and all sequences of realized private beliefs and preference types.*

In order to assess the inefficiency of Bayesian rational social learning a benchmark needs to be derived. It is given by the *(constrained) social optimum* – the maximum of the social welfare function under the restriction that players cannot directly reveal their private information but have to communicate it through their actions. That is players $i = 1, \ldots, n$ are restricted to choosing strategies $\sigma_i : [\underline{b}, \overline{b}] \times H_i \times T \to \Delta(A)$. The social optimum can be interpreted as the outcome of the maximization of a social planner. However, Smith and Sørensen (2008a, Lemma 1) show that any social optimum is given by a *team equilibrium* (Radner, 1962), the equilibrium of the distorted social learning game $\mathscr{G}' = \langle n, H, \Theta, \mathbf{p}, (G_0, G_1), T, \pi, u' \rangle$ given by

$$u_i'(a, \delta) = E\left[(1-\delta) \sum_{j=i}^{\infty} \delta^{j-i}\, u\left(\tilde{a}_j, \tilde{\theta}, \tilde{t}_j\right)\right].$$

In the social learning game with symmetric binary signals, Chamley (2004b, Section 4.5) establishes that in the team equilibrium cascades start at more extreme beliefs than in the Bayesian equilibrium.[12] Smith and Sørensen (2008a) extend this result to the general social learning game with a single preference type. They show that the team equilibrium requires players to lean against the public belief, i.e. rely more on their private information. In summary, if players internalize the information externality their decisions will rely more strongly on private information thereby revealing more information to others.

Existing theoretical studies on social learning do not quantify the inefficiencies of the equilibrium outcome nor do they relate these inefficiencies to observable variables. The medium-run properties of the equilibrium outcome are the key to assess more accurately social welfare and they are discussed in the next section.

## I.4. Medium-Run Properties of Bayes-Rational Social Learning

We take the perspective of a social planner who aims at increasing equilibrium welfare by releasing public information. Since the communication of information to the market is costly, the planner's policy is only warranted when equilibrium inefficiencies are sufficiently large. The welfare level of the equilibrium outcome should therefore be assessed. Moreover, welfare levels have to be assessed by relying on observable aspects of the social learning environment. A natural candidate is the sequence of actions.

Indeed, if the speed of convergence of players' actions differs between two social learning environments which are undistinguishable from the outside then this intuitively suggests that private belief distributions are distinct. And the social learning environment which generates slow convergence is likely to have more diverse private beliefs, a good property for the aggregation of information. Thus, the degree of uniformity observed in the medium-run of the equilibrium outcome is likely to inform the planner about the richness of private beliefs. The latter has been considered by the economic literature as the main determinant of the equilibrium welfare. Note also that policy costs are positively related to the precision of the information necessary to shift behavior. Consequently, the fragility of the equilibrium outcome is another desirable property.

[12] In a similar vein, Sgroi (2002) shows that when $\delta = 1$ it is profitable for the social planner to force a subset of the players to decide simultaneously in the first period.

In this section, we investigate (i) how the degree of inefficiencies generated by Bayesian rational social learning depends on the shape of the distribution of private information, (ii) how this is reflected in the uniformity of the equilibrium outcome, and (iii) how the fragility of the equilibrium outcome is related to the two.

### I.4.1. Measures of Uniformity and Fragility

As indicated in Section I.3.2 the economic literature on social learning has defined uniformity of behavior in various ways. However, the distinction between informational cascades, herds and also long but finite sequences of similar choices is hardly relevant for our purposes. We prefer to focus on the self-defeating property according to which the weight of history induces actions to *cluster* and slows down social learning. Interestingly enough, Chamley (2004b, Chapter 3) shows that this even applies in a model with a continuum of actions where others' actions perfectly reveal their private information. Our measure of uniformity tries to capture this robust property of social learning.

**Definition I.7.** *In the social learning game $\mathcal{G}$, the **degree of clustering** for strategy profile $\sigma$ is given by*

$$\mathcal{C}_n(\sigma) = E\left[\frac{2}{n\,(n-1)}\sum_{i=2}^{n}\sum_{j=1}^{i-1}\mathbb{1}_{\{a_i=a_j\}}(\tilde{h}_{n+1})\right] \tag{I.10}$$

*where $\mathbb{1}_{\{a_i=a_j\}}$ denotes the indicator function of the set $\{a_i = a_j\} = \{h_{n+1} : h_{n+1}(i) = h_{n+1}(j)\}$ and the expectation is taken over (complete) histories $h_{n+1}$, states of Nature $\theta \in \Theta$, and sequences $(\tilde{t}_i)_{i=1}^{n}$ of preference types.*

The degree of clustering measures the expected similarity for a finite sequence of choices. More precisely, it measures the average fraction of pairs of players which make the *same* choice. Hence, $\mathcal{C}_n(\sigma) \in [0,1]$ for each $\sigma \in \Sigma$ (for a related measure, see Gul and Lundholm, 1995).

Some properties of $\mathcal{C}_n$ are noteworthy. First, the definition expresses a very restrictive notion of uniformity. A more permissive notion would count as clustering any bundling of *similar* actions. However the latter requires a notion of similarity or distance between actions. Absent a topological action space and in the presence of preference heterogeneity no such notion naturally exists. Second, rather than defining uniform behavior in the long run, we provide a quantification of uniformity for any finite number of players. Our measure enables us to study the evolution of uniformity over time. Finally, the degree of clustering is valid independent of the distribution of private beliefs. In particular the measure does not distinguish between cascades, herds, and temporary uniformity. As will become clear later the measure focuses on the key aspect of the clustering of decisions – its detrimental effect on the accumulation of information.

To demonstrate the usefulness of the measure, we characterize uniformity in the specific model of Bikhchandani, Hirshleifer, and Welch (1992) (the proof is given in Appendix C).

**Lemma I.2.** *Assume that the quality of private beliefs equals $1/2 < q < 1$.*

***(i)*** *If all players follow private information then $\mathcal{C}_n(\sigma_{PI}) = q^2 + (1-q)^2$.*

***(ii)*** *In equilibrium, $\mathcal{C}_n(\sigma^*)$ tends to $1$ like $1/n$.*

***(ii)*** *If the first player is endowed with a private belief of higher quality ($q_1 > q$) then $\mathscr{C}_n(\sigma^*) = 1$.*

Though notions of uniformity are abundant in the literature, the complementary notion of fragility has rarely been discussed except in the seminal paper of Bikhchandani, Hirshleifer, and Welch (1992). Our measure of fragility is inspired by this seminal study.

**Definition I.8.** *Let*

$$\bar{q}(\mathscr{G}) = \frac{1}{2} + \frac{1}{2}\sum_{\theta\in\Theta}\int_{\underline{b}}^{\bar{b}}\left|s - \frac{1}{2}\right| dG_\theta(s)$$

*denote the **average signal precision** in the social learning game $\mathscr{G}$ and define for each $h \in H$ and each rational preference type $t \in supp(\pi)$ the following bounds*

$$\begin{aligned}
\underline{\beta}(h,t) &= \min\left\{\beta > 0 : \left(\frac{1-\bar{q}(\mathscr{G})}{\bar{q}(\mathscr{G})}\right)^{\beta} < \frac{1-\bar{b}}{\bar{b}} * \frac{1-b(\emptyset,h)}{b(\emptyset,h)} * \frac{\bar{r}\left(a^1(t)\mid t\right)}{1-\bar{r}\left(a^1(t)\mid t\right)}\right\},\\
\bar{\beta}(h,t) &= \min\left\{\beta > 0 : \left(\frac{\bar{q}(\mathscr{G})}{1-\bar{q}(\mathscr{G})}\right)^{\beta} > \frac{\underline{b}}{1-\underline{b}} * \frac{1-b(\emptyset,h)}{b(\emptyset,h)} * \frac{\underline{r}\left(a^{m(t)}(t)\mid t\right)}{1-\underline{r}\left(a^{m(t)}(t)\mid t\right)}\right\}.
\end{aligned}$$

*The **degree of robustness** for strategy profile $\boldsymbol{\sigma}$ is given by*

$$\mathscr{R}_n(\sigma) = E\left[\frac{1}{2n}\sum_{i=1}^{n}\left(\underline{\beta}(\tilde{h}_i,\tilde{t}_i) + \bar{\beta}(\tilde{h}_i,\tilde{t}_i)\right)\right] \tag{I.11}$$

*where the expectation is taken over histories, preference types, and states of Nature.*

To interpret this measure note that

$$\underline{\beta}(h\mid t) = \min\left\{\beta > 0 : \frac{\bar{b}\,b(\emptyset,h)\,(1-\bar{q})^{\beta}}{\bar{b}\,b(\emptyset,h)\,(1-\bar{q})^{\beta} + (1-\bar{b})\,(1-b(\emptyset,h))\,\bar{q}^{\beta}} < \bar{r}(a^1(t)\mid t)\right\}.$$

Hence, $\underline{\beta}(h \mid t)$ measures the minimal amount of information necessary to induce a player of type $t$ at history $h$ to choose her most preferred action under state 0 with probability one. Equivalently $\bar{\beta}(h \mid t)$ measures the minimal amount of information necessary to induce (with probability one) the action most preferred under state 1. Accordingly, $\mathscr{R}_n(\sigma)$ measures the average amount of information necessary to induce any player of any preference type to choose either of her most extreme actions. The higher $\mathscr{R}_n$, the more *robust* (the less fragile) the equilibrium outcome. Two remarks are in order. First, we measure the amount of information via the number of signals of average precision which implicitly assumes that the planner cannot acquire information more easily than the average player. Second, we assume that the planner tries to shift behavior *with probability one*. Related to this assumption is our restriction on extreme actions. Indeed, *secure actions* may be such that there always exists some extreme private belief which renders them unprofitable meaning that secure actions cannot be induced with probability one. Note however that shifting equilibrium behavior to a state where some secure action is chosen with maximal probability does require a smaller amount of information than shifting equilibrium behavior such that an extreme action is chosen with probability one. Accordingly, our measure overestimates the degree of fragility at worst which is of low concern since we are interested in comparing social learning environments.

The following result holds almost by definition.

**Lemma I.3.** *If private beliefs are unbounded then* $\mathscr{R}_n(\sigma^*) = \infty$.

*Proof.* This follows straightforwardly from $\underline{b} = 1 - \overline{b} = 0$. In words, if infinitely strong private beliefs are possible then no single action can be induced to be chosen with probability one. Unbounded private beliefs render policy intervention impossible. □

To demonstrate the usefulness of the measure, we again turn to the specific model of Bikhchandani, Hirshleifer, and Welch (1992).

**Lemma I.4.** *Assume that the quality of private beliefs equals* $1/2 < q < 1$. *The degree of fragility is given by*

$$\begin{aligned}\mathscr{R}_n(\sigma^*) &= 1 + \frac{1}{2}\frac{\log((1+q)/(2-q))}{\log(q/(1-q))} \\ &\quad - \frac{1}{n(1-q+q^2)}\frac{\log((1+q)/(2-q))}{\log(q/(1-q))} * \begin{cases} 1-(q-q^2)^{(n-1)/2}\,\frac{1+q(1-q)}{2} & \text{if } n \text{ odd} \\ 1-(q-q^2)^{n/2} & \text{if } n \text{ even} \end{cases}.\end{aligned}$$

*Proof.* See Appendix C. □

Since $(1+q)/(2-q) < q/(1-q)$, this degree is smaller than $3/2$. Any action can be induced by less then two signals of average quality.

We now discuss the relationships between the shape of private belief distributions, our measures of uniformity and fragility, and equilibrium welfare.

### I.4.2. Private Beliefs, Uniformity, Fragility, and Equilibrium Welfare

We now study the impact of private belief distributions on equilibrium welfare in various social learning games. In addition, we relate equilibrium welfare to our measures of uniformity and fragility. We restrict ourselves to the following class of discrete private belief distributions. We assume that players $i = 1, \dots, n$ receive private signals whose distribution is characterized by signal precision $q_i$ where $q_i$ is the realization of a random variable $\tilde{q}$ which has a discrete distribution on the interval $(1/2, 1)$. We represent the distribution of $\tilde{q}$ by the vector $\boldsymbol{\pi} = (q_1, Pr(q_1); \dots; q_k, Pr(q_k))$. This class of distributions allows us to easily characterize the properties of private beliefs. Vector $\boldsymbol{\pi}$ corresponds to the *information quality distribution*. Obviously, we fix the set of actions and the distribution of preference types.

The results we present have been obtained numerically using a dynamic programming procedure. In each case, a series of three figures shows the evolution of the equilibrium welfare, the degree of clustering, and the degree of fragility as a function of the number of players. The equilibrium welfare is normalized according to the maximal social welfare which is given by the benchmark where all players choose the most profitable action in each state. In each figure, different colors correspond to different social learning environments.

We start with the simplest case of two actions and a single rational preference type.

#### Two Actions: Heterogeneity is Efficiency-Enhancing

We let $A = \{0, 1\}$, $\delta = 0.9$ and fix the utility function at $u(1, \theta) = \theta - 1/2$ and $u(0, \theta) = 0$. Figure I.1 compares $\boldsymbol{\pi}_1 = (0.7, 1)$ to $\boldsymbol{\pi}_2 = \left(0.6, \frac{1}{5}; 0.7, \frac{3}{5}; 0.8, \frac{1}{5}\right)$, $\boldsymbol{\pi}_3 = \left(0.6, \frac{2}{5}; 0.7, \frac{1}{5}; 0.8, \frac{2}{5}\right)$ and

$\pi_4 = \left(0.6, \frac{1}{2};\ 0.8, \frac{1}{2}\right)$. For each private beliefs distribution, the average quality of information equals $\bar{q} = 0.7$. However, private beliefs distributions differ according to the weight they put on the boundaries of the support $[0.6, 0.8]$. We therefore investigate how equilibrium welfare, uniformity and fragility evolve as the "hill-shaped" information quality distribution ($\pi_1$) transforms into a "U-shaped" information quality distribution ($\pi_4$).

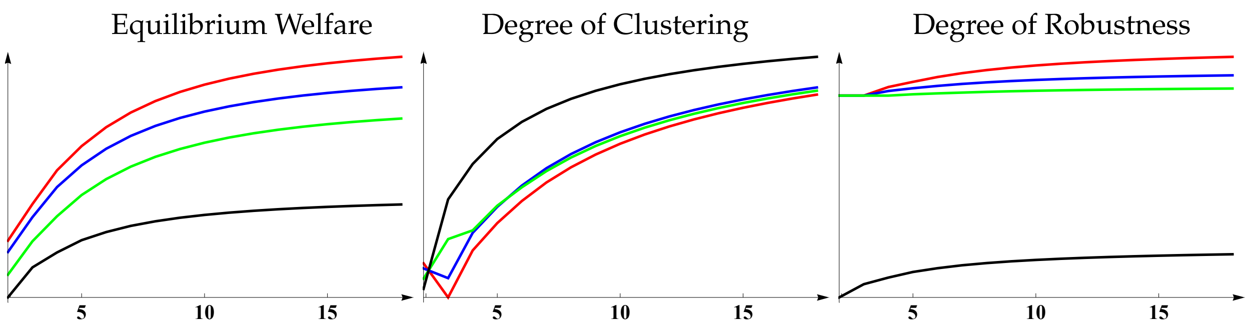


Figure I.1.: Welfare, Uniformity, and Fragility – 2 Actions and Constant Average Quality of Information. Private belief distributions are given by $\pi_1 = (0.7, 1)$ (black lines), $\pi_2 = \left(0.6, \frac{1}{5};\ 0.7, \frac{3}{5};\ 0.8, \frac{1}{5}\right)$ (green lines), $\pi_3 = \left(0.6, \frac{2}{5};\ 0.7, \frac{1}{5};\ 0.8, \frac{2}{5}\right)$ (blue lines), and $\pi_4 = \left(0.6, \frac{1}{2};\ 0.8, \frac{1}{2}\right)$ (red lines).

The highest equilibrium welfare is obtained in the setting where the private beliefs distribution puts most weight on the boundaries of the (common) support, i.e. the most U-shaped information quality distribution. Additionally, the more U-shaped the information quality distribution the less clustering and the less fragile is the equilibrium outcome. Result 1 summarizes our initial findings.

**Result 1.** *Assume that the average quality of private information remains constant across social learning games. The equilibrium achieves higher welfare and exhibits less uniformity and less fragility the more U-shaped the information quality distribution.*

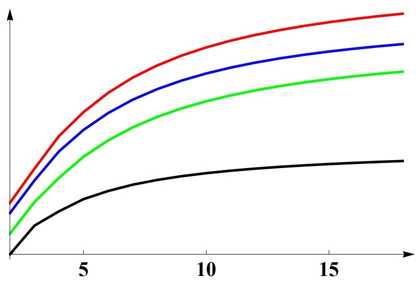


Figure I.2.: Welfare, Uniformity, and Fragility – 2 Actions and Constant Average Quality of Information (ctd). Social welfare for the same private belief distributions as in Figure I.1 given $\delta = 0.99$

We subsequently proceed to validate the robustness of this result. As a first step we establish that it does not depend on the discount factor $\delta$. Figure I.2 shows social welfare for the same group of private belief distributions if $\delta = 0.99$. Apparently the discount factor is not crucial to our basic claim. Notice that this is not at odds with the effects of the discount factor on the social optimum[13]

[13]See for instance Section 7 in Smith and Sørensen (2008a).

since we compare the effects across social learning settings.

In the example above we have kept constant the average quality of information. However, by increasing the weight on the boundaries of the support we have also increased the amount of players with the highest quality of information which may provide a different explanation for the result on equilibrium welfare. In response to this problem we consider a second group of private belief distributions given by $\pi'_1 = \left(0.6, \frac{1}{10}; 0.7, \frac{4}{10}; 0.8, \frac{1}{2}\right)$, $\pi'_2 = \left(0.6, \frac{2}{10}; 0.7, \frac{3}{10}; 0.8, \frac{1}{2}\right)$, and $\pi'_3 = \left(0.6, \frac{1}{2}; 0.8, \frac{1}{2}\right)$. For each of these distributions the fraction of players with the highest quality of information equals $\pi_{0.8} = 1/2$. Again, we investigate how equilibrium welfare, uniformity and fragility evolve as a hill-shaped information quality distribution ($\pi'_1$) transforms into a U-shaped information quality distribution ($\pi'_3$). Notice that at the same time the average quality of information *decreases*. Intuitively higher average quality of information favors equilibrium welfare. Figure I.3 summarizes our results.

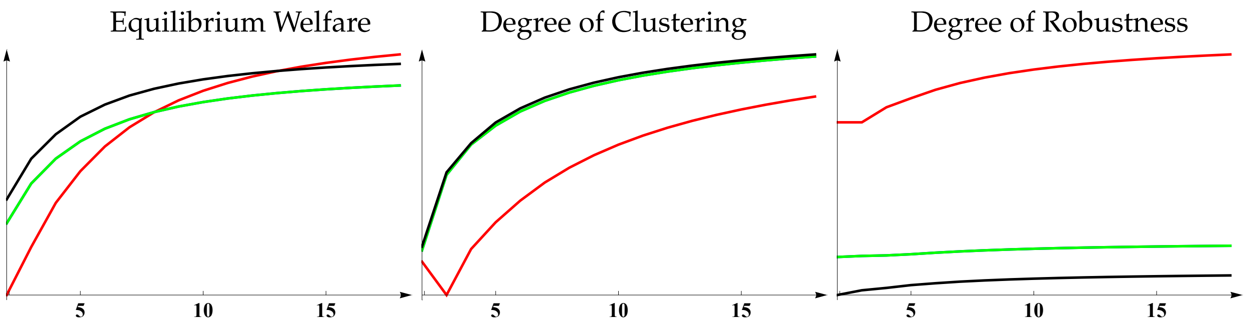


Figure I.3.: Welfare, Uniformity, and Fragility – 2 Actions with Constant Proportion of Well-Informed Individuals. Private belief distributions are given by $\pi'_1 = \left(0.6, \frac{1}{10}; 0.7, \frac{4}{10}; 0.8, \frac{1}{2}\right)$ (black lines), $\pi'_2 = \left(0.6, \frac{2}{10}; 0.7, \frac{3}{10}; 0.8, \frac{2}{5}\right)$ (green lines), and $\pi'_3 = \left(0.6, \frac{1}{2}; 0.8, \frac{1}{2}\right)$ (red lines).

As before the more U-shaped the information quality distribution the less clustering and the less fragile is the equilibrium outcome. Additionally, the highest equilibrium welfare is obtained in the setting where the private beliefs distribution puts most weight on the boundaries of the (common) support *provided n is sufficiently large*. We therefore corroborate Result 1. Result 2 summarizes this finding.

**Result 2.** *Assume that the fraction of player with the highest quality of private information remains constant across social learning games. The equilibrium exhibits less uniformity and less fragility and, provided n is sufficiently larger, achieves higher welfare the more U-shaped the information quality distribution.*

We strengthen this result even further by providing similar results for another group of distributions given by $\pi''_1 = (0.7, 1)$, $\pi''_2 = \left(0.6, \frac{1}{6}; 0.7, \frac{5}{6}\right)$, and $\pi''_3 = \left(0.6, \frac{1}{5}; 0.7, \frac{4}{5}\right)$. We therefore investigate how equilibrium welfare, uniformity and fragility evolve as the hill-shaped information quality distribution with highest average quality of information and highest fraction of players with the largest quality of information ($\pi''_1$) transforms into a U-shaped information quality distribution with lower average quality of information and lower fraction of players with the largest quality of information ($\pi''_3$).

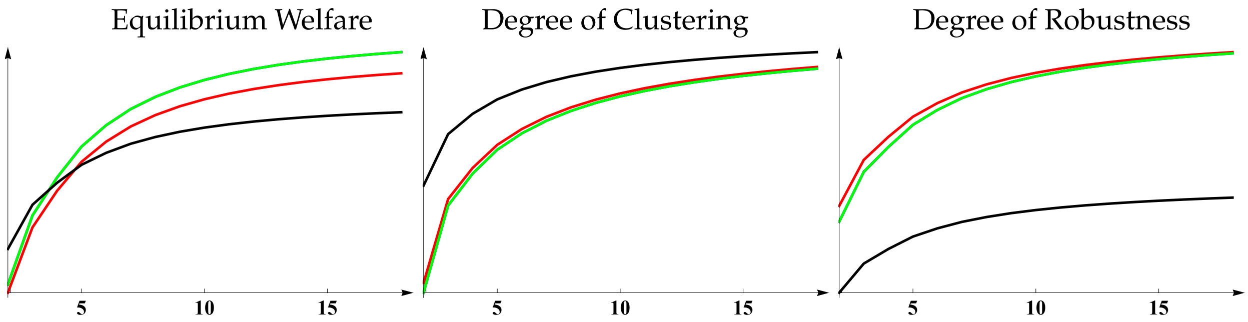


Figure I.4.: Welfare, Uniformity, and Fragility – 2 Actions with Constant Proportion of Well-Informed Individuals (ctd). Private belief distributions are given by $\pi_1'' = (0.7, 1)$ (black lines), $\pi_2'' = \left(0.6, \frac{1}{6}; \; 0.7, \frac{5}{6}\right)$ (green lines), and $\pi_3'' = \left(0.6, \frac{1}{5}; \; 0.7, \frac{4}{5}\right)$ (red lines).

Again, the highest equilibrium welfare is obtained in the setting with the most U-shaped information quality distribution (accompanied by similar effects on clustering and fragility).[14]

Having established robustness of the result for 2 actions and a single, rational preference type, we now move on to more actions and more complicated distributions of preferences.

### Noise fosters Heterogeneity

Our first aim is to extend the result to allowing for heterogeneity in preferences. In particular we introduce for each action $a = 1, 2$ a fraction $\xi_a$ of noise types for which this action is the dominant choice. Figure I.5 presents results for the cases of $\xi_a = 0.05$ (top row) and $\xi_a = 0.1$ (bottom row) and the distributions of the previous example $\pi_1''$, $\pi_2''$, and $\pi_3''$.

Surprisingly in either case highest equilibrium welfare is obtained in the setting where the private beliefs distribution puts *least* weight on the boundaries of the (common) support, i.e. the most *hill-shaped* information quality distribution.Additionally the more U-shaped the information quality distribution the *more* clustering and (for $n$ sufficiently large) the *more* fragile is the equilibrium outcome. This seems to contradict Result 1 for a setting with noise types. However, we argue that the addition of noise types is similar to a change in the distribution of information qualities. More precisely, introducing noise types is comparable to endowing a strictly positive fraction of players with uninformative signals of precision $q = 0.5$. Like noise types these players will act independently of the realized state of Nature. In particular uninformed players always follow the public belief. Hence, any decision in line with the public belief yields the same updating of beliefs regardless of whether noise types or uniformed players are present. On the contrary beliefs are updated differently in the two scenarios following a decision against the public belief. Still a reinterpretation of the results in this spirit helps explain the findings from Figure I.5: Introducing a strictly positive fraction of uninformed players changes the (common) support of the information quality distributions $\pi_1''$, $\pi_2''$, and $\pi_3''$ from $[0.6, 0.7]$ to $[0.5, 0.7]$. Therefore while with support $[0.6, 0.7]$, $\pi_1''$ is the most hill-shaped distribution, it is least hill-shaped for support $[0.5, 0.7]$.In this sense the results with noise types confirm Result 1.

[14]Notice that $\pi_2''$ still generates higher social welfare than $\pi_3''$. Since the difference in shape between these two is negligible, we attribute this to the higher average information quality of $\pi_2''$.

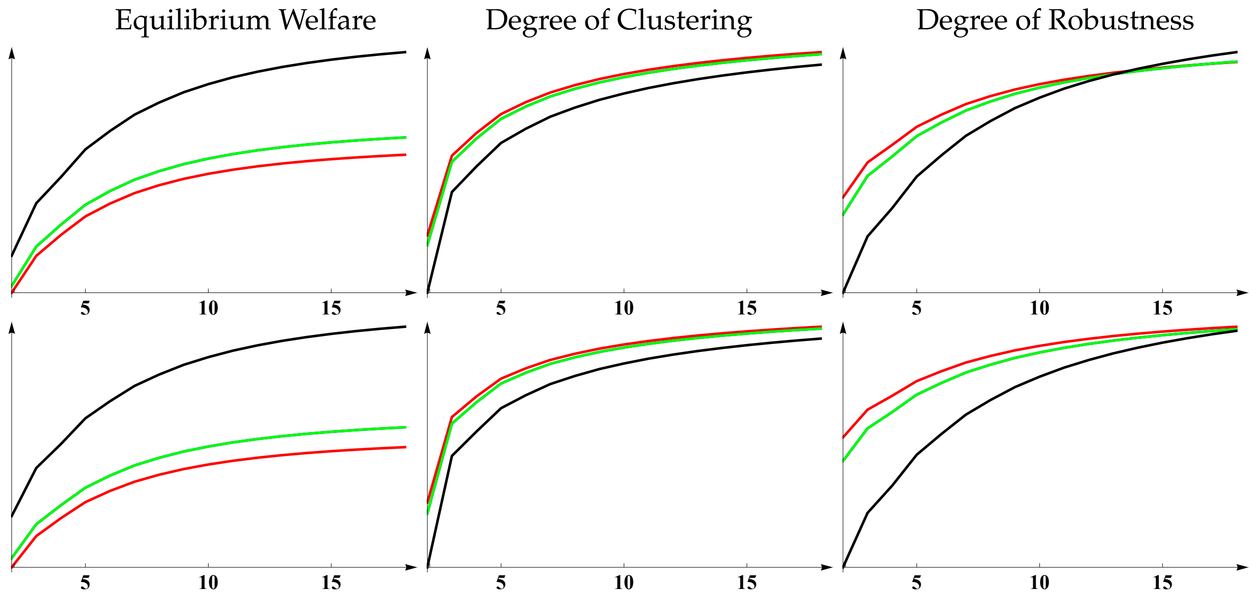


Figure I.5.: Welfare, Uniformity, and Fragility – 2 Actions with Noise. Private belief distributions are given by $\pi_1''$ (black lines), $\pi_2''$ (green lines), and $\pi_3''$ (red lines) as in Figure I.4. Top row – $\xi_a = 0.05$, bottom row – $\xi_a = 0.1$.

#### More Actions: Robustness of the Result

We finally investigate whether Result 1 extends to settings with more actions.

We start with the case of $m = 3$. We let $u(1, \theta) = 0$, $u(3, \theta) = 2\theta - 1$ and $u(2, \theta) = \theta - 1/3$. Hence, action 1 is optimal for posteriors $b(s, h) \in [0, 1/3$, the secure action 2 is optimal provided $b(s, h) \in [1/3, 2/3]$, and action 3 is optimal at $b(s, h) \in [2/3, 0]$. Consider first an information $\boldsymbol{\pi} = (q, 1)$ for some $1/2 < q < 1$. Apparently if $q < 2/3$, an informational cascade on the secure action $a = 2$ starts in the first period. Consequently, it holds $\mathscr{C}_n = 1$, $W_n/W_n^{max} = 1/6$ and $\mathscr{R}_n = 1 + \frac{\log(2)}{\log(q/(1-q))}$. Now compare this to an information quality distribution $\hat{\boldsymbol{\pi}} = \left(\underline{q}, \frac{1}{2}; \overline{q}, \frac{1}{2}\right)$ such that $q = (\underline{q} + \overline{q})/2$ and $\overline{q} > 2/3$ (which is possible provided $q > 7/12$). With strictly positive probability the first player does not choose action $a = 2$. Therefore a cascade does not arise in the first period and uniformity and fragility of behavior decrease. Figure I.6 summarizes how equilibrium welfare, uniformity and fragility evolve in the case of the U-shaped information quality distribution $\hat{\boldsymbol{pi}} = \left(0.55, \frac{1}{2}; 0.75, \frac{1}{2}\right)$ in comparison to the hill-shaped information quality distribution $(0.65, 1)$.

This simple example confirms Result 1: The highest equilibrium welfare is obtained in the setting with the U-shaped information quality distribution $\hat{\boldsymbol{\pi}}$ and this setting generates less clustering and less fragility of the social equilibrium. To establish robustness of this example we revisit our first example $(\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_4)$ for the case of 3 actions and the utility function introduced above. Figure I.7 summarizes the results establishing that Result 1 robustly generalizes to the case of $m = 3$ actions.

Our last example investigates the robustness of Result 1 in a setting with a continuum of actions. W.l.o.g. we let $A = [0, 1]$ and $u(a, \theta) = 1 - (a - \theta)^2$. Hence, a player maximizes expected utility by choosing $a(s_i, h_i) = b(s_i, h_i)$, her posterior given realized private belief $b(s_i, \emptyset)$ and realized history $h_i$. Therefore players' choices perfectly reveal their private information.

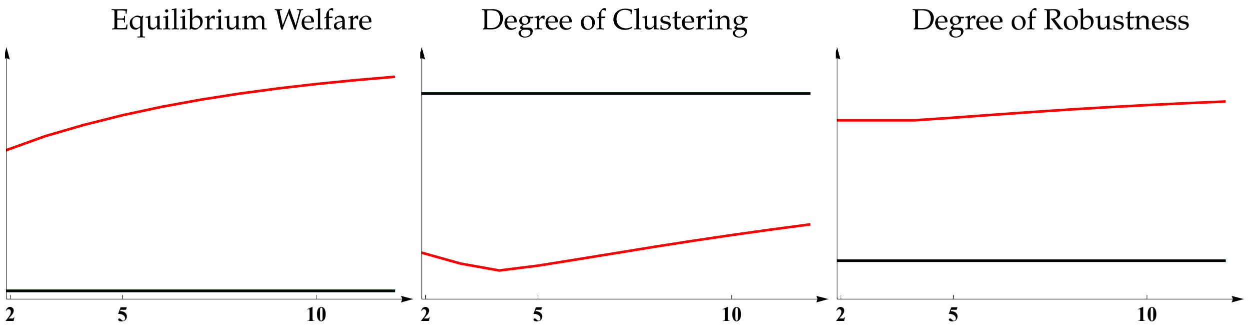


Figure I.6.: Welfare, Uniformity, and Fragility – 3 Actions and Constant Average Quality of Information. Private belief distributions are given by $\pi = (0.65, 1)$ (black lines), and $\hat{\pi} = \left(0.55, \frac{1}{2};\ 0.75, \frac{1}{2}\right)$ (red lines).

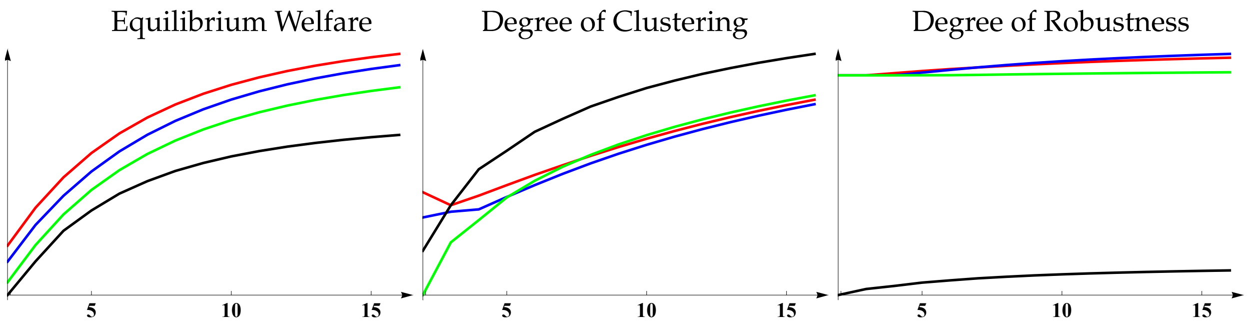


Figure I.7.: Welfare, Uniformity, and Fragility – 3 Actions and Constant Average Quality of Information. Private belief distributions are given by $\pi_1 = (0.7, 1)$ (black lines), $\pi_2 = \left(0.6, \frac{1}{5};\ 0.7, \frac{3}{5};\ 0.8, \frac{1}{5}\right)$ (green lines), $\pi_3 = \left(0.6, \frac{2}{5};\ 0.7, \frac{1}{5};\ 0.8, \frac{2}{5}\right)$ (blue lines), and $\pi_1 = \left(0.6, \frac{1}{2};\ 0.8, \frac{1}{2}\right)$ (red lines).

In order to investigate validity of Result 1 in this setting, we first need to extend our measures $\mathscr{C}_n$ and $\mathscr{F}_n$. Notice first that fragility is not well defined in the settings above since players are able to perfectly adapt their choices to their information. Thus no action can be induced from the outside. We therefore restrict ouselves to analyzing the relationship between the shape of the information quality distribution, equilibrium welfare and uniformity. Secondly, we have to admit that our previous notion of uniformity as choices of exactly the same action is not well suited to a continuous action space. However, we can amend it using a notion of distance on the set of actions. Accordingly, we let

$$\mathscr{C}_n'(\sigma) = E\left[\frac{2}{n\,(n-1)} \sum_{i=2}^{n} \sum_{j=1}^{i-1} \left[1 - \left(a_i - a_j\right)^2\right]\right] = 1 - E\left[\frac{2}{n\,(n-1)} \sum_{i=2}^{n} \sum_{j=1}^{i-1} \left(a_i - a_j\right)^2\right]. \tag{I.12}$$

Apparently on $A = [0, 1]$, $0 \leq \mathscr{C}_n' \leq 1$. Furthermore $\mathscr{C}_n$ and $\mathscr{C}_n'$ coincide on $A = \{0, 1\}$. Hence, $\mathscr{C}_n'$ is the natural extension of $\mathscr{C}_n$.

We investigate how equilibrium welfare and uniformity evolve as the information quality distribution transforms from a hill-shaped distribution to a U-shaped distribution using the information quality distributions $\pi_1 = (0.7, 1)$, $\pi_2 = \left(0.6, \frac{1}{5};\ 0.7, \frac{3}{5};\ 0.8, \frac{1}{5}\right)$, $\pi_3 = \left(0.6, \frac{2}{5};\ 0.7, \frac{1}{5};\ 0.8, \frac{2}{5}\right)$ and $\pi_4 = \left(0.6, \frac{1}{2};\ 0.8, \frac{1}{2}\right)$ or our first example. Figure I.8 shows that Result 1 extends to the setting with continuous action space.

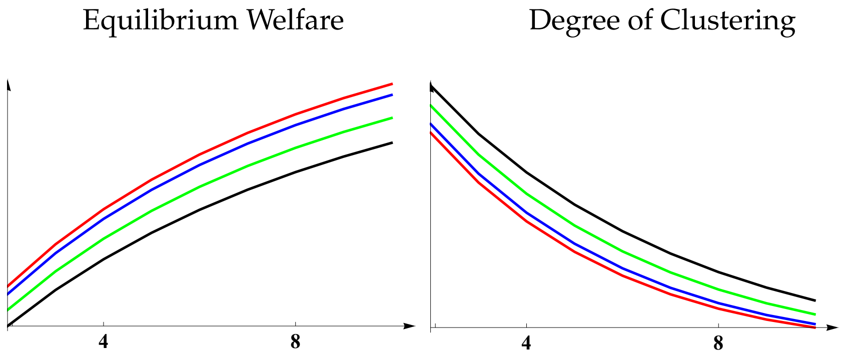


Figure I.8.: Welfare, Uniformity, and Fragility – Continuum of Actions and Constant Average Quality of Information. Private belief distributions are given by $\pi_1 = (0.7, 1)$ (black lines), $\pi_2 = \left(0.6, \frac{1}{5};\ 0.7, \frac{3}{5};\ 0.8, \frac{1}{5}\right)$ (green lines), $\pi_3 = \left(0.6, \frac{2}{5};\ 0.7, \frac{1}{5};\ 0.8, \frac{2}{5}\right)$ (blue lines), and $\pi_1 = \left(0.6, \frac{1}{2};\ 0.8, \frac{1}{2}\right)$ (red lines).

From the examples above we conclude

**Result 3.** *The equilibrium robustly achieves higher welfare and exhibits less uniformity and less fragility the more U-shaped the information quality distribution.*

## I.5. Conclusion

This chapter provides an overview of the economic literature on social learning with a particular emphasis on the relation between the properties of the equilibrium outcome and the primitives of the environment. In contexts where the set of investment options is discrete, Bayes-rational social learning is inefficient but rich distributions of private signals are conducive to complete learning and investors choosing asymptotically the most profitable investment option. Dissatisfied with the emphasis of the literature on asymptotic learning, the chapter introduces new measures of uniformity and fragility to discuss the medium-run properties of the equilibrium outcome. This change of focus seems a necessary step to foster the development of positive models of social learning.

In the next chapter we question more fundamentally the Bayes-rational view of social learning. We show that rational herding lacks learning-theoretic foundations. Accordingly, the predictions of economic models of social learning are likely to be misleading in many applications. A new direction for the theoretical modeling of social learning is then offered in our third chapter. There we extend the classic model of social learning by assuming that individuals update their beliefs in a non-Bayesian way. Individuals either underweigh or overweigh their public information relative to the public information revealed by the decisions of others and each individual's updating rule is private information. We show that introducing heterogeneous belief updating rules into the social learning setting is equivalent to increasing the support of private beliefs. In particular a social learning setting with sufficiently rich updating rules corresponds to a setting with a more U-shaped distribution of information qualities. In addition we show that introducing heterogeneous updating rules in a social learning setting reconciles equilibrium predictions with evidence from the laboratory and the field. Accordingly, in real-world social learning settings higher social welfare may be obtained than predicted by existing models of social learning.

## Appendix I.A Probabilistic Foundations

Economic models of social learning usually assume the existence of an underlying probability space $(\Omega, \mathfrak{F}, P)$ without going into details. Though implicitly relying on a strong mathematical background is quite common, we feel that providing a more formal discussion of the probabilistic assumptions in this case is not only a matter of completeness. Indeed it helps identifying the process structure of social learning and sheds light on the tools useful and required to perform research in this field.

We start by fixing some $\theta \in \Theta$. As noted, private beliefs are i.i.d. across players and distributed according to the probability measure $\nu_\theta$ on the measurable space $([0,1], \mathfrak{B}([0,1]))$. Here, $\mathfrak{B}(A)$ denotes the $\sigma$–algebra of Borel sets of the set $A \subseteq \mathbb{R}$. In addition preference types $\tilde{t}_i$ constitute an i.i.d. sequence of random variables distributed according to the probability measure $\pi$ on the space $T = \mathbb{R}^{2\times m}$ which we simply endow with its $\sigma$–algebra of Borel sets $\mathfrak{B}(T)$ as well. The probability spaces $([0,1], \mathfrak{B}([0,1]), \nu_\theta)$ and $(T, \mathfrak{B}(T), \pi)$ are the premises. Write $J = [0,1] \times T$. For every set $A \times B \in J$ we can write down its probability by $Q_\theta(A \times B) = \nu_\theta(A) * \pi(B)$ given independence of preference types and private beliefs. Hence, for every $i \in \mathbb{N}$, the *private information set* $\tilde{\mathscr{I}}_i = (b(\tilde{s}_i, \emptyset), \tilde{t}_i)$ is a random variable taking values in the product measure space $(J, \mathfrak{B}(J), \nu_\theta \times \pi)$ where $\mathfrak{B}(J) = \mathfrak{B}([0,1]) \times \mathfrak{B}(T)$ (see for instance Athreya and Lahiri, 2006, Chapter 5.1). Let $\Omega = J^n$ and let $\mathfrak{F} = \bigtimes_{i=1}^{n} \mathfrak{B}(j)$. By the independence of private beliefs and preference types for any finite collection of private information sets $\left(\hat{\mathscr{I}}_i\right)_{i=1}^n \in \Omega$ we can write $Q_\theta^{(n)}\left(\bigtimes_{i=1}^{n} \hat{\mathscr{I}}_i\right) = \prod_{i=1}^{n} Q_\theta(\hat{\mathscr{I}}_i)$. Accordingly for any state of Nature we have constructed a probability space $(\Omega, \mathscr{F}, Q_\theta^{(n)})$. We finally need to incorporate the state. Given that $\Theta$ is a binary (finite) set, we may take the set of probability measures $Q_\theta^{(n)}$ together with its power set and endow this measurable space with the measure determined by the prior $p$. Then, Nature first draws a probability measure on the space $(\Omega, \mathfrak{F}_\infty)$ and subsequently a sequence of realized private beliefs and preference types. Notice that a similar argumentation applies for an infinite sequence of players by the existence theorem of Kolmogorov (see e.g. Athreya and Lahiri, 2006, Theorem 6.3.1).

## Appendix I.B Properties of Private Beliefs

**Lemma I.B.1.** *A private belief distribution* $(G_0, G_1)$ *satisfies*

*(i)* $G_1$ *is a cdf with* $G_1(0) = 0$ *and* $G_1(1) = 1$.

*(ii)* $\int_0^1 \frac{1}{s}\, dG_1(s) = 2$.

*Conversely, if a function* $G_1$ *satisfies these conditions, there exists some* $G_0$ *such that* $(G_0, G_1)$ *is a private belief distribution.*

**Remark:** This result has been stated in the first version of Acemoglu, Munther, Lobel, and Ozdaglar (2010).

*Proof.* Let $(G_0, G_1)$ be a private belief distribution. Then clearly $G_1$ is a cdf with $G_1(1) = 1$ and by absolute continuity $G_1(0) = 0$. Furthermore

$$1 = \int_0^1 dG_0(s) = \int_0^1 \frac{1-s}{s} dG_1(s) = \int_0^1 \left(\frac{1}{s} - 1\right) dG_1(s) = \int_0^1 \frac{1}{s}\, dG_1(s) - 1.$$

For the converse define $G_0(s) = \int_0^1 \frac{1-r}{r} dG_1(r)$ and let private signals be distributed on $[0, 1]$ according to the measure induced by the cdf $G_\theta$ conditional on $\theta$ being the realized state of Nature. □

**Lemma I.B.2.** *For a private belief distribution $(G_0, G_1)$, if $G_1$ and $G_0$ are continuous and have densities $g_0, g_1$, there exists a density $g$ on $[0, 1]$ such that $g_0(s) = 2(1 - s)\, g(s)$ and $g_1(s) = 2s\, g(s)$.*

**Remark:** Chamley (2004b, p.71) and Smith and Sørensen (2008a) denote this property of continuous private belief distributions.

*Proof.* Let $\nu_\theta$ be the probability measure associated with the cdf $G_\theta$. Define $\nu = \frac{\nu_0+\nu_1}{2}$. Then by the no introspection condition, $\frac{d\nu_1}{d_\nu}(b) = 2s$ and $\frac{d\nu_0}{d\nu}(s) = 2(1 - s)$. If $\nu_0, \nu_1$ have densities, so does $\nu$ and the two LHS are $\frac{g_1}{g}(s)$ and $\frac{g_0}{g}(s)$ respectively. □

**Lemma I.B.3.** *Any private belief distribution $(G_0, G_1)$ satisfies*

*(i)* $G_1(s) \le \frac{s}{1-s}\, G_0(s)$ *for each* $s \in [0, 1]$ *with strict inequality for* $s > \underline{b}$,

*(ii)* $(1 - G_1(s)) \ge \frac{s}{1-s}\, (1 - G_0(s))$ *for each* $s \in [0, 1]$ *with strict inequality for* $s < \overline{b}$,

*(iii)* $G_1(s) < G_0(s)$ *for each* $s \in (0, 1)$ *(first–order stochastic dominance),*

*(iv)* $\frac{d}{ds}\frac{G_1}{G_0}(s) > 0$ *for each* $s \in (\underline{b}, 1]$,

*(v)* $\frac{d}{ds}\frac{1-G_1}{1-G_0}(s) > 0$ *for each* $s \in [0, \overline{b})$,

*(vi)* $\int_0^1 \frac{1-2s}{1-s}\, dG_0(s) = 0 = \int_0^1 \frac{1-2s}{s}\, dG_1(s)$.

*Proof.* The function $\frac{1-r}{r}$ is strictly decreasing in $r$. Consequently

$$\underline{(i):} \quad G_0(s) = \int_0^s d\,G_0(r) = \int_0^s \tfrac{1-r}{r}\, d\,G_1(r) \ge \tfrac{1-s}{s} \int_0^s d\,G_1(r) = \tfrac{1-s}{s}\, G_1(s)$$

$$\underline{(ii):} \quad 1 - G_0(s) = \int_s^1 d\,G_0(r) = \int_s^1 \tfrac{1-r}{r}\, d\,G_1(r) \le \tfrac{1-s}{s} \int_s^1 d\,G_1(r) = \tfrac{1-s}{s}\, (1 - G_1(s))$$

The strict inequalities hold, if the integration interval has probability strictly greater than 0 under $G_0$ and $G_1$, i.e. if $s > \underline{b}$ in the first case and if $s < \overline{b}$ in the second case. For (iii) $G_0(s)\, (1 - G_1(s)) > G_1(s)\, (1 - G_0(s))$ which follows from (i) and (ii) gives the desired result. For (iv) and (v)

$$\begin{aligned} \frac{d}{ds}\frac{G_1}{G_0}(s) &= \frac{\frac{dG_1}{ds}(s)\, G_0(s) - \frac{dG_0}{ds}(s)\, G_1(s)}{G_0^2(s)} = dG_1(s)\, \frac{G_0(s) - \frac{1-s}{s}\, G_1(s)}{G_0^2(s)}, \\ \frac{d}{ds}\frac{1-G_1}{1-G_0}(s) &= \frac{-\frac{dG_1}{ds}(s)\, [1 - G_0(s)] + \frac{dG_0}{ds}(s)\, [1 - G_1(s)]}{G_0^2(s)} = dG_1(s)\, \frac{\frac{1-s}{s}\, [1 - G_1(s)] - [1 - G_0(s)]}{G_0^2(s)} \end{aligned}$$

where in each case the last equality follows from the no introspection condition (I.1). By (i) and (ii) both are strictly greater than 0. Finally (vi) immediately follows from

$$\int_0^1 \frac{1-s}{s}\, dG_1(s) = \int_0^1 dG_0(s) = 1 = \int_0^1 dG_1(s) = \int_0^1 \frac{s}{1-s}\, dG_0(s).$$

□

Some specifications of private beliefs (private signals respectively) merit special attention. First assume that $S_i = \{0, 1\}$ for each $i \in N$ and the distribution of $\tilde{s}_i$ is given by probabilities

$$Pr\left(\tilde{s}_i = 1|\tilde{\theta} = 1\right) = Pr\left(\tilde{s}_i = 0|\tilde{\theta} = 0\right) = q > 1/2.$$

This is usually referred to as the *symmetric binary signal (SBS)* model and has first been introduced in the seminal paper by Bikhchandani, Hirshleifer, and Welch (1992). Given the prior $p = 1/2$ and a realized private signal $s_i = 1$, player $i$ holds the realized private belief $b(1, \emptyset) = q$ while his realized private belief given a realized private signal $s_i = 0$ is $b(0, \emptyset) = 1 - q$. Hence, private beliefs are binary as well. It is thus an example of bounded private beliefs. The interpretation of this specific model is that each player receives a signal indicating that either investing is profitable or not which is correct on average in a fraction $q$ of all cases.

The second specification of private beliefs is a generalization of this model introduced by Smith and Sørensen (2000, Section 3.1.A and 2008b, Section 5.2). It is referred to as the *uniform signal quality (USQ)* model. As in the example above, each player gets a binary signal. However, any player's (privately) realized signal precision $q_i$ is the realization of a random variable $\tilde{q}$ which is uniformly distributed on the interval $[1/2, a]$ for some $1/2 < a \leq 1$. Therefore private beliefs are atomlessly distributed on $[1 - a, a]$ with conditional densities

$$g_0(s) = \frac{2\,(1-s)}{2\,a\, -\, 1} \qquad \text{and} \qquad g_1(s) = \frac{2\,s}{2\,a\, -\, 1}.$$

Consequently $G_0(s) = \frac{a^2 - (1-s)^2}{2a-1}$ and $G_1(s) = \frac{s^2 - (1-a)^2}{2a-1}$. In particular private beliefs in this model are unbounded, iff $a = 1$.

The last example was studied by Chamley (2004b) and is known as the *gaussian* model. It is characterized by signals which are normally distributed with mean $\theta$ and standard deviation $\sigma$ for any realization $\theta \in \{0, 1\}$ of $\tilde{\theta}$. Consequently, given a realized private signal $s \in \mathbb{R}$, a player entertains the realized private belief

$$b(s, \emptyset) = \left[1 + \frac{d\,\Phi_{0,\sigma}}{d\,\Phi_{1,\sigma}}(s)\right]^{-1} = \left[1 + \exp\left(-\frac{2\,x - 1}{2\,\sigma^2}\right)\right]^{-1}.$$

Thus a realized private belief $b$ is one–to–one related to the realized private signal

$$s = 1/2 - \sigma^2\, \log\left(\frac{1-b}{b}\right)$$

and conditional on the realization $\theta$ of $\tilde{\theta}$ private beliefs are distributed according to the density

$$g_\theta(b) = \phi_{\theta,\sigma}\left(1/2 - \sigma^2 \log\left((1-b)/b\right)\right)$$

where $\phi_{\mu,\sigma}(\cdot)$ denotes the density of the normal distribution with mean *mu* and standard deviation $\sigma$.

All of these examples satisfy the following notion of *symmetry*.

**Definition** (Symmetry of Private Beliefs)**.** *Private Beliefs are symmetric, iff* $G_0(s) = 1 - G_1(1-s)$ *for any* $s \in \left(\underline{b}, \overline{b}\right)$.

**Lemma I.B.4.** *Private beliefs in the SBS model, the USQ model and the gaussian model are symmetric.*

*Proof.* In the SBS model it holds $G_0(s) = q$ and $G_1(1-s) = 1-q$ for $s \in (1-q, q)$. For the USQ model the assertion follows straightforwardly via

$$1 - G_1(1-s) = 1 - \frac{(1-s)^2 - (1-a)^2}{2a-1} = \frac{2a - 1 - (1-s)^2 + 1 - 2a + a^2}{2a-1} = \frac{a^2 - (1-s)^2}{2a-1} = G_0(s).$$

Finally, in the gaussian model

$$\begin{aligned} g_1(1-s) &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2\sigma^2}\left(1/2 - \sigma^2 \log(s/(1-s)) - 1\right)^2\right] \\ &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2\sigma^2}\left(-1/2 + \sigma^2 \log((1-s)/s)\right)^2\right] \\ &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{1}{2\sigma^2}\left(1/2 - \sigma^2 \log((1-s)/s)\right)^2\right] = g_0(s). \end{aligned}$$

Therefore,

$$\begin{aligned} G_0(s) + G_1(1-s) &= \int_0^s g_0(r)\,dr + \int_0^{1-s} g_1(r)\,dr = \int_0^s g_0(r)\,dr + \int_0^{1-s} g_0(1-r)\,dr \\ &= \int_0^s g_0(r)\,dr - \int_1^s g_0(r')\,dr' = \int_0^s g_0(r)\,dr + \int_s^1 g_0(r')\,dr' = 1. \end{aligned}$$

□

## Appendix I.C Omitted Proofs

***Proof of Lemma I.1.*** Fix rational type $t$ and $1 \le j < k \le m(t)$ and denote in an abuse of notation $u(j,\theta,t) = u(a_j(t),\theta,t)$. It can be straightforwardly seen that

$$\begin{aligned} & r\,u(k,1,t) + (1-r)\,u(k,0,t) > r\,u(j,1,t) + (1-r)\,u(j,0,t) \\ \Leftrightarrow\quad & r > r(j,k \mid t) = \frac{u(j,0,t) - u(k,0,t)}{u(j,0,t) - u(k,0,t) + u(k,1,t) - u(j,1,t)}. \end{aligned}$$

Define

$$d(j,k \mid t) = \frac{u(k,1,t) - u(j,1,t)}{u(j,0,t) - u(k,0,t)}.$$

Then $r(j,k \mid t) = 1/[1+d(j,k \mid t)]$ and for any $1 \leq j \leq m(t)$ it must hold that $d(j-1,j \mid t) > d(j,j+1 \mid t)$, since otherwise $r(j-1,j \mid t) \geq r(j,j+1 \mid t)$ in which case $a_j(t)$ would not be optimal for an open set of posteriors contradicting $a_j(t) \in A(t)$. Therefore for any $j < k-1$

$$\begin{aligned} d(j,k \mid t) = \frac{u(k,1,t) - u(j,1,t)}{u(j,0,t) - u(k,0,t)} &= \frac{\sum_{z=j}^{k-1} (u(z+1,1,t) - u(z,1,t))}{\sum_{z=j}^{k-1} (u(z,0,t) - u(z+1,0,t))} \\ &= \frac{\sum_{z=j}^{k-1} d(z,z+1 \mid t)\ (u(z,0,t) - u(z+1,0,t))}{\sum_{z=j}^{k-1} (u(z,0,t) - u(z+1,0,t))} > d(k-1,k \mid t). \end{aligned}$$

By a similar argument $d(j,k \mid t) \leq d(j,j+1 \mid t)$ for any $k > j+1$. Hence, $r(j,k \mid t) < r(k-1,k \mid t)$ for any $j < k-1$ and $r(j,k \mid t) > r(j,j+1 \mid t)$ for any $k > j+1$. □

***Proof of Proposition I.1***. Notice first that we need only care about rational preference types since noise types have only a single undominated action $a \in A(t)$. Accordingly fix player $i$ with rational preference type $t_i$, realized private belief $b(s_i, \emptyset)$, and realized history $h_i$. Let the equilibrium be given by $\sigma^*$. For readability we omit the argument henceforth. If $Pr(h_i \mid \tilde{\theta} = \theta) > 0$ for some $\theta \in \Theta$ we may rewrite

$$\begin{aligned} U(a \mid s_i, h_i, t_i) &= Pr(h_i \mid \tilde{\theta} = 1)\, dG_1(s_i)\, u(a,1,t) \ +\ Pr(h_i \mid \tilde{\theta} = 0)\, dG_0(s_i)\, u(a,0,t) \\ &= \Big[Pr(h_i \mid \tilde{\theta} = 1)\, dG_1(s_i) \ +\ Pr(h_i \mid \tilde{\theta} = 0)\, dG_0(s_i)\Big] \\ &\quad * \Bigg[\frac{Pr(h_i \mid \tilde{\theta} = 1)}{Pr(h_i \mid \tilde{\theta} = 1) \ +\ \frac{dG_0}{dG_1}(s_i)\, Pr(h_i \mid \tilde{\theta} = 0)}\, u(a,1,t_i) \\ &\quad + \frac{\frac{dG_0}{dG_1}(s_i)\, Pr(h_i \mid \tilde{\theta} = 0)}{Pr(h_i \mid \tilde{\theta} = 1) \ +\ \frac{dG_0}{dG_1}(s_i)\, Pr(h_i \mid \tilde{\theta} = 0)}\, u(a,0,t_i)\Bigg]. \end{aligned}$$

The no introspection condition (I.1) implies that $dG_0/dG_1(s_i) = s_i/(1-s_i) = b(s_i,\emptyset)/[1-b(s_i,\emptyset)]$. Therefore the second bracket equals $b(s_i,h_i)\, u(a,1,t_i) \ +\ [1-b(s_i,h_i)]\ u(a,0,t_i)$ provided $b(s_i,h_i)$ is formed using Bayes' rule. Consequently Lemma I.1 implies that expected payoff of type $t_i$ at history $h_i$ given private belief $b(s_i,\emptyset)$ is maximized by choosing $\sigma_i(a \mid s_i,h_i,t_i) = \sigma_i^*(a \mid s_i,h_i,t_i) = 1$ for $a \in A$ satisfying $\underline{r}(a \mid t_i) < b(s_i,h_i) < \overline{r}(a \mid t_i)$. However, if $b(s_i,h_i) = \underline{r}(a \mid t_i) = \overline{r}(a' \mid t_i)$ a player is indifferent between action $a$ and action $a'$. Hence, any assignment $\sigma(s_i,h_i,t_i)$ which assigns strictly positive probability only to actions $a$ and $a'$ maximizies expected payoff. But such cases occur with strictly positive probability only if the distribution of private beliefs has an atom at

$$\underline{s}(a \mid t_i, h_i) = \frac{\underline{r}(a \mid t_i)\, Pr(h_i \mid \tilde{\theta} = 0)}{\underline{r}(a \mid t_i)\, Pr(h_i \mid \tilde{\theta} = 0) \ +\ \Big[1 - \underline{r}(a \mid t_i)\Big]\ Pr(h_i \mid \tilde{\theta} = 1)}.$$

Apparently this is a non-generic property of the social learning settings (according to Smith and Sørensen, 2000, a property is generic if it holds for an open and dense subset of parameters of a game).

On the other hand the assignment $\sigma^*(s_i, h_i, t_i)$ maximizing expected payoff is also not uniquely defined at histories satisfying $Pr(h_i \mid \tilde{\theta} = \theta) = 0$ for each $\theta \in \Theta$ as in this case $U(a \mid s_i, h_i, t_i, \sigma) = 0$ for each $a \in A$. Yet, such histories have no bearing on the outcome of the game which can therefore be shown to be unique except in non-generic social learning settings.

Finally, to see this we show that there exists a unique iteratively undominated hence rationalizable strategy profile (again abstracting from non-generic settings). Notice first that the sequential structure and the absence of payoff externalities jointly imply that each player need only care about the strategies of his predecessors. Moreover the first player need not care about others' strategies. By the argumentation above she has a dominant strategy given by $\sigma(a \mid s_i, \emptyset, t_i) = 1$ if $\underline{r}(a \mid t_i) < b(s_i, \emptyset) < \overline{r}(a \mid t_i)$ for each rational $t_i \in supp(\pi)$, each $a \in A$, and each $b(s_i, \emptyset) \in [\underline{b}, \overline{b}]$. Knowledge of this strategy is sufficient for the second player to derive conditional probabilities of histories $Pr(h_2 \mid \tilde{\theta} = \theta)$ and thereby his set of dominant strategies. The latter uniquely defines behavior at all histories occurring a strictly positive fraction of the time under some state of Nature while not fixing behavior at other histories. A similar argumentation inductively applies to players $2 < i \leq n$ completing the proof. □

***Proof of Lemma I.2.*** If all players follow private information since private signals are independent $Pr(\tilde{a}_i = \tilde{a}_j \mid \tilde{\theta} = \theta) = q^2 + (1-q)^2$ for each $\theta \in \Theta$ and each $i \neq j$. Consequently

$$E[\sum_{i=2}^{n} \sum_{j=1}^{i-1} \mathbb{1}_{\{a_i = a_j\}}(\tilde{h}_{n-1})] = [q^2 + (1-q)^2] \sum_{i=2}^{n} (i-1) = [q^2 + (1-q)^2] \frac{n\,(n-1)}{2}.$$

On the other hand if the first player has the highest signal precision all players follow the first decision. Consequently, $\tilde{a}_i = \tilde{a}_j$ for each $i \neq j$ with probability one.

To proof the statement for the Bayesian equilibrium recall that a cascade starts once one action has been chosen two times more often than the other action for the first time. Conversely a cascade does not start if players 1 and 2, 3 and 4, …, $2k+1$ and $2k+2$, … choose different actions. Given the tie-breaking rule that players when indifferent choose either action with equal probability each of this events happens with probability $q(1-q)$. On the other hand if no cascade has started after $2k$ decisions, it will start if the next two decisions match. This is the case with probability $1 - q(1-q)$. If cascade starts after $2k$ decisions a number of $k + n - 2k = n - k$ players choose one action (the cascade action) and $k$ players choose the other action. Consequently, along any such history the measure of clustering is given by $\frac{n\,(n-1)}{2} - k\,(n-k)$. We thus obtain for an even number $n$ of player (we omit the argument $\boldsymbol{\sigma}^*$)

$$\begin{aligned}
C_n^{even} &= \frac{2}{n\,(n-1)} \left\{ \sum_{j=0}^{(n-2)/2} [q\,(1-q)]^j \, [1-q\,(1-q)] \left( \frac{n\,(n-1)}{2} - k\,(n-k) \right) \right. \\
&\qquad \left. + \; [q\,(1-q)]^{n/2} \, \frac{n}{2} \left( \frac{n}{2} - 1 \right) \right\} \\
&= \frac{2}{n\,(n-1)} \left\{ [q\,(1-q)]^{n/2} \, \frac{n\,(n-2)}{4} + \frac{n\,(n-1)}{2} \left[ 1 - (q-q^2)^{n/2} \right] \right. \\
&\qquad \left. - \, [1-q(1-q)] \left[ n \sum_{j=0}^{(n-2)/2} j\,(q-q^2)^j + \sum_{j=0}^{(n-2)/2} j^2\,(q-q^2)^j \right] \right\}.
\end{aligned}$$

By induction one establishes

$$\begin{aligned}
\sum_{j=0}^{z} j\,p^j &= \frac{p \, - \, p^{z+1}\,(1+z\,(1-p))}{(1-p)^2}, \\
\sum_{j=0}^{z} j^2\,p^j &= \frac{p\,(1+p) \, - \, p^{z+1} \left( p^2\,z^2 \, + \, (z+1)^2 \, + \, p\,(1-2z-2z^2) \right)}{(1-p)^3}.
\end{aligned}$$

Therefore by straightforward calculations

$$C_n^{even} = 1 - \frac{2\,q\,(1-q)}{(n-1)\,(1-q+q^2)} + \frac{2q(1-q)\,(1+q-q^2)}{[1-q(1-q)]^2} * \frac{1-(q-q^2)^{n/2}}{n\,(n-1)}.$$

On the other hand for $n$ odd

$$C_n^{odd} = \frac{n-2}{n}\,C_{n-1}^{even} + \frac{2}{n(n-1)} \left[ n - 1 - q\,(1-q)\,\frac{1-(q-q^2)^{(n-1)/2}}{1-q(1-q)} \right].$$

It is thus easy to see that

$$C_n = 1 \, - \, \frac{2\,q\,(1-q)}{(n-1)\,(1-q+q^2)} \, + \, \mathcal{O}\left( \frac{1}{n^2} \right)$$

$\square$

***Proof of Lemma I.4***. First, $\overline{b} = q = 1 - \underline{b}$ and $\bar{q} = q$. Second, beliefs in a cascade on $a = 1$ ($a = 0$) are given by $\frac{q\,(1-q)}{2-2\,q\,(1-q)}$ $\left(\frac{(1-q)\,(2-q)}{2-2\,q\,(1-q)}\right)$. Third apparently in a cascade on $a$, the amount of information necessary to induce $a$ is zero.

Consider first the case of a cascade which is the same for player $2j+1$ and $2j$. Unconditionally either of these is in a cascade on action $a = 1$ with probability $\frac{1}{2}\left[1-(q-q^2)^j\right]$ in which case the amount of information required to induce action 1 is zero and the amount to induce $a = 0$ is given by

$$\begin{aligned}
\underline{\beta}(0) &= \frac{\log\left(\overline{b}/(1-\overline{b})\right) \, + \, \log\left(b(\emptyset,h)/(1-b(\emptyset,h))\right)}{\log\left(\bar{q}/(1-\bar{q})\right)} \\
&= \frac{\log\left(q/(1-q)\right) \, + \, \log\left([q\,(1+q)]/[(1-q)\,(2-q)]\right)}{\log\left(q/(1-q)\right)} \\
&= 2 + \frac{\log\left((1-q)/(2-q)\right)}{\log\left(q/(1-q)\right)}
\end{aligned}$$

where we have used in the last line that $\log(ab) = \log(a) + \log(b)$. A similar resoning applies to the case of a cascade on $a = 0$ with roles of $a$ and $b$ switched. In conclusion in a cascade the average (over actions) degree of fragility for players $2j + 1$ and $2j + 2$ is given by $1 + \frac{1}{2}\frac{\log((1+q)/(2-q))}{\log(q/(1-q))}$.

On the other hand if player $2j + 1$ is not in a cascade her public belief must be given by $b(\emptyset, h) = 1/2$. It is then easy to see that $\underline{\beta}(h) = \overline{\beta}(h) = 1$, i.e. the average degree of clustering at such histories equals one as well. If player $2j + 2$ is not in a cascade her public belief is given by either $b(\emptyset, h) = q$ or $b(\emptyset, h) = 1 - q$. In the first case $\overline{\beta}(h) = 0$ and $\underline{\beta}(h) = 2$ while in the second case $\underline{\beta}(h) = 0$ and $\overline{\beta}(h) = 2$. Hence, the average degree of clustering is one. In conclusion for an even number of players (omitting the Bayesian equilibrium strategy $\boldsymbol{\sigma}^*$ as an argument)

$$
\begin{aligned}
\mathscr{R}_n &= \frac{1}{n}\sum_{j=0}^{(n-2)/2}\left\{2\,\frac{1-(q-q^2)^j}{2}\left[1 + \frac{1}{2}\,\frac{\log((1+q)/(2-q))}{\log(q/(1-q))}\right] + (q-q^2)^j\right\} \\
&= \frac{1}{n}\sum_{j=0}^{(n-2)/2}\left\{1 + \frac{1}{2}\,\frac{\log((1+q)/(2-q))}{\log(q/(1-q))}\left[1-(q-q^2)^j\right]\right\} \\
&= 1 + \frac{1}{2}\,\frac{\log((1+q)/(2-q))}{\log(q/(1-q))}\left[1 - \frac{1}{n}\,\frac{1-(q-q^2)^{n/2}}{1-q+q^2}\right]
\end{aligned}
$$

while for an odd number of players

$$
\mathscr{R}_n = \frac{n-1}{n}\,\mathscr{R}_{n-1} + \frac{1}{n}\left\{1 + \frac{1}{2}\,\frac{\log((1+q)/(2-q))}{\log(q/(1-q))}\left[1-(q-q^2)^{(n-1)/2}\right]\right\}
$$

yields the desired result. □

# Chapter II.

# Adaptive Social Learning

## II.1. Introduction

As our previous chapter shows, economic models of social learning have the potential to deepen our understanding of real-world phenomena like social epidemics. However, we fear that the Bayesian rational view of social learning in its current form reaches unsound conclusions which in turn limits its behavioral relevance. First, rational herds are fragile meaning that even after having observed one thousand identical decisions the average individual is not extremely confident about the appropriateness of the chosen action. Second, because Bayesian rational individuals are aware of the fragility of herds they overturn the accumulated evidence from many previous actions when endowed with a sufficiently precise signal which points in a different direction. According to the overturning principle (Smith and Sørensen, 2000), the belief of the average individual is drastically revised after the observation of such a contrarian action which implies that society eventually learns the truth in rich-enough signal spaces. Both conclusions stand in sharp contradiction to the experimental evidence on social learning (Ziegelmeyer, Koessler, Bracht, and Winter, 2010; Weizsäcker, 2010) and they have been criticized by recent models of boundedly rational social learning (Eyster and Rabin, 2010; Guarino and Jehiel, 2009).

Though we are sympathetic to the bounded rationality approach and admit that rational inferences probably go beyond what real decision makers are capable of, we are also convinced that experience can substitute for knowledge and that passive observation and active experimentation can reduce uncertainty. In other words, we agree with the prevalent idea in economics that rational play usually arises as the outcome of some adaptation process.[1] In this chapter, we investigate the learning foundations of economic models of social learning which enables us to indirectly assess the whole set of assumptions underlying these models. After all, these equilibrium models not only assume that individuals are Bayesian rational but they also suppose that all individuals agree on how to assess the uncertainty surrounding the parameters of the model. In particular, the distribution of private signals is assumed to be mutual knowledge among the individuals even in the case of arbitrarily strong signals. The learning perspective offers us the possibility to clarify when and why equilibrium is likely to capture observed regularities in the field. More precisely it enables us to rigorously evaluate the assumption of mutual knowledge of the signal distributions.

We consider an adaptive process which allows players to learn during repeated play of a social learning game how to derive information from others' decisions. The process is in the

[1] Walliser (1998) provides a nice classification of equilibration processes. See also Hart (2005).

spirit of fictitious play (Brown, 1951): In each repetition players myopically best respond to their assessments formed by employing a frequentist method. Our first main result provides sufficient conditions for the learned beliefs and choices to approach those predicted by the Bayesian equilibrium of the game. We term these conditions, given by an infinite number of repetitions of the same stable social learning game, *perfect learning opportunities*. A straightforward consequence is that under perfect learning opportunities the Bayes-rational strategy, i.e. combining information derived from assessments about others and private information in a Bayesian way and best responding to the resulting posteriors, maximizes expected payoff and uniquely so in some social learning environments.

Limit results are fairly standard in the theoretical literature on learning in games. Yet, many real-world strategic environments do not remain unchanged in the long run. Hence, players will rarely encounter the same game a great number of times. The outcome of the learning process after a finite number of repetitions may thus be more relevant to behavior. As a response to this argument various authors have argued that any sort of learning involves extrapolation across environments which are considered similar (Fudenberg, 2006; Fudenberg and Levine, 1998). Accordingly what matters is how often players have played similar games. This may justify retaining the long-run perspective.

We check the validity of this justification by explicitly modeling learning across social learning games. In fact our assumption is stronger: We have players learning in a world of several social learning environments which differ *only* in the distribution of private beliefs. In each repetition of the learning process one of these environments is selected randomly and *unobserved* by the players. We then argue that absent strong assumptions regarding prior knowledge about the space of possible distributions the settings are not just similar but *indistinguishable* for players. Our results demonstrate that adapting across such environments is highly problematic. In particular Bayes-rationally responding to assessments and private information ceases to be the optimal strategy. More precisely since players cannot disentangle environments they adapt to the artificial mixed setting given by the linear combination of private belief distributions. The Bayes-rational strategy in the latter generically does not coincide with the optimal benchmark strategy for the mixed world. Therefore players who do not respond to assessments (and private information) in a Bayes-rational fashion may ultimately achieve a higher expected payoff in the mixed world. In conclusion the results cast doubts upon the assumption of mutual knowledge of signal distributions.

In order to corroborate and complement these results and to investigate further the structure of alternative strategies we extend our model to account for other imperfections of the individual learning environment. In particular we directly model learning within a finite number of repetitions. Under this assumption assessments must remain noisy estimates of their true counterparts. We show in a result of separate interest that noise in assessments causes Bayesian posteriors to in expected terms *underweight* private information. This indicates but by no means proves the sub-optimality of the Bayes-rational strategy. Yet, we verify again that the combination of Bayesianism and best response is outperformed by other strategies. In particular we indicate that *overweighting of private information* tends to imply a higher expected payoff.

Our results on imperfect learning opportunities demonstrate how more complex social learning worlds may favor non-Bayes-rational strategies. We thus provide a rationale for the persistence of these strategies. In particular in the spirit of the work on rule rationality we shed a new light on the experimental findings.[2] If adaptation takes place across environments, subjects' behavior in simple laboratory social learning settings may be an artifact of their adjustment to more complex real-world environments. Especially the framework of mixed worlds explicitly formalizes how learning across similar environments may ultimately induce suboptimal behavior in each environment separately.

A common perspective in previous work on rule rationality asserts that rules are shaped by an evolutionary process, instead of being consciously derived. The results of this chapter are easily extended in this direction. While this may have provided more precise results regarding the exact shape of alternative strategies, two considerations have guided our choice not to become more specific on this point. On the one hand we did not want to tie our hands with concrete behavioural models or concepts of evolutionary stability. On the other hand we believe this to be as much an empirical question. We thus frequently discuss the model of *heterogeneous updating rules* which takes into account the overweighting of private information and which is shown to potentially capture the experimental regularities in the final chapter (chapter III) of the thesis.

This chapter is structured as follows. We start with a basic illustration of our main results in Section II.2. In Section II.3 we introduce the general social learning setting. We set up the adaptive process in Section II.4. Section II.5 presents the results on perfect learning opportunities. In Section II.6 we establish how mixed worlds and other learning imperfections may shift the optimal strategy. We discuss our results in Section II.7. An appendix contains various additional material and some of the proofs.

## II.2. A Basic Illustration

In this section, we illustrate the main results of our chapter with the help of two simple examples.

### Example 1: Social Learning in a Single Context

Consider a setting where two players, Anna and Bob, face similar investment decisions under uncertainty and have noisy perceptions about the payoff of the investment which are private information. Players decide in sequence with Anna deciding first and Bob deciding second after having observed Anna's investment decision. Payoffs from investing and rejecting are the same for both players. The investment payoff is denoted by the random variable $\tilde{\theta}$ with possible realizations 0 and 1 which are equally likely, and $\theta$, the realization of $\tilde{\theta}$, denotes the true value of the investment payoff. The cost of an investment is $c = 1/2$ and the payoff to rejecting is zero. Players' perceptions about the payoff of the investment are modeled via *private beliefs* which are probability estimates about the profitability of the investment decision. Concretely, player $i$'s private belief, $i \in \{A(nna), B(ob)\}$, is given by $b_i$, an estimated probability that the true payoff of the investment is 1.[3] We assume that a player is more likely to hold a higher private belief if the

[2] This is much in line with the recent survey of Al-Najjar and Weinstein (2009), who argue similarly about uncertainty aversion.

[3] Thus, $b_i = Pr\left(\tilde{\theta} = 1 \mid b_i\right)$.

true payoff of the investment is 1 than if it is 0.[4] This reflects the presumption that players' private beliefs are not merely opinions but are (at least partially) informative. Without loss of generality, we also assume that players' private beliefs are independent conditional on the state.

In the following, we first characterize players' behavior as a result of eductive learning. The eductive process converges to the unique rationalizable outcome (Bernheim, 1984; Pearce, 1984). Next we show that adaptive learning leads to the same outcome.

### Eductive Social Learning

In an eductive process, Anna and Bob have complete prior information on the social learning context. They both know the payoff structure and the structure of perceptions (simply information structure). In fact, the social learning context is not only assumed to be mutual knowledge but common knowledge (in this illustration it is obviously enough to assume that Bob knows that Anna knows). In other words, both players are endowed with *structural knowledge*. Additionally, players' preferences and degree of rationality are also assumed to be common knowledge. We abstract from bounded rationality considerations and assume that both players are endowed with instrumental and cognitive rationality in a strong sense. Since knowledge is represented by probabilities, cognitive rationality means that both players form beliefs about the profitability of the investment by using the laws of conditional probability, where the conditioning is on all available information (in this illustration it is obviously enough to assume that Bob is cognitively rational). Instrumental rationality means that once a belief is formed a decision is made which maximizes expected payoff.[5] In a nutshell, Anna and Bob are Bayesian rational and Bayesian rationality is assumed to be common knowledge which means that both players are endowed with *strategic knowledge*.

Anna's belief is straightforwardly given by her private belief $b_A$. Therefore Anna's expected payoff from investing is given by $E[\tilde{\theta} \mid b_A] - 1/2 = b_A - 1/2$. Accordingly, Anna invests if $b_A > 1/2$, i.e. if she believes it more likely that the true payoff of the investment is 1, and she rejects if $b_A < 1/2$, i.e. if she believes it more likely that the true payoff of the investment is 0.[6]

Bob, the second player, in addition to holding a private belief observes Anna's investment decision. Concretely, Bob observes a history $h_B \in \{\text{invest, reject}\}$. Since Anna's decision strictly depends on her private belief and since the latter comprises at least some information about the true payoff of the investment, Anna's decision constitutes a second source of information about the payoff of the investment for Bob. Hence, Bob's belief is based on his private belief and Anna's investment decision. Concretely, $Pr\left(\tilde{\theta} = 1 \mid b_B, h_B\right)$ is Bob's belief that the true value of the

[4]Formally, we assume that $b_i$ is the realization of a random variable $\tilde{b}_i$ whose conditional distributions satisfy $Pr\left(\tilde{b}_i = b_i \mid \tilde{\theta} = 1\right)/Pr\left(\tilde{b}_i = b_i' \mid \tilde{\theta} = 0\right) \leq Pr\left(\tilde{b}_i = b_i' \mid \tilde{\theta} = 1\right)/Pr\left(\tilde{b}_i = b_i' \mid \tilde{\theta} = 0\right)$ if $b_i < b_i'$ for each $i \in \{A, B\}$. This assumption is weaker than the *proportional property* $Pr\left(\tilde{b}_i = b_i \mid \tilde{\theta} = 1\right)/Pr\left(\tilde{b}_i = b_i \mid \tilde{\theta} = 0\right) = b_i/(1 - b_i)$ (Chamley, 2004b, p.31) which must hold if players' private beliefs result solely from updating the prior with private information in a Bayesian way.

[5]The same conclusion holds if one relies on the weaker assumption of probabilistic sophistication in the sense of Machina and Schmeidler (1992).

[6]If $b_A = 1/2$, Anna is indifferent between investing and rejecting. By neglecting this indifference case, we strengthen our conclusion.

investment is 1 given his private belief $b_B$ and the history $h_B$. Since Bob's expected payoff from investing is given by $E\left[\tilde{\theta} \mid b_B, h_B\right] - 1/2 = Pr\left(\tilde{\theta} = 1 \mid b_B, h_B\right) - 1/2$, Bob invests if $Pr\left(\tilde{\theta} = 1 \mid b_B, h_B\right) > 1/2$ or equivalently if his likelihood ratio $\lambda(b_B, h_B) = Pr\left(\tilde{\theta} = 1 \mid b_B, h_B\right) / Pr\left(\tilde{\theta} = 0 \mid b_B, h_B\right)$ is strictly greater than 1. Bob's likelihood ratio is given by

$$\lambda(b_B, h_B) = \frac{b_B}{1 - b_B} * \frac{Pr\left(\tilde{h}_B = h_B \mid \tilde{\theta} = 1\right)}{Pr\left(\tilde{h}_B = h_B \mid \tilde{\theta} = 0\right)}.$$

How can Bob assess the probabilities $Pr\left(\tilde{h}_B = h_B \mid \tilde{\theta} = \theta\right)$ $(\theta = 0, 1)$? Recall that Bob does not observe (or know) Anna's private belief. In order to assess the probabilities, Bob must reason about Anna's decision. Since Bob is endowed with strategic knowledge, he knows that Anna invests provided $b_A > 1/2$ and rejects provided $b_A < 1/2$. Moreover, Bob is endowed with structural knowledge which implies that he knows the likelihood of the events $\{b_A > 1/2\}$ and $\{b_A < 1/2\}$ conditional on the realized payoff of the investment. Therefore, Bob's likelihood ratio is given by

$$\lambda(b_B, h_B) = \begin{cases} \frac{b_B}{1-b_B} \frac{Pr\left(b_A > 1/2 \mid \tilde{\theta} = 1\right)}{Pr\left(b_A > 1/2 \mid \tilde{\theta} = 0\right)} & \text{if Anna invests} \\ \frac{b_B}{1-b_B} \frac{Pr\left(b_A < 1/2 \mid \tilde{\theta} = 1\right)}{Pr\left(b_A < 1/2 \mid \tilde{\theta} = 0\right)} & \text{if Anna rejects} \end{cases}.$$

It can be shown that $Pr\left(b_A > 1/2 \mid \tilde{\theta} = 1\right) > Pr\left(b_A > 1/2 \mid \tilde{\theta} = 0\right)$ and equivalently $Pr\left(b_A < 1/2 \mid \tilde{\theta} = 1\right) < Pr\left(b_A < 1/2 \mid \tilde{\theta} = 0\right)$.[7] On the one hand, if Bob's private belief confirms Anna's decision (Anna invests and $b_B > 1/2$ or Anna rejects and $b_B < 1/2$) then Bob follows his private belief or equivalently imitates Anna's decision. On the other hand, if Bob's private belief contradicts Anna's decision (Anna invests and $b_B < 1/2$ or Anna rejects and $b_B > 1/2$) then his decision depends upon the strength of his private belief and the information contained in Anna's decision. Concretely, if

$$1 < \frac{1 - b_B}{b_B} < \frac{Pr\left(b_A > 1/2 \mid \tilde{\theta} = 1\right)}{Pr\left(b_A > 1/2 \mid \tilde{\theta} = 0\right)}$$

then Bob follows Anna's decision and invests despite his pessimistic private belief $b_B < 1/2$ about the true payoff of the investment. Equivalently, if

$$1 > \frac{1 - b_B}{b_B} > \frac{Pr\left(b_A < 1/2 \mid \tilde{\theta} = 1\right)}{Pr\left(b_A < 1/2 \mid \tilde{\theta} = 0\right)}$$

then Bob follows Anna's decision and rejects despite his private belief $b_B > 1/2$ indicating that investing is superior to rejecting.

In conclusion, eductive learning leads to the rationalizable outcome which is unique and is therefore also the equilibrium outcome.[8]

### Adaptive Social Learning

Contrary to the eductive justification for equilibrium, the adaptive approach does not assume that players are endowed with individual and interactive knowledge before the adjustment process

[7] See e.g. Milgrom (1981).

[8] Needless to say, eductive learning is a collective mental process which runs in "virtual" time.

starts. We assume epistemic learning (or beliefs-based learning) which means that players monitor others' past decisions and form extrapolative expectations about their strategies completed and corrected through real time. To simplify exposition, we do not provide much details on the learning-theoretic model. Instead, we simply assume that players are Bayesian rational and that they do not entertain any repeated play considerations but decide myopically each time they interact in the social learning context. Finally, to give best chances to the emergence of the equilibrium outcome, we abstract from learning costs i.e. players maximize the sum of their per-interaction expected payoffs. The adaptive process is fully exposed in Section II.4.

Concretely, Anna and Bob repeatedly interact in the social learning context with the investment payoff being revealed after both Anna and Bob have decided. In each repetition Anna follows her private belief (i.e. invests if $b_A > 1/2$ and rejects if $b_A < 1/2$) which implies that she invests more often in repetitions with an investment payoff of 1. More precisely, the conditional relative frequency of Anna investing given an investment payoff of 1 (0) eventually approaches $Pr\left(b_A > 1/2 \mid \tilde{\theta} = 1\right)$ ($Pr\left(b_A > 1/2 \mid \tilde{\theta} = 0\right)$) and equivalently for Anna rejecting. If across repetitions Bob keeps track of Anna's choices conditional on the payoff of the investment and if there are sufficiently many repetitions Bob will eventually learn these relative frequencies. In other words, Bob eventually learns to infer the same information from Anna's decision than he would deduce by being endowed with strategic and structural knowledge.

In conclusion, adaptive learning also leads to the rationalizable outcome.

Our first example suggests that Bayesian rational social learning can be viewed as the long-run outcome of a dynamic process of adjustment. This example is however misleading. Indeed, learning-theoretic models build on the assumptions that the learning horizon is infinite and that there is a large number of players who interact relatively anonymously (to prevent repeated play considerations). The validity of these two assumptions is questionable when learning takes place in a single context. As argued by Fudenberg and Levine (1998, p. 4), "our presumption that players *do* extrapolate across games they see as similar is an important reason to think that learning models have some relevance to real-world situations." Accordingly, we further investigate the learning-theoretic foundations of Bayesian rational social learning in a second example where players adapt across social learning contexts.

### Example 2: Social Learning in Multiple Contexts

We assume that there are two different settings which differ only in the distributions of private beliefs. Additionally, we stick to our previous assumptions i.e. players are Bayesian rational who maximize the sum of their per-interaction expected payoffs and they do not entertain any repeated play considerations.

In the *low information setting* the distributions of Anna's private belief satisfy $Pr\left(b_A > 1/2 \mid \tilde{\theta} = 1, L\right) = \pi_A^L > 1/2$ and $1/2 > Pr\left(b_A > 1/2 \mid \tilde{\theta} = 0, L\right) = 1 - \pi_A^L$.[9] Bob's private belief takes one of two pos-

[9] We assume symmetry of private belief distributions only for the ease of exposition. In particular $\pi_A^L$ denotes the probability that Anna chooses the more profitable action given the true payoff of the investment. The results straightforwardly extend to the asymmetric case.

sible values $0 < 1 - b_B^L < 1/2 < b_B^L < 1$ and $b_B^L > \pi_A^L$.[10] Accordingly, if Anna and Bob would only interact in the low information setting then Bob would eventually learn to always follow his private belief irrespective of Anna's decision.[11]

In the *high information setting* the condition on Anna's private belief distributions is given by $Pr\left(b_A > 1/2 \mid \tilde{\theta} = 1, H\right) = \pi_A^H > b_B^L$ and $Pr\left(b_A > 1/2 \mid \tilde{\theta} = 0, H\right) = 1 - \pi_A^H$. The possible values of Bob's private belief are $0 < 1 - b_B^H < 1/2 < b_B^H < 1$ and satisfy $b_B^H > \pi_A^H$. Again, if Anna and Bob would only interact in the high information setting then Bob would eventually learn to always follow his private belief irrespective of Anna's decision.

As before our premise is that Bob's strategic and structural knowledge is acquired over time. Anna and Bob interact with each other repeatedly and Bob keeps track of the same information as before (the relationship between Anna's investment decision and the realized payoff of the investment). A new twist arises: The settings keeps switching across repetitions. Concretely, we assume that each repetition is equally likely to take place in either the low information setting or the high information setting.[12] If Bob were able in each repetition to identify the setting then he could simply keep track of Anna's choices in each setting separately. But in general this would require him to possess a significant degree of structural knowledge before the adaptive process starts.[13] In particular, Bob would have to know that there exists a correlation between the strength of his private belief and the strength of the information that he can derive from Anna's decision. Absent this knowledge, Bob *cannot* distinguish between the two settings.[14]

The fact that Bob cannot distinguish between the two settings implies that he learns *across* them. Concretely, Bob keeps track of Anna's investment decisions conditional on the realized investment payoff across all repetitions. Since Anna always follows her private belief and since each setting is equiprobable, Bob eventually infers the following conditional relative frequencies

$$\begin{aligned} Pr\left(h_B = \text{invest} \mid \tilde{\theta} = 1\right) = Pr\left(h_B = \text{reject} \mid \tilde{\theta} = 0\right) &= \overline{\pi}_A = \frac{\pi_A^L + \pi_A^H}{2}\text{, and} \\ Pr\left(h_B = \text{invest} \mid \tilde{\theta} = 0\right) = Pr\left(h_B = \text{reject} \mid \tilde{\theta} = 1\right) &= 1 - \overline{\pi}_A. \end{aligned}$$

Based on these frequencies, Bob adopts the following strategy: In the high information setting, Bob follows his private belief independent of Anna's decision ($\pi_A^L < \overline{\pi}_A < \pi_A^H < b_B^H$); in the low information setting, Bob follows his private belief irrespective of Anna's decision if $\overline{\pi}_A < b_B^L$ but he imitates Anna's decision irrespective of his private belief if $\overline{\pi}_A > b_B^L$. In summary, if $\overline{\pi}_A > b_B^L$ then Bob suffers an expected loss in the low information setting as he would be better off by following

[10] We discuss the opposite case at the end of the section.

[11] Since $\lambda(b_B,\ \text{invest}) = \frac{b_B}{1-b_B} * \frac{\pi_A^L}{1-\pi_A^L} > 1 \Leftrightarrow b_B > 1 - \pi_A^L \Leftrightarrow b_B > 1/2$ and $\lambda(b_B,\ \text{reject}) = \frac{b_B}{1-b_B} * \frac{1-\pi_A^L}{\pi_A^L} < 1 \Leftrightarrow b_B < \pi_A^L \Leftrightarrow b_B < 1/2$.

[12] Our results straightforwardly extend to the non-symmetric case.

[13] By identification of the setting we mean the identification of its information structure. Obviously, social learning contexts differ not only according to their information structure. For instance, the nature of the investment is idiosyncratic to the social learning context. However, unless one is willing to assume that such physical cues are correlated with the information structure our arguments remain valid.

[14] Strictly speaking, Bob can discriminate the two settings according to the strength of his private belief. Our results straightforwardly extend to a more complex setting where this is not possible (see the robustness section).

his private belief in this setting as well.

Our second example shows that the use of Bayes' rule to form beliefs about the profitability of the investment is detrimental to Bob's fitness, measured by the sum of his per-interaction expected payoffs. Obviously, our assumption that Bob is Bayesian rational was made only for convenience. A more reasonable learning-theoretic framework allows for the adjustment of updating rules since players resort to active experimentation in order to test unusual strategies. We conclude our second example by showing that non-Bayesian rational learning emerges in such an extended framework.

Remember that Bayesian rational Bob invests if the following likelihood ratio

$$\lambda(b_B, h_B) = \frac{b_B}{1 - b_B} * \frac{Pr\left(\tilde{h}_B = h_B \mid \tilde{\theta} = 1\right)}{Pr\left(\tilde{h}_B = h_B \mid \tilde{\theta} = 0\right)}$$

is strictly greater than 1. Accordingly, Bob invests provided

$$[\log(b_B) - \log(1 - b_B)] + \left[\log\left(Pr(\tilde{h}_B = h_B \mid \tilde{\theta} = 0)\right) - \log\left(Pr(\tilde{s}_h = h_B \mid \tilde{\theta} = 1)\right)\right] > 0.$$

Bayesian rational Bob weights equally his private belief and the information he derives from Anna's decision. Alternatively, non-Bayesian rational Bob invests provided

$$\beta * [\log(b_B) - \log(1 - b_B)] + \left[\log\left(Pr(\tilde{h}_B = h_B \mid \tilde{\theta} = 0)\right) - \log\left(Pr(\tilde{s}_h = h_B \mid \tilde{\theta} = 1)\right)\right] > 0$$

where $\beta > 0$ captures the weight of the private belief relative to the weight of the information inferred from Anna's decision. Non-Bayesian rational Bob's strategy is characterized as follows: In each setting $k \in \{L, H\}$, there exists $\beta_k^*$ such that Bob follows his private belief irrespective of Anna's decision provided $\beta > \beta_k^*$ and Bob imitates Anna's decision irrespective of his private belief provided $\beta < \beta_k^*$.[15] Since $\beta_L^* > \beta_H^*$, Bob follows his private belief irrespective of Anna's decision in either setting provided $\beta > \beta_L^*$. Therefore, a non-Bayesian rational Bob whose relative weight on his private belief is strictly greater than $\beta_L^*$ achieves higher fitness than a Bayesian rational Bob.

In conclusion, if Bob cannot distinguish between the two settings then he is not able to infer the correct magnitude of information that Anna's decision conveys in each setting. In other words, the complete resolution of uncertainty (strategic and structural) is not possible in the presence of large structural uncertainty.[16] We conjecture that our result holds for most adaptive processes though a proof of this assertion is left for future research.

### Robustness of Example 2

In our second example, Bob has the possibility to discriminate the two settings according to the strength of his private belief. We now show that this is not possible in an extended framework.

[15] $\beta_k^* = \log(\overline{\pi}_A / (1 - \overline{\pi}_A)) / \log\left(b_B^k / (1 - b_B^k)\right)$.

[16] Large structural uncertainty prevails in field environments since the generating process of players' private beliefs is rarely commonly known (see Dekel and Gul, 1997, p. 101). This being said, real-world situations with low structural uncertainty also exist. Assume for instance that the payoff of the investment is determined by the uncertain amount of oil in some tract and that players' private beliefs result from each of them taking a soil sample (Hendricks and Kovenock, 1989). Published experiments provide a thorough understanding of both the prior likelihood of oil and the distribution of samples as a function of the oil in the tract. In this situation, (common) knowledge of the distribution of private beliefs is easily justified.

Assume that Bob can be weakly ($b_B \in \{1-b_B^L, b_B^L\}$) or well informed ($b_B \in \{1-b_B^H, b_B^H\}$) in *both* settings but that he is more likely to be weakly informed in the low information setting. Concretely, assume that in the low (high) information setting the probability that Bob is weakly (well) informed is given by $\alpha > 1/2$. In this case it is optimal for Bob to follow an extreme private belief $b_B \in \{1-b_B^H, b_B^H\}$ since $\pi_A^L < b_B^L < \pi_A^H < b_B^H$ and it is optimal for him to follow a less extreme private belief $b_B \in \{1-b_B^L, b_B^L\}$ provided $b_B^L > \alpha\, \pi_A^L + (1-\alpha)\, \pi_A^H$.[17] In particular $\alpha\, \pi_A^L + (1-\alpha)\, \pi_A^H < \overline{\pi}_A$ since $\alpha > 1/2$ and there exists a set of parameters such that $\alpha\, \pi_A^L + (1-\alpha)\, \pi_A^H < b_B^L < \overline{\pi}_A$, i.e. such that Bob given a less extreme private belief imitates Anna's decision although he should optimally follow his private belief. Hence, we reproduce the result that learning across settings may lead to suboptimal behavior for Bayesian rational players.[18]

Finally, if players rely on the adaptive process to assess both their own information and uncertainty then Bayesian rationality eventually yields the strategy with the highest fitness. We find this argument unconvincing. Social learning occurs in situations where informed players (though imperfectly) try to learn from observing others' decisions.

## II.3. Preliminaries

### II.3.1. The Social Learning Game

A finite number of players $i = 1, 2, \dots, n$ choose an action, in that exogenous order, from the set $A = \{0, 1\}$. Each player's payoff depends on the realization of an underlying state of Nature and the chosen action. The state of Nature is given by the random variable $\tilde{\theta}$ distributed on $\Theta = \{0, 1\}$, over which players share a common prior belief. Without loss of generality, the prior is assumed to be flat, with both states equally likely. Players' payoffs are given by the mapping $u : A \times \Theta \to \mathbb{R}$ where $u(1, \theta) = \theta - \frac{1}{2}$ and $u(0, \theta) = 0$ for each $\theta \in \Theta$. In the following, player $i$'s action $a_i = 1$ (resp. $a_i = 0$) is sometimes referred to as "invest" (resp. "reject") and the cost of the investment is set equal to $\frac{1}{2}$ merely to simplify the exposition (our results straightforwardly extend to any cost in the interval $(0, 1)$). In a similar vein, both the underlying state of Nature and the action set are binary to simplify the exposition (our results extend to any finite number of actions and states but at significant algebraic cost).

The realized state $\theta$ is unknown and each player is endowed with some (imperfect) private information about the realized state. Before any action is taken, player $i$'s imperfect knowledge about the realized state is called her *private belief* and is denoted by $b(\tilde{s}_i, \emptyset)$. This endowment of player $i$ is a probability estimate of the state of Nature which can be interpreted as resulting from the prior probability of the state and player $i$'s private signal $\tilde{s}_i$. In the following, we often identify $\tilde{s}_i$ (resp. $s_i$) with the private belief (resp. realization of the private belief) $b(\tilde{s}_i, \emptyset)$ (resp. $b(s_i, \emptyset)$). Conditional on the realized state, $(\tilde{s}_i)_{i=1}^n$ is an i.i.d. sequence generated according to the

[17] $Pr\left(\tilde{\theta} = 1 \mid b_B^L, h_B\right) = \frac{\alpha\, b_B^L\, Pr\left(h_B|\tilde{\theta}=1,L\right) + (1-\alpha)\, b_B^L\, Pr\left(h_B|\tilde{\theta}=1,H\right)}{\alpha\, b_B^L\, Pr\left(h_B|\tilde{\theta}=1,L\right) + (1-\alpha)\, b_B^L\, Pr\left(h_B|\tilde{\theta}=1,H\right) + \alpha\left(1-b_B^L\right) Pr\left(h_B|\tilde{\theta}=0,L\right) + (1-\alpha)\left(1-b_B^L\right) Pr\left(h_B|\tilde{\theta}=0,H\right)}$ and equivalently for $Pr\left(\tilde{\theta} = 1 \mid 1 - b_B^L, h_B\right)$.

[18] Bayesian rational Bob could insist on distinguishing between the two settings and he would eventually arrive at the optimal strategy. We doubt the relevance of such cases and abstract from them by assumption but we leave it to the reader to evaluate the restrictiveness of this assumption.

c.d.f. $G_\theta(s)$. $G_0$ and $G_1$ satisfy the standard assumptions meaning that they have common support whose convex hull is given by $\left[\underline{b}, \overline{b}\right]$ and their Radon-Nikodym derivative is such that $\frac{dG_1}{dG_0}(s) = \frac{s}{1-s}$. We assume that $\underline{b} > 0$ and $\overline{b} < 1$ i.e. private beliefs are bounded. Note that the property on the Radon-Nikodym derivative implies that $G_0$ and $G_1$ satisfy the property of first-order stochastic dominance: $G_1(s) < G_0(s)$ for each $s \in (\underline{b}, \overline{b})$.

Though player $i$ cannot observe the private belief of any other player, she observes the complete history $h_i = (a_1, \ldots, a_{i-1}) \in H_i = A^{i-1}$ of previous actions with $h_1 \equiv \emptyset$. In fact, we assume that history $h_i$ is observed by all players $i, i+1, \ldots, n$ and that this knowledge is common to all players. $H_{n+1} = A^n$ denotes the set of complete histories with element $h_{n+1} = (a_1, \ldots, a_n)$ and $H = \bigcup_{i=1}^n H_i$. Given a sequence of actions $(a_1, \ldots, a_{i-1})$, the probability estimate of the state of Nature that is based solely on the public information is called the *public belief* and is given by $b(\emptyset, h_i) = Pr\left(\tilde{\theta} = 1 \mid h_i\right) = Pr\left(h_i \mid \tilde{\theta} = 1\right) / \left(Pr\left(h_i \mid \tilde{\theta} = 1\right) + Pr\left(h_i \mid \tilde{\theta} = 0\right)\right)$ with $b(\emptyset, h_1) = 1/2$.

We denote by $\langle n, A, u, \Theta, (G_0, G_1)\rangle$ the *social learning game*. Smith and Sørensen (2000) study a non-straightforward generalization of the social learning game where players have heterogeneous preferences which are private information and some of them have state independent preferences with a single dominant action. In subsection II.6.2, we discuss how our results extend to this generalized social learning game.

We conclude this subsection with a definition (Smith and Sørensen, 2000).

**Definition II.1.** *A property is **generic** or **robust** if it holds for an open and dense subset of parameters of the social learning game.*

### II.3.2. Bayesian Rational Play

We now define rational play in the social learning game. In this subsection, we assume that the social learning game is common knowledge among players. Given the considerations in the first chapter we rely on the concept of rationalizability.

Recall that players are characterized by *beliefs* $b : \left[\underline{b}, \overline{b}\right] \times H \to [0,1]$ where $b(s,h) = Pr\left(\tilde{\theta} = 1 \mid b(\tilde{s}, \emptyset) = s, \tilde{h} = h\right)$ and associated behavioral strategies $\sigma : \left[\underline{b}, \overline{b}\right] \times H \to [0,1]$ where $\sigma(s_i, h_i) = Pr(a_i = 1 \mid s_i, h_i)$ denotes the probability of investment given realized private belief $s_i$ and history $h_i$.

**Definition II.2.** *Rational play in the social learning game $\langle n, A, u, \Theta, (G_0, G_1)\rangle$ is given by a profile of behavioral strategies $\{\sigma^*\}_{i=1}^n$ and a system of beliefs $\{b^*\}_{i=1}^n$ satisfying*

*(i) beliefs are formed according to **Bayes' rule**, i.e.*

$$b^*(s,h) = \frac{b(s,\emptyset)\, Pr\left(h \mid \tilde{\theta} = 1, \sigma^*\right)}{b(s,\emptyset)\, Pr\left(h \mid \tilde{\theta} = 1, \sigma^*\right) + [1 - b(s,\emptyset)]\; Pr\left(h \mid \tilde{\theta} = 0, \sigma^*\right)}$$

*whenever $Pr\left(h \mid \tilde{\theta} = \theta, \sigma^*\right) > 0$ for some $\theta \in \Theta$ where $Pr\left(h_i \mid \tilde{\theta} = \theta, \sigma^*\right) = \prod_{j<i} \int_{\underline{b}}^{\overline{b}} \sigma^*(a_j \mid s', h_j)\, dG_\theta(s')$, $a_j = h_i(j)$ and $h_j \subset h_i$,*

(ii) *behavioral strategies are* ***sequentially rational****, i.e.*

$$\sigma^*(1 \mid s, h_i) = \begin{cases} 1 & \text{if } b^*(s, h_i) > \frac{1}{2} \\ 0 & \text{if } b^*(s, h_i) < \frac{1}{2} \end{cases} .$$

In the following, we denote by $\sigma^*$ the strategy induced by Bayesian updating and sequential rationality. Note that $\sigma^*$ is undefined whenever $b(s, \emptyset) = 1 - b(\emptyset, h)$ i.e. in case of a tie. The occurrence of such ties with strictly positive probability is a non-generic property of the social learning game which implies that there is no need to commit to a specific tie-breaking rule. Proposition I.1 in the previous chapter shows that rational play as defined above characterizes the (absent ties) unique rationalizable outcome of the game. This obviously conincides with the unique Bayesian equilibrium outcome of the game. As in the previous chapter we stick with the latter notion.

An equivalent representation of $\sigma^*$ is given by

$$\sigma^*(s, h) = \begin{cases} 1 & \text{if } b(s, \emptyset) > 1 - b(\emptyset, h) \\ 0 & \text{if } b(s, \emptyset) < 1 - b(\emptyset, h). \end{cases}$$

Accordingly, $\sigma^*$ is a *cutoff-strategy* such that the player invests if her private belief (strictly) exceeds a certain history-dependent threshold and rejects if the private belief falls (strictly) below this threshold. This alternative representation is especially useful for the graphical illustration of strategies.

Finally, for the sake of completeness

$$U_i = \frac{1}{2} \sum_{\theta \in \Theta} \sum_{h_i \in H_i} Pr\left(h_i \mid \tilde{\theta} = \theta, \sigma\right) \int_{\underline{b}}^{\bar{b}} \sum_{a \in A} \sigma\left(a \mid s, h_i\right) \, u(a, \theta) \, dG_\theta(s)$$

denotes the ex-ante expected payoff of player $i$ (respectively in period $i$).

## II.4. The Adaptive Process

In this section we describe a learning process according to which myopic players who do not possess *a priori* a model of others' decisions might learn to play rationally the social learning game. First, we discuss the properties of the learning environment. Second, we detail the rules guiding players' learning by linking our approach with fictitious play. Finally, we discuss how our approach relates to other learning approaches.

### II.4.1. The Learning Environment

We consider an *extended social learning game* $\langle n, A, u, \Theta, (G_0, G_1), R \rangle$ where players play the social learning game repeatedly in rounds $r = 1, 2, \ldots, R$. In each round a player's position in the sequence is determined randomly such that the player is equally likely to act in any period $i \in \{1, \ldots, n\}$. For the sake of clarity, we assume that the social learning game is known by all players (this assumption is relaxed in section II.5.2). Though we assume mutual knowledge of the social learning game, we refrain from assuming any higher-order interactive knowledge. Payoffs are realized immediately at the end of each round and the state of Nature, a player's position in the sequence as well as her private belief are assumed to be i.i.d. across rounds with distributions given by $p = 1/2$, $\mathcal{U}(\{1, \ldots, n\})$ and $G_{\theta(r)}$ respectively. Since the elements of the social learning game, most importantly the distribution of private beliefs, remain unchanged across rounds we call this environment *stable*.

With regard to the state of Nature our assumptions guarantee that no learning *about the state* is possible across rounds. For instance if to the contrary the state were fixed once and for all in round 1, a player receiving new private information each round would eventually learn this state as $R \to \infty$. More generally any correlation of states across rounds confounds learning about the state and about other players' strategic behavior. Since our focus is on the latter we abstract from such possibilities entirely.

The assumption on the state straightforwardly induces independence of private beliefs across rounds.[19] On the other hand repeated play raises a difficulty concerning the distribution of private beliefs. Clearly a player's quality of information is endogeneously determined e.g. via her cognitive or material resources. Therefore heterogeneity of players' information is a reasonable assumption. While in the standard one-shot social learning game this is readily captured by a single distribution, repeated draws average out these differences. However, a player's information quality may decisively influence her learning opportunities. For instance a player with constantly very informative private beliefs will create more deviations from herds than a player with constantly less informative beliefs. In order to take this into account we propose the following extension on the distribution of private beliefs applying the *signal quality structure* model of Smith and Sørensen (2008b, section 5.2): Each player is told one of two signals $\tilde{s} \in \{0, 1\}$ where $Pr\left(\tilde{s} = \theta \mid \tilde{\theta} = \theta\right) = \tilde{q}$ denotes the signal precision. We let $\tilde{q}$ be a random variable, distributed on the set $\left(\frac{1}{2}, 1\right)$ according to the measure $Q$. For instance the Dirac measure on some $q \in (0.5, 1)$ yields the specific model of Bikhchandani, Hirshleifer, and Welch (1992). We now assume that each individual is characterized by her measure $Q$ which is private information and fixed across rounds. Accordingly, the population is characterized by a (finite discrete) probability distribution $\mathcal{Q}$ on the set $\Delta\left((1/2, 1)\right)$ of all measures $Q$ satisfiying

$$\int\limits_{Q \in \Delta((1/2,1))} \int\limits_{q=1/2}^{1} \sum_{s \in \{0,1\}} Pr(s \mid \theta, q) \, \mathbb{1}_A \left( \frac{Pr(s \mid 1, q)}{Pr(s \mid 1, q) + Pr(s \mid 0, q)} \right) Q(dq) \, \mathcal{Q}(dQ) = \int\limits_{s \in A} dG_\theta(s)$$

for each $A \subseteq [\underline{b}, \overline{b}]$ with $\mathbb{1}_A$ the indicator function of the set $A$. This construct permits us to make the following distinction: The learning process has *homogeneous support of private beliefs* if $\mathcal{Q}$ is a Dirac measure on some $Q \in \Delta((0.5, 1))$. Otherwise we say that the support of private beliefs is *heterogeneous*.

[19] A third possibility would fix both the state and every player's private belief once and for all at the beginning of the game. However, this does not permit feedback on the state of Nature. Given that others' private belief realizations are not observed either no strategic learning is possible in such a setting.

### Feedback

In line with much of the literature we will consider learning within a large population of players (see Fudenberg and Levine, 1998, section 1.2 for an overview) governed by a random matching mechanism. That is in each round a player is matched with $(n-1)$ other players to play the game. At the end of the round the player receives only feedback generated in her own match. Since a player's quality of information (see above), her position in the sequence, and her choice (see below) may significantly influence feedback we consider this a reasonable assumption.[20]

Concerning feedback generated in each round of the learning process it has been noted that in extensive form games (unlike in simultaneous-move games) players will not be able to observe complete strategies of their opponents. Rather it is usually assumed that agents observe the terminal nodes of the game i.e. the unique path through the tree. With incomplete information even the latter assumption is rigid. For instance in a social learning game observation of the terminal node requires revelation of every player's realized private information. We view this as inconsistent with the notion of private beliefs. However, given the specific payoff structure of the social learning game players are interested in others' strategies only insofar as these reveal information about the realized state of Nature. We thus assume that players' feedback in period $r$ is comprised of a subpart of $(h_{n+1}/r), \theta(r))$, the (complete) sequence of actions and the state of Nature in period $r$. (Furthermore apparently, $s_i(r) \in y_i(r)$). Along each dimension we distinguish two cases: (i) whether or not the state of Nature is observed regardless of a player's choice, and (ii) whether or not choices made after a player's own decision are revealed to the player. Regarding the first point notice that given the payoff function a player who invests will eventually get to know the state when payoffs are realized at the end of a round. On the other hand players who do not invest cannot distinguish states from payoffs alone. Still, it might be possible for them to learn the state of Nature in a round for instance by observing success or failure of others. Accordingly, if observation of the state depends on (does not depend on) a player's choice we say that feedback on the state is *(un)conditional*. The second point concerns the question whether players "stay on the market" after having decided. In particular feedback on choices is *complete* if $h_{n+1}(r) \in y_i(r)$ for each $i = 1, \ldots, n$ where $y_i(r)$ denotes the feedback player $i$ receives at the end of round $r$. On the contrary feedback on choices is *partial* if for each player $i$, $a_j(r) \in y_i(r)$ if and only if $j < i$.[21]

Note that one could assume a learning process in which players observe not the complete sequence of choices but only statistics (or signals respectively, see Dekel, Fudenberg, and Levine, 2004) generated by this sequence. For instance players may only observe relative investment frequencies conditional on the state. In this case the associated learning process cannot converge to Bayesian equilibrium but may be expected to approach the coarse analogy-based expectations equilibrium (ABEE) studied by Guarino and Jehiel (2009).

---

[20] Different matching protocols provide players with feedback of all matches which circumvents these problems. However, we consider active experimentation a more reasonable way to cope with limited feeback; see appendix II.A.

[21] Compare this to the study of Esponda (2008) on learning in auctions. Esponda distinguishes (i) whether players who do not win the object receive information about its value, and (ii) whether players who win the object receive information about the second highest bid. He finds that non-Nash equilibria and inefficiencies may arise if information is not complete.

We conclude this subsection with a definition.

**Definition II.3.** *If $R \to \infty$ (resp. $R < \infty$) then the learning horizon is said to be* ***infinite*** *(resp.* ***finite****).*

### II.4.2. Frequentist Assessments and Myopic Responses

We now detail the rules governing players' learning. We stick closely to the ideas underlying the concept of *fictitious play* (Brown, 1951).

Denote by $\mathbf{y}(r)$ the vector of feedback players receive at the end of round $r$. In each round $r$, $\mathbf{y}(r) \in Y$ where for instance in the case of unconditional state feedback and complete observations $Y = A^n \times \Theta$. Let $\zeta_r = (\mathbf{y}(1), \ldots, \mathbf{y}(r-1))$ denote the collection of feedback at the beginning of round $r$. Clearly, $\zeta_r \in \mathscr{Z}_r = Y^{r-1}$ where $\mathscr{Z}_1 = \emptyset$. We call $\zeta_r$ a realized *super-history* in round $r$ while $\zeta \in \mathscr{Z} = Y^\infty$ denotes an infinite super-history.

As discussed above a player never gets to know other players' private belief realizations. Therefore it is not possible for her to learn others' complete behavioral strategies $\sigma : \left[\underline{b}, \overline{b}\right] \times H \to [0,1]$. At least this holds true absent strong assumptions as for instance common knowledge of the distribution of private information, payoff functions, belief updating and response to beliefs which we would like to abstain from. Yet, we have argued already that players care about others' strategies only insofar as to be able to draw inferences from observed choices about the realized state of Nature. More precisely players care about conditional probabilities of histories $Pr\left(h \mid \tilde{\theta} = \theta\right)$. It is thus sufficient for players to learn these probabilities. Accordingly for some fixed representative player with feedback $y$ we let $\varphi(h \mid \theta)$ ($\varphi : H \times \Theta \to [0,1]$) denote the probability assigned to history $h$ conditional on state $\theta$. $\varphi$ consequently satisfies $\sum_{h_i \in H_i} \varphi(h_i \mid \theta) = 1$ for each $i = 1, \ldots, n$ and each $\theta \in \Theta$. Sometimes it is more convenient to work with probabilities of the sort $Pr\left(a_i \mid h_i, \tilde{\theta} = \theta\right)$ for $a_i \in A$ and $i = 1, \ldots, n$. These are one-to-one related to probabilities $Pr(h_i \mid \theta)$ via $Pr((h_i, a_i) \mid \theta) = Pr(h_i \mid \theta) * Pr\left(a_i \mid h_i, \tilde{\theta} = \theta\right)$ given independence *across periods*. Notice that the latter is a weak assumption as the extensive-form structure of the game[22] already allows for correlation across periods. Therefore in a slight abuse of notation we denote by $\varphi(a \mid h, \theta)$ the probability assigned to observing action $a$ after history $h$ conditional on state $\theta$. Clearly, $\varphi(h_i \mid \theta) = \prod_{j=1}^{i-1} \varphi(a_j \mid h_j, \theta)$ and $\varphi(a_i \mid h_i, \theta) = \varphi((h_i, a_i) \mid \theta)/\varphi_i(h_i \mid \theta)$, i.e. either mapping uniquely defines the other.

As in fictitious play we have players learn via a frequentist approach. That is players keep track of frequencies of observations conditional on history and state. Formally, let $\kappa(a, h, \theta \mid \zeta_r)$ ($\kappa : A \times H \times \Theta \times \bigcup_{r=1}^{\infty} \mathscr{Z}_r \to \mathbb{N}$) denote the number of times in rounds $\rho = 1, \ldots, r$ that action $a \in A$ occurred after history $h$ when the observed state was $\theta$. Hence, $\kappa(a, h, \theta \mid \zeta_r) = \left|\{1 \leq \rho \leq r : \{h, a, \theta\} \subseteq y(\rho)\}\right|$. Players then use these frequencies to update their state-contingent probabilities of histories. Accordingly, we define *assessments* $\hat{\varphi} : H \times \Theta \times \bigcup_{r=1}^{\infty} \mathscr{Z}_r \to \Delta(A)$ via

$$\hat{\varphi}(a \mid h, \theta; \zeta_r) = \begin{cases} (1-\epsilon) * \frac{\kappa(a,h,\theta \mid \zeta_r)}{\sum_{a' \in A} \kappa(a',h,\theta \mid \zeta_r)} & \text{if } \kappa(a, h, \theta \mid \zeta_r) > 0 \\ \epsilon & \text{if } \kappa(a, h, \theta \mid \zeta_r) = 0 \end{cases} \tag{II.1}$$

[22]This considers the games the way they are introduced with complete observation of the past. Under different assumptions on the observation structure, this no longer holds true in general.

for each $r = 1, \ldots, R$, $\zeta_r \in \mathscr{Z}_r$, $\theta \in \Theta$, $h \in H$ and $a \in A$ while $\hat{\varphi}(h_i \mid \theta; \zeta_r) = \prod_{j<i} \hat{\varphi}(a_j \mid h_j, \theta; \zeta_r)$. At this point we deviate from the original fictitious play idea in that we do not require players to have initial weights $\kappa(a, h, \theta \mid \emptyset) > 0$ for each $a$, $h$ and $\theta$. The main reason is that we want to also study the process with finite repetitions where the influence of such weights is not negligible. However, without initial weights assessments and derived posteriors may not always be well-defined. We deal with this problem using the idea of trembles: as long as a player cannot rely on accumulated evidence she attaches a small probability $\epsilon > 0$ to each observation she cannot rule out a priori. We will then study the limit as $\epsilon \to 0$. This approach implies that players attach probability zero to histories never observed before and therefore do not take them into account in calculating expected payoffs. On the other hand beliefs at such histories are well-defined: Observations made for the first time are assumed to carry no information about the state.

We return to the fictitous play approach in defining players' responses to assessments. In particular we assume that players in any round take into account their current observations in a myopic way. In other words players do not engage in strategic repeated play considerations, an assumption justified for instance in large population models. To determine myopic response to observations fix an assessment $\hat{\varphi}$ and denote by $\hat{\sigma} : \bigcup_{r=1}^{\infty} \mathscr{Z}_r \to \Sigma$ a *strategic response*. $\hat{\sigma}$ is *myopic Bayes-rational* iff

$$\hat{\sigma}(s, h \mid \zeta_r) = \begin{cases} 1 & \text{if } b(s,h) > 1 - \hat{b}(\emptyset, h \mid \zeta_r) \\ 0 & \text{if } b(s,h) < 1 - \hat{b}(\emptyset, h \mid \zeta_r) \end{cases}$$

where

$$\hat{b}(\emptyset, h \mid \zeta_r) = \frac{\hat{\varphi}(h \mid 1; \zeta_r)}{\hat{\varphi}(h \mid 1; \zeta_r) + \hat{\varphi}(h \mid 0; \zeta_r)}$$

is the assessed public belief for history $h$ at super-history $\zeta_r$. Assuming myopic Bayes-rational responses also means that we abstract from learning costs i.e. players maximize the sum of their per-interaction expected payoffs. We do this to give best chances to the emergence of the equilibrium outcome. In subsection II.6.2 we return briefly to the issue of learning costs.

We have now collected all elements of the learning process. Obviously, our aim is to study the limit outcome of this process in terms of beliefs and choices as $r \to R$. At any super-history $\zeta_R$ both are uniquely determined by the matrix of assessments $(\hat{\varphi}(h \mid \theta; \zeta_R))_{h \in H, \theta \in \Theta}$ at all histories satisfying $\sum_{\theta \in \Theta} \hat{\varphi}(h \mid \theta; \zeta_R) > 0$. On the other hand uniqueness of the Bayesian equilibrium outcome guarantees that if $\hat{\varphi}(\zeta_R)$ coincides with equilibrium assessments $\varphi^*(h \mid \theta) = Pr(h \mid \tilde{\theta} = \theta, \sigma^*)$ at all histories on the equilibrium path associated strategies and public beliefs must coincide (on the path) as well. Consequently we say that the learning process *approaches the equilibrium outcome* along super-history $\zeta_R$, if

$$\lim_{r \to R} \hat{\varphi}(\zeta_r) = \varphi^*$$

where for each $r \leq R$, $\zeta_r \subseteq \zeta_R$. However, the limit outcome of the learning process along a particular super-history is not very informative. Our focus is therefore on the probabilistic limit of the stochastic process $\hat{\varphi}(\tilde{\zeta}_r)$. Accordingly, denoting by $\mathbf{P}$ the distribution of final super-histories $\tilde{\zeta}_R$, we say that the learning process *eventually approaches the equilibrium outcome*, if

$$\mathbf{P}\Big(\zeta_R \in \mathscr{Z}_R : \lim_{r \to \infty} \hat{\varphi}(\zeta_r) = \varphi^*\Big\} = 1.$$

### II.4.3. Other Approaches to Learning in Games

Learning in games has been extensively studied in the past (for an overview see e.g. Fudenberg and Levine, 1998, 2009). However, most of this work has focused on normal-form games with complete information where players can perfectly observe strategies of others. This feature is usually inappropriate in extensive-form games of incomplete information. Much of the literature on such games has therefore examined the relationship between the feedback players receives about others' strategies and characteristics of possible limit outcomes of the associated learning process. This has led to the development of the weaker concept of a subjective/conjectural/self-confirming equilibrium (see Battigalli, 1997; Kalai and Lehrer, 1993; Fudenberg and Levine, 1993a; Fudenberg and Kreps, 1995; Dekel, Fudenberg, and Levine, 2004). Only few studies on learning in extensive-form games study the question whether the learning process converges at all. Most notably among those are the study of Fudenberg and Levine (1993b) which examine a Bayesian learning process and more recently Beggs (1993) in Bayesian games with binary actions. In such settings optimal strategies are specified by a single threshold. Beggs shows that players can learn this optimal strategy by adjusting employed thresholds over time via a simple adaptive rule given limited feedback.

In general learning models can be classified into three classes – rational, epistemic and behavioral learning (see Walliser, 1998; Hart, 2005). The fictitious play approach we pursue in this chapter belongs to the class of epistemic learning models in which players best respond to conjectures about others' future strategies. Conjectures are revised over time in a simple, unsophisticated, mechanical way in response to the (perhaps imperfect) observation of others' actions taken previously. Epistemic learning is thus more sophisticated than behavioral learning where players no longer maximize but choose strategies probabilistically according to their success in past rounds. Strategies which fare well in terms of payoff are reinforced while strategies which fare worse are inhibited. Yet, while conceptionally epistemic and behavioral learning are very different they are close from a mathematical point of view (see for instance Hopkins, 2002 or Camerer and Ho, 1999).[23] Given that we are interested in conditions which prevent players from approaching equilibrium behavior, choice of a more sophisticated epistemic learning model would not seem to work to our advantage.

On the other hand epistemic learning is still rather naive in the way conjectures are updated and responded to. In the direction of higher sophistication alternative models permit active experimentation with suboptimal strategies or have players update beliefs in a Bayesian way (rational learning). In this sense our model is still restrictive. However, we do not think that our main result – the fact that the presence of both strategic and structural uncertainty may prevent players from learning to behave optimally – is sensitive to the choice of the adaptive process. For instance assume that players are endowed with a prior on the strategy space and a prior on the space of the unknown parameters of the model (the space of private belief distributions) and across repetitions update these priors in a Bayesian way in the light of the feedback received. Assume that a history becomes less informative. This change may either be attributed to a different distribution of strategies (more noise players) or a different distribution of private beliefs (less

[23] Indeed, the model of Beggs can be interpreted as a combination of the two.

informative). The feedback we have discussed does not permit players to distinguish between these effects. Therefore some uncertainty must prevail and by similar arguments as below this may lead to systematic and severe mistakes in the outcome of the learning process. We leave a thorough investigation of this conjecture for future research. We do however address the idea of experimentation in an extension of our learning process presented in Appendix II.A.

## II.5. Perfect Learning Opportunities

This section presents the first of our key results: If the environment is stable and the learning horizon is infinite then the learning process eventually approaches the equilibrium outcome. We also argue that this does not depend on knowledge about the primitives of the model. A simple consequence is that under such "perfect" learning opportunities, best responding to posteriors formed in a Bayesian way maximizes the (ex-ante) expected payoff.

### II.5.1. Learning Conditions for Rational Play

In social learning a player's task comprises two challenges. First, the player has to draw correct inferences from previous choices and combine this information with her private information in an optimal way. Second, the player must best respond to the resulting belief given her payoff function. However, in period 1 no inferences from observed choices need to be drawn, players are left with the task of responding to private information. Therefore by myopia of the learning process every player will always follow the Bayesian equilibrium strategy in the first period. This on the other hand implies that given the stability of the environment inferences from the first player's decision eventually approach the correct ones provided first period decisions have been observed sufficiently often. Thus, players learn to make correct inferences in the second period. Accordingly their behavior approaches the equilibrium behavior in the second period as well. More generally, by a simple inductive argument this holds for all situations occurring with strictly positive probability as $R$ increases without bounds. In other words the learning process unravels. Consequently as $R \to \infty$ the unique equilibrium outcome is attained. Notice finally that this happens independent of feedback and private belief support as with $R \to \infty$ each player infinitely often decides at position $n$ where even with heterogeneous private belief support she will infinitely often make an observation which leads her to invest.

**Proposition II.1.** *Under perfect learning opportunities, i.e. a stable learning environment and an infinite learning horizon, the learning process eventually approaches the unique Bayesian equilibrium outcome regardless of the support of private beliefs and the feedback on state and choices.*

Perfect learning opportunities correspond to the necessary[24] and sufficient conditions that permit the learning process to eventually approach equilibrium.

<u>Remark:</u> The proof of this proposition (provided in Appendix II.D) is basically a formalization of the arguments presented above. A slight complication arises in dealing with situations off the

[24] Necessity is shown in the next section.

equilibrium path. Therefore what we can show is that the learning process (randomly) selects a Bayesian equilibrium. In section II.A of the appendix we return to the issue of equilibrium selection for a more sophisticated version of the learning process. It turns out that we can sharpen the result only insofar as with the sophisticated process *any* Bayesian equilibrium may possibly be selected.

The following result is a straightforward consequence of the proposition above.

**Corollary.** *With perfect learning opportunities, for each period $i = 1, \ldots, n$ the Bayes-rational strategy $\sigma^*$ maximizes the (ex-ante) expected payoff $U_i(\sigma)$ and uniquely so in some social learning games.*

*Proof.* This follows by the same argumentation as applied in deriving the Bayesian equilibrium of the standard social learning game: Given correct assessments $\varphi^*$ of other players' strategies, in each period expected payoff is maximized by best responding to Bayesian posteriors. Since with perfect learning opportunities the assessments in the learning process eventually converge to equilibrium assessments for histories on the path and since furthermore off-path-histories have no bearing on the expected payoff, Bayesian updaters will eventually obtain the highest expected payoff. While with discrete private belief distribution the maximizer (on $\Sigma$) is not necessarily unique, this is guaranteed in settings with rich (continuous) private belief distribution. □

### II.5.2. Ex-ante Knowledge

So far we have assumed that the primitives of the game are known by the players *a priori*. In this subsection we discuss the robustness of the result from the previous subsection with regard to this assumption. Throughout this subsection we maintain the assumption of perfect learning opportunities. Furthermore, given the considerations in Proposition II.1 we restrict attention to the learning process with homogeneous private belief support and complete and unconditional feeedback. Finally, we will not question that players possess knowledge about their opportunities (action space). Under these assumptions we argue that indeed little has to be known as long as the truth is not ruled out a priori.

Consider first the structure of interactions. Assume that the truth is not ruled out a priori. That is no player believes with certainty that she makes her decision in complete isolation. Said differently every player attaches strictly positive probability to the existence of both payoff externalities and information externalities. Then on one hand repeated payoff information will eventually rule out payoff externalities. On the other hand players' awareness of the possibility that others' choices convey information about the state of Nature is sufficient to have them take it into account. More precisely if a certain history is observed much more often in combination with one rather than the other state of Nature players will realize this and update beliefs accordingly. Similar considerations imply that knowledge about others' action spaces is superfluous provided respective priors are non-degenerate. Since players will eventually observe any action occurring with strictly positive probability they will eventually explore any player's action space. The assumption on priors rules out the necessity to update probability-zero-events. Finally, it is obvious that the adaptive process we have set up does not require players to have information about others' payoffs. Yet, such information might accelerate the learning substantially. Note however that in the literature it is rarely assumed that players do receive feedback about others' payoffs.

Second, suppose that players lack knowledge about the distribution of Nature's moves, i.e. the prior on the state space and the distribution of their own private information. In this case players are faced with a second task, the classical statistical learning problem of making inferences from i.i.d. data. Formally, for some fixed representative player with feedback $y$ define $\kappa(\theta \mid \zeta_R) = \left|\{1 \leq r \leq R : \theta \in y(r)\}\right|$ and $\kappa_i(A, \theta \mid \zeta_r) = \left|\{1 \leq r \leq R : \theta \in y(r) \text{ and } s \in A\}\right|$ where $s$ is the player's private belief in round $r$. Let $\hat{\phi} : \bigcup_{r=1}^{\infty} \mathscr{Z}_r \to \Delta(\Theta \times [0,1])$ denote the player's *assessment on Nature* where $\hat{\phi}(\theta; \zeta_r) = \hat{\phi}(\theta, [0,1] \mid \zeta_r)$ is her assessment of the prior and $\hat{\phi}(s \mid \theta; \zeta_r) = \hat{\phi}(\theta, ds \mid \zeta_r)$ is her assessment on the distribution $G_\theta$. In a similar fashion as for assessments regarding observations we define

$$\begin{aligned} \hat{\phi}(\theta; \zeta_r) &= \frac{\kappa(\theta \mid \zeta_r)}{r}, \\ \hat{\phi}(A \mid \theta; \zeta_r) &= \frac{\kappa(\theta, A \mid \zeta_r)}{\kappa(\theta \mid \zeta_r)}. \end{aligned}$$

With uncertainty about Nature's moves a player computes her private belief from the given $s$ in a Bayesian way using her assessment $\hat{\phi}$:

$$b(s, \emptyset \mid \zeta_r) = \frac{\hat{\phi}(1; \zeta_r)\, \hat{\phi}(s \mid 1; \zeta_r)}{\hat{\phi}(1; \zeta_r)\, \hat{\phi}(s \mid 1; \zeta_r) \;+\; \hat{\phi}(0; \zeta_r)\, \hat{\phi}(s \mid 0; \zeta_r)}.$$

The law of large numbers guarantees in most settings[25] that along any $\zeta$, for any $\theta \in \Theta$ and any $s \in [\underline{b}, \overline{b}]$, $\hat{\phi}(\theta; \zeta_r) \to \frac{1}{2}$ and $\hat{\phi}(s \mid \theta; \zeta_r) \to dG_\theta(s)$. Therefore the results of the previous subsection are robust to uncertainty about Nature's move. Notice that this even extends to uncertainty about state and private belief space: If players don't rule out any possible state or private belief a priori, the support of the probabilities will asymptotically be revealed.

Finally, if players are uncertain about their own payoffs then having them engage in a prepended trial-and-error stage of learning will eventually resolve this uncertainty. More precisely if each action is tried sufficiently often in any circumstance, players should be able to learn how payoffs depend (or do not depend) upon own choices, others' choices and Nature's draw.

In conclusion, the results from the previous subsection are robust with regard to the assumptions on ex-ante knowledge.

## II.6. Imperfect Learning Opportunities

### II.6.1. Mixed Worlds

Our convergence result (Proposition II.1) requires that players play the same social learning game a great (infinite) number of times. In the real world this will rarely be the case. It has however been

[25] In a recent paper, Al-Najjar (2009) shows that successfully learning from i.i.d. data might be hard in some settings. More precisely if the underlying space is discrete infinite and the $\sigma$-algebra of all subsets has to be learned, learning the correct probability is impossible. Since in our case the state space is binary and the private belief space bounded, his results do not apply to this setting.

argued that "any sort of learning involves extrapolation from past observations to settings that are deemed (implicitly or explicitly) to be similar, so what matters is how often agents have played 'similar' games" (Fudenberg, 2006, p. 701).[26] However, while the idea is prevalent in economics[27] few studies explicitly model learning across similar games.

We extend our adaptive process in order to investigate which strategic behavior emerges if adaptation takes place across contexts. In particular we analyze whether players which respond in a Bayes-rational fashion to assessments continue to eventually learn the optimal strategy. Furthermore we discuss the behavioral differences which arise. Throughout this subsection we maintain the assumption of an infinite learning horizon ($R \to \infty$).

Let $E_k = \left\langle n, \Theta, A, u, \left(G_0^{(k)}, G_1^{(k)}\right)\right\rangle$ denote a social learning *environment* with $\left[\underline{b}_k, \overline{b}_k\right]$ the convex hull of the common support of $(G_0^k, G_1^k)$. A collection of social learning environments $\mathscr{E} = \{E_k : k = 1, \ldots, K\}$ with associated probability vector $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_K) \in \Delta(\{1, \ldots, K\})$[28] will be called a (finite) ***social learning world***. Notice that environments in $\mathscr{E}$ differ only in their distribution of private information. In principle one could assume more generally worlds whose environments differ in either element of the basic structure. However, non-revelation of the environment is a key to our argumentation (see below) and we decided not to question players' abilities (or willingness) to distinguish games according to the number of players, the state or action space, or the payoff function.[29]

We amend our learning process by assuming that in each round $r$, the relevant social learning environment $\tilde{E}(r)$ is drawn from $\mathscr{E}$ according to $\boldsymbol{\alpha}$ independently across rounds. Furthermore for any round $r$, $E(r) \notin y(r)$, i.e. the environment drawn is not revealed to players. The latter assumption presumes that players are not able to identify the environment they are acting in. Consequently our argumentation is much stronger than considered in the literature: Rather than across *similar* games we have players learn across games which are *indistinguishable* to them. We claim that especially with regard to the distribution of private information this presumption is justified. First and foremost we have serious doubts about players knowing the distribution from which other players draw their private beliefs respectively even the space of possible distributions. To be clear we do not question the ability of players to correctly interpret their *own* private information. Rather we claim that assuming (mutual) knowledge of the distribution of private beliefs is heroic.[30] In particular it would presuppose an amount of prior information which as argued above is prohibitively costly. Notice that we therefore challenge the Harsanyi doctrine (Harsanyi, 1967)

---

[26]Fudenberg and Levine (1998) denote that the "presumption that players *do* extrapolate across games they see as similar is an important reason to think that learning models have some relevance to real-world situations" (p.4).

[27]Examples include Gilboa and Schmeidler's (1995) case-based decision theory and Jehiel's (2005) analogy-based expectations equilibrium.

[28]That is $\alpha_k > 0$ for each $k = 1, \ldots, K$ and $\sum_{k=1}^{K} \alpha_k = 1$.

[29]This does not reflect a conviction, that players possess such abilities. Indeed, it is not clear, how a player might identify the correct environment when she is about to act. For a generalization of this idea, see the next subsection.

[30]If a player were to know the set of possible distributions and furthermore these would not have common support, her own private belief may provide the player with sufficient information to identify the environment. But with common support even knowledge of $\mathscr{E}$ is in general not sufficient.

which subjugates structural uncertainty by assuming common knowledge of a prior probability distribution on the types of players (and Nature). In other words we assume a more *fundamental* structural uncertainty.[31] Second, we claim that our assumptions on the feedback players receive (see section II.4.1) do not permit players to derive this distribution in the course of the learning process. Absent feedback about other players' private belief realizations a player's sole direct information about the distribution of private beliefs is the sequence of her own realizations across rounds. While these permits players to eventually learn the *average* distribution of private beliefs in each state, nothing allows them to disentangle the different environments that make up for this average. For instance the world with the two environments given by $G^1_\theta$ and $G^2_\theta$, $\theta \in \{0, 1\}$, drawn according to $(\alpha, 1 - \alpha)$ generates the same limiting frequencies of private beliefs conditional on either state as the world comprised of the single environment $\overline{G}_\theta = \alpha\, G^1_\theta + (1 - \alpha)\, G^0_\theta$ for each $\theta \in \{0, 1\}$. Finally, we abstract from extreme worlds where players might be able to easily identify single environments from their observations.

For a given social learning world $\mathscr{E} = \{E_k : k = 1, \dots, K\}$ we denote by $\overline{E}$ the mixed environment given by private belief distributions $\overline{G}_\theta = \sum_{k=1}^{K} \alpha_k\, G^k_\theta$. Furthermore we let $\sigma^*_k$ (respectively $\overline{\sigma}^*$) denote the Bayesian equilibrium strategy in environment $E_k$ (respectively $\overline{E}$) as characterized in Definition II.2 with associated equilibrium assessments $\varphi^*_k$ (respectively $\overline{\varphi}^*$). Our first result establishes that players learning in a social learning world without feedback about the environment chosen in each round adapt to the mixed environment $\overline{E}$.

**Proposition II.2.** *Players adapt to the mixed environment $\overline{E}$: Beliefs and choices in the adaptive process on social learning world $\mathscr{E}$ eventually approach beliefs and choices (on the path) of the unique Bayesian equilibrium outcome of the mixed environment $\overline{E}$, i.e.*

$$\mathbf{P}\left(\lim_{r\to\infty} \hat{\varphi}(a \mid h, \theta; \tilde{\zeta}_r) = \overline{\varphi}^*(a \mid h, \theta)\right) = 1$$

*for each $a \in A$, $\theta \in \Theta$, and $h \in H$ s.t. $\overline{\varphi}^*(h \mid \theta) > 0$.*

*Proof.* See appendix II.D. □

Subsequently we call $\overline{\sigma}^*$ the *learned strategy* in social learning world $\mathscr{E}$. A trivial implication of the result above is the possible discrepancy between learned behavior and optimal behavior in single environments.

**Corollary.** *Generically for some social learning environment $E_k$ ($k \in \{1, \dots, K\}$), the learned strategy $\overline{\sigma}^*$ does not coincide with the optimal strategy $\sigma^*_k$ at all histories reached a non-vanishing fraction of the time.*[32]

*Proof.* We provide a generic example: For $K = 2$ let $G^k_\theta$, $k = 1, 2$ be continuous and satisfy $G^1_0(1/2) * G^2_0(1/2) < G^1_1(1/2) * G^2_1(1/2)$. Furthermore $0 < \alpha = \alpha_1 = 1 - \alpha_2 < 1$.

---

[31] Indeed a better term for structural uncertainty about which is resolved by endowing players with common knowledge of a distribution over fundamentals would be *structural risk*.

[32] That is $\overline{\sigma}^*(s, h) \neq \sigma^*_k(s, h)$ for some $h \in H$ and each private belief $b(s, \emptyset)$ in some subset $B \subseteq supp(G^k_\theta)$ where $\int_B dG^k_\theta(s) > 0$ and $\overline{\varphi}^*(h \mid \theta) + \varphi^*_k(h \mid \theta) > 0$ for some $\theta \in \Theta$.

Since players follow private information in the first period a rejection causes second period assessments $\varphi_k^*(0 \mid \emptyset, \theta) = G_\theta^k(1/2)$ and $\overline{\varphi}*(0 \mid \theta) = \alpha\, G_\theta^1(1/2) + (1-\alpha)\, G_\theta^2(1/2)$ respectively. Consequently the associated public belief satisfies

$$\begin{aligned} b(\emptyset, (0) \mid \overline{\varphi}*) &= \frac{\alpha\, G_1^1(1/2) + (1-\alpha)\, G_1^2(1/2)}{\alpha \left(G_1^1(1/2) + G_0^1(1/2)\right) + (1-\alpha) \left(G_1^2(1/2) + G_0^2(1/2)\right)} \\ &< \frac{G_1^1(1/2)}{G_1^1(1/2) + G_0^1(1/2)} = b(\emptyset, (0) \mid \varphi_1^*). \end{aligned}$$

Accordingly by continuity of private belief distributions, players eventually learn to imitate the first player's rejection too often in environment $E_1$. □

The corollary ascertains that players which are not able to distinguish social learning environments cannot achieve maximal expected payoff. However, the Bayes-rational strategy may still be the optimal choice under the restriction imposed by the environment. Our main result below invalidates this claim. Players who do not respond to assessments in a Bayes-rational way may achieve a higher expected payoff.

To establish this result let $\overline{\boldsymbol{\varphi}}^* = \left(\overline{\varphi}_1^*, \ldots, \overline{\varphi}_K^*\right)$ denote the vector of assessments associated with strategy $\overline{\sigma}^*$.[33] Expected payoff in period $i$ is then given by[34]

$$U_i^{\mathscr{E}}(\sigma \mid \overline{\boldsymbol{\varphi}}^*) = \frac{1}{4} \sum_{k=1}^{K} \alpha_k \sum_{h_i \in H_i} \int_{\underline{b}_k}^{\overline{b}_k} \sigma(s, h_i) \left[\overline{\varphi}_k^*(h_i \mid 1)\, dG_1^k(s) - \overline{\varphi}_k^*(h_i \mid 0)\, dG_0^k(s)\right].$$

**Lemma II.1.** *The benchmark strategy $\sigma_{\mathscr{E}}^*$ given by*

$$\sigma_{\mathscr{E}}^*(s, h) = \begin{cases} 1 & \text{if } \sum_{k=1}^{K} \alpha_k\, \overline{\varphi}_k^*(h \mid 1)\, dG_1^K(s) > \sum_{k=1}^{K} \alpha_k\, \overline{\varphi}_k^*(h \mid 0)\, dG_0^k(s) \\ 0 & \text{if } \sum_{k=1}^{K} \alpha_k\, \overline{\varphi}_k^*(h \mid 1)\, dG_1^k(s) < \sum_{k=1}^{K} \alpha_k\, \overline{\varphi}_k^*(h \mid 0)\, dG_0^k(s) \end{cases}$$

*maximizes expected payoff $U_i^{\mathscr{E}}(\sigma \mid \overline{\boldsymbol{\varphi}}^*)$ on $\Sigma$ for each $i = 1, \ldots, n$.*

*Proof.* See appendix. □

Generically (across social learning worlds) this benchmark strategy does NOT coincide with the Bayes-rational strategy.

**Proposition II.3.** *Generically the learned strategy $\overline{\sigma}^*$ does not coincide with the optimal strategy $\sigma_{\mathscr{E}}^*$ at all histories reached a non-vanishing fraction of the time.*[35]

---

[33]That is for each $k = 1, \ldots, K$, $h \in H$, $\theta \in \Theta$, and $a \in A$, $\overline{\varphi}_k^*(a \mid h, \theta) = \int_{\underline{b}_k}^{\overline{b}_k} \overline{\sigma}^*(s, h)\, dG_\theta^k(s)$.

[34]Notice that since $n$ and $A$ are the same across environments in $\mathscr{E}$, all sets $H_i$ are the same in each environment.

[35]That is $\overline{\sigma}^*(s, h) \neq \sigma_{\mathscr{E}}^*(s, h)$ for some $h \in H$ and each $b(s, \emptyset)$ in some $B \subseteq \bigcup_{k=1}^{K} \left[\underline{b}_k, \overline{b}_k\right]$ where $\sum_{k=1}^{K} \alpha_k \int_{s \in B} dG_\theta^k(s) > 0$ and $\sum_{k=1}^{K} \alpha_k\, \overline{\varphi}_k^*(h \mid \theta) > 0$ for some $\theta \in \Theta$.

*Proof.* We provide a second generic example with $K = 2$. For $k = 1, 2$, $G^k_\theta$ is continuously distributed on $supp(G^k_\theta) = [1 - a_k, a_k]$ according to conditional densities $g^k_0(s) = 2(1-s)/2a_k - 1$ and $g^k_1(s) = 2s/(2a_k - 1)$. W.l.o.g. $a_1 < a_2$. The optimal benchmark $\sigma^*_{\mathscr{E}}$ can be straightforwardly derived as

$$\sigma^*_{\mathscr{E}}(s,h) = \begin{cases} 1 & \text{if } s > 1 - \check{b}(\emptyset, h \mid \overline{\boldsymbol{\varphi}}^*) \\ 0 & \text{if } s < 1 - \check{b}(\emptyset, h \mid \overline{\boldsymbol{\varphi}}^*) \end{cases}$$

where the *benchmark public belief* is given by

$$\check{b}(\emptyset, h \mid \overline{\boldsymbol{\varphi}}^*) = \frac{\sum_{k=1}^{K} \frac{\alpha_k}{2a_k - 1} \overline{\varphi}^*_k(h \mid 1)}{\sum_{k=1}^{K} \frac{\alpha_k}{2a_k - 1} \left[\overline{\varphi}^*_k(h \mid 1) + \overline{\varphi}^*_k(h \mid 0)\right]}.$$

On the other hand if responding to assessments in a Bayes-rational fashion players eventually invest if $b(s, \emptyset) > 1 - b(\emptyset, h \mid \overline{\boldsymbol{\varphi}}^*)$ and reject if $b(s, \emptyset) < 1 - b(\emptyset, h \mid \overline{\boldsymbol{\varphi}}^*)$ where the *learned public belief* is given by

$$b(\emptyset, h \mid \overline{\boldsymbol{\varphi}}^*) = \frac{\sum_{k=1}^{K} \alpha_k \overline{\varphi}^*_k(h \mid 1)}{\sum_{k=1}^{K} \alpha_k \left[\overline{\varphi}^*_k(h \mid 1) + \overline{\varphi}^*_k(h \mid 0)\right]}.$$

As before continuity of private beliefs implies that optimal and learned behavior differ unless at all histories occurring with strictly positive probability either learned and optimal public belief coincide or both induce the same action with probability one (i.e. both lie strictly within the same cascade set). By straightforward algebraic transformation coincidence of the beliefs is equivalent to

$$\frac{2\,(a_2 - a_1)}{(2a_2 - 1)(2a_1 - 1)} * \left[\overline{\varphi}^*_2(h \mid 1)\overline{\varphi}^*_1(h \mid 0) - \overline{\varphi}^*_1(h \mid 1)\overline{\varphi}^*_2(h \mid 0)\right] = 0.$$

Therefore learned behavior is optimal if and only if for each history occurring a strictly positive fraction of the time either $\overline{\varphi}^*_1(h \mid 1) * \overline{\varphi}^*_2(h \mid 0) = \overline{\varphi}^*_2(h \mid 1) * \overline{\varphi}^*_1(h \mid 0)$, or learned and optimal public belief lie strictly within a cascade set. Generically this is not satisfied.

For instance in period 2

$$\begin{aligned} \overline{\varphi}^*_k(1 \mid 1) &= 1 - G^k_1(1/2) = \frac{a_k}{2} + \frac{1}{4} > 0 \\ \text{and } \overline{\varphi}^*_k(1 \mid 0) &= 1 - G^k_0(1/2) = \frac{3}{4} - \frac{a_k}{2} > 0 \end{aligned}$$

jointly imply

$$\begin{aligned} \overline{\varphi}^*_2(1 \mid 1)\overline{\varphi}^*_1(1 \mid 0) - \overline{\varphi}^*_1(1 \mid 1)\overline{\varphi}^*_2(1 \mid 0) &= \frac{a_2 - a_1}{2} > 0, \\ b(\emptyset, h \mid \overline{\boldsymbol{\varphi}}^*) &= \frac{1}{4} + \frac{\alpha a_1 + (1-\alpha)a_2}{2} \in (1 - a_1, a_2) \end{aligned}$$

where $1 - a_1 < \frac{1}{4} + \frac{\alpha a_1 + (1-\alpha) a_2}{2}$ follows from $a_1 < \alpha a_1 + (1-\alpha) a_2 < a_2$ and $a_2 > a_1 > 1/2$. A similar result holds for $h_2 = (0)$. □

Proposition II.3 demonstrates how Bayes-rationally responding to assessments ceases to be the optimal strategy when players decide in a world of multiple indistinguishable social learning environments. Said differently the optimal benchmark shifts[36] when players adapt across social learning games. This idea has become prominent by the term *rule rationality* (Aumann, 1997, 2008, the opposite notion is *act-rationaliy*). In single decision tasks rule rational choices may lead to severe and systematic errors as demonstrated for instance in the corollary to Proposition II.2. Rule rationality is thus frequently employed to explain the divergence of theoretical predictions and experimental observations in both individual and strategic choice contexts.[37] Hence, our results may provide a new interpretation of the experimental results on social learning.

We are however far from claiming that players consciously derive the optimal strategy $\sigma_{\mathscr{E}}$. This would presume a degree of knowledge on the part of players which we have ruled out before. Rather we interpret the results as an indication that (seemingly) non-Bayes-rational strategies may arise in an evolutionary process of mutation and survival-of-the-fittest where we equate fitness with expected payoff. The evolutionary perspective is standard in models of rule-rationality. For instance we may assume that the rules which are inherited are directly given by strategies $\sigma \in \Sigma$. Our results then suggest that in worlds comprised of several indistinguishable environments the strategy $\sigma^*_{\mathscr{E}}$ is the most successful. Since by the result above this strategy does not involve responding to learned assessments in a Bayes-rational way, this provides an explanation for the persistence of such strategies. However, rules are more commonly interpreted as simple "rules of thumb" (Aumann, 1997) or "heuristics" (Tversky and Kahnemann, 1974) as opposed to the complex strategies $\sigma \in \Sigma$. In the remainder of this subsection we thus demonstrate how less complex strategies derived from a simple generalization of the Bayesian rule for updating probabilities may be more successful than the Bayes-rational strategy $\sigma^*$.[38]

### Simple Non-Bayesian Strategies

We consider the heterogeneous updating model studied more extensively in Chapter III. In this model players sequentially best respond to beliefs formed according to the generalized updating rule

$$b(s,h,\beta) = \frac{[b(s,\emptyset)]^\beta \; b(\emptyset,h)}{[b(s,\emptyset)]^\beta \; b(\emptyset,h) \; + \; [1-b(s,\emptyset)]^\beta \; [1-b(\emptyset,h)]}. \tag{II.2}$$

The parameter $\beta \in (0,\infty)$ is player-specific and determines how much weight a player puts on her private information compared to the public information. Apparently for $\beta = 1$ the expression reduces to the standard Bayesian updating formula. If $\beta > 1$ a player overweights her private information while for $\beta < 1$ a player overweights public information.

For a given mixed world $\mathscr{E}$ we investigate whether types $\beta \neq 1$ can have a higher expected payoff than Bayesian players. In general expected payoff also depends on the distribution of

[36] This is reminiscent of the work on robustness to model misspecification in macroeconomic risk models. In these models having decision-makers take into account possible misspecifications of their model of the world induces model uncertainty premia on equilibrium prices of risk. Hansen (2007) for instance provides a nice overview of these results and discusses furthermore how model uncertainty can arise as the outcome of imperfect learning about a complex statistical problem.

[37] Proponents of this idea are e.g. Tversky and Kahnemann (1974), Myerson (1991), and more recently Al-Najjar and Weinstein (2009). A formalization has been provided by Samuelson (2001).

[38] More illustrations can be found in the appendix (part II.C).

updating rules in the population. For simplicity we assume that the population is composed of Bayesian players ($\beta = 1$) and a single negligible player satisfying $\beta \neq 1$. Therefore our results will show whether non-Bayesian players can invade a population of Bayesian players.

The social learning world we consider is a composition of the specific model of Bikhchandani, Hirshleifer, and Welch (1992) and a variant thereof (hence $K = 2$). We maintain $A = \Theta = \{0, 1\}$. In both environments each player $i$ receives one of two possible private signals (either high, H, or low, L) which is the realization of a random variable $\tilde{s}_i$. If the realized state of Nature is 1 (0) then the probability that the signal is H (L) is equal to $1 > q > 1/2$, the signal precision. Furthermore signals are drawn independently across players. In environment $E_1$ each player's signal precision is given by $1/2 < q_1 < 1$. Hence, the private belief distribution in $E_1$ is given by $supp(G^1_\theta) = \{1 - q_1, q_1\}$ with $Pr\left(b(\tilde{s}_i, \emptyset) = q_1 \mid \tilde{\theta} = 1\right) = Pr\left(b(\tilde{s}_i, \emptyset) = 1 - q_1 \mid \tilde{\theta} = 0\right) = q_1$.[39] On the other hand in environment $E_2$ every player's signal precision is drawn independently from the set $\left\{q_2, q_2^2/\left(q_2^2 + (1 - q_2)^2\right)\right\}$ with each element equally likely. In other words every player is equally likely to be informed by a signal of single (double) precision. The associated distribution of private beliefs can be derived straightforwardly.[40] Finally, we assume that $\alpha_1 = \alpha_2 = 1/2$.[41]

In the first period all players follow private information. Since players cannot distinguish social learning environments they learn to derive an average amount of information from the first player's decision; more precisely

$$b(\emptyset, (1)) = \frac{q_1}{2} + \frac{q_2}{4} + \frac{1}{2}\frac{q_2^2}{q_2^2 + (1 - q_2)^2}$$

and $b(\emptyset, (0)) = 1 - b(\emptyset, (1))$. Assume $q_1 > \frac{q_2^2}{q_2^2+(1-q_2)^2} > q_2$. Then in environment $E_1$ players follow private information also in the second period while in $E_2$ players with private beliefs $b(s_2, \emptyset) \in \{1 - q_2, q_2\}$ imitate the first player's decision. These choices each maximize expected payoff in each environment separately. On the other hand the decision of players in $E_2$ with a signal of double precision depends on the values of $q_1$ and $q_2$. In particular there exists $\overline{q}_2$ such that for all environmens $E_2$ satisfying $q_2 < \overline{q}_2$, players with doubly informative private signals imitate the first player's decision as well. This is not optimal: Since with strictly positive probability the first player's decision relies on a less informative signal and no player can receive a more informative one, players with a signal of double precision are better off following their private information in the second period as well.

[39] Strictly speaking, $E_1$ is a non-generic social learning environment in which ties can occur. However, slight perturbations may solve this problem. For instance assuming that the signal precision is continuously distributed on a small intervall $[q_1 - \epsilon, q_1 + \epsilon]$ around $q_1$ for some $\epsilon > 0$ will resolve ties without changing the basic properties of the environment.

[40] $supp(G^2_\theta) = \left\{\frac{(1-q_2)^2}{q_2^2+(1-q_2)^2}, 1 - q_2, q_2, \frac{q_2^2}{q_2^2+(1-q_2)^2}\right\}$, and $Pr\left(b(\tilde{s}_i, \emptyset) = \frac{q_2^c}{q_2^c+(1-q_2)^c} \mid \tilde{\theta} = 1\right) = Pr\left(b(\tilde{s}_i, \emptyset) = \frac{(1-q_2)^c}{q_2^c+(1-q_2)^c} \mid \tilde{\theta} = 0\right) = \frac{1}{2} * \frac{q_2^c}{q_2^c+(1-q_2)^c}$ for $c = 1, 2$.

[41] Notice that a similar argument as in the example of our basic illustration (section II.2) applies: Players could identify the environment via their realized private beliefs. Again assuming that each private belief arises in each environment with strictly positive albeit negligible probability solves this problem.

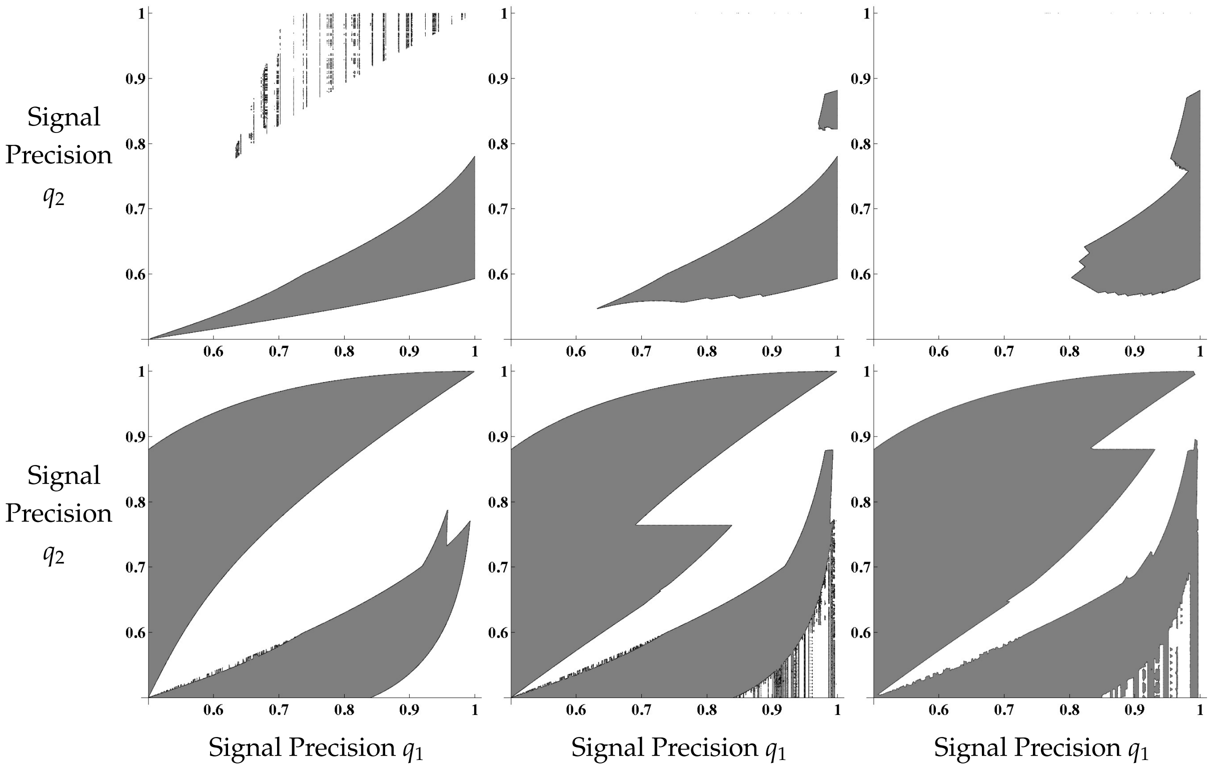


Figure II.1.: Values of $q_1$ and $q_2$ for which the non-Bayesian type has the higher expected payoff – left column: $n = 3$, middle column: $n = 6$, right column: $n = 10$; top row: $\beta = 2$, bottom row: $\beta = 1/2$.

It is then obvious that overweighting of private information is potentially beneficial. That is there robustly exist $\beta > 1$ which induce optimal choices in period 2. Since $b(s, h, \beta) > 1/2$ if and only if $[b(s, \emptyset)]^{\beta}\ b(\emptyset, h) > [1 - b(s, \emptyset)]^{\beta}\ [1 - b(\emptyset, h)]$ a player of type $\beta$ follows her doubly precise signal provided

$$\beta > \underline{\beta} = \frac{\log(b(\emptyset, (1))) - \log(1 - b(\emptyset, (1)))}{2\ \log(q_2) - 2\ \log(1 - q_2)}.$$

where $q_2 > \overline{q}_2$ induces the RHS to be strictly larger than 1. On the other hand players with low precision signal imitate provided

$$\beta < \overline{\beta} = \frac{\log(b(\emptyset, (1))) - \log(1 - b(\emptyset, (1)))}{\log(q_2) - \log(1 - q_2)}.$$

It is easily shown that $\underline{\beta} < \overline{\beta}$ and all players satisfying $\underline{\beta} < \beta < \overline{\beta}$ are strictly better off in $\mathscr{E}$ with $n = 2$ than Bayesian players.

To extend this result to $n > 2$ we rely on a computerized numerical procedure. Figure II.1 illustrates the results for $n = 3$, $n = 6$, and $n = 10$. The top row depicts the values of $q_1$ and $q_2$ such that players with $\beta = 2$, i.e. who give twice as much weight to their private information obtain a higher expected payoff in $\mathscr{E}$ than Bayesian players. The bottom row depicts analogous values for players with $\beta = 1/2$, i.e. who give twice as much weight to the public information. As one can see in either case there robustly exist parameter constellations such that the non-Bayesian types have the better strategy. In particular overweighting of private information is beneficial whenever

$q_1 \gg q_2$ while underweighting is beneficial for the majority of pairs $(q_1, q_2)$.[42]

### II.6.2. Other Imperfections

In the previous subsection we have argued that the external value of results in the long run is disputable. We have also addressed one caveat to this line of argumentation: Adaptation across similar games. In this subsection we complement these results in two ways. First, we directly assess the impacts of a finite number of repetitions. Second, we discuss how in more complex social learning environments, e.g. with heterogeneity in preferences or an endogeneous timing of decisions, the effects of learning imperfections may be exacerbated.

As discussed before several authors (see e.g. Roth, 1996) have argued that the *medium run* of a learning process is a more relevant predictor of behavior. On one hand many strategic environments do not remain unchanged in the very long term. Hence, the environment may offer only a finite time for adaptation. On the other hand active learning is costly. A player will only engage in further repetitions dedicated to learning if the learning costs are outweighed by the expected benefits measured in terms of the increase in expected payoff from future interactions due to improvement of the player's strategy. The particular extensive form structure of social learning games however induces this expected benefit to be strongly decreasing. A simple mathematical exercise shows that given $R$ rounds of play, precision of assessments at position $i$ of the social learning game is of the order $\sqrt{2^i/R}$.[43] Thus particularly for assessments at positions which are farther to the back of the social learning sequence the expected benefit of further repetitions is negligible. Finally, cognitive abilities are costly as well. Therefore individuals may be better characterized by finite memory and information processing capabilities (see for instance Jehiel, 1995; Rubinstein, 1986; Young, 1993). For instance Samuelson (2004) argues that endowing humans with better information processing capabilities is prohibitively costly-[44]

Since we have addressed learning across games in the previous subsection we fix the social learning environment $E$. We model learning within a finite number of repetitions as follows. For any history $h \in H_{-1} = H \backslash H_1$[45] we assume that each players holds a noisy assessments of the public belief. More precisely at history $h$ each player bases her decision on a distorted public belief $b'(\emptyset, h) = b(\emptyset, h) + \tilde{\epsilon}_h$. The distortions $\tilde{\epsilon}_h$ are assumed to be independent across histories $h \in H_{-1}$ and normally distributed with mean zero and standard deviations given by the vector $\boldsymbol{\eta} = (\eta_h)_{h \in H_{-1}}$ which is such that $Pr\,(0 < b'(\emptyset, h) < 1) \approx 1$.

To defend this formalization notice that public beliefs satisfy $b(\emptyset, h) = E\left[\tilde{\theta} \mid h\right]$. This relation justifies a reinterpretation of the learning process. In particular players can learn public beliefs by

[42] Notice that the graphs do not provide information about general optimality of $\beta = 2$ or $\beta = 1/2$.

[43] At any history-state-pair $(h, \theta)$ players attempt to learn the frequency of investments observing realizations of Bernoulli draws. Assuming for simplicity a stationary environment with investment probability $\pi$, the standard deviation of the r.v. $\kappa(1, h, \theta)/\kappa(h, \theta)$ is given by $\sqrt{\pi\,(1-\pi)/\kappa(h, \theta)}$. Furthermore, at position $i$ of the social learning game there exist $2^i$ different history-state-pairs, i.e. $\kappa(h_i, \theta)$ is of the order of $1/2^i$.

[44] Samuelson (2006) explicates that humans are already outliers in Nature regarding for instance the energy required to maintain the brain or the length of postnatal development.

[45] Obviously, there is no room for noisy assessments in the first period.

calculating *sample means*

$$b'(\emptyset, h) = \bar{\theta}(h, \zeta_R) = \frac{1}{\kappa(h \mid \zeta_R)} \sum_{r \leq R; h \in y(r)} \theta(r)$$

where $\kappa(h, \zeta_R)$ denotes the number of occurrences of history $h$ in super-history $\zeta_R$. By the central limit theorem (see for instance Shao, 2003) $\bar{\theta}(h, \zeta_R)$ is asymptotically normally distributed with mean $b(\emptyset, h)$ and variance $\eta_h^2 = b(\emptyset, h)\,(1 - b(\emptyset, h))/\kappa(h \mid \zeta_R)$. Accordingly, with a finite number of repetitions assessments of public beliefs remain noisy ($\eta_h^2 > 0$).

The existence of simple white noise in assessments remains our only assumption. Our first result which is of separate interest shows how this affects the formation of Bayesian posterior beliefs.

**Proposition II.4.** [46] *For any history $h \in H_{-1}$ and private belief $b(s, \emptyset)$ , if $\eta_h > 0$,*

$$E_{\tilde{\epsilon}_h}\left[b'(s, h)\right] = \begin{cases} > b(s, h) & \text{if } b(s, \emptyset) < \frac{1}{2} \\ < b(s, h) & \text{if } b(s, \emptyset) > \frac{1}{2} \end{cases}$$

In words white noise induces players (in expected terms) to lean more towards the public information, i.e. to underweight private information. This result rests upon the convexity (resp. concavity) of the Bayesian updating rule in the domain $b(s, \emptyset) <$ (resp. $>$) $\frac{1}{2}$. It is such that contradictory information is weighed more heavily than confirmatory information. Since with white noise assessed public beliefs are equally likely to differ slightly from correct ones in either direction, resulting posteriors are biased in the direction of the public information. The formalization of this argument is straightforward but tedious and is thus relegated to the appendix.

In a world with rich action space were players are rewarded for stating posteriors, the result of the Lemma above implies that noise in assessments favors overweighting of private information.[47]. With coarse action set $A = \{0, 1\}$, this implication is not true. Our main result of this subsection more generally proves that Bayes-rationally responding to noisy assessments is not the optimal strategy.

In order to establish this result let $E_\eta = \langle n, \Theta, A, u, (G_0, G_1), \boldsymbol{\eta} \rangle$ denote the *noisy social learning game*. $E_\eta$ differs from the standard game by the restriction that players $i = 2, \ldots, n$ instead of histories $h_i \in H_i$ observe realizations of $b'(\emptyset, h_i) = b(\emptyset, h_i) + \tilde{\epsilon}_{h_i}$. Accordingly in the noisy social learning game behavioral strategies of players are given by mappings $\sigma_\eta : [\underline{b}, \overline{b}] \times \{\emptyset\} \cup [0, 1] \to [0, 1]$ where $\sigma_\eta(s, \emptyset) = Pr(a = 1 \mid b(s, \emptyset), h = h_1 = \emptyset)$ governs behavior in period 1 and behavior in later periods is determined by $\sigma_\eta(s, x) = Pr\,(a = 1 \mid b(s, \emptyset), b'(\emptyset, h) = x)$. Said differently we presume that a player's public information is completely captured by her (distorted) public belief. In particular if a player's perceived public belief is the same at different histories in (possibly) different periods she has to employ the same strategy mapping private beliefs into actions. Accordingly in the noisy social learning game taking into account that players are randomly assigned to positions in the sequence expected payoff to strategy $\sigma_\eta$ is given by

[46] Appendix II.B derives a similar result for a setting with normally distributed state and private signals which is more closely related to the literature on learning in rational expectations.

[47] See Appendix II.B.

$$U_\eta\left(\sigma_\eta \mid \varphi^*\right) = \frac{1}{4n} \int_{\underline{b}}^{\overline{b}} \sigma_\eta(s,\emptyset)\ [dG_1(s) - dG_0(s)]$$

$$+ \frac{1}{4n} \int_{\underline{b}}^{\overline{b}} \sum_{h \in H_{-1}} \int_{-b(\emptyset,h)}^{1-b(\emptyset,h)} \sigma_\eta(s, b(\emptyset,h)+\epsilon)\ [\varphi(h \mid 1)\, dG_1(s) - \varphi(h \mid 0)\, dG_0(s)]\ \phi_h(\epsilon)\, d\epsilon.$$

Here $\varphi^*$ denotes assessments associated with the Bayes-rational strategy $\sigma^*(s,x) = 1$ of $b(s,\emptyset) > 1-x$ (resp. $\sigma^*(s,\emptyset) = 1$ if $b(s,\emptyset) > 1/2$) and $\sigma^*(s,x) = 0$ if $b(s,\emptyset) < 1-x$ (resp. $\sigma^*(s,\emptyset) = 0$ if $b(s,\emptyset) < 1/2$) and $\phi_h(\epsilon)$ is the density of the normal distribution with mean 0 and standard deviation $\eta_h$.

As in the previous subsection we show that generically the Bayes-rational strategy does not maximize expected payoff.

**Lemma II.2.** *Fix assessment $\varphi^*$ and let*

$$t(x \mid \boldsymbol{\eta}, \varphi^*) = \frac{\sum_{h \in H_{-1}} \varphi^*(h \mid 0)\, \phi_h\left(x - b(\emptyset, h \mid \varphi^*)\right)}{\sum_{h \in H_{-1}} \left[\varphi^*(h \mid 0) + \varphi^*(h \mid 1)\right] \phi_h\left(x - b(\emptyset, h \mid \varphi^*)\right)},$$

*and in an abuse of notation $t(\emptyset \mid \boldsymbol{\eta}, \varphi^*) = 1/2$. The strategy $\sigma^*_{\boldsymbol{\eta}}$ given by*

$$\sigma^*_{\boldsymbol{\eta}}(s,x) = \begin{cases} 1 & \text{if } b(s,\emptyset) > t(x \mid \boldsymbol{\eta}, \varphi^*) \\ 0 & \text{if } b(s,\emptyset) < t(x \mid \boldsymbol{\eta}, \varphi^*) \end{cases}$$

*maximizes expected payoff $U_{\boldsymbol{\eta}}$ in the noisy social learning game.*

*Proof.* See appendix. □

**Proposition II.5.** *Generically for sufficiently large $\boldsymbol{\eta}$, the Bayes-rational strategy $\sigma^*$ does not coincide with the optimal strategy $\sigma^*_{\boldsymbol{\eta}}$ at all histories reached a non-vanishing fraction of the time.*[48]

*Proof.* See appendix. □

As argued before we do not consider the strategy $\sigma^*_{\boldsymbol{\eta}}$ to be consciously derived by players. It is the benchmark of an evolutionary process (again with expected payoffs representing fitness) taking place directly on the strategy space $\Sigma_{\boldsymbol{\eta}}$ or indirectly on some subspace of less complicated rules. The difference between $\sigma^*_{\boldsymbol{\eta}}$ and the Bayes-rational strategy then provides a foundation for disagreement between the latter and behavior of real decision-makers. As before we illustrate in an example how simple rules can do better than the Bayes-rational strategy.

### Simple Non-Bayesian Strategies (ctd.)

As in the previous subsection we consider the heterogeneous belief updating model of Chapter III driven by the generalized updating rule (II.2). Again we assume a Bayesian population and a single player of type $\beta \neq 1$. We study the specific social learning model of Bikhchandani, Hirshleifer, and Welch (1992) (environment $E_1$ in the illustration of the previous subsection) with fixed signal

[48] That is $\sigma^*(s,x) \neq \sigma^*_{\boldsymbol{\eta}}(s,x)$ for each $x \in X \subseteq [0,1]$ and each $b(s,\emptyset) \in B \subseteq [\underline{b}, \overline{b}]$ where $X$ and $B$ satisfy $\int_B dG_\theta(s) > 0$ and $\int_X \sum_{h \in H} \varphi^*(h \mid \theta)\, \phi_h\left(x - b(\emptyset, h \mid \varphi^*)\right) dx > 0$ for some $\theta \in \Theta$.

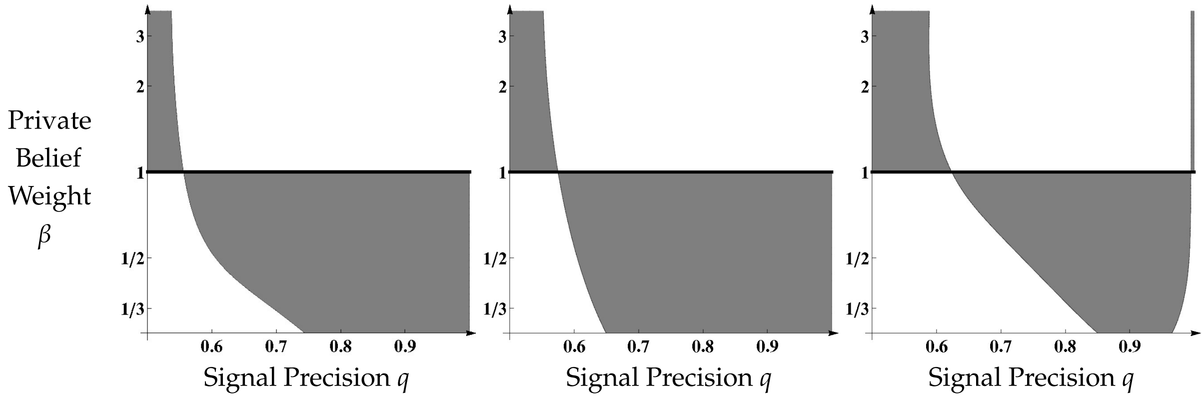


Figure II.2.: Expected Payoff Difference of Non-Bayesian updaters versus Bayesian players in dependence upon signal precision $q$ – left column: $n = 3$, middle column: $n = 6$, right column: $n = 10$. Orange line: $\beta = 1/2$, blue line: $\beta = 2/3$, red line: $\beta = 3/2$, green line: $\beta = 2$.

precision $1/2 < q < 1$.

Noise does not affect choices in the first period, accordingly players of each type follow private information. In the second period correct public beliefs are given by $b(\emptyset,(1)) = 1 - b(\emptyset,(0)) = q$. Recall

$$U_\eta(\sigma) = \frac{1}{4n} \sum_{h\in H} \int_{-\infty}^{\infty} \sum_{s\in\{1-q,q\}} \sigma_\eta\,(s, b(\emptyset,h)+\epsilon)\; U_\eta(s,h)\,\phi_h(\epsilon)\,d\epsilon$$

where $U_\eta(s,h) = \varphi^*(h \mid 1) * Pr(b(\tilde{s},\emptyset) = s \mid \tilde{\theta} = 1) - \varphi^*(h \mid 0) * Pr(b(\tilde{s},\emptyset) = s \mid \tilde{\theta} = 0)$ denotes the expected payoff given private belief $s$ and history $h$ and a population of Bayesian decision-makers. It can be straightforwardly derived that $U_\eta(q,(1)) = -U_\eta(1-q,(0)) = 2q-1$ while $U_\eta(q,(0)) = U_\eta(1-q,(1)) = 0$. On the other hand given private belief $s$, history $h$, and realized public belief shock $\epsilon$ a player of type $\beta \in (0,\infty)$ invests provided $\epsilon > \frac{(1-s)^\beta}{s^\beta+(1-s)^\beta} - b(\emptyset,h)$ and rejects otherwise (abstracting from ties). Hence, her probability of investment is given by $1 - \Phi_h\left(\frac{(1-s)^\beta}{s^\beta+(1-s)^\beta} - b(\emptyset,h)\right)$. Notice that $\Phi_h$ is symmetric around zero. Accordingly the player's expected payoff can be written as

$$U_\eta(\beta) = \frac{2q-1}{4n}\left[2 - \Phi_h\left(\frac{(1-q)^{\beta+1} - q^{\beta+1}}{q^\beta + (1-q)^\beta}\right)\right]$$

and is strictly increasing in $\beta$. Hence, the more weight a player puts on private information the higher her expected payoff in the noisy social learning game with $n = 2$.

For $n > 2$ we resort to numerical approximations. Recall that for each period $i = 2,\ldots,n$ the assessed private belief at history $h_i$ is normally distributed around its correct counterpart with standard deviation $\eta_{h_i}$. We make the following simplifying assumption: for each $i = 2,\ldots,n$ and each $h_i \in H_i$, $\eta_{h_i}$ is given by

$$\eta_{h_i} = \eta_i = \eta_2 * 2^{(i-2)/2}.$$

This is motivated by the considerations in the first paragraph of this subsection where we have derived the standard deviation in period $i$ to be of the order $\sqrt{2^i/R}$. We further assume that $\eta_2 = 0.03$ which corresponds to a value of $R \approx 100$. Figure II.2 represents pairs of signal precision $q$

and private information weight $\beta$ such that $U_\eta(\beta) > U_\eta(1)$ i.e. such that the non-Bayesian type has a more profitable strategy than the Bayesian players. As one can see if private signals convey few information, the noisy environment favors overweighting of private information. On the other hand underweighting is favorable for more informative private signals. However the larger $n$ the larger the range of signal precisions for which overweighters are better off. We consequently conjecture that in very long social learning games with noisy assessments a successful strategy must involve overweighting of private information. This provides a novel explanation for the experimental results.

#### Extensions

We have so far restricted ourselves to a simple social learning game. In the remainder of this subsection we argue that the considered learning process is even less likely to lead to rational play in generalized social learning games.

Consider first a world with heterogeneity in preferences as introduced in chapter I. That is we distinguish *rational* types whose payoff decisively depends upon the realized state of Nature from *noise* types who always choose the same action and thereby introduce noise into the social learning process. The standard assumption in such a model (see e.g. Smith and Sørensen, 2000) is common knowledge of the distribution of preference types. As before from an individual learning perspective this assumption is disputable. In particular it would require players to receive feedback about payoffs of other players. Without this information identification of single environments is not possible and adaptation takes place across settings. Therefore presence of multiple social learning settings which differ in their distribution of preferences may trivially explain deviations from act-rational behavior in single environments as well. However, it can be shown that absent differences in private belief distributions the Bayes-rational strategy remains optimal if no player can distinguish settings.

Yet, differences in preferences may reinforce suboptimality of the Bayes-rational strategy in the presence of multiple private belief distributions. To see this consider a world composed of two simple 2-by-2-by-2 environments occurring with equal probability: $E_1$ is characterized by signal precision $q_1$ and all players having the same standard preferences we have maintained so far (i.e. $u(1, \theta) = \theta - 1/2$ and $u(0, \theta) = 0$). On the other hand in environment $E_2$ with signal precision $q_2 > q_1$ only a fraction $(1 - 2\,\xi_2)$ of players have this preferences while a fraction of each $\xi_2 > 0$ are noise types which always invest or always reject respectively. In particular assume that $\xi_2$ is sufficiently large such that players should optimally follow private information even after observing the first two players making the same choice. If players learn across these environments absence of noise in environment $E_1$ leads Bayes-rational players in $E_2$ to infer too much from the first players' decisions and may cause them to suboptimally imitate in period 3. On the contrary it can be shown that there robustly exist settings where overweighters of private information outperform Bayes-rational players by avoiding this mistake (see appendix II.D).

While noise players are an extreme form of preference heterogeneity similar arguments apply to the case of several rational types as has been most impressively demonstrated by the confounding outcomes of Smith and Sørensen (2000). In general heterogeneity in preferences complicates learning from others' actions by significantly reducing the amount of information single decisions convey. This brings into focus the problem of imperfect learning within a finite number of repetitions much more than in the simple settings studied so far. If heterogeneity confounds the opportunities to learn by observing others and imperfect learning opportunities add another source of noise, the process of social learning may break down alltogether. Moreover as discussed above the inability of players to distinguish social learning environments may spread such complications across settings. Even if a single environment gives best chances to learning from others, these may be thwarted by player' adaptation across settings and their limited time to do so.

In conclusion imperfect learning opportunities may have much more drastic influence in settings where players must learn from subtle differences since these are most vulnerable to the introduction of any kind of noise. Accordingly such settings may be particularly favorable to non-Bayesian strategies.

A different complication is introduced if players need to choose the timing of their decisions. In this case the very simple structure of the game which for instance drives our convergence result (Proposition II.1) is lost. More precisely endogeneous timing of decisions requires players to not only derive information from past choices but also foresee future social learning opportunities. By choosing the timing of decisions players may then significantly influence their possibilities for individual learning. Moreover the importance of further restrictions of learning opportunities such as heterogeneous support of private beliefs, partial feedback on choices or conditional feedback on the state of Nature increases.

Consider for instance a setting where the costs of investing are large such that only sufficiently strong favorable information may induce a player to invest (e.g. $u(1,\theta) = \theta - 2/3$). Assume further conditional feedback on the state of Nature, i.e. a player receives feedback about the profitability of the investment (the state of Nature) if and only if she invests herself. Suppose finally that some players are restricted to a limited quality of information such that neither of them will ever invest conditional on private information alone. In the simple adaptive process this group of players will not be able to learn how to infer information from others' choices. Consequently neither of them will ever invest. Overweighting of private information may then be profitable not only because sometimes investing must yield a positive payoff but also because it would allow players to learn how to take others' actions into account in future interactions.

While this constitutes a rather extreme example, it fascilitates how an endogeneous timing of decisions complicates the task of learning from others and players' opportunities to adapt to this task. At least it shows how individual experimentation and self-confirming equilibria are much more important than in the present context.

## II.7. Discussion

This chapter constitutes a first attempt to discuss the learning foundations of rational play in social learning games. Our results suggest that too much structural knowledge has been assumed in standard economic models of social learning. Though in the absence of **fundamental** structural uncertainty and with an infinite learning horizon epistemic learning leads to Bayesian rational play, the same learning process favors non-Bayesian play whenever players do not know the distribution of private beliefs. As a consequence, further economic models of social learning should allow for the presence of **fundamental** structural uncertainty and learning models should inform those economic models about the nature and scope of the structural uncertainty.

Assuming that **fundamental** structural uncertainty is unavoidable, our theoretical study also suggests that *rule rationality* rather than *act rationality* is the appropriate benchmark for discussing rational social learning. If players lack structural knowledge then they are likely to develop rational *rules* of social learning rather than trying to act in a Bayesian rational way in each single social learning environment. This, in turn, seriously questions the informativeness of the existing experimental evidence on rational social learning. Indeed, the validity of the rational view of social learning should be tested in contexts that are familiar to economic actors, i.e. contexts with **fundamental** structural uncertainty. As far as we know, all existing experimental studies on social learning have considered laboratory settings which correlate strongly with the standard economic models of social learning. Though we understand that such simple settings with full control of information flows were the natural candidates for testing the existing economic models, we fear that the field environments in which social learning mainly takes place differ substantially from those laboratory settings. As a consequence, existing laboratory settings might be perceived as artificial by subjects and the fact that laboratory behavior systematically deviates from rational play in those settings might not come as a surprise and does not constitute conclusive evidence against rational social learning.

Our results therefore suggest a re-evaluation of both the experimental and the theoretical research in social learning. First, new laboratory experiments should be designed to test the rational view of social learning in more familiar contexts. Aspects of the field environment which are likely to strongly influence social learning (e.g. structural uncertainty) have to be incorporated in those laboratory settings. Second, theoretical models with players lacking structural knowledge are likely to provide a better understanding of real-world social learning. In chapter III we show that this does not necessarily increase the complexity of these models.

## Appendix II.A More Sophisticated Learning

In this addendum we extend the learning process to allow for players who are more sophisticated. We follow a general approach suggested for instance by Fudenberg and Kreps (1995, FK henceforth). Since most of the concepts apply primarily to the long run we assume an infinite number of repetitions throughout.

In the basic learning process we have assumed that each player holds a single assessment $\hat{\varphi}(\tilde{\zeta}_r)$ which she updates over time given evidence $\zeta_r$. More generally one may assume that a player considers possible several models of others. Hence, let $\Phi = \{\varphi : H \times \Theta \to [0,1]\}$ denote the set of possible assessments. Define a *general assessment rule*[49] as $\hat{\gamma} : \bigcup_{r=1}^{\infty} \mathscr{Z}_r \to \Delta(\Phi)$ attaching to each super-history $\zeta_r$ a probability distribution on $\Phi$. We do not explicitly model the updating of general assessments over time. In staying at an abstract level our concept entails various more specific models, most importantly it includes Bayesian inference as shown by FK (Section 3.2). We do require that once enough evidence has accumulated players respect this evidence as follows. For each $r = 1, 2, \ldots$, each $\zeta_r \in \mathscr{Z}_r$ and each $h$ s.t. $\kappa(h, \theta \mid \zeta_r) = \sum_{x=0}^{n-1} \kappa(a, h, \theta \mid \zeta_r) > 0$ define

$$\bar{\varphi}(a \mid h, \theta; \zeta_r) = \frac{\kappa(a, h, \theta \mid \zeta_r)}{\kappa(a, \theta \mid \zeta_r)}$$

and $\bar{\varphi}(h \mid \theta; \zeta_r)$ accordingly.

**Definition II.4.** *The general assessment rule $\hat{\gamma}$ is **asymptotically empirical** if for every $\epsilon > 0$, every $\zeta$ and for every $\theta \in \Theta$ and $h \in H$ such that $\liminf_{r\to\infty} \frac{\kappa(h,\theta|\zeta_r)}{r} > 0$,*

$$\lim_{r\to\infty} \hat{\gamma}(\zeta_r)\left(\left\{\varphi : \left\|\varphi(h \mid \theta) - \bar{\varphi}(h \mid \theta; \zeta_r)\right\| < \epsilon\right\}\right) = 1$$

In the basic learning process we have assumed that players myopically respond to their assessments in each round. Apart from repeated play considerations this is also restrictive because it rules out the possibility that players experiment. For instance one can argue (see FK for an exhaustive discussion of this issue) that in extensive-form games a player's action affects what she learns about others' behavior. We take this into account following the idea of FK. For general assessment $\gamma$ and some $i \in \{1, \ldots, n\}$, let $U_i(\sigma \mid \gamma) = \sum_{h_i \in H_i} \int_{\underline{b}}^{\bar{b}} \sigma(s, h_i)\, U(s, h_i \mid \gamma)$ where $U(s, h \mid \gamma) = \int_{\varphi\in\Phi} [\varphi(h \mid 1)\, dG_1(s) - \varphi(h \mid 0)\, dG_0(s)]\, \gamma(d\varphi)$. The strategic response $\hat{\sigma}$ is **asymptotically myopic** with regard to $\hat{\gamma}$ if there exists a sequence of non-negative numbers $\{\epsilon_r\}_{r=1}^{\infty}$ such that $\lim_{r\to\infty} \epsilon_r = 0$ and, for each $r$ and $\zeta_r$, and each $i = 1, \ldots, n$

$$U_i\left(\hat{\sigma}(\zeta_r) \mid \hat{\gamma}(\zeta_r)\right) + \epsilon_r \geq \max_{\sigma\in\Sigma} U_i\left(\sigma, \hat{\gamma}(\zeta_r)\right).$$

Asymptotic myopia permits players to choose suboptimal strategies where the suboptimality vanishes over time. In particular players may choose slightly suboptimal strategies with larger probabilities or grossly suboptimal strategies with small probabilities (see (Fudenberg and Kreps, 1993)). Thus while at early dates depending on the sequence $\{\epsilon_r\}_{r=1}^{\infty}$ the player may consciously

[49] This is not to be confused with the "rules" we discuss in connection with rule rationality.

experiment with suboptimal strategies, she eventually has to confine herself to *random experimentation* with decreasing overall probability. In this regard asymptotic myopia it still very restrictive. A more general notion is the following

**Definition II.5.** *The strategic response $\hat{\sigma}$ is* ***asymptotically myopic with calendar-time limitations on experimentation*** *with respect to $\hat{\gamma}$ if there exist*

*(i) a sequence $\{\epsilon_r\}_{r=1}^{\infty}$ s.t. $\epsilon_r > 0$ for each $r = 1, 2, \ldots$ and $\lim_{r \to \infty} \epsilon_r = 0$,*

*(ii) a sequence $\{\delta_r\}_{r=1}^{\infty}$ s.t. $\delta_r \geq 0$ and $\delta_{r+1} \geq \delta_r$ for each $r = 1, 2, \ldots$ and $\lim_{r \to \infty} \delta_r / r = 0$,*

*(iii) an asymptotically myopic strategic response $\hat{\sigma}^{opt}$,*

*(iv) a strategic response $\hat{\sigma}^{exp}$ s.t. for every $\zeta$, $r$ and every $s$ and $h$, $\hat{\sigma}^{exp}(a \mid s, h; \zeta_r) > 0$ if and only if $\kappa(a, h \mid \zeta_r) < \delta_r$ where $\kappa(a, h \mid \zeta_r)$ is the number of times the player associated with $\hat{\sigma}$ chose action $a$ at history $h$ along $\zeta_r$,*

*(v) a mapping $\hat{\alpha} : \bigcup_{r=1}^{\infty} \mathscr{Z}_r \times H \to [0, 1]$ s.t. $\hat{\alpha}(\zeta_r, h) < 1$ only if $\kappa(a, h \mid \zeta_r) < \delta_r$ for some $a \in A$*

*and if for each $r = 1, 2, \ldots$, $\zeta_r \in \mathscr{Z}_r$, $s$ and $h \in H$*

$$\hat{\sigma}(s, h \mid \zeta_r) = \hat{\alpha}(\zeta_r, h) * \hat{\sigma}^{opt}(s, h \mid \zeta_r) \; + \; [1 - \hat{\alpha}(\zeta_r, h)] * \hat{\sigma}^{exp}(s, h \mid \zeta_r).$$

An *individual learning model* for the social learning stage game is an array of assessment rules and strategic responses, one each for each player. It is *conforming*, if each player's assessment rule is asymptotically empirical and each player's strategic response is asymptotically myopic with calendar-time limitations on experimentation with respect to the assessment rule.

**Definition II.6.** *A strategy profile $\sigma^{**}$ is* ***locally stable****, if there exists some conforming learning model, such that $\mathbf{P}\left(\lim_{r \to \infty} \hat{\sigma}^{opt}(\zeta_r) = \sigma^{**}\right) > 0$ where $\hat{\sigma}^{opt}$ denotes the non-experimental part of the array of strategic responses.*

**Proposition II.6.** *In the social learning game independent of feedback and private belief support a strategy profile is locally stable if and only if it is a Bayesian equilibrium*

*Proof.* First a similar argumentation as in the proof for the more basic learning process applies, i.e. every player infinitely often observes the complete sequence of choices together with the state of Nature whatever the specification of feedback and private belief support. We thus concentrate on the case with complete feedback on choices, unconditional feedback on the state and homogeneous private belief support.

The proof follows the same ratio as in the more basic case. We fix some Bayesian equilibrium $\sigma^*$ with associated public beliefs $b^*(\emptyset, h)$ for $h \in H$ and define $\varphi^*$ as in the proof of Proposition II.1. Also we restrict ourselves to analyzing the learning process of a single representative player characterized by an asymptotically empirical assessment rule $\hat{\gamma}$ and an associated strategic response $\hat{\sigma}$ which is asymptotically myopic with calender-time limitations on experimentation (asymptotic myopia wCTLE).

For general assessment $\gamma$ best reponse is given by

$$\sigma(s,h \mid \gamma) = \begin{cases} 1 & \text{if } \int\limits_{\varphi\in\Phi} \varphi(h \mid 1)\,\gamma(d\varphi)\,dG_1(s) > \int\limits_{\varphi\in\Phi} \varphi(h \mid 0)\,\gamma(d\varphi)\,dG_0(s) \\ 0 & \text{if } \int\limits_{\varphi\in\Phi} \varphi(h \mid 1)\,\gamma(d\varphi)\,dG_1(s) < \int\limits_{\varphi\in\Phi} \varphi(h \mid 0)\,\gamma(d\varphi)\,dG_0(s) \end{cases}.$$

A similar induction argument applies: In the first period no inferences about others are necessary and asymptotic myopia wCTLE guarantees that in the long run the player's strategy obeys equilibrium. In later periods asymptotic empiricism guarantees that generalized assessments at histories reached a non-vanishing fraction of the time eventually put probability 1 on equilibrium assessments. Hence, assessments eventually match equilibrium assessments at such histories. Asymptotic myopia wCTLE in turn implies that the player's strategy eventually comes to match her equilibrium strategy (with a similar argument regarding ties). Therefore we turn to histories reached a vanishing fraction of the time.

Fix an out-of-equilibrium history $h_i$. $b^*(\emptyset, h_i)$ is not specified by $\varphi^*$. We construct explicitly an asymptotically empirical assessment rule and an associated *myopic* behavior rule following the proof of Fudenberg and Kreps (1993, Proposition 6.3). Assume that the player's initial assessment is given by $\gamma^* = \delta_{\varphi^*}$ the Dirac measure on $\varphi^*$. Clearly, myopically best response to $\gamma^*$ is given by $\sigma^*$. Now assume that the assessment rule is such that players only change their assessments, if sufficient evidence against the current assessment accumulates. Formally, given $\sigma^*$ we can create a probability space over complete sequences of actions $(a_1, \ldots, a_n) \in A^n$. We may then construct a non-decreasing sequence $n_r$ with $\lim\limits_{r\to\infty} n_r = \infty$ such that the event

$$\left\{\left\|\bar{\varphi}(\zeta_r) - \varphi^*\right\| < 1/n_r \text{ for } r = 1, 2, \ldots\right\}$$

has probability at least $\frac{1}{2}$ for an appropriate norm $\left\|\bar{\varphi}(\zeta_r) - \varphi^*\right\|$[50] (see Fudenberg and Kreps, 1993, for details). Let the assessment rule be given by

$$\hat{\gamma}(\zeta_r) = \begin{cases} \gamma^* & \text{if } \left\|\bar{\varphi}(\zeta_r) - \varphi^*\right\| < 1/n_t \\ \bar{\gamma}(\zeta_r) & \text{else} \end{cases}$$

where $\bar{\gamma}(\zeta_r)\,(\bar{\varphi}(\zeta_r)) = 1$. Then the players begin with $\gamma^*$ and best respond with $\sigma^*$ and by the construction above with probability at least $\frac{1}{2}$ stick with both.

Notice that as long as players behave according to $\sigma^*$ assessments for off-path-behavior are never changed, because these situations are reached with probability zero and hence no new evidence is accumulated. However, if assessments on the path change, this may imply different myopic best responses which may in turn yield strictly positive probabilities for some out-of-equilibrium choices. Therefore the construction above is necessary.

In conclusion, the first case shows that any locally stable strategy profile must satisfy Bayes' rule and sequential rationality on the equilibrium path. On the other hand the second case proves that any off-path-behavior can be locally stable. The proof is complete.

□

[50]For instance $\left\|\bar{\varphi}(\zeta_r) - \varphi^*\right\| = \sup_{a,h,\theta} \left|\bar{\varphi}(a \mid h, \theta; \zeta_r) - \varphi^*(a \mid h, \theta)\right|$.

The Proposition shows that even in the more sophisticated learning process the set of strategic outcomes of the learning process is large (infinite) though of course all are outcome equivalent to the unique Bayesian equilibrium outcome. Yet, many more refined equilibrium concepts (sequential equilibrium, trembling-hand perfect equilibrium) do not significantly reduce the set of equilibria. However under mild additional conditions a unique equilibrium is selected as the limit of regular quantal response equilibria as players become more and more inclined to choose the best response (the payoff disturbances approach zero). This equilibrium is given by extending Bayesian updating to probability-zero observations by assuming that such events indicate those private belief under which the loss in expected payoff is smallest. We feel that under mild additional conditions on the individual learning process, this unique equilibrium could as well be the unique stable outcome of individual learning. The reason is as follows. Within the (individual) learning process, out-of-equilibrium choices may occur for two reasons: (i) mistaken assessments yielding wrong beliefs, and (ii) experimentation. In the first case mistaken assessments imply wrong inferences from previous choices, not from private beliefs. Hence, if assessments are such that the out-of-equilibrium choice is optimal for some private belief, it must be optimal for its most favorable private belief $\hat{s}$ as well. If assessments (and therefore beliefs) change in a smooth way, the deviation should remain optimal for $\hat{s}$ longer than for any other private belief. Consequently, over time it should be observed most frequently if $\hat{s}$ occurred. Secondly, random experimentation is such that least costly mistakes can be made most often. If the possibilities for random experimentation vanish sufficiently slow, deviations should again be observed most frequently, if the player's private belief is $\hat{s}$. Finally, if players take into account such information (notice that this is not required by asymptotic empiricism, as only state and history contingent action distributions matter and these should in the limit put probability zero on any deviation) the unique equilibrium seems most reasonable.

## Appendix II.B The Gaussian-Quadratic Model

In social learning settings more closely related to the classical model of learning in rational expectations the following setting has become very popular (see for instance Vives, 1993). The state of Nature $\tilde{\theta}$ is normally distributed with mean $\mu(\theta)$ and variance $\sigma^2(\theta)$. Furthermore for realized state $\theta$ private signals are given by $\tilde{s}_i = \theta + \tilde{\epsilon}$ where $\tilde{\epsilon}$ is a noise independent of $\theta$ and normally distributed around mean 0 with variance $\sigma_\epsilon^2 > 0$. This information structure usually is combined with action space $A = [0,1]$ and the quadratic payoff (loss) function $u(a,\theta) = -(a-\theta)^2$.

The beauty of the Gaussian model stems from the fact that conditional on a signal realization $s$ by Bayes'rule $\tilde{\theta}$ is still normally distributed with updated mean and variance given by

$$\begin{aligned}\sigma^2(\theta \mid s) &= \frac{\sigma^2(\theta)\,\sigma_\epsilon^2}{\sigma^2(\theta) + \sigma_\epsilon^2},\\ \mu(\theta \mid s) &= \frac{\sigma^2(\theta)}{\sigma^2(\theta) + \sigma_\epsilon^2}\, s + \frac{\sigma_\epsilon^2}{\sigma^2(\theta) + \sigma_\epsilon^2}\,\mu(\theta).\end{aligned}$$

This follows from the Bayesian formula $f(\theta \mid s) = f(\theta)\, g_\theta(s) / \int_{\theta'} f(\theta')\, g_{\theta'}(s)\, d\theta'$ by straightforward calculations. Furthermore expected payoff is maximized by setting $a(s) = \mu(\theta \mid s)$. Hence, if $\mu(\theta)$,

$\sigma^2(\theta)$ and $\sigma_\epsilon^2$ are known $a$ perfectly reveals $s$ and $f(\theta \mid a) = f(\theta \mid s(a))$ where $s(a)$ is uniquely defined.

Assume as above that players attempt to learn $f(\theta \mid a)$ for each $a \in A$ in a statistical way.[51] That is players assume that conditional on $a$, $\tilde{\theta}$ is still normally distributed[52] and attempt at determining $\mu(\theta \mid a)$ and $\sigma^2(\theta \mid a)$ via the sample mean $\bar{\theta}(a) = \frac{1}{\kappa(a)} \sum_{r:a \in y(r)} \theta(r)$ and the sample variance $S^2(\theta \mid a) = \frac{1}{\kappa(a)-1} \sum_{r:a \in y(r)} [\theta(r) - \bar{\theta}(a)]^2$ respectively. Then $\bar{\theta}(a)$ is normally distributed around mean $\mu(\theta \mid a)$ with variance $\sigma^2(\theta \mid a)/\kappa(a)$ and $S^2(\theta \mid a)$ is asymptotically normally distributed around $\sigma^2(\theta \mid a)$ with variance $2\,\sigma^4(\theta \mid a)/\sqrt{n}$. Furthermore by Cochran's Theorem (Cochran, 1934) $\bar{\theta}(a)$ and $S^2(\theta \mid a)$ are independent. Using these measures for determining mean and variance of the updated normal distribution of $\tilde{\theta}$ conditional on private signal $s$ and observation $a$ players arrive at

$$\begin{aligned} \hat{\mu}(\theta \mid s, a) &= \frac{S^2(\theta \mid a)}{S^2(\theta \mid a) + \sigma_\epsilon^2}\, s + \left[1 - \frac{S^2(\theta \mid a)}{S^2(\theta \mid a) + \sigma_\epsilon^2}\right] \bar{\theta}(a) \\ \text{and } \hat{\sigma}^2(\theta \mid s, a) &= \frac{\sigma_\epsilon^2\, S^2(\theta \mid a)}{\sigma_\epsilon^2 + S^2(\theta \mid a)}. \end{aligned}$$

The mapping $\alpha(S) = \frac{S}{\sigma_\epsilon^2 + S}$ is strictly concave in $S$. Therefore by Jensen's inequality $E\left[\alpha\left(S^2(\theta \mid a)\right)\right] \leq \alpha\left(E\left[S^2(\theta \mid a)\right]\right) = \alpha\left(\sigma^2(\theta \mid a)\right)$ provided $\kappa(a) < \infty$. Independence of $\bar{\theta}(a)$ and $S^2(\theta \mid a)$ then implies that $\hat{\mu}(\theta \mid s, a)$ is adjusted too little towards the signal $s$.

We adapt the specific model of heterogeneous belief updating studied in chapter III to this setting. That is

$$f_\beta(\theta \mid s) = \frac{[g_\theta(s)]^\beta\ f(s)}{\int_{\theta'} [g_{\theta'}(s)]^\beta\ f(\theta')\, d\theta'}$$

for some $\beta > 0$. Then the posterior distribution of $\tilde{\theta}$ given signal $s$ is given by a normal distribution with biased mean and variance

$$\begin{aligned} \hat{\mu}_\beta(\theta \mid s, a) &= \frac{\beta\, S^2(\theta \mid a)}{\sigma_\epsilon^2 + \beta\, S^2(\theta \mid a)}\, s + \left[1 - \frac{\beta\, S^2(\theta \mid a)}{\sigma_\epsilon^2 + \beta\, S^2(\theta \mid a)}\right] \bar{\theta}(a), \\ \sigma_\beta^2(\theta \mid s, a) &= \frac{\sigma_\epsilon^2\, S^2(\theta \mid a)}{\sigma_\epsilon^2 + \beta\, S^2(\theta \mid a)}. \end{aligned}$$

Hence, if $\kappa(a) < \infty$ there exists $\beta > 1$ such that $E\left[\hat{\mu}_\beta(\theta \mid s, a)\right] = \mu(\theta \mid s, a)$.

Finally, we show that with the quadratic payoff function this bias in posteriors implies non-optimal choices of Bayesians and therefore fitness benefits of overweighters. Notice first that expected payoff $U(a \mid s, x) = E_{\tilde{\theta}}\left[u(a, \theta) \mid s, x\right]$ where $x$ is the observed action is strictly concave in $a$. Thus, it holds $E_{\tilde{\epsilon}}\left[U(a \mid s, x)\right] < U\left(E_{\tilde{\epsilon}}[a(s, x)] \mid s, x\right)$. Furthermore $U(a \mid s, x)$ is maximized at $a^* = E[\tilde{\theta} \mid s, x] = \mu(\theta \mid s, x)$. The result above straightforwardly implies that $E_{\tilde{\epsilon}}[a_1(s, x)] = E_{\tilde{\epsilon}}\left[\hat{\mu}_1(\theta \mid s, x)\right] \neq$

[51] Clearly, since $A$ is an interval this requires players to learn an uncountably infinite number of values. However, little is lost by assuming that players partition $A$ into a finite set of intervals $A_j$ and learn $f(\theta \mid A_j)$ for each $j$. If the partition is sufficiently fine since $\mu(\theta \mid s(a))$ is continuous in $a$, $f(\theta \mid A_j)$ approximates $f(\theta \mid a)$ for each $a \in A_j$.

[52] Equivalently, they could apply an appropriate statistical test which will be confirmed eventually.

$\mu(\theta \mid s,x)$ while for some $\beta > 1$ $E_{\tilde{\epsilon}}[a_\beta(s,x)] = \mu(\theta \mid s,x)$. Therefore $U\left(E_{\tilde{\epsilon}}[a_1(s,x)] \mid s,x\right) < U\left(E_{\tilde{\epsilon}}[a_\beta(s,x)]\right)$.

# Appendix II.C Illustration: The Uniform Signal Quality Model

To illustrate further our results regarding imperfect learning opportunities, we consider a specific social learning model introduced in a slightly simpler version by Smith and Sørensen (2000, Section 3.1.A and 2008b, Section 5.2). It is given by private beliefs distributed continuously on the interval $[1-a, a]$ according to densities $g_0(s) = 2\,(1-s)/(2a-1)$ and $g_1(s) = 2\,s/(2a-1)$ where $1/2 < a \leq 1$. This private belief distribution can be interpreted as an extension of the example from the basic illustration (section II.2) where every player's information quality is uniformly distributed on the interval $\left(\frac{1}{2}, a\right)$. Each player knows the realized precision of her own signal but not those of others' signals. We have used such distributions already in the proof of Proposition II.3. Their simple form permits us to express the benchmark strategies in closed form and illustrate them graphically. As in the illustrations in the main part we will also discuss expected payoff advantages of simple generalized belief updating strategies.

## II.C.1 Mixed Worlds

We restrict ourselves to the case $K = 2$. Private belief distributions $(G_0^k, G_1^k)$, $k = 1, 2$, are characterized by parameters $a_k$ where w.l.o.g. $a_1 < a_2$. Furthermore we fix $\alpha_1 = \alpha_2 = 1/2$. The simple form of the distributions $(G_0^k, G_1^k)$, $k = 1, 2$ permits us to express the benchmark strategy $\sigma^*_{\mathscr{R}}$ via a private belief threshold function. Recall the (naive) Bayes-rational strategy $\sigma^*(s,h)$ given by

$$\sigma^*(s,h) = \begin{cases} 1 & \text{if } b(s,\emptyset) > 1 - b(\emptyset, h \mid \overline{\varphi}_1^*, \overline{\varphi}_2^*) \\ 0 & \text{if } b(s,\emptyset) < 1 - b(\emptyset, h \mid \overline{\varphi}_1^*, \overline{\varphi}_2^*) \end{cases}$$

where

$$b(\emptyset, h \mid \overline{\varphi}_1^*, \overline{\varphi}_2^*) = \frac{\overline{\varphi}_1^*(h \mid 1) + \overline{\varphi}_2^*(h \mid 1)}{\overline{\varphi}_1^*(h \mid 1) + \overline{\varphi}_1^*(h \mid 0) + \overline{\varphi}_2^*(h \mid 1) + \overline{\varphi}_2^*(h \mid 0)}$$

and

$$\overline{\varphi}_k^*(h_i \mid \theta) = \prod_{j<i} \int_{1-a_k}^{a_k} \sigma^*(s, h_j)\, dG_\theta^k(s).$$

Then $\sigma^*_{\mathscr{E}}$ may be expressed as

$$\sigma^*_{\mathscr{E}}(s,h) = \begin{cases} 1 & \text{if } b(s,\emptyset) > t^*_{\mathscr{E}}(h \mid \overline{\varphi}_1^*, \overline{\varphi}_2^*) \\ 0 & \text{if } b(s,\emptyset) < t^*_{\mathscr{E}}(h \mid \overline{\varphi}_1^*, \overline{\varphi}_2^*) \end{cases}$$

where

$$t^*_{\mathscr{E}}(h \mid \overline{\varphi}_1^*, \overline{\varphi}_2^*) = \frac{\overline{\varphi}_1^*(h \mid 1)/(2a_1 - 1) + \overline{\varphi}_2^*(h \mid 1)/(2a_2 - 1)}{\left[\overline{\varphi}_1^*(h \mid 1) + \overline{\varphi}_1^*(h \mid 0)\right]/(2a_1 - 1) + \left[\overline{\varphi}_2^*(h \mid 1) + \overline{\varphi}_2^*(h \mid 0)\right]/(2a_2 - 1)}.$$

Threshold strategies permit a simple graphical representation in public belief – private belief – space ($[0,1]^2$) via the threshold function. In particular for each history $h$ we plot the threshold

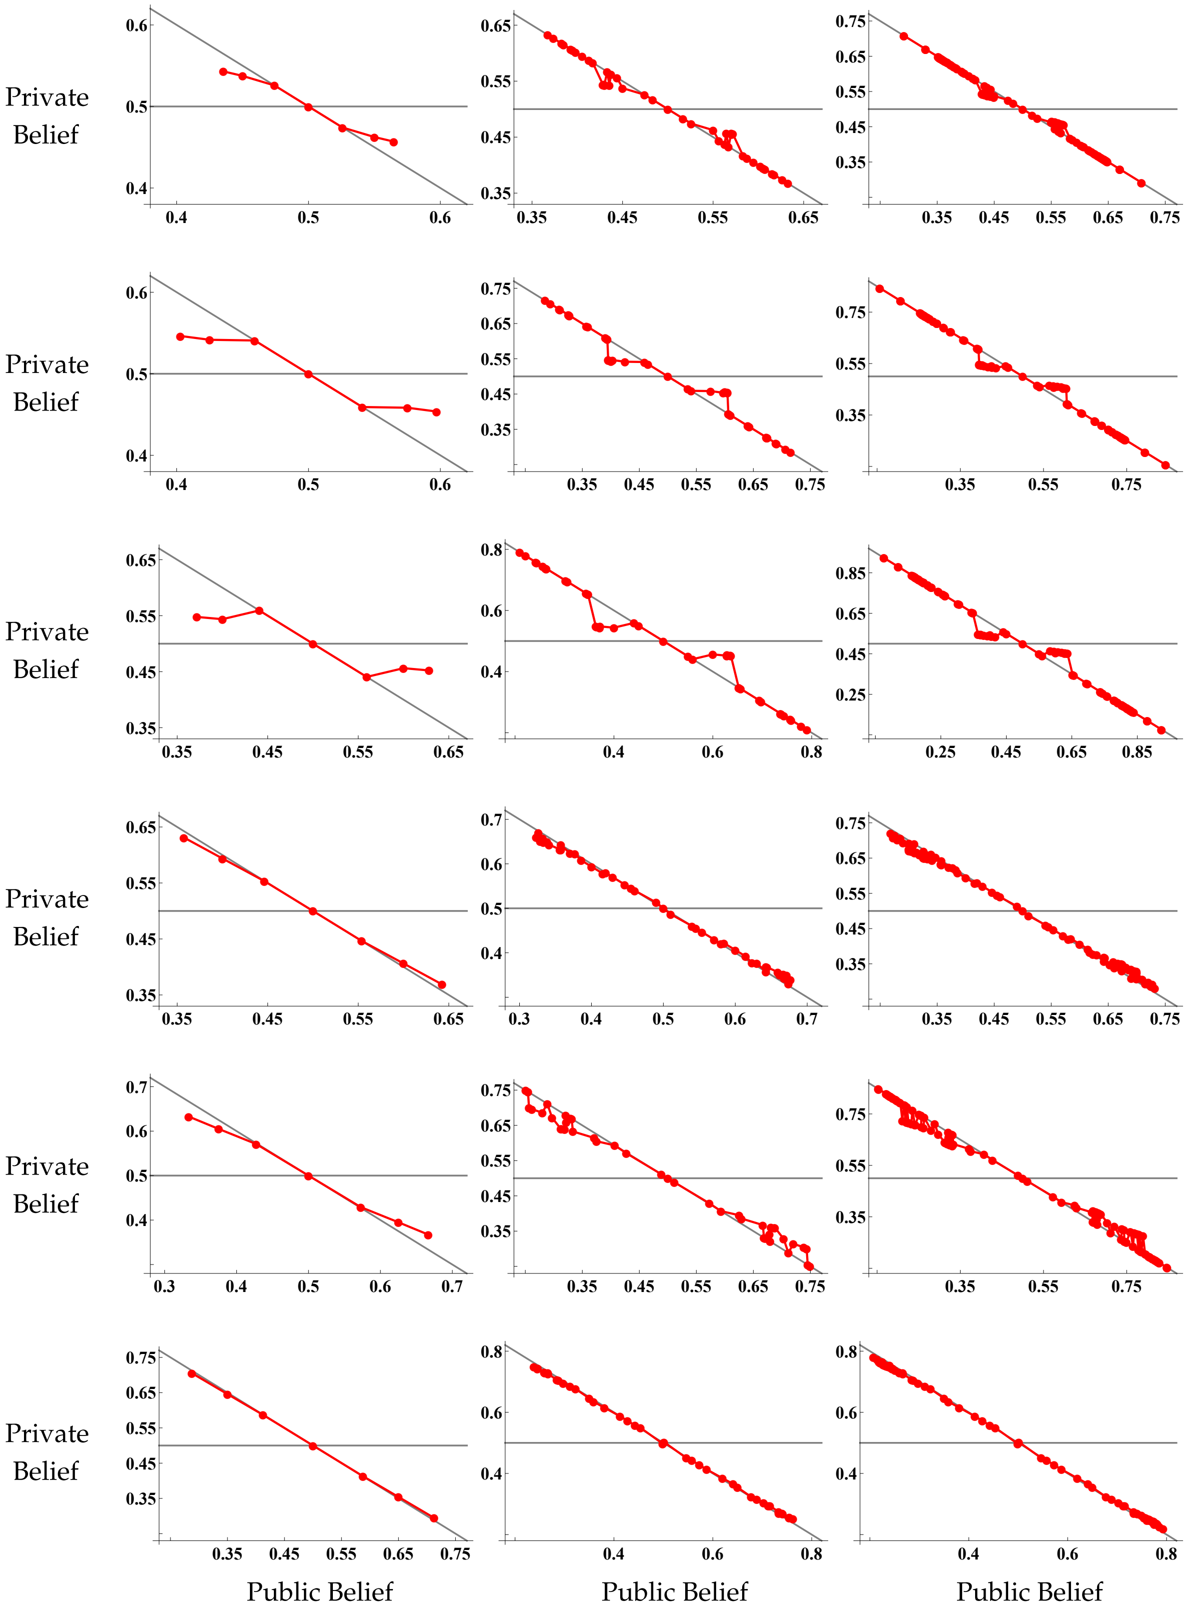


Figure II.3.: Benchmark strategy (red) and learned Bayes-rational strategy (gray) in a world of two uniform signal quality environments: Columns from left to right – $n = 3$, $n = 6$, and $n = 10$. Rows from top to bottom – $(a_1, a_2) = (0.55, 0.65)$, $(a_1, a_2) = (0.55, 0.75)$, $(a_1, a_2) = (0.55, 0.85)$, $(a_1, a_2) = (0.65, 0.75)$, $(a_1, a_2) = (0.65, 0.85)$, and $(a_1, a_2) = (0.75, 0.85)$. Markers indicate histories. For readability we restrict ourselves to histories occurring at least 2% of the time in a population of naive Bayes-rational learners.

functions $t^*_{\mathscr{E}}$ and $t^*(h \mid \overline{\varphi}_1^*, \overline{\varphi}_2^*) = 1 - b(\emptyset, h \mid \overline{\varphi}_1^*, \overline{\varphi}_2^*)$ against the public belief $b(\emptyset, h \mid \overline{\varphi}_1^*, \overline{\varphi}_2 *)$. The latter separate the space in two parts such that players invest for public belief-private belief-pairs above the threshold and reject for the remaining combinations. The closer a threshold function lies to the constant function $f(b) \equiv 1/2$, the stronger the associated strategy relies on the private belief. Figure II.3 presents numerical calculations of the threshold $t^*_{\mathscr{E}}$ associated with the benchmark strategy $\sigma^*_{\mathscr{E}}$ (red lines) and the threshold $t^*$ of the (naive) Bayes-rational strategy $\sigma^*$ (gray lines) for different values of $a_1$, $a_2$ and $n$. Histories are marked (on the red lines). For readability we restrict ourselves to histories occurring with probability at least 1/50 under $\alpha_1 \overline{\varphi}_1^* + \alpha_2 \overline{\varphi}_2^*$. As one can see while the Bayes-rational strategy approximates well the benchmark at many histories, there do exist histories such that overweighting of private information is profitable. Furthermore this is more pronounced the more different the distributions $(G_0^1, G_1^1)$ and $(G_0^2, G_1^2)$ (i.e. the larger $a_2 - a_1$).

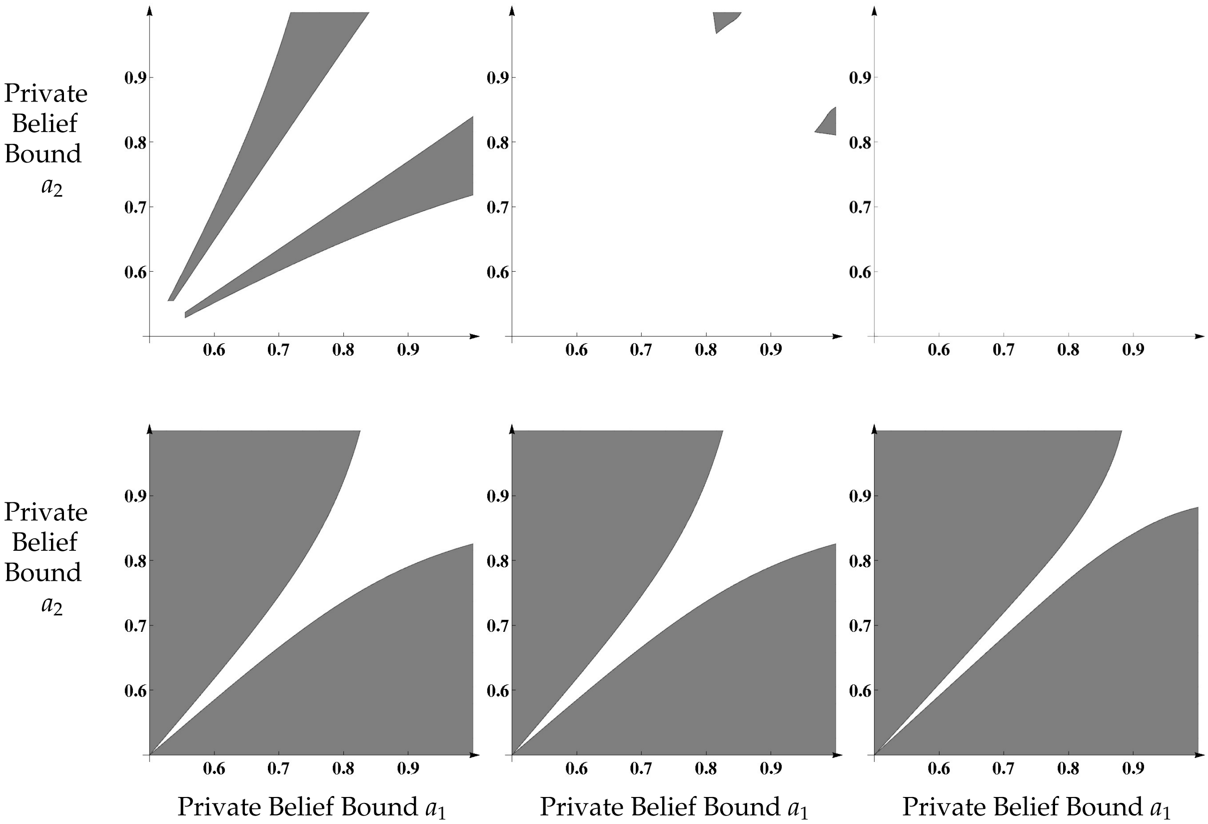


Figure II.4.: Uniform signal quality worlds for which the non-Bayesian type has the higher expected payoff: Columns from left to right – $n = 3$, $n = 6$, and $n = 10$. Rows – $\beta = 1.1$ (top), and $\beta = 0.9$ (bottom).

As in the main part of the text we now illustrate for this specific social learning world how simple strategies may outperform the naive Bayes-rational strategy $\sigma^*$ using the example of heterogeneous belief updating rules (Chapter III). To wit in this model each player is characterized by a private belief weight $\beta \in (0, \infty)$ which can be shown to induce the strategy

$$\sigma_\beta(s, h) = \begin{cases} 1 & \text{if } b(s, \emptyset) > t_\beta(h \mid \overline{\varphi}_1^*, \overline{\varphi}_2^*) \\ 0 & \text{if } b(s, \emptyset) < t_\beta(h \mid \overline{\varphi}_1^*, \overline{\varphi}_2^*) \end{cases}$$

where

$$t_\beta(h \mid \overline{\varphi}_1^*, \overline{\varphi}_2^*) = \frac{\left[\overline{\varphi}_1^*(h \mid 0) + \overline{\varphi}_2^*(h \mid 0)\right]^{1/\beta}}{\left[\overline{\varphi}_1^*(h \mid 1) + \overline{\varphi}_2^*(h \mid 1)\right]^{1/\beta} + \left[\overline{\varphi}_1^*(h \mid 0) + \overline{\varphi}_2^*(h \mid 0)\right]^{1/\beta}}.$$

This straightforwardly permits the computation of expected payoff $U_{\mathscr{E}}(\beta) = U_{\mathscr{E}}(\sigma_\beta \mid \overline{\varphi}_1^*, \overline{\varphi}_2^*)$ assuming as before a population of naive Bayes-rational ($\beta = 1$) players and a single player of type $\beta \neq 1$.[53] For given $\beta$ and $n$ the difference $U_{\mathscr{E}}(\beta) - U_{\mathscr{E}}(1)$ can easily be calculated numerically. Figure II.4 depicts regions in $(a_1, a_2)$-space where this difference is positive for $\beta \in \{0.9, 1.1\}$ and different values of $n$. The figures indicate that for this model few constellations favor overweighting of private information. On the other hand slight underweighting of private information seems profitable for many pairs $(a_1, a_2)$. In general however, the naive Bayes-rational strategy ($\beta = 1$) apparently provides a good approximation of the benchmark strategy $\sigma_{\mathscr{E}}^*$.

### II.C.2 Finite Repetitions

We consider a single uniform signal quality distribution parametrized by $a$. Throughout this part we furthermore maintain the assumption introduced in the main part that for each $i = 2, \dots, n$ and each $h_i \in H_i$, $\eta_{h_i} = \eta_i = \eta_2\, 2^{(i-2)/2}$ where $\eta_2 = 0.03$. We start with an illustration of the benchmark strategy $\sigma_\eta^*$. Recall that $\sigma_\eta^*(s, \emptyset) = 1$ if $b(s, \emptyset) > 1/2$, $\sigma_\eta^*(s, \emptyset) = 0$ if $b(s, \emptyset) < 1/2$, and

$$\sigma_\eta^*(s, x) = \begin{cases} 1 & \text{if } b(s, \emptyset) > t_\eta^*(x \mid \varphi^*) \\ 0 & \text{if } b(s, \emptyset) < t_\eta^*(x \mid \varphi^*) \end{cases}$$

where

$$t_\eta^*(x \mid \varphi^*) = \frac{\sum_{h \in H_{-1}} \varphi^*(h \mid 0)\, \phi_h\left(x - b(\emptyset, h \mid \varphi^*)\right)}{\sum_{h \in H_{-1}} \left[\varphi^*(h \mid 1) + \varphi^*(h \mid 0)\right] \phi_h\left(x - b(\emptyset, h \mid \varphi^*)\right)},$$

and for $h \in H_i$, $\phi_h$ is the density of the normal distribution with mean zero and standard deviation $\eta_i = \eta_2 * 2^{(i-2)/2}$. Furthermore $\varphi^*$ denotes occurrence probabilities of histories under the naive Bayes-rational strategy determined by the threshold function $t^*(x) = 1 - x$ (and $\sigma^*(s, \emptyset) = \sigma_\eta^*(s, \emptyset)$). As before we represent strategies $\sigma^*$ and $\sigma_\eta^*$ graphically via their threshold functions in public belief – private belief – space.

Graphs are given in Figure II.5. We consider $n \in \{3, 6, 10\}$ and $a \in \{0.55, 0.65, 0.75, 0.85\}$. Notice that since with noise a set of public beliefs is likely to occur at any history, we do not mark histories in the graphs. One can see from the graphs that noise favors overweighting of private beliefs since thresholds of the benchmark strategy are closer to the constant function $f(x) \equiv 1/2$ than those of the naive Bayes-rational strategy $\sigma^*$. Furthermore this becomes more pronounced the smaller $a$, the upper bound of private beliefs.

We finally turn again to strategies derived from the heterogeneous updating rule model. A player of type $\beta \in (0, \infty)$ invests at history $h$ provided $[b(s, \emptyset)]^\beta\ [b(\emptyset, h) + \tilde{\epsilon}_h] > [1 - b(s, \emptyset)]\ [1 - b(\emptyset, h) - \tilde{\epsilon}_h]$

[53] We calculate expected payoff for a slightly extended version of the game were players are randomly assigned to positions at the beginning of time.

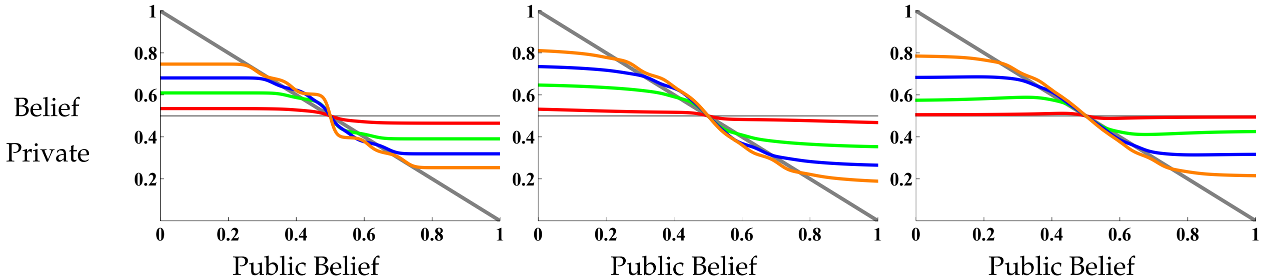


Figure II.5.: Benchmark strategy (colored) and learned Bayes-rational strategy (gray) in a noisy uniform signal quality social learning game: Columns from left to right – $n = 3$, $n = 6$, and $n = 10$. Colors capture different parameters of the private belief distributions – $a = 0.55$ (red), $a = 0.65$ (green), $a = 0.75$ (blue), and $a = 0.85$ (orange).

which is equivalent to

$$\tilde{\epsilon}_h > \frac{[1 - b(s, \emptyset)]^{\beta}}{[b(s, \emptyset)]^{\beta} + [1 - b(s, \emptyset)]^{\beta}}.$$

Accordingly,

$$U_{\eta}(\beta) = U_{\eta}(\sigma_{\beta}) =$$

$$\frac{1}{4n}\left\{G_0(1/2) - G_1(1/2) + \sum_{h \in H_{-1}} \int_{1-a}^{a} \left[1 - \Phi_h\left(\frac{(1-s)^{\beta}}{s^{\beta} + (1-s)^{\beta}} - b(\emptyset, h \mid \varphi^*)\right)\right] [\varphi^*(h \mid 1)\, dG_1(s) \; - \; \varphi^*(h \mid 0)\, dG_0(s)]\right\}.$$

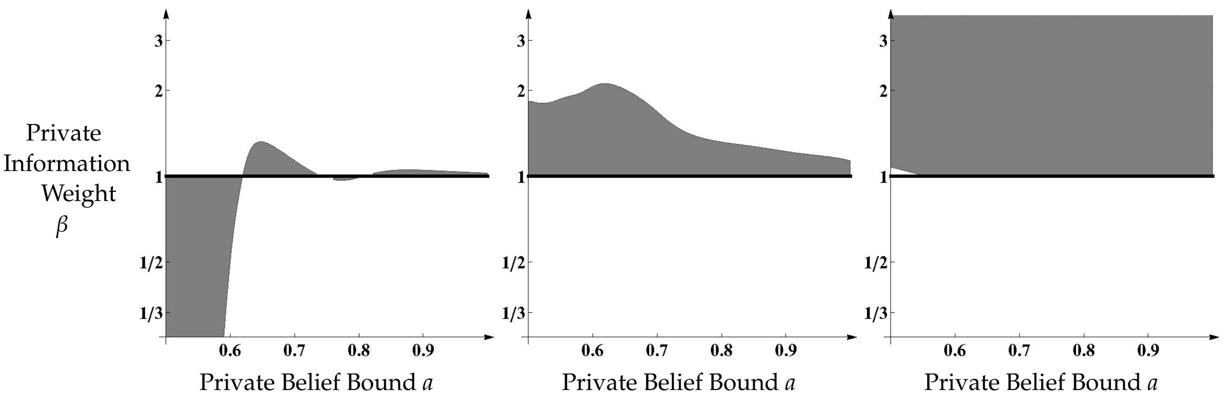


Figure II.6.: Expected Payoff Advantages of Non-Bayesian updaters versus Bayesian players in the uniform signal quality model: The graphs depicts values of $a$ and $\beta$ such that in the noisy social learning game with uniform signal quality structure parametrized by $a$, a player of type $\beta$ has a higher expected payoff than a Bayes-rational player ($\beta = 1$). From left to right: $n = 3$, $n = 6$, and $n = 10$.

In figure II.6 we illustrate for different lengths $n$ of the social learning sequence values of $a$ and $\beta$ such that $U(\beta) > U(1)$, i.e. such that players of type $\beta$ have a higher expected payoff than Bayes-rational players with $\beta = 1$ in the associated noisy social learning game. As one can see for small $n$ small $a$ favor underweighting of private information while overweighting of private information is rarely profitable. For larger $n$ however the picture shifts: For almost all upper bounds $a$ on the distribution of private beliefs, overweighting of private information is more profitable than Bayesian updating.

## Appendix II.D Omitted Proofs

***Proof of Proposition II.1.*** We first show that the result is independent of private belief supports and feedback. Let $R \to \infty$. First, as each round a player's position is drawn uniformly randomly, each player will infinitely often decide in period $n$. Hence, even with partial feedback on choices players observe the complete sequence of choices infinitely often. Second, even with the least informative support of private beliefs each round each player is endowed with a private belief $b(s, h) > \frac{1}{2}$ favoring investment with strictly positive probability. Therefore a finite sequence of players are all endowed with such beliefs with strictly positive probability each round as well. Hence, independent of her private belief support each player will infinitely often invest at period $n$ in which case she observes both, the state of Nature and the complete sequence of choices. Consequently, we may restrict ourselves to the case where each player has complete feedback on choices and unconditional feedback on the state and the private belief support is homogeneous.

Fix some Bayesian equilibrium $\sigma^*$ and let $(b^*(\emptyset, h))_{h \in H}$ denote the associated system of public beliefs. We define equilibrium assessments $\varphi^*$ via $\varphi^*\left(a_j \mid h_j, \theta\right) = \int_{\underline{b}}^{\bar{b}} \sigma^*\left(a_j \mid s, h_j\right) dG_\theta(s)$ and $\varphi^*(h_i \mid \theta) = \prod_{j<i} \varphi^*(a_j \mid h_j, \theta)$. Throughout the proof we restrict ourselves to analyzing the learning process of a single representative player.

In the first period, $h_1 = \emptyset$ w. Pr. 1 which clearly implies $\hat{\varphi}(h_1 \mid \theta; \tilde{\zeta}_r) = \varphi^*(h_1 \mid \theta) = 1$ with probability 1 for each $r \geq 1$, hence $\mathbf{P}(\{\zeta : \hat{\varphi}(h_1 \mid \theta; \zeta) \to \varphi^*(h_1 \mid \theta) = 1 \text{ for each } \theta \in \Theta\}) = 1$.

Assume now that the latter holds for each $j \leq i$, $h_j \in H_j$ and $\theta \in \Theta$, i.e.

$$\mathbf{P}\left(\left\{\zeta : \lim_{r \to \infty} \hat{\varphi}(h_j \mid \theta; \zeta) = \varphi^*(h_j \mid \theta) \text{ for each } j \leq i, h_j \in H_j \text{ and } \theta \in \Theta\right\}\right) = 1.$$

In period $i$ we distinguish two cases:

**Case 1 – $h_i$ lies <u>on</u> the equilibrium path:** In this case $\varphi^*(h_i \mid \theta)$ for some $\theta \in \Theta$ and thus $b^*(\emptyset, h_i)$ is well-defined. By induction assumption with $\mathbf{P}$-probability 1, $\hat{\varphi}(h_i \mid \theta; \tilde{\zeta}_r) \to \varphi^*(h_i \mid \theta)$. Thus, since $\hat{b}(\emptyset, h_i \mid \tilde{\zeta}_r)$ is continuous in $\hat{\varphi}(h_i \mid \theta; \tilde{\zeta}_r)$ for any $\epsilon > 0$ there exists $R > 1$ s.t. for each $r > R$, $|b^*(\emptyset, h_i) - \hat{b}(\emptyset, h_i \mid \tilde{\zeta}_r)| < \epsilon$ with probability 1 on $\mathscr{Z}_r$. Thus, unless the distribution of private beliefs has an atom at $1 - b^*(\emptyset, h_i)$ (in which case there would be a tie at $h_i$ – as noted before a non-generic event) there exists $R > 1$ s.t. for each $r > R$, $\hat{\sigma}(s, h_i \mid \tilde{\zeta}_r) = \sigma^*(s, h_i)$ since responses are myopic Bayes-rational. Accordingly, by the law of large numbers it must hold that

$$\hat{\varphi}\left(a_i \mid h_i, \theta; \tilde{\zeta}_{R+r_2}\right) \to \int_{\underline{b}}^{\bar{b}} \sigma^*(s, h_i)\, dG_\theta(s) = \varphi^*(a_i \mid h_i, \theta)$$

eventually as $r_2 \to \infty$ for each $\theta \in \Theta$ and $a_i \in \{0, 1\}$. Consequently, $\hat{\varphi}((h_i, a_i) \mid \theta; \tilde{\zeta}_r) \to \varphi^*((h_i, a_i) \mid \theta)$ eventually as $r \to \infty$.

**Case 2 – $h_i$ lies <u>off</u> the equilibrium path:** In this case necessarily $\varphi^*(h_i \mid \theta) = 0$ for each $\theta \in \Theta$ and by the induction assumption along any $\zeta \in \mathscr{Z}$, $\hat{\varphi}(h_i \mid \theta; \zeta_r) \to 0$ for any $\theta \in \Theta$. Hence, technically $h_i$ has no bearing on whether the learning process eventually approaches equilibrium. Still, we will subsequently show that under the current assumptions on the learning process, limit behavior at $h_i$ is not restricted. Clearly, there exists $a_j \in h_i$ s.t.

$$\varphi^*(a_j \mid h_j, \theta) = \int_{\underline{b}}^{\bar{b}} \sigma^*(a_j \mid b, h_j)\, dG_\theta(b) = 0$$

for each $\theta \in \Theta$. If also $\kappa(a_j, h_j, \theta \mid \zeta) = 0$ for each $\theta \in \Theta$, $a_j$ conveys no information by assumption. However, depending on the realized super-history of play $\zeta$ it might be that $\kappa(a_j, h_j, \theta \mid \zeta) > 0$ for some $\theta \in \Theta$ if $a_j$ has been the myopic response of some player to a belief derived from wrong expectations before convergence took place. In this case there exists $\zeta' \in \mathscr{Z}$ with $\mathbf{P}(\zeta') > 0$ and $r' \geq 1$ such that for each $r$ with $r' < r < \infty$, $\hat{\varphi}(a_j \mid h_j, \theta; \zeta'_r) > 0$ and posterior beliefs after $a_j$ are well-defined for any $r > r'$ and hence in the limit as $r \to \infty$. Moreover in the limit the observation of $a_j$ may result in a change of posterior beliefs unlike in the case where $\kappa(a_j, h_j, \theta \mid \zeta) = 0$ for each $\theta \in \Theta$. Therefore different super-histories $\zeta$ may induce different public beliefs $\hat{b}(\emptyset, h_i \mid \zeta)$. Consequently, the strategic limit $\hat{\sigma}(\mu, h_i \mid \tilde{\zeta})$ need not be a Dirac measure on some $z \in [0,1]$.

In conclusion the strategic outcome of the learning process agrees with the equilibrium strategy profile along all histories on the equilibrium path. Off the equilibrium path agreement occurs with some probability $\pi \in [0,1]$. Hence, depending on the specificities of the game several Bayesian equilibria are possibly selected by the learning process. But since the Bayesian equilibrium outcome is unique (outside non-generic frameworks) the conclusion of the proposition follows.

□

***Proof of Proposition II.2***. We employ a similar induction argument as in the proof of Proposition II.1. Fix $\mathscr{E}$. In period 1 in each round of the learning process players invest provided $b(\tilde{s}_1, \emptyset) > 1/2$, i.e. with probability $1 - G_\theta^k(1/2)$ in environment $E_k$ given realized state $\theta$. Absent feedback about the realized environment in a round players employ a single assessment $\hat{\varphi}$. By the law of large numbers along any $\zeta$, $\hat{\varphi}(1 \mid \emptyset, \theta; \zeta_r) \to \sum_{k=1}^{K} \alpha_k \left[1 - G_\theta^k(1/2)\right] = 1 - \sum_{k=1}^{K} \alpha_k\, G_\theta^k(1/2) = \overline{\varphi}^*(1 \mid \emptyset, \theta)$ as $r \to \infty$ for any $\theta \in \Theta$.

Assume $\mathbf{P}\left(\lim_{r\to\infty} \hat{\varphi}(a \mid h_j, \theta; \tilde{\zeta}_r) = \overline{\varphi}^*(a \mid h_j, \theta)\right) = 1$ for any $j < i$, $h_j \in H_j$ s.t. $\overline{\varphi}^*(h_j \mid \theta) > 0$ for some $\theta \in \Theta$, $a \in A$ and $\theta \in \Theta$. Hence, eventually

$$b(\emptyset, h_i \mid \hat{\varphi}) = \frac{\prod_{j<i} \hat{\varphi}(a_j \mid h_j, 1; \tilde{\zeta}_r)}{\prod_{j<i} \hat{\varphi}(a_j \mid h_j, 1; \tilde{\zeta}_r) \;+\; \prod_{j<i} \hat{\varphi}(a_j \mid h_j, 0; \tilde{\zeta}_r)} \to b(\emptyset, h_i \mid \overline{\varphi}^*)$$

for any $h_i$ such that $\overline{\varphi}^*(h_i \mid \theta) > 0$ for some $\theta \in \Theta$. Then absent ties there exists $R_1$ such that eventually for any $r > R_1$

$$\hat{\sigma}(s, h_i \mid \tilde{\zeta}_r) = \begin{cases} 1 & \text{if } b(s,\emptyset) > 1 - b(\emptyset, h_i \mid \overline{\varphi}^*) \\ 0 & \text{if } b(s,\emptyset) < 1 - b(\emptyset, h_i \mid \overline{\varphi}^*) \end{cases}$$

for any $h_i$ as described above and any $b(s,\emptyset) \in \bigcup_{k=1}^{K} supp(G_\theta^k)$. Consequently,

$$Pr(\tilde{a}_i = 1 \mid h_i, \theta; \tilde{\zeta}_{R_1+r}) = 1 - Pr(\tilde{a}_i = 0 \mid h_i, \theta; \tilde{\zeta}_{R_1+r}) \to 1 - \sum_{k=1}^{K} \alpha_k\, G_\theta^k \left(1 - b(\emptyset, h_i \mid \overline{\varphi}^*)\right) = \overline{\varphi}^* * (1 \mid h_i, \theta)$$

eventually as $r \to \infty$ for each $\theta \in \Theta$. Therefore by another application of the law of large numbers,

$$\mathbf{P}\left(\lim_{r\to\infty} \hat{\varphi}(a \mid h_i, \theta; \tilde{\zeta}_r) = \overline{\varphi} * (a \mid h_i, \theta)\right) = 1$$

for each $a \in A$, each $\theta \in \Theta$ and each $h_i \in H_i$ s.t. $\overline{\varphi} * (h_i \mid \theta) > 0$ for some $\theta \in \Theta$. On the other hand each $h_i$ s.t. $\overline{\varphi} * (h_i \mid \theta) = 0$ for each $\theta \in \Theta$ lies off the equilibrium path and hence does not have any bearing upon the (unique) equilibrium outcome.

□

***Proof of Lemma II.1***. Fix period $i$. We rewrite expected payoff as

$$\begin{aligned} U_i^{\mathscr{E}}(\sigma \mid \overline{\boldsymbol{\varphi}}^*) &= \frac{1}{4} \sum_{k=1}^{K} \alpha_k \sum_{h_i \in H_i} \int_0^1 \sigma(s, h_i) \left[\overline{\varphi}_k^*(h_i \mid 1)\, dG_1^k(s) - \overline{\varphi}_k^*(h_i \mid 0)\, dG_0^k(s)\right] \\ &= \frac{1}{4} \sum_{h_i \in H_i} \int_0^1 \sigma(s, h_i) \sum_{k=1}^{K} \alpha_k \left[\overline{\varphi}_k^*(h_i \mid 1)\, dG_1^k(s) - \overline{\varphi}_k^*(h_i \mid 0)\, dG_0^k(s)\right]. \end{aligned}$$

Obviously, this expression is maximized on $\Sigma$ by selecting $\sigma(s,h_i) = 1$ whenever $\overline{U}(s,h_i) = \sum_{k=1}^{K} \alpha_k \left[\overline{\varphi}_k^*(h_i \mid 1)\, dG_1^k(s) - \overline{\varphi}_k^*(h_i \mid 0)\, dG_0^k(s)\right] > 0$ and $\sigma(s,h_i) = 0$ whenever $\overline{U}(s,h_i) < 0$. This straightforwardly yields the optimal strategy $\sigma_{\mathscr{E}}^*$.

□

***Proof of Proposition II.4***. We calculate the second derivative of

$$\check{b}(s,h,\epsilon) = \frac{b(s,\emptyset)\ [b(\emptyset,h) + \epsilon]}{b(s,\emptyset)\ [b(\emptyset,h) + \epsilon]\ +\ [1 - b(s,\emptyset)]\ [1 - b(\emptyset,h) - \epsilon]}$$

with respect to $\epsilon$. We obtain

$$\frac{\partial^2\, \check{b}(s,h,\epsilon)}{\partial\, \epsilon^2} = \frac{2\, b(s,\emptyset)\ [1 - b(s,\emptyset)]\ [1 - 2\, b(s,\emptyset)]}{[b(s,\emptyset)\ [b(\emptyset,h) + \epsilon]\ +\ [1 - b(s,\emptyset)]\ [1 - b(\emptyset,h) - \epsilon]]^3}.$$

Obviously, $\partial^2\, \check{b}(s,h,\epsilon)/\partial\, \epsilon^2 > 0$ if $b(s,\emptyset) < 1/2$ and $\partial^2\, \check{b}(s,h,\epsilon)/\partial\, \epsilon^2 < 0$ if $b(s,\emptyset) > 1/2$. Therefore $\check{b}(s,h,\epsilon)$ is strictly convex in $\epsilon$ if $b(s,\emptyset) < 1/2$ and strictly concave in $\epsilon$ if $b(s,\emptyset) > 1/2$. Furthermore clearly, $b'(s,h) = \check{b}(s,h,\tilde{\epsilon}_h)$. Thus if $\eta_h > 0$ by Jensen's inequality for $b(s,\emptyset) < 1/2$,

$$E_{\tilde{\epsilon}_h}\,[b'(s,h)] > \check{b}\left(s,h,E_{\tilde{\epsilon}_h}[\tilde{\epsilon}_h]\right) = \check{b}(s,h,0) = b(s,h)$$

and for $b(s,\emptyset) > 1/2$

$$E_{\tilde{\epsilon}_h}\,[b'(s,h)] < \check{b}\left(s,h,E_{\tilde{\epsilon}_h}[\tilde{\epsilon}_h]\right) = \check{b}(s,h,0) = b(s,h).$$

□

***Proof of Lemma II.2.*** Let $\sigma \in \Sigma_\eta$. We rewrite expected payoff in the noisy social learning game as

$$U_\eta(\sigma \mid \varphi^*) = \frac{1}{4n} \int_{\underline{b}}^{\bar{b}} \left\{ \sigma(s,\emptyset)\, [dG_1(s) - dG_0(s)] + \int_0^1 \sigma(s,x)\, [f(x \mid 1)\, dG_1(s) - f(x \mid 0)\, dG_0(s)]\, dx \right\}$$

where

$$f(x \mid \theta) = \sum_{h \in H_{-1}} \varphi^*(h \mid \theta)\, \phi_h\, (x - b(\emptyset, h \mid \varphi^*))$$

is the probability density function of the random variable $b'(\emptyset, \tilde{h})$. The same argument we have used before then implies, that the optimal strategy $\sigma^*_\eta \in \Sigma_\eta$ is given by

$$\sigma^*_\eta(s,\emptyset) = \begin{cases} 1 & \text{if } b(s,\emptyset) > 1/2 \\ 0 & \text{if } b(s,\emptyset) < 1/2 \end{cases}$$

and

$$\sigma^*_\eta(s,x) = \begin{cases} 1 & \text{if } b(s,\emptyset) > \frac{\sum_{h\in H_{-1}} \varphi^*(h|0)\, \phi_h\left(x - b(\emptyset,h|\varphi^*)\right)}{\sum_{h\in H_{-1}} \left[\varphi^*(h|1) + \varphi^*(h|0)\right] \phi_h\left(x - b(\emptyset,h|\varphi^*)\right)} \\ 0 & \text{if } b(s,\emptyset) < \frac{\sum_{h\in H_{-1}} \varphi^*(h|0)\, \phi_h\left(x - b(\emptyset,h|\varphi^*)\right)}{\sum_{h\in H_{-1}} \left[\varphi^*(h|1) + \varphi^*(h|0)\right] \phi_h\left(x - b(\emptyset,h|\varphi^*)\right)} \end{cases}.$$

□

***Proof of Proposition II.5.*** We show that generically the first derivative of

$$t(x \mid \boldsymbol{\eta}, \varphi^*) = \frac{\sum_{h\in H_{-1}} \varphi^*(h \mid 0)\, \phi_h\, (x - b(\emptyset, h \mid \varphi^*))}{\sum_{h\in H_{-1}} \left[\varphi^*(h \mid 1) + \varphi^*(h \mid 0)\right] \phi_h\, (x - b(\emptyset, h \mid \varphi^*))}$$

with respect to $\eta_h^2$ is different from zero provided $\eta_{h'}$ is sufficiently large for some $h' \in H$. This in turn implies that for generic $\boldsymbol{\eta}$, $t(x \mid \boldsymbol{\eta}, \varphi^*) \neq 1 - x$.

As a first step we calculate

$$\begin{aligned} \frac{\partial\, \phi_h\, (x - b(\emptyset, h \mid \varphi^*))}{\partial\, \eta_h^2} &= \frac{1}{\sqrt{2\,\pi}} \exp\left( - \frac{(x - b(\emptyset, h \mid \varphi^*))^2}{2\,\eta_h^2} \right) \left( - \frac{1}{2\,\eta_h^2\, \sqrt{\eta_h^2}} \right) \\ &\quad + \frac{1}{\sqrt{2\,\pi\,\eta_h^2}} \exp\left( - \frac{(x - b(\emptyset, h \mid \varphi^*))^2}{2\,\eta_h^2} \right) \frac{(x - b(\emptyset, h \mid \varphi^*))^2}{2\,\eta_h^2\,\eta_h^2} \\ &= \frac{1}{\sqrt{2\,\pi\,\eta_h^2}} \exp\left( - \frac{(x - b(\emptyset, h \mid \varphi^*))^2}{2\,\eta_h^2} \right) \frac{1}{2\,\eta_h^2} \left( \frac{1}{\eta_h^2} - 1 \right) \\ &= \phi_h\, (x - b(\emptyset, h \mid \varphi^*))\, \frac{1 - \eta_h^2}{2\left[\eta_h^2\right]^2}. \end{aligned}$$

This implies that $\partial\,\phi_h\,(x - b(\emptyset, h \mid \varphi^*))\,/\partial\,\eta_h^2 \neq 0$ whenever $\eta_h^2 \neq 1$ for each $x$ satisfying $\phi_h\,(x - b(\emptyset, h \mid \varphi^*)) > 0$. The latter is approximately satisfied for each $x$ such that

$$\left|x - b(\emptyset, h \mid \varphi^*)\right| < 3\,\eta_h. \tag{II.3}$$

Denote by $N = \sum_{h \in H_{-1}} [\varphi^*(h \mid 1) + \varphi^*(h \mid 0)]\,\phi_h\,(x - b(\emptyset, h \mid \varphi^*))$. The derivative of $t(x \mid \boldsymbol{\eta}, \varphi^*)$ with respect to $\eta_h^2$ is then given by

$$\begin{aligned}
\frac{\partial\,t(x \mid \boldsymbol{\eta}, \varphi^*)}{\partial\,\eta_h^2} &= \frac{1}{N^2}\Bigg\{\varphi^*(h \mid 0)\,\frac{\partial\,\phi_h\,(x - b(\emptyset, h \mid \varphi^*))}{\partial\,\eta_h^2} * \sum_{h' \in H_{-1}} [\varphi^*(h' \mid 1) + \varphi^*(h' \mid 0)]\,\phi_{h'}\,(x - b(\emptyset, h' \mid \varphi^*)) \\
&\qquad - [\varphi^*(h \mid 1) + \varphi^*(h \mid 0)]\,\frac{\partial\,\phi_h\,(x - b(\emptyset, h \mid \varphi^*))}{\partial\,\eta_h^2} \sum_{h' \in H_{-1}} \varphi^*(h' \mid 0)\,\phi_{h'}\,(x - b(\emptyset, h \mid \varphi^*))\Bigg\} \\
&= \frac{1}{N^2}\,\frac{\partial\,\phi_h\,(x - b(\emptyset, h \mid \varphi^*))}{\partial\,\eta_h^2} \sum_{h' \in H_{-1}} \phi_{h'}\,(x - b(\emptyset, h \mid \varphi^*))\,[\varphi^*(h' \mid 1)\,\varphi^*(h \mid 0) - \varphi^*(h \mid 1)\,\varphi^*(h' \mid 0)]\,.
\end{aligned}$$

From the argumentation above $\partial\,t(x \mid \boldsymbol{\eta}, \varphi^*)/\partial\,\eta_h^2 \neq 0$ can only hold for $x$ satisfying (II.3). On the other hand the summand at $h$ of the sum in the last line is zero. Hence, a necessary condition for the sum not to be zero is existence of at least one $h' \neq h$ such that $\phi_{h'}\,(x - b(\emptyset, h')) \neq 0$ for at least one $x$ satisfying (II.3). Since generically $b(\emptyset, h') \neq b(\emptyset, h)$ this requires that $\eta_{h'}$ is sufficiently large. Existence of such a history $h'$ is however generically also sufficient for the sum to be different from zero. To see this assume first that there exists exactly one such history $h' \neq h$. In this case $\partial\,t(x \mid \boldsymbol{\eta}, \varphi^*)/\partial\,\eta_h^2 \neq 0$ requires $\varphi^*(h' \mid 1)\,\varphi^*(h \mid 0) \neq \varphi^*(h' \mid 0)\,\varphi^*(h \mid 1)$. Not being an identity, this is generically satisfied (see Smith and Sørensen, 2000, p.389). More generally if more than one $h'$ satisfy the condition the sum is one of several generically non-zero terms each weighted by $\phi_{h'}\,(x - b(\emptyset, h \mid \varphi^*))$. Clearly, this sum must generically be different from zero.

□

***Benefits of Overweighting with Heterogeneous Preferences (Section II.6.2)***. We have $K = 2$ and $\alpha_1 = \alpha_2 = 1/2$. The distribution of private beliefs in environment $E_k$, $k = 1, 2$, is given by $supp(G_\theta^k) = \{1 - q_k, q_k\}$ and $Pr\left(b(\tilde{s}, \emptyset) = q_k \mid \tilde{\theta} = 1\right) = Pr\left(b(\tilde{s}, \emptyset) = 1 - q_k \mid \tilde{\theta} = 0\right) = q_k$ where $1/2 < q_1 < q_2 < 1$. Furthermore in $E_2$ only a fraction $(1 - 2\,\xi_2)$ has standard preferences while there exist $\xi_2$ noise players each which always invest or always reject respectively. In both environments in the first two periods players optimally follow private information. (In $E_1$ we make this assumption for simplicity. It can be justified for instance by assuming a small positive fraction $0 < \xi_1 \ll \xi_2$ of noise players in this environment as well.) In the third period clearly imitating two similar decisions in period 3 is optimal in $E_1$ while in $E_2$ following private information is optimal provided $q_2$ and $\xi_2$ jointly satisfy

$$[\xi_2 + (1 - 2\,\xi_2)\,q_2]^2\,(1 - q_2) < [\xi_2 + (1 - 2\,\xi_2)\,(1 - q_2)]^2\,q_2 \;\Leftrightarrow\; \left(\frac{\xi_2}{1 - 2\,\xi_2}\right)^2 > q_2\,(1 - q_2).$$

On the other hand in the mixed environment assessments in the third period are given by

$$\begin{aligned}
\overline{\varphi}^*((1,1) \mid \tilde{\theta} = 1) = \overline{\varphi}^*((0,0) \mid \tilde{\theta} = 0) &= \left[\frac{q_1}{2} + \frac{\xi_2}{2} + \frac{1 - 2\,\xi_2}{2}\,q_2\right]^2, \\
\overline{\varphi}^*((1,1) \mid \tilde{\theta} = 0) = \overline{\varphi}^*((0,0) \mid \tilde{\theta} = 1) &= \left[\frac{1 - q_1}{2} + \frac{\xi_2}{2} + \frac{1 - 2\,\xi_2}{2}\,(1 - q_2)\right]^2.
\end{aligned}$$

Consider the heterogeneous updating rule model of chapter I. A player of type $\beta \in (0, \infty)$ in environment $E_k$ imitates the first two decisions in period 3 iff $\overline{\varphi}^*((1,1) \mid 1)\,(1-q_k)^\beta > \overline{\varphi}^*((1,1) \mid 0)\, q_k^\beta$ and follows private information if the inequality is reversed. Accordingly provided

$$\frac{\log\left(\overline{\varphi}^*((1,1) \mid 1)\right) - \log\left(\overline{\varphi}^*((1,1) \mid 0)\right)}{\log(q_2) - \log(1-q_2)} < \beta < \frac{\log\left(\overline{\varphi}^*((1,1) \mid 1)\right) - \log\left(\overline{\varphi}^*((1,1) \mid 0)\right)}{\log(q_1) - \log(1-q_1)}$$

the player correctly imitates in $E_1$ and correctly follows private information in $E_2$. Since $q_1 < q_2$ this intervalis well-defined. Furthermore the lower bound strictly exceeds 1 provided

$$q_2 < \frac{q_1^2}{q_1^2 + (1-q_1)^2}$$

and

$$\frac{\sqrt{q_2\,(1-q_2)} - 2\,q_2\,(1-q_2)}{(2\,q_2-1)^2} < \xi_2 << \frac{2\,\sqrt{q_2\,(1-q_2)} - 2\,q_2\,(1-q_2) - q_2\,(1-q_1) - q_1\,(1-q_2)}{(2\,q_2-1)^2}.$$

Hence, for these values of the parameters, overweighting is profitable with $n = 3$.

□

# Chapter III.

# Behavioral Social Learning

## III.1. Introduction

Our learning-theoretic critique of rational herding developed in the previous chapter suggests that individuals who combine private and public information in a non-Bayesian way are common in social learning environments. Non-Bayesian updating of beliefs results from adaptation to the significant structural uncertainty which prevails in the field. This finding therefore indicates that introducing heterogeneous updating rules in economic models of social learning is likely to improve their explanatory and predictive power.

In this chapter we follow this indication and 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. Introducing heterogeneous updating rules in a simple model of social learning leads to equilibrium predictions which are more in line with the laboratory evidence on rational herding 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.

This chapter is structured as follows. We first provide a short review of the laboratory evidence on rational herding in Section III.2. Section III.3 illustrates the implications of heterogeneous updating rules for the process of social learning in the specific model of Bikhchandani, Hirshleifer, and Welch (1992). In Section III.4 we extend the previous analysis to a general social learning environment. Section III.5 discusses related work. We conclude in Section III.6. Appendix III.A contains omitted proofs.

## III.2. Laboratory Evidence on Bayesian Rational Herding

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, 2010). Most of these economic experiments on social learning implement a simple environment which is based on the specific model of Bikhchandani, Hirshleifer, and Welch (1992) (simply 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.

As already mentioned, the fact that laboratory behavior systematically deviates from Bayesian rational play in the considered experimental settings comes as no surprise. Field environments in which social learning mainly takes place are likely to differ substantially from those experimental settings. If, as suggested in our second chapter, subjects develop rational rules of social learning then their behavior will deviate from Bayesian rational play in a setting which lacks structural uncertainty. The next section shows that incorporating subjects' adaptive response to structural uncertainty into an economic model of social learning drastically reduces the discrepancy between laboratory data and theoretical predictions.

## III.3. 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 the 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 the 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 the 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 the original 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.

### III.3.1. 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 $\theta$, 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 $i$, $i = 1, 2, \ldots$, observes a private signal (either high, $H$, or low, $L$) about $\theta$ 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 Bikhchandani, Hirshleifer, and Welch (1992) in their specific model.

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. $Pr\left(\tilde{\theta} = 1 \mid s_i, h_i\right)$ is player $i$'s belief that the true value of the investment payoff is one and $Pr\left(\tilde{\theta} = (-1) \mid s_i, h_i\right) = 1 - Pr\left(\tilde{\theta} = 1 \mid s_i, h_i\right)$. Since the expected payoff of investment is given by $E\left[\tilde{\theta} \mid s_i, h_i\right] = 2\,Pr\left(\tilde{\theta} = 1 \mid s_i, h_i\right) - 1$, player $i$ invests if $Pr\left(\tilde{\theta} = 1 \mid s_i, h_i\right) > 1/2$. Said differently, player $i$ invests if her likelihood ratio $\lambda(s_i, h_i) = Pr\left(\tilde{\theta} = 1 \mid s_i, h_i\right) / Pr\left(\tilde{\theta} = (-1) \mid s_i, h_i\right)$ is strictly greater than 1. Player $i$'s likelihood ratio is given by

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

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

In the 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 the 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 overweighters* ($\beta_i = 3$) (i.e. conformists, weak and strong overweighters 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.[1]

### 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.

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

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

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 overweighter 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 overweighter then he follows his private signal. The first weak overweighter who imitates after a sequence of identical investment decisions is Emma, the fifth player (assuming that $q$ is not too large).[2] Clearly, a strong overweighter will only imitate a long sequence of identical investment decisions since little information can be inferred from any herding decision.

[2]If $q$ 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.

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.[3] An investment cascade (rejection cascade) starts and lasts forever but, unlike in the 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 the 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. 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 overweighters in the 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 the specific model (see proof 2.1.2 in the Appendix). Figure III.1 shows the lower (upper) bound of the

[3]Conformists 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) overweighters 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 $(1-q)^2/q^2$ ($(1-q)^3/q^3$).

probability that a correct (incorrect) cascade arises in comparison to the probabilities in the 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 overweighters who follow their private signal when Bayes' rule predicts to imitate).

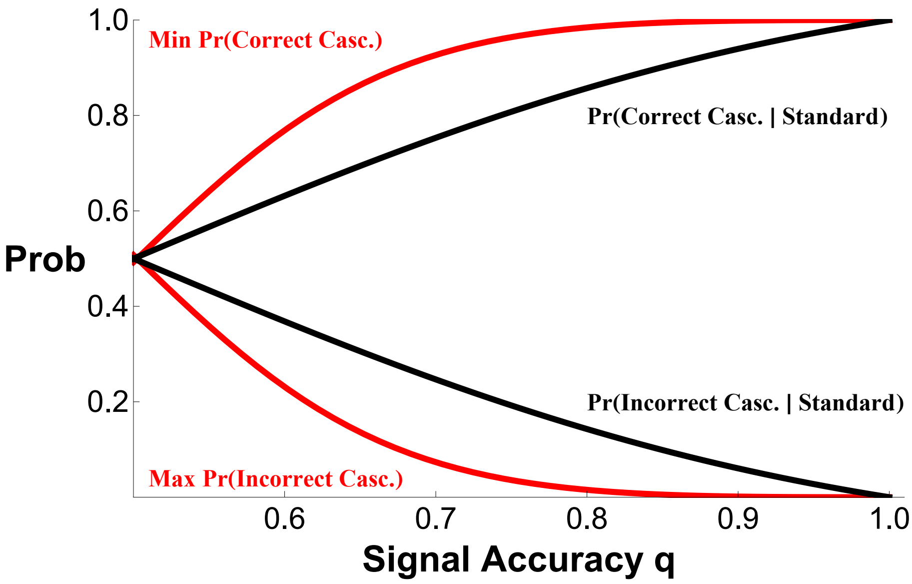


Figure III.1.: Lower (upper) bound upon the probability that a correct (incorrect) cascade arises compared to probabilities in the standard model.

### The Fragility of Information Cascades

Information cascades are less fragile compared to the original predictions 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 overweighters need at least 2 contradictory private signals.

### Summary

We have analyzed a straightforward extension of the 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 the specific model with heterogeneous updating rules nicely organize the experimental evidence on social learning. Compared to the original model, information cascades start latter 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.

### III.3.2. 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 the 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 the 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 the specific model.The existence and unicity of the ABEE in the social learning environment considered by the specific model with heterogeneous updating rules has been established with the help of simulations.[4] Figure III.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 (fictitious) private signal is of a slightly larger accuracy than $q$.

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

[4] In 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.

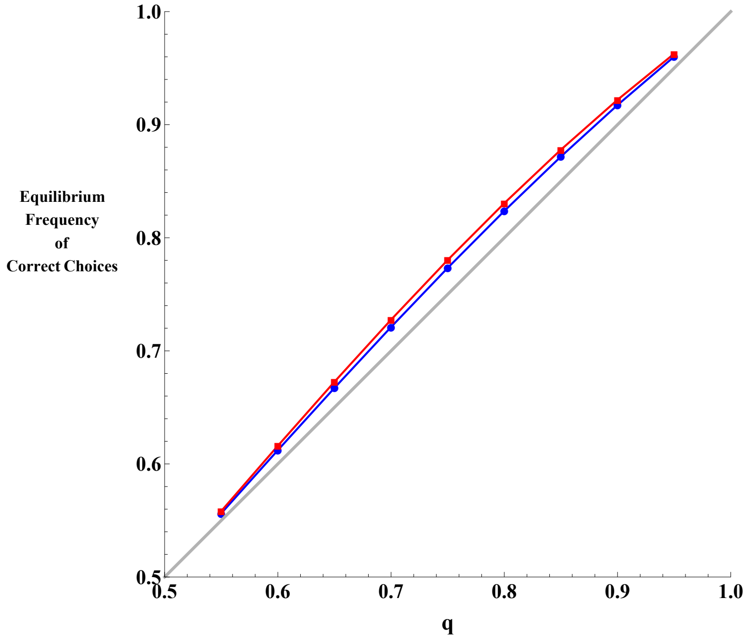


Figure III.2.: Frequencies of correct choices in ABEE for sequence lengths of 6 (blue line) and 10 players (red line).

(respectively reject) regardless of their private signal.

Like in the 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 the 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 III.3 illustrates the evolution of the public belief with the depth of a herd (length of identical investment decisions).

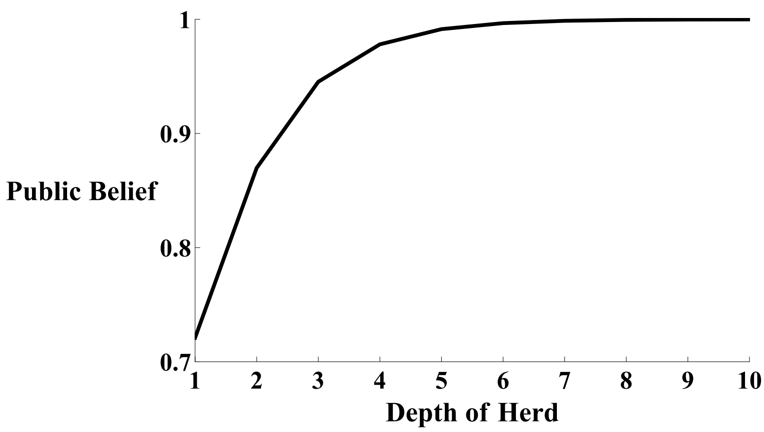


Figure III.3.: Evolution of public belief in a herd.

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 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.

## III.4. 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 overweighters. 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 overweighters on societal welfare in this second setting is much weaker.

### III.4.1. 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 $Pr\left(\tilde{\theta} = 1\right) = \frac{1}{2}$. Players $i = 1, 2, \ldots$ sequentially choose from the set of actions $A = \{0, 1\}$. Payoffs are given by the mapping $u : A \times \theta \to \mathbb{R}$ 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 $\theta$ 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 $G_\theta$. $G_{(-1)}$ 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 $supp(G_1) = supp(G_{(-1)}) = \left[\underline{b}, \overline{b}\right]$ where $\underline{b} > 0$ and $\overline{b} < 1$ meaning that private beliefs are bounded. Finally, each player $i$ observes the history $h_i = (a_1, \ldots, a_{i-1})$ of actions of all preceding players where $h_i \in H_i = A^{i-1}$.

Players aim at maximizing their expected utility $U(a) = E\left[u\left(a, \tilde{\theta}\right)\right]$. 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: $b(s_i, h_i) = Pr\left(\tilde{\theta} = 1 | b(\tilde{s}_i, \emptyset) = b(s_i, \emptyset), \tilde{h}_i = h_i\right)$. 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

$$\lambda(s_i, h_i) = (\lambda(s_i, \emptyset))^{\beta_i} \ \lambda(\emptyset, h_i) \tag{III.1}$$

where $\beta_i$ 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 $\beta_i$ 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 $\left(\tilde{\beta}_i\right)_{i\in N}$ 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 $\Theta$ together with the flat prior, the action set $A$, the utility function $u$, the private belief distributions $\left(G_{(-1)}, G_1\right)$ and the cumulative distribution function $W$ are commonly known among the players.

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

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

*(i) Beliefs are updated according to (III.1) where* $\lambda(s_i,\emptyset) = \frac{b(s_i,\emptyset)}{1-b(s_i,\emptyset)}$ *and*
$\lambda(\emptyset,h_i) = \prod_{j<i} \frac{Pr\left(\tilde{a}_j=h_i(j)|\tilde{h}_j=\left(h_i(1)\dots,h_i(j-1)\right),\tilde{\theta}=1\right)}{Pr\left(\tilde{a}_j=h_i(j)|\tilde{h}_j=\left(h_i(1)\dots,h_i(j-1)\right),\tilde{\theta}=(-1)\right)}$,

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

### 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 $\beta_i$. Using the updating rule (III.1) and the posterior LR threshold, the following lemma describes the player's strategy in terms of her realized private belief.

**Lemma III.1.** *For player* $i \in \mathbb{N}$ *let* $\lambda_i = \lambda(\emptyset,h_i)$ *be the realized public likelihood ratio and let* $\beta_i$ *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* $\sigma_i$ *is given by*

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

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

*(i)* $t(\lambda,1) = 1/(1+\lambda)$, $\lim_{\beta\to\infty} t(\lambda,\beta) = \frac{1}{2}$, $\lim_{\beta\searrow 0} t(\lambda,\beta) = \begin{cases} 1 & \text{if } \lambda < 1 \\ 0 & \text{if } \lambda > 1 \end{cases}$;

*(ii)* $\lim_{\lambda \to 0} t(\lambda, \beta) = 1$, $\lim_{\lambda \to \infty} t(\lambda, \beta) = 0$, $t(1, \beta) = \frac{1}{2}$;

*(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 III.4 shows the graph of the private belief thresholds as a function of the realized public belief for different values of the decision weight $\beta_i$. 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 $\beta_i$, the more pronounced this behavior becomes. As $\beta_i = 0$, players rely only on the public information except when it is indecisive $\left(b(\emptyset, h_i) = \frac{1}{2}\right)$.

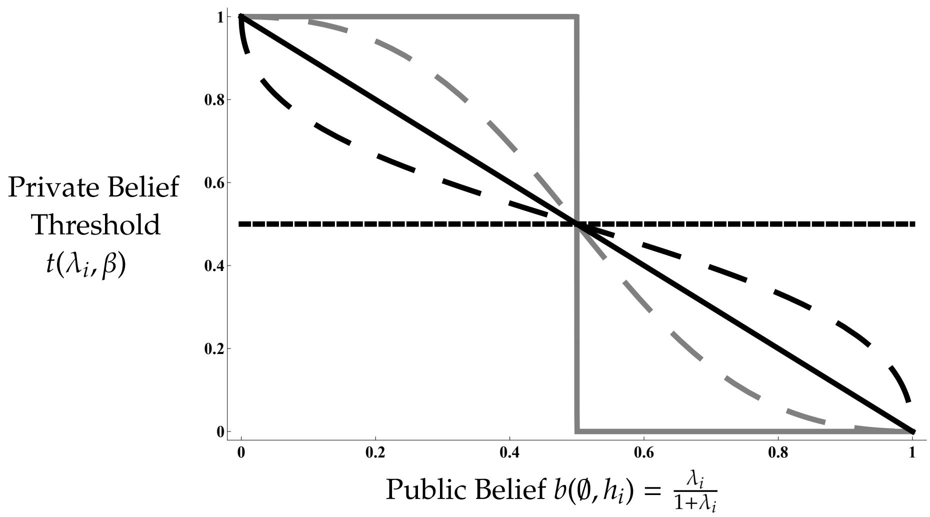


Figure III.4.: Private belief thresholds as a function of the realized public belief $b(\emptyset, h_i)$ for weights $\beta_i = 0$ (solid grey line), $\beta_i = \frac{1}{2}$ (dashed grey line), $\beta_i = 1$ (solid black line), $\beta_i = 2$ (dashed black line) and $\beta_i \to \infty$ (dotted black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

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

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

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

$\beta_i$ is private information, players have to form an average of these probabilities over the possible types using the cumulative distribution function $W$. Formally

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

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

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

### An Equivalence Result

Let us go back to the social learning environment considered in the 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 III.2.** *Let the private belief distributions be given by $G_{(-1)}$ and $G_1$ with support $\left[\underline{b}, \overline{b}\right] \subseteq [0,1]$. The equilibrium of $\left\langle \mathbb{N}, \Theta, A, u, \left(G_{(-1)}, G_1\right), W\right\rangle$ is identical to the equilibrium of $\left\langle \mathbb{N}, \Theta, A, u, \left(G'_{(-1)}, G'_1\right), W'\right\rangle$ 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 III.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 (III.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 III.3.** *Let $[\underline{\beta}, \overline{\beta}] \subseteq [0, \infty)$ 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 $\left[\underline{b}', \overline{b}'\right]$ where*

$$\underline{b}' = \frac{\underline{b}^{\overline{\beta}}}{\underline{b}^{\overline{\beta}} + (1-\underline{b})^{\overline{\beta}}} \qquad \textit{and} \qquad \overline{b}' = \frac{\overline{b}^{\overline{\beta}}}{\overline{b}^{\overline{\beta}} + \left(1-\overline{b}\right)^{\overline{\beta}}}.$$

*Furthermore it holds $G_1'(b) < G_{(-1)}'(b)$ for each $b \in \left(\underline{b}', \overline{b}'\right)$.*

*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 $\overline{\beta} < \infty$. 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.** *In the social learning game $\left\langle \mathbb{N}, \Theta, A, u, \left(G_{(-1)}', G_1'\right), W' \right\rangle$, $G_{(-1)}'$ and $G_1'$ are unbounded if and only if $W$ is unbounded in the social learning game with rich updating rules $\left\langle \mathbb{N}, \Theta, A, u, \left(G_{(-1)}, G_1\right), W \right\rangle$.*

### 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 III.2.

### Learning

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

**Lemma III.4.** *With bounded private information weights, the larger $\overline{\beta}$, the more extreme (i.e. the farther from $\frac{1}{2}$) beliefs are in the limit.*

*Proof.* From the equivalence result (Lemma III.2), the larger $\overline{\beta}$, the larger the support of the modified private beliefs. The boundaries of the cascade regions satisfy $\underline{\lambda} = \frac{1-\overline{b}}{\overline{b}} < 1$ and $\overline{\lambda} = \frac{1-\underline{b}}{\underline{b}} > 1$. As $\overline{\beta}$ 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 $\overline{\beta}$ the farther away this limit. As mentioned in the previous section, information cascades still emerge in the 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 $\beta$, 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 $\beta$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.** *Uniform behavior eventually arises in finite time.*

**Lemma III.5.** *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 $\overline{\beta} \to \infty$.*

*Proof.* Given Lemma III.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 $\overline{b'}$ which in turn increases in $\overline{\beta}$. □

### 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 $\beta$, 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 $\beta_i$ choose accordingly. We have shown above that as $\overline{\beta} \to \infty$, 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 III.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 $\left(\check{\beta}_i\right)_{i \in \mathbb{N}}$ of mappings $\check{\beta}_i : [0,1]^2 \to [0,\infty)$ 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.

### III.4.2. 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 $\bar{\sigma}(a \mid \theta)$ for each $(\theta,a) \in \Theta \times A$. Hence, the information value $\bar{\sigma}(a \mid 1)/\bar{\sigma}(a \mid (-1))$ 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 III.6.** *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 $\pi_\theta$ 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 $\beta_i$ 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 III.7.** $\pi_1 < 1$ *and* $\pi_{(-1)} > 0$ *if*

*(i) Private information weights are bounded* ($\overline{\beta} < \infty$);

*(ii) Private information weights are unbounded* ($\overline{\beta} = \infty$) *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. □

## III.5. Related Work

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).[5]

Eyster and Rabin (2010) considers a social learning environment where individuals choose actions from a continuum and receive arbitrarily informative signals. Rational social learning

[5] Bernardo 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.

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 equilibrium 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.

## III.6. Conclusion

In this chapter, we revisit 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.

## Appendix III.A 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 $\left[0,(1-q)^3/q^3\right)$ and $\left(q^3/(1-q)^3,\infty\right)$. 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 overweighters 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

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

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

□

***Proof 2.1.2.*** Consider the equation

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

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 $\pi$ 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 $\pi \ge \frac{q^3\,(5+q)\left(q^3-(1-q)^3\right)}{q^6\,(5+q)-(1-q)^6\,(6-q)}$. The RHS of the latter exceeds $\frac{q\,(1+q)}{2\,(1-q+q^2)}$, the probability in the specific model of Bikhchandani, Hirshleifer, and Welch (1992), 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 $\lambda_{IC}$. In the specific model of Bikhchandani, Hirshleifer, and Welch (1992), $\lambda_{IC}=\frac{q\,(1+q)}{2\,(1-q+q^2)}$ 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 $\beta_i$, $\beta_i\in\{0,1,2,3\}$, 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 the specific model of Bikhchandani, Hirshleifer, and Welch (1992), $\beta_i = 1$ 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 the analysis of Bikhchandani, Hirshleifer, and Welch, 1992). For weak overweighters we obtain less fragility compared to the standard model since $z'_{min}(2) > 2$. For strong overweighters 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 III.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

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

from which the properties follow immediately. □

***Proof of Lemma III.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 standard equilibrium each player's decision is characterized by two assumptions. First, each player updates his belief about the state of Nature $\tilde{\theta}$ 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 $G = \left\langle \mathbb{N}, \Theta, A, u, \left(G_{(-1)}, G_1\right), B\right\rangle$ 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 $h : [0,1] \times [0,\infty) \to [0,\infty)$ be given by $h(b,\beta) = \frac{b^\beta}{b^\beta+(1-b)^\beta}$. Define further $\tilde{\mu}_i := h\left(\tilde{b}_i, \tilde{\beta}\right)$ 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 $\theta$ being the realization of $\tilde{\theta}$. We have

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

We check first whether $G'_\theta$ 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 III.1, $G'_\theta(0) = \int_0^\infty G_\theta(0)\, dW(\beta) = 0$. Similarly $G'_\theta(1) = 1$. Furthermore by Lemma III.1 (iv) $t(\lambda, \beta)$ is decreasing in $\lambda$. As $x/(1-x)$ is increasing in $x$ and $G_\theta$ is weakly increasing, the integrand is weakly increasing in $x$ and hence $G'_\theta$ is. Finally, continuity from the right follows straightforwardly. Thus $G'_\theta$ is a proper cdf. 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 standard equilibrium update beliefs the same way as players in the heterogeneous decision weight equilibrium. As further the second requirement of sequential rationality is satisfied in either equilibrium this proves the Lemma. □

***Proof of Lemma III.3.*** We first characterize the support of the new private belief distribution. Let $\left[\underline{\beta}, \overline{\beta}\right]$ be the support of $W$, the distribution of decision weights. First we have $G'_\theta(b) = 0$ if for each $\beta \in \left[\underline{\beta}, \overline{\beta}\right]$, $\frac{b^{\frac{1}{\beta}}}{b^{\frac{1}{\beta}}+(1-b)^{\frac{1}{\beta}}} < \underline{b}$. This holds exactly when $b \leq \frac{\underline{b}^\beta}{\underline{b}^\beta+\left(1-\underline{b}\right)^\beta}$ for each $\beta \in \left[\underline{\beta}, \overline{\beta}\right]$. Given that $\underline{b} < \frac{1}{2}$ the RHS is strictly decreasing in $\beta$ and thus attains its minimum at $\overline{\beta}$. Hence, the condition holds for all $\underline{\beta} \leq \beta \leq \overline{\beta}$ provided it holds for $\overline{\beta}$. Hence $G'_\theta(b) = 0$ if and only if $b \leq \frac{\underline{b}^{\overline{\beta}}}{\underline{b}^{\overline{\beta}}+\left(1-\underline{b}\right)^{\overline{\beta}}}$.

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

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

□

***Proof of Lemma III.6.*** We fix $T \in \mathbb{N}$ and for $\theta \in \{(-1), 1\}$ let $\pi_\theta = \bar{\sigma}(1|\theta)$ denote investment frequencies under state $\theta$. The proof follows along the same lines as the one by Guarino and Jehiel (2009) (Proposition 7). Define the function $\Phi : [0,1]^2 \to [0,1]^2$ that for a given vector $\left(\pi_{(-1)}, \pi_1\right)$ where $\pi_\theta = \bar{\sigma}(1 \mid \theta)$ 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 $\beta_i$ distributed according to $W$. We show that for continuous distribution $W$, $\Phi$ is continuous and furthermore $\Phi$ maps the upper triangle of the unit interval $\Delta = \{(x,y) : 0 \le x \le 1 \wedge x \le y \le 1\}$ into itself. Then by Brouwer's fixed point theorem there exists $\left(\pi^*_{(-1)}, \pi^*_1\right) = \Phi\left(\pi^*_{(-1)}, \pi^*_1\right)$.

We denote by $\Phi_\theta(\pi_{(-1)}, \pi_1)$ the investment frequency $\Phi$ determines for state $\theta$, 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

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

and we have that

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

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

First, continuity of $G'_\theta$ implies continuity of $Pr\left(\tilde{a}_i = 1 \mid \tilde{h}_i = h_i, \tilde{\theta} = \theta\right)$ and also of $Pr\left(\tilde{h}_i = h_i \mid \tilde{\theta} = \theta\right)$[6] with respect to the vector $(\pi_{(-1)}, \pi_1)$. Hence, $\Phi_\theta$ is continuous if $G'_\theta$ is. To show the latter notice first that trivially if $G_\theta$ is continuous, so is $G'_\theta$. On the other hand $G_\theta$ can only possess countably many jumps. For each $x \in (0,1)$ there thus exist countably many $\beta > 0$ such that $G_\theta$ 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 $G'_\theta$ and thus $\Phi_\theta$ is continuous for each $\theta \in \{(-1), 1\}$.

Second, we show that $\Phi : \Delta \to \Delta$. As $0 \le \Phi_\theta(\pi_{(-1)}, \pi_1) \le 1$ by definition the only thing we need to show is that $\pi_1 > \pi_{(-1)}$ implies $\Phi_1(\pi_{(-1)}, \pi_1) > \Phi_0(\pi_{(-1)}, \pi_1)$. To proof the latter showing $Pr\left(a_i = 1 \mid \tilde{\theta} = 1\right) > Pr\left(a_i = 1 \mid \tilde{\theta} = 0\right)$ for any $i = 1, 2, \ldots, T$ is sufficient. For each $i = 1, \ldots, T$ we turn to the distribution of the differences $\Delta\Sigma\left(\tilde{h}_i\right) = \Sigma_1\left(\tilde{h}_i\right) - \Sigma_0\left(\tilde{h}_i\right)$ conditional on the state of the world. We first collect the following properties:

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

(i) Investment probabilities are constant across histories

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

and given by

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

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

Given these properties it suffices to show that the distribution of differences $\Delta\Sigma\left(\tilde{h}_i\right)$ 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. $K_{i,\theta}(x) = Pr\left(\Delta\Sigma(\tilde{h}_i) < x \mid \tilde{\theta} = \theta\right)$. We show that $K_{i,(-1)}(x) > K_{i,1}(x)$ for any $x \in \{-(i-1), i-3\}$ and any $i = 1, 2, \ldots, T$ 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 $z \in \{-(i-1), -(i-1)+2, \ldots, i-3\}$ it holds

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

Then

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

and rearranging terms yields

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

where

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

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 III.7.*** 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 players 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

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

First assume (a) $\mu_\beta = E\left[\tilde{\beta}_i\right] < \infty$ and $\sigma_\beta^2 = Var\left[\tilde{\beta}_i\right] < \infty$. 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 $\log(x) < x$

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

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 $\tilde{\theta} = 0$

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

As $\sum_{i=1}^{\infty} \frac{1}{(ai+b)^2} < \infty$ provided $b/a > -1$,[7] $\mu_\beta$ and $\sigma_\beta^2$ are finite and as $\frac{\pi_1}{\pi_{(-1)}} = \frac{1-\epsilon}{\epsilon} \to \infty$ as $\epsilon \to 0$ the argument of the exponential function converges to zero as $T \to \infty$ and $\epsilon \to 0$. Hence, the

[7] The 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 $x > 0$.

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

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

Therefore by the same calculations as above

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

Now, as $T \to \infty$ 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 $T \to \infty$ and $\epsilon \to 0$. With strictly positive probability, all players invest in state 0.

□

# Conclusion

Economic models of social learning have contributed to the understanding of market frenzies and crashes by showing that idiosyncratic and fragile herding can be the result of imperfectly informed rational investors who try to learn from observing others' investment decisions. In this thesis we elaborate our concern that the Bayesian rational view employed in these studies in its current form limits the capacity of the models to explain social epidemics. In particular we express a dissatisfaction with the eductive justification of equilibrium in social learning, more precisely its assumption that fundamentals of the social learning setting are not only shared knowledge but common knowledge. Concretely in a game of incomplete information the latter requires the prior probability distribution of fundamentals to be commonly known (Harsanyi, 1967). As soon as the setting is not a "small world"(Savage, 1954) and gets complex, that is in many economically relevant situations in which social learning takes place, this assumption seems highly disputable. In addition we argue that some of the conclusions reached by the rational view of social learning are unsound or of minor relevance from a welfare point of view. We therefore suggest a new perspective on economic models of social learning which focusses on the positive aspect of these models with the aim to provide a better understanding of real-world social learning.

Chapter 1 of the thesis provides an overview of the economic literature on social learning with a particular emphasis on the relation between the properties of the equilibrium outcome and the primitives of the environment. In situations where the investment space is discrete the equilibrium outcome is always inefficient and a herd almost surely arises. Dissatisfied with the emphasis of the literature on the long-run properties of the equilibrium outcome the chapter offers a new perspective on social learning which relies on new medium-run measures of uniformity and fragility. Comparing, in line with previous studies, social learning settings which differ only in the distribution of private information but focussing on observables to assess social welfare we find that less uniform investment decisions lead to higher welfare levels of the equilibrium outcome but lower fragility of the herds.

Chapter 2 questions more thoroughly the Bayes-rational view of social learning by investigating the learning foundations of economic models of social learning. The learning perspective offers us the possibility to indirectly assess the whole set of assumptions underlying these models in particular regarding mutual knowledge of the distribution of private information and to clarify when and why equilibrium is likely to capture observed regularities in the field. Our results suggest that too much structural knowledge has been assumed in standard economic models of social learning. While in the absence of fundamental (not resolved by the Harsanyi doctrine) structural uncertainty and with an infinite learning horizon epistemic learning leads to Bayesian rational play, the same learning process favors non-Bayesian play whenever players do not know the distribution of private beliefs. We therefore suggest that further economic models of social learning should allow for the presence of fundamental structural uncertainty.

Chapter 3 offers one simple possibility to incorporate fundamental structural uncertainty into a model of social learning. We extend the classic model of social learning by assuming that individuals update their beliefs in a non-Bayesian way. Individuals either underweigh or overweigh their public information relative to the public information revealed by the decisions of others and each individual's updating rule is private information. In a setting with perfectly rational individuals and commonly known distribution of updating rules we show that introducing heterogeneous updating rules reconciles equilibrium predictions with evidence from the laboratory and the field. In particular individuals decisions' are less uniform which is why in a herd individuals tend to more quickly become extremely confident about the appropriateness of the chosen action and herds exhibit less fragility than in the standard model of social learning. We also consider a variant of the setting 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.

The thesis constitutes a first attempt to investigate learning foundations of rational play in social learning in order to assess the nature, scope and behavioral consequences of structural uncertainty in these models. Our results indicate that this kind of analysis is likely to enhance our understanding of both experimental and real-world social learning. The implications of our results for future research on social learning are twofold: On the one hand theoretical studies should thoroughly investigate modeling assumptions in comparison to the aspects of the field environments they intend to model. In particular accommodating the fact that field environments offer limited learning opportunities to players promises to generate models whose predictions are more accurate. Furthermore in more elaborate settings with continuous actions, flexible prices, or endogeneous sequencing learning opportunities and assumptions on the structural knowledge of individuals are likely to be of even greater importance. On the other hand new laboratory experiments should be designed which test the rational view of social learning in more familiar contexts. To the best of our knowledge all existing experimental studies have considered laboratory settings which correlate strongly with the standard economic models of social learning. While such simple settings permit full control of information flows they might be perceived as artificial by subjects if they differ substantially from the field environments in which social learning mainly takes place. Accordingly deviation of observed behavior from rational play might not come as a surprise and might not constitute conclusive evidence against rational social learning.

Viewed from a broader perspective, our results also offer new insights for behavioral economists.[8] Despite a regular exchange between experimentalists and theorists over the past two decades, there

[8] In recent years the field of economics has witnessed an increased emergence of behavioral models, many of them designed to explain experimental phenomena not captured by standard notions of equilibrium or rationality (see e.g. *Advances in Behavioral Economics* Camerer, Loewenstein, and Rabin, 2004). While these models usually better accommodate the evidence, a common critique concerns their lack of foundations (see e.g. Fudenberg, 2006).

is no satisfactory behavioral model of social learning. Previous attempts have imported psychological insights (e.g. judgmental biases, limited depth of reasoning) into existing economic models of social learning. These behavioral models acknowledge the cognitive limitations of economic actors by relaxing the assumption of Bayesian rationality in the direction of greater psychological realism. Though we are sympathetic to this approach, we show that a thorough investigation of the modeling assumptions may straightforwardly yield an alternative model of identical complexity but with increased explanatory power. We believe that our approach is likely to be fruitful not only in social learning. Economic models which accommodate the fact that field environments provide limited learning opportunities to players are likely to generate more accurate predictions without diminished tractability and in this sense they complement other models which incorporate more realistic psychological foundations.

# List of Figures

# Bibliography


ACEMOGLU, D., D. MUNTHER, I. LOBEL, AND A. OZDAGLAR (2010): "Bayesian Learning in Social Networks," mimeo, MIT.

AL-NAJJAR, N. I. (2009): "Decision Makers as Statisticians: Diversity, Ambiguity, and Learning," *Econometrica*, 77, 1339–1369.

AL-NAJJAR, N. I., AND J. WEINSTEIN (2009): "The Ambiguity Aversion Literature: A Critical Assessment," *Economics and Philosophy*, 25, 249–284.

ANDERSON, L., AND C. HOLT (1997): "Information cascades in the laboratory," *American Economic Review*, 87, 847–862.

ATHREYA, K., AND S. LAHIRI (2006): *Measure Theory and Probability Theory*. Springer, New York.

AUMANN, R. (1997): "Rationality and Bounded Rationality," *Games and Economic Behavior*, 21, 2–14.

AUMANN, R. J. (2008): "Rule-Rationality versus Act-Rationality," mimeo, The Hebrew University of Jerusalem.

BANERJEE, A. V. (1992): "A simple model of herd behavior," *The Quarterly Journal of Economics*, 107, 797–817.

BATTIGALLI, P. (1997): "Comportamento Razionale ed Eqiulibrio nei Giochi e nelli Situatzioni Sociali," Ph.D. thesis, Bocconi University.

BEGGS, A. (1993): "Learning in Bayesian Games with Binary Actions," *The B.E. Journal of Theoretical Economics*, 9: Iss. 1 (Advances), Article 33.

BERNARDO, A., AND I. WELCH (2001): "On the Evolution of Overconfidence and Entrepeneurs," *Journal of Economics & Management Strategy*, 10, 301–330.

BERNHEIM, B. D. (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007–1028.

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.

——— (1998): "Learning from the behavior of others: Conformity, fads, and informational cascades," *Journal of Economic Perspectives*, 12, 151–170.

BROWN, G. W. (1951): "Iterative Solution of Games by Fictitious Play," in *Activity Analysis of Production and Allocation*, ed. by T. Koopmans, pp. 374–376. Wiley, New York.

CAMERER, C. F., AND T.-H. HO (1999): “Experience-weighted attraction learning in normal-form games,” *Econometrica*, 67, 827–874.

CAMERER, C. F., G. LOEWENSTEIN, AND M. RABIN (eds.) (2004): *Advances in Behavioral Economics*. Russell Sage Foundation and Princeton University Press.

ÇELEN, B̆., AND S. KARIV (2004): “Observational Learning under Imperfect Observation,” *Games and Economic Behavior*, 47, 72–86.

CHAMLEY, C. (2004a): “Delays and Equilibria with Large and Small Information in Social Learning,” *European Economic Review*, 48, 477–501.

——— (2004b): *Rational herds: Economic models of social learning*. Cambridge, Cambridge University Press.

CHAMLEY, C. P. (2003): “Cascades and Slow Social Learning,” Working paper, Boston University.

COCHRAN, W. G. (1934): “The distribution of quadratic forms in a normal system, with applications to the analysis of covariance,” *Math Proc of the Cambridge Philosophical Society*, 30, 178–191.

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.

CYERT, R. M., AND M. H. DEGROOT (1970): “Bayesian Analysis and Duopoly Theory,” *The Journal of Political Economy*, 78, 1168–1184.

——— (1974): “Rational Expectations and Bayesian Analysis,” *The Journal of Political Economy*, 82, 521–536.

DEKEL, E., D. FUDENBERG, AND D. K. LEVINE (2004): “Learning to Play Bayesian Games,” *Games and Economic Behavior*, 46, 282–303.

DEKEL, E., AND F. GUL (1997): “Rationality and Knowledge in Game Theory,” in *Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress*, ed. by D. M. Kreps, and K. F. Wallis, vol. 1 of *Econometric Society Monographs, Vol. 26*, pp. 87–172. Cambridge: Cambridge University Press.

DELAVANDE, A. (2008): “Measuring Revisions to Subjective Expectations,” *Journal of Risk and Uncertainty*, 36, 43–82.

DIACONIS, P., AND D. FREEDMAN (1986): “On the Consistency of Bayes Estimates,” *Annals of Statistics*, 11, 1–26.

EL-GAMAL, M., AND D. GRETHER (1995): “Are People Bayesian? Uncovering Behavioral Strategies,” *Journal of the American Statistical Association*, 90, 1137–1145.

ESPONDA, I. (2008): “Information Feedback in First Price Auctions,” *RAND Journal of Economics*, 39, 491–508.

Eyster, E., and M. Rabin (2010): "Naive Herding in Rich-Information Settings," *American Economic Journal: Microeconomics, forthcoming*.

Fudenberg, D. (2006): "Advancing Beyond "Advances in Behavioral Economics"," *Journal of Economic Literature*, 44, 694–711.

Fudenberg, D., and D. M. Kreps (1993): "Learning Mixed Equilibria," *Games and Economic Behavior*, 5, 320–367.

——— (1995): "Learning in Extensive-Form Games, I: Self-Confirming Equilibria," *Games and Economic Behavior*, 8, 20–55.

Fudenberg, D., and D. K. Levine (1993a): "Self-Confirming Equilibrium," *Econometrica*, 61, 523–545.

——— (1993b): "Steady State Learning and Nash Equilibrium," *Econometrica*, 61, 547–573.

——— (1998): *The Theory of Learning in Games*. Cambridge, MA: MIT Press.

——— (2009): "Learning and Equilibrium," *Annual Review of Economics*, 1, 385–419.

Gale, D. (1996): "What have we learned from social learning?," *European Economic Review*, 40, 617–628.

Gilboa, I., and D. Schmeidler (1995): "Case-Based Decision Theory," *The Quarterly Journal of Economics*, 110, 605–639.

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

Goeree, J. K., T. R. Palfrey, and B. W. Rogers (2006): "Social Learning with Private and Common Values," *Economic Theory*, 28, 245–264.

Guarino, A., and P. Jehiel (2009): "Social Learning with Coarse Inference," Working paper.

Gul, F., and R. Lundholm (1995): "Endogeneous Timing and the Clustering of Agents' Decisions," *The Journal of Political Economy*, 103, 1039–1066.

Hansen, L. P. (2007): "Beliefs, Doubts and Learning: Valueing Macroeconomic Risk," *The American Economic Review*, 97, 1–30.

Harsanyi, J. C. (1967): "Games with Incomplete Information played by Bayesian Players," *Management Science*, 14, 159–182, 320–334, 486–502.

Hart, S. (2005): "Adaptive Heuristics," *Econometrica*, 73, 1401–1430.

Hendricks, K., and D. Kovenock (1989): "Asymmetric Information, Information Externalities, and Efficiency: The Case of Oil Exploration," *The RAND Journal of Economics*, 20, 164–182.

Hendricks, K., R. H. Porter, and B. Boudreau (1987): "Information and Returns in OCS Auctions: 1954–69," *Journal of Industrial Economics*, 35, 517–542.

Hopkins, E. (2002): "Two Competing Models of How People Learn in Games," *Econometrica*, 70, 2141–2166.

Jackson, M. O., E. Kalai, and R. Smorodinsky (1999): "Bayesian Representation of Stochastic Processes under Learning: De Finetti Revisited," *Econometrica*, 67, 875–893.

Jehiel, P. (1995): "Limited Horizon Forecast in Repeated Alternate Games," *Journal of Economic Theory*, 67, 497–519.

——— (2005): "Analogy-Based Expectation Equilibrium," *Journal of Economic Theory*, 123, 81–104.

Kalai, E., and E. Lehrer (1993): "Rational Learning Leads to Nash Equilibrium," *Econometrica*, 61, 1019–1046.

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.

Lee, I. H. (1993): "On the Convergence of Informational Cascades," *Journal of Economic Theory*, 61, 395–411.

Lohmann, S. (1994): "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989–91," *World Politics*, 47, 42–101.

Machina, M. J., and D. Schmeidler (1992): "A More Robust Definition of Subjective Probability," *Econometrica*, 60, 745–780.

McKelvey, R. D., and T. R. Palfrey (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6–38.

——— (1998): "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1(1), 9–41.

Milgrom, P. R. (1981): "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12, 380–391.

Moscarini, G., M. Ottaviani, and L. Smith (1998): "Social Learning in a Changing World," *Economic Theory*, 11, 657–665.

Myerson, R. (1991): *Game Theory: Analysis of Conflict*. Harvard University Press.

Odean, T. (1998): "Volume, Volatility, Price, and Profit When All Traders Are Above Average," *Journal of Finance*, 53, 1887–1934.

Pearce, D. G. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029–1050.

Radner, R. (1962): "Team Decision Problems," *Annals of Mathematical Statistics*, 33, 857–881.

REINHART, C. M., AND K. S. ROGOFF (2009): *This Time Is Different: Eight Centuries of Financial Folly*. Princeton, NJ: Princeton University Press.

ROTH, A. (1996): "Adaptive Behaviour and Strategic Rationality: Evidence from the Laboratory and the Field," in *The Rational Foundations of Economic Behaviour: Proceedings of the IEA Conference held in Turin, Italy*, ed. by K. J. Arrow, E. Colombatto, M. Perlman, and C. Schmidt, pp. 255–273. London: Macmillan.

RUBINSTEIN, A. (1986): "Finite Automata Play the Repeated Prisoner's Dilemma," *Journal of Economic Theory*, 39, 83–96.

SAMUELSON, L. (2001): "Analogies, Adaptation, and Anomalies," *Journal of Economic Theory*, 97, 320–366.

——— (2004): "Information-Based Relative Consumption Effects," *Econometrica*, 72, 93–118.

——— (2006): "Information, Evolution and Utility," *Theoretical Economics*, 1, 119–142.

SAVAGE, L. (1954): *The Foundations of Statistics*. Wiley, New York.

SGROI, D. (2002): "Optimizing Information in the Herd: Guinea Pigs, Profits, and Welfare," *Games and Economic Behavior*, 39(1), 137–166.

SHAO, J. (2003): *Mathematical Statistics*. Springer, New York, 2nd edn.

SMITH, L., AND P. SØRENSEN (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68(2), 371–398.

——— (2008a): "Informational Herding and Optimal Experimentation," Working paper.

——— (2008b): "Rational Social Learning with Random Sampling," 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.

VIVES, X. (1993): "How fast do rational agents learn?," *Review of Economic Studies*, 60, 329–347.

——— (1996): "Social Learning and Rational Expectations," *European Economic Review*, 40, 589–601.

——— (2008): *Information and Learning in Markets: The Impact of Market Microstructure*. Princeton University Press.

WALLISER, B. (1998): "A Spectrum of Equilibration Processes in Game Theory," *Journal of Evolutionary Economics*, 8, 67–87.

WEIZSÄCKER, G. (2010): "Do we follow others when we should? A simple test of rational expectations," *forthcoming in American Economic Review*.

Young, H. P. (1993): "The Evolution of Conventions," *Econometrica*, 61, 57–84.

Ziegelmeyer, A., F. Koessler, J. Bracht, and E. Winter (2010): "Fragility of information cascades: An experimental study using elicited beliefs," *Experimental Economics*, 13, 121–145.