

March, Christoph

Working Paper

Adaptive Social Learning

CESifo Working Paper, No. 5783

Provided in Cooperation with:

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

Suggested Citation: March, Christoph (2016) : Adaptive Social Learning, CESifo Working Paper, No. 5783, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/130412>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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



Working Papers

www.cesifo.org/wp

Adaptive Social Learning

Christoph March

CESIFO WORKING PAPER NO. 5783

CATEGORY 13: BEHAVIOURAL ECONOMICS

FEBRUARY 2016

An electronic version of the paper may be downloaded

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

ISSN 2364-1428

CESifo

Center for Economic Studies & Ifo Institute

Adaptive Social Learning

Abstract

The paper investigates social-learning when the information structure is not commonly known. Individuals repeatedly interact in social-learning settings with distinct information structures. In each round of interaction, they use their experience gained in past rounds to draw inferences from their predecessors' current decisions. Such adaptation yields rational behavior in the long-run if and only if individuals distinguish social-learning settings and receive rich feedback after each round. Limited feedback may lead individuals to imitate uninformed predecessors. Moreover, adaptation across social-learning settings renders Bayes' rule payoff-inferior compared to non-Bayesian belief updating rules and suggests that belief-updating rules are heterogeneous in the population.

JEL-Codes: C730, D820, D830.

Keywords: informational herding, adaptation, analogy-based expectations equilibrium, Non-Bayesian updating.

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

This version: February 1, 2016

The paper is based on chapter II of my PhD thesis. I thank in particular Anthony Ziegelmeyer, Christophe Chamley, and Philippe Jehiel for advice, suggestions and support. I also thank Thomas Daske, Christian Feilcke, Florian Inderst, Frédéric Koessler, Dorothea Kübler, Marco Ottaviani, Marco Sahm, Georg Weizsäcker, and seminar audiences at Berlin, Boston, Jena, Munich, Paris and Stony Brook for helpful comments. Finally, I gratefully acknowledge financial support from the ERC (advanced grant GTAPCL) and hospitality of Paris School of Economics during my time as a postdoctoral researcher.

1 Introduction

In many economic settings with observable actions, individuals with limited information about a payoff-relevant state of nature benefit from learning from the actions of other individuals. Such *social-learning* has been identified among others in financial and microloan markets and with respect to consumption of experience goods like movies or restaurant meals (Cipriani and Guarino, 2014; Zhang and Liu, 2012; Moretti, 2011; Cai et al., 2009).

Beginning with Bikhchandani et al. (1992) and Banerjee (1992), an extensive research program has investigated *rational social-learning*. This literature assumes that individuals are rational, form beliefs using Bayes' rule, and possess common strategic and structural knowledge meaning that Bayesian rationality and the structure of the social-learning setting are commonly known. If the set of actions is discrete, rational social-learning leads individuals to eventually *herd* on an action. As a consequence, information aggregation fails spectacularly and the economic outcome is inefficient.¹

Despite the many important insights it delivers, rational social-learning also reaches some unsound conclusions. For instance, even in a long herd, individuals never (or only very slowly) become extremely confident in the correctness of the chosen action. Moreover, if an individual with very precise information goes against the herd, her decision overturns it since her successors correctly infer the precise information (Smith and Sørensen, 2000). Many of these conclusions hinge upon the models' strong underlying assumptions.

Though (commonly known) rationality can be criticized on many grounds, the present paper mainly investigates the implications of assuming that individuals possess structural knowledge. In particular, I question the assumption that the information structure, i.e. the distribution of private information, is commonly known. Common knowledge of the information structure is innocuous only if a “physical” ex ante procedure generates private information (Dekel and Gul, 1997). For example, if the state of nature is the amount of oil in some tract and firms' private information results from taking soil samples, 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 (Hendricks and Kovenock, 1989). On the other hand, fueled by the rapid diffusion of Internet technologies, there are many important economic settings where interaction is anonymous and information sources are wide and dispersed. For example, lenders in online peer-to-peer lending markets successfully use “soft” information – such as written statements of borrowers about the

¹Social-learning is efficient only if individuals choose from a continuum of actions and are rewarded according to the proximity of their action to the most profitable one (Lee, 1993). If the set of actions is discrete and private signals are of bounded strength, rational social-learning quickly stops, and individuals herd on a wrong action with positive probability (Banerjee, 1992; Bikhchandani et al., 1992). With a discrete set of actions and unbounded private signals, the correct action is chosen asymptotically, but information aggregation is often extremely slow and long herds still emerge (Smith and Sørensen, 2000). The results carry over to social-learning settings with general observation structures (Smith and Sørensen, 2008; Acemoglu et al., 2011) and settings with an informationally efficient price process and a sufficiently rich information structure (Avery and Zemsky, 1998; Park and Sabourian, 2011).

reasons for the loan application – to predict the creditworthiness of borrowers (Iyer et al., 2015; Lin et al., 2013).² Furthermore, some decision settings are characterized by an abundance of diverging information (Heal and Millner, 2014).

In the absence of common knowledge of the information structure, individuals are unable to *deduce* the informational content of the actions they observe. My main assumption is that repeated interaction enables individuals to acquire an understanding of the information conveyed by observed actions. Specifically, individuals repeatedly interact in social-learning settings and they get *feedback* about the state of nature and the actions of others after each repetition. They use this feedback to assess conditional probabilities of a given history of actions, conditional on each possible state of nature. For a given history and state, the conditional probability is assessed by the relative frequency with which this history occurred across past repetitions in which the given state was realized. I further assume that relative frequencies for each state are combined according to Bayes' rule with private information, and that individuals pick the action which maximizes their expected payoff in the current repetition. Accordingly, individuals are myopic because they ignore repeated-game considerations.

The paper analyzes the long-run outcome of the above defined *adaptive process*. Yet, in real-world environments individuals are unlikely to encounter *exactly* the same strategic setting many times. I therefore assume that repeated interactions take place in several different social-learning settings. I focus on settings which differ only with respect to the information structure.

The first result of the paper establishes sufficient conditions under which long-run behavior mimics rational social learning: First, individuals are able to distinguish settings and second, the state of nature is revealed after each repetition (Proposition 1). The subsequent analysis proves that these assumptions are also necessary. First, I analyze how limited feedback regarding the state of nature affects adaptation. In particular, the state may not be revealed unless a certain action is taken (e.g. a good is bought or an investment is realized). As a consequence, individuals are likely to excessively imitate actions whose payoffs cannot reveal the state (Proposition 2). Such imitation may be based on no information at all. This explains why people are susceptible to following false prophets, joining cults, or relying on anecdotal reasoning. Second, I study the long-run outcome of *adaptation across settings*. Individuals adapt across settings if they use the feedback from past repetitions regardless of the respective social-learning setting. This may stem from an inability to distinguish settings or a desire to rely on a larger amount of feedback. I show (in Proposition 3) that the long-run outcome of adaptation across settings is captured by an *analogy-based expectations equilibrium* (Jehiel, 2005). In equilibrium, individuals bundle the decision situations of others into analogy classes and have a correct understanding of average behavior in each class. In the present setting, bundling leads

²See Oberlechner and Hocking (2004), Kulkarni et al. (2012) for further examples.

to systematically biased inferences from observed actions. Therefore, long-run behavior does not maximize individuals' (ex ante) expected payoffs (Proposition 4), and herding may spill-over from one social-learning setting to another. Moreover, adaptation across settings renders Bayes' rule payoff inferior compared to non-Bayesian belief updating rules: Individuals with the most precise private information benefit from overweighting private information relative to others' actions. Conversely, individuals with the least precise private information benefit from underweighting private information (Proposition 5).

The results of the paper provide several new insights: First, the results clarify when and why the assumptions underlying rational social-learning are justified. Second, the results suggest that belief updating rules are likely to be heterogeneous in the population when individuals are unable to distinguish social-learning settings. This provides an explanation for corresponding empirical findings (see e.g. Palfrey and Wang, 2012). Third, the results straightforwardly lead to a behavioral model of social-learning with heterogeneous belief updating rules which is able to accommodate the experimental regularities on social-learning. Numerous laboratory studies have established that herds emerge later than predicted by rational social-learning and that the length and strength of herds are positively correlated (e.g. Kübler and Weizsäcker, 2004, 2005; Goeree et al., 2007). In particular, individuals become extremely confident even in wrong actions. Though the experiments have stimulated an active behavioral literature (see below), none of the alternative theories organizes well the bulk of the experimental evidence. A behavioral model of social-learning with flexible belief updating rules accommodates the experimental regularities and is also able to capture more recent evidence of a 'social-confirmation'-bias among participants in a social-learning experiment with a richer information structure (March and Ziegelmeyer, 2015).

The paper relates to a growing literature on social-learning with bounded rationality. It complements the analysis of Guarino and Jehiel (2013) who assume that individuals only understand the relation between the aggregate distribution of actions and the state of nature. This correct understanding is assumed to emerge from an adaptive process with limited feedback. The two papers therefore characterize the long-run outcome of adaptation under different feedback regimes. The present paper also suggests a re-interpretation of Bohren (2015) who assumes that (i) a fraction of individuals are socially uninformed and decide based only on their private signals, and (ii) other individuals misperceive the exact proportion of uninformed predecessors. Based on the results presented here, Bohren's (2015) model may be re-interpreted as a social-learning model with (a specific form of) heterogeneous belief updating rules and a non-common prior.³ Finally, the naïve inference model of Eyster and Rabin (2010), in which individuals believe that the action of each previous individual reveals that individual's private information, can be seen as

³See also Bernardo and Welch (2001) and Kariv (2005) both of which introduce individuals who overweight their private information into the standard model of social-learning.

an alternative model of initial behavior.

A recurrent issue in the literature has been the emergence of extreme confidence in wrong beliefs (Eyster and Rabin, 2010; Guarino and Jehiel, 2013). This confidence cannot emerge in models of rational social-learning, but it has been repeatedly observed in the laboratory and in the field. A social-learning model with belief updating heterogeneity is able to predict extreme, false beliefs alongside a delayed formation of herds, which has been consistently found in the lab as well.

The paper is structured as follows. Section 2 outlines the results of the paper with the help of a simple example. Section 3 introduces the analytical framework, discusses rational social-learning, and formalizes adaptation. Section 4 characterizes the long-run outcome of adaptation under different assumptions regarding individuals' feedback. Section 5 introduces a model of social learning based on heterogeneous belief updating and discusses its relation to the adaptation results. Section 6 presents extensions of the model towards heterogeneous preferences and an endogenous timing of decisions. Section 7 concludes. The appendix contains the proofs and some additional analyses.

2 A Simple Example

Consider the following social-learning game with two players: Anna and Bob decide in sequence whether or not to invest. Anna decides first, and Bob decides after having observed Anna's decision. Payoffs from investing and, respectively, rejecting are identical for both players. The investment payoff θ takes values 0 and 1 with equal probability, the cost of an investment is $c = 1/2$, and the payoff from rejection is zero. Before deciding, each player $i \in \{A, B\}$ observes a symmetric, binary private signal $s_i \in S = \{0, 1\}$ where $\Pr(s_i = 0 \mid \theta = 0) = \Pr(s_i = 1 \mid \theta = 1) = q_i \in (0.5, 1)$ denotes the *signal precision* of player i . Signals are independent conditional on the state θ .

I will assume throughout that Anna is Bayes-rational which implies that her dominant strategy is to reject if $s_A = 0$ and to invest if $s_A = 1$.

2.1 Rational Social-Learning

In this example, the assumptions of rational social-learning entail that (i) not only Anna but also Bob is Bayes-rational, (ii) Bob knows that Anna is Bayes-rational, i.e. he possesses strategic knowledge, and (iii) both Anna and Bob know the payoff function and the signal precisions q_A and q_B , i.e. they possess structural knowledge. Bob's knowledge implies that he identifies Anna's decision to invest (reject) with the signal $s_A = 1$ ($s_A = 0$). Accordingly, Bob's *likelihood ratio*, i.e. the ratio of probabilities he assigns to the investment payoff being 1 versus 0 given signal $s_B \in \{0, 1\}$ and Anna's action

$x_A \in \{\text{invest}, \text{reject}\}$ is given by

$$\lambda(s_B, x_A) = \frac{\Pr(\tilde{\theta} = 1 \mid s_B, x_A)}{\Pr(\tilde{\theta} = 0 \mid s_B, x_A)} = \frac{\Pr(s_B \mid \theta = 1)}{\Pr(s_B \mid \theta = 0)} \cdot \begin{cases} \frac{q_A}{1-q_A} & \text{if Anna invests} \\ \frac{1-q_A}{q_A} & \text{if Anna rejects} \end{cases}.$$

Bob invests (rejects) if $\lambda(s_B, x_A) > (<) 1$. Therefore, Bob follows his private signal if it confirms Anna's decision (Anna invests and $s_B = 1$, or Anna rejects and $s_B = 0$). On the other hand, if Anna's decision contradicts his private signal, Bob follows his private signal if $q_B > q_A$ and *imitates* Anna's decision if $q_B < q_A$.⁴

2.2 Adaptive Social-Learning

Contrary to rational social-learning, adaptive social-learning assumes that individuals do not possess strategic or structural knowledge, but interact repeatedly. Since Bayes-rational Anna has the dominant strategy to follow her private signal, only Bob's behavior differs between the two approaches. Concretely, I make the following assumptions: Anna and Bob repeatedly play the social-learning game in rounds $r = 1, 2, \dots$. In each round, a new investment payoff is determined randomly and independently from the investment payoff in previous rounds. Players learn the realized investment payoff at the end of a round. This enables Bob to track the relationship between Anna's decision and the investment payoff, i.e. he keeps counts of how often Anna invests and rejects when the investment payoff is 1 and when it is 0. Finally, both players are Bayes-rational and myopically maximize their expected payoff in the current round.

The long-run outcome of the above defined *adaptive process* is easily derived: Myopic Anna follows her private signal in each round. Therefore, the relative frequency with which Anna invests when the investment payoff equals 1 (respectively 0) approaches q_A (respectively $1 - q_A$). As Bob tracks this relative frequency, he eventually infers the same information from Anna's decision as he could deduce when possessing strategic and structural knowledge. Hence, Bob eventually plays his unique rationalizable strategy.

In summary, the adaptive social-learning outcome coincides with the rational social-learning outcome.

2.3 Adaptive Social-Learning Across Games

Real-world social-learning is likely to take place in a multitude of settings and individuals are unlikely to distinguish those settings in their finest details. I therefore investigate whether adaptive social-learning also leads to rational social-learning when players must simultaneously adapt to multiple games.

I assume the following: Anna and Bob repeatedly play two different social-learning

⁴I omit the case $q_A = q_B$ for which rational social-learning makes no clear-cut prediction.

games. The games differ only with respect to the private signal precisions. In game $k \in \{H, L\}$ the signal precision of player $i \in \{A, B\}$ is $q_i^k \in (0.5, 1)$ where $q_A^L < q_A^H$ without loss of generality. In each round, the game to be played is determined randomly and independently of previous rounds; both games are equally likely. As before, Bob tracks the relation between Anna's decision and the realized investment payoff across rounds. However, Bob does not distinguish the games, i.e. he keeps single counts of how often Anna invests and rejects when the investment payoff is 1 and when it is 0, and he uses them to learn from Anna's decision in both games.

As before, Anna follows her private signal in each round and each game. Because each game occurs on average in half of the rounds, the relative frequency with which Anna invests when the investment payoff equals 1 (resp. 0) approaches $\bar{q}_A = \frac{1}{2} (q_A^L + q_A^H)$ (resp. $1 - \bar{q}_A$) *across the two games*. Since Bob does not distinguish the two games, this relative frequency eventually guides his behavior. Accordingly, Bob eventually follows his private signal in game k if $q_B^k > \bar{q}_A$, and he eventually imitates Anna's decision if $q_B^k < \bar{q}_A$.

Obviously, Bob's long-run behavior when he adapts across games is not optimal in game L (resp. H) if $q_A^L < q_B^L < \bar{q}_A$ (resp. $\bar{q}_A < q_B^H < q_A^H$). Moreover, Bob's long-run behavior may not maximize his *ex ante expected payoff* from the randomly selected social-learning game. Consider for instance two games such that signal precisions are low (high) in game L (H) since information is scarce (abundant). If Bob's signal precision is higher than Anna's in each game ($q_B^k > q_A^k$ for each k), but Anna's average signal precision is higher than Bob's signal precision in game L ($\bar{q}_A > q_B^L$), Bob suboptimally imitates Anna's decision in game L in the long-run. Equivalently, Bob suboptimally follows his private signal in game H in the long-run, if $q_B^k < q_A^k$ for each k and $q_B^H > \bar{q}_A$.

So far, I have assumed that players update beliefs according to Bayes' rule. A more general adaptive process allows players to also grope for the optimal response to beliefs. I therefore extend the adaptive process by considering flexible belief updating rules where Bob's generalized likelihood ratio given signal s_B and Anna's decision x_A is given by

$$\hat{\lambda}(s_B, x_A) = \left(\frac{\Pr(s_B \mid \theta = 1)}{\Pr(s_B \mid \theta = 0)} \right)^{\beta_B} \cdot \begin{cases} \frac{\bar{q}_A}{1 - \bar{q}_A} & \text{if Anna invests} \\ \frac{1 - \bar{q}_A}{\bar{q}_A} & \text{if Anna rejects} \end{cases}.$$

$\beta_B > 0$ denotes Bob's *private information weight*. Bob is Bayesian for $\beta_B = 1$, overweights private information for $\beta_B > 1$, and underweights private information for $\beta_B < 1$.

If Bob has a higher signal precision than Anna in each game, adaptation across games may lead him to eventually imitate Anna's decision in setting L . Bob's strategy will therefore improve if he overweights his private signal. Equivalently, underweighting his private signal improves Bob's strategy, if adaptation across games leads him to suboptimally follow his private signal in game H . Figure 1 shows all possible combinations of signal precisions $(q_A^L, q_A^H, q_B^L, q_B^H) \in (\frac{1}{2}, 1)^4$ for which over- or underweighting his private signal improves Bob's strategy. Settings with signal qualities in the blue (light blue) area

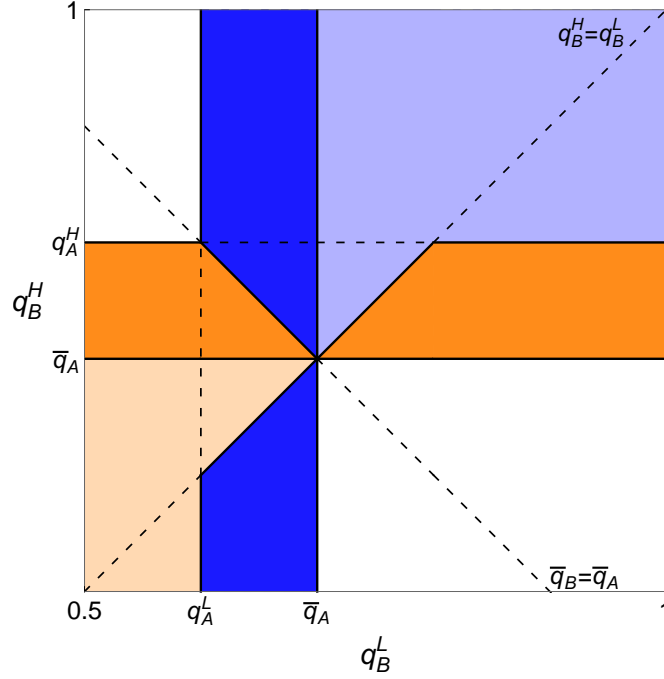


Figure 1: OPTIMAL BELIEF-UPDATING RULES IN THE SIMPLE EXAMPLE.

strictly (weakly) favor overweighting of private information, while settings with signal qualities in the orange (light orange) area strictly (weakly) favor underweighting of private information.⁵ The result indicates that optimal belief updating rules are likely to be heterogeneous in the population.

2.4 Discussion

There are two main reasons for Bob to adapt across games. First, Bob may be unable to distinguish games since he lacks or neglects relevant information. For instance, Anna's signal precision may not be known to Bob, and he may fail to recognize that his own and Anna's signal precision are correlated. Second, Bob's experience is likely restricted to a finite number of interactions in which case it may be optimal to adapt across games. Appendix B provides numerical results establishing that adaptation across games is optimal for Bob even if it leads to systematically mistaken long-run beliefs as long as the number of repetitions is not too large.

A second remark concerns the feedback Bob receives after each round. Indeed, the investment payoff may not be revealed to Bob unless he invests. Assuming this changes the long-run outcome of the adaptive process. After any finite number of rounds, Bob believes with strictly positive probability that imitating Anna's decision to reject regardless of his private signal maximizes his expected payoff. At this point however, the investment payoff is no longer revealed to Bob whenever Anna rejects. Bob is therefore unable to ever revise his wrong inference from a rejection by Anna and he imitates this decision in the long-run.

⁵See Proposition A.1 in the appendix for a formal statement and the proof.

Most importantly, this happens with strictly positive probability even if $q_B > q_A$, i.e. if it is optimal for Bob to never imitate Anna's decision, and if $q_A \approx 0.5$, i.e. if Anna does not possess any valuable information.

3 General Analytic Framework

3.1 A Social-Learning Game

There is a finite sequence of players $t = 1, 2, \dots, T$ who each choose an action a_t from the set $A = \{0, 1\}$. The sequential order of players is exogeneously, randomly determined and players are indexed according to their position. While payoff externalities are absent, player t 's payoff from her action depends on the realization of the state of Nature (henceforth state) $\tilde{\theta} \in \Theta = \{0, 1\}$.⁶ *Ex ante* the two states are equally likely. More precisely, player t 's payoff for each $t = 1, \dots, T$ is determined by the vN-M payoff function

$$u_t(a_t, \theta) = \begin{cases} \theta - \frac{1}{2} & \text{if } a_t = 1 \\ 0 & \text{if } a_t = 0 \end{cases}.$$

In the following, action $a = 1$ ($a = 0$) is sometimes referred to as “invest” (“reject”) and the costs of the investment are set equal to $\frac{1}{2}$ merely to simplify the exposition.⁷

Each player receives a private signal $\tilde{s}_t \in [0, 1]$ about the state. Conditional on the state, signals are independent and identically distributed. When the true state is θ , the signal distribution is given by the cumulative distribution function G_θ . G_0 and G_1 are mutually absolutely continuous and have common support $[\underline{b}, \bar{b}] \subseteq [0, 1]$. Therefore, a positive, finite Radon-Nikodym derivative exists and satisfies $f(s) = \frac{s}{1-s}$ (Smith and Sørensen, 2000). To ensure that some signals are informative, I rule out $f = 1$ almost surely. The assumptions imply that G_0 first-order stochastically dominates G_1 and that $s_t = \Pr(\tilde{\theta} = 1 \mid s_t)$. I assume that $\underline{b} > 0$ and $\bar{b} < 1$, i.e. private signals are bounded.

While player t only knows her own private signal, she observes additionally the complete history of previous actions denoted by $h_t = (a_1, \dots, a_{t-1}) \in H_t = A^{t-1}$ (and $h_1 \equiv \emptyset$). Subsequently, $H = \bigcup_{t=1}^T H_t$ denotes the complete set of histories, and $H_{T+1} = A^T$ denotes the set of final histories with element $h_{T+1} = (a_1, \dots, a_T)$.

The social-learning game is summarized by the collection $\langle T, A, \{u_t\}_{t=1}^T, \Theta, (G_\theta)_{\theta \in \Theta} \rangle$.

3.2 Rational Social-Learning

Rational social-learning relies on four main assumptions. First, players form beliefs about the state of nature by combining all the available information using Bayes' rule. Second,

⁶Throughout, tilded letters ($\tilde{\theta}$) denote random variables and standard letters (θ) denote realizations.

⁷Similarly, the state and the action set are binary to simplify the exposition. The results extend to any finite number of actions and states but at significant algebraic cost.

players maximize their expected utility given these beliefs. I refer to the first two assumptions as *Bayesian rationality* of players. Third, Bayesian rationality is commonly known; in other words players possess (common) *strategic knowledge*. Fourth, the social-learning game is commonly known meaning that players also possess (common) *structural knowledge*.⁸ The distinction between the two forms of common knowledge is important. Indeed, the next section investigates which adaptive process generates long-run outcomes equivalent to rational outcomes of social-learning when players possess strategic knowledge but are deprived from structural knowledge.

Without loss of generality, I focus on *behavioral strategies* $\sigma_t : [\underline{b}, \bar{b}] \times H_t \rightarrow \Delta(A)$, $t = 1, \dots, T$.⁹ In a slight abuse of notation, $\sigma_t(s_t, h_t)$ denotes player t 's probability of investment given her private signal s_t and the history h_t . Let Σ_t denote the strategy set of player t , $\Sigma = \times_{t=1}^T \Sigma_t$ the set of strategy profiles, and $\Sigma_{-t} = \times_{\tau \neq t} \Sigma_\tau$ the set of profiles of strategies for all players excluding t . The *ex-ante expected payoff* of player t following strategy $\sigma_t \in \Sigma_t$ for given $\sigma_{-t} \in \Sigma_{-t}$ is given by

$$U_t(\sigma_t, \sigma_{-t}) = \frac{1}{2} \sum_{\theta \in \Theta} \sum_{h_t \in H_t} \Pr(h_t \mid \tilde{\theta} = \theta, \sigma) \int_{\underline{b}}^{\bar{b}} \sum_{a \in A} \sigma_t(a \mid s_t, h_t) u(a, \theta) dG_\theta(s_t). \quad (1)$$

As shown by Tan and Werlang (1988), the assumptions above restrict players to iteratively undominated strategies. Strategy $\sigma_t \in \Sigma_t$ is *strictly dominated* if there exists $\sigma'_t \in \Sigma_t$ such that $U_t(\sigma'_t, \sigma_{-t}) > U_t(\sigma_t, \sigma_{-t})$ for every $\sigma_{-t} \in \Sigma_{-t}$. A strategy profile $\sigma = (\sigma_1, \dots, \sigma_T)$ is *iteratively undominated* if each strategy σ_t survives the iterated elimination of strictly dominated strategies. Lemma 1 characterizes all iteratively undominated strategies in terms of sequential best responses to Bayesian consistent *beliefs* $b_t : [\underline{b}, \bar{b}] \times H_t \rightarrow [0, 1]$ where $b_t(s_t, h_t) = \Pr(\tilde{\theta} = 1 \mid \tilde{s}_t = s_t, \tilde{h}_t = h_t)$.

Lemma 1. *To any iteratively undominated strategy profile σ there exists a belief system $\{b_t^*\}_{t=1}^T$ such that for each $t = 1, \dots, T$, each $s_t \in [\underline{b}, \bar{b}]$, and each $h_t \in H_t$*

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

$$b_t^*(s_t, h_t) = \left[1 + \frac{1 - s_t}{s_t} \frac{\Pr(h_t \mid \tilde{\theta} = 0, \sigma^*)}{\Pr(h_t \mid \tilde{\theta} = 1, \sigma^*)} \right]^{-1}$$

if $\Pr(h_t \mid \tilde{\theta} = \theta, \sigma^) = \prod_{\tau < t} \int_{\underline{b}}^{\bar{b}} \sigma_\tau^*(a_\tau \mid s_\tau, h_\tau) dG_\theta(s_\tau) > 0$ for each $\theta \in \Theta$ where $a_\tau = h_\tau(\tau)$ and $h_\tau \subset h_t$ for each $\tau < t$,*

⁸See Brandenburger (1996) for a similar distinction between strategic and structural *uncertainty*.

⁹For a given set M , $\Delta(M)$ is the set of probability distributions over M .

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

$$\sigma_t^*(s_t, h_t) = \begin{cases} 1 & \text{if } b_t^*(s_t, h_t) > \frac{1}{2} \\ 0 & \text{if } b_t^*(s_t, h_t) < \frac{1}{2} \end{cases}.$$

The iteratively undominated outcome is almost always unique.¹⁰

In the following, σ^* denotes the iteratively undominated strategy profile characterized in the lemma. While iterated dominance does not restrict behavior in case of a tie ($b_t^*(s_t, h_t) = 1/2$), ties occur with probability zero almost always. I therefore assume henceforth that the social-learning game does not allow for ties which is why there is no need to commit to a specific tie-breaking rule. Absent ties, iterated elimination of dominated strategies yields a unique outcome which is also the unique (Perfect) Bayesian equilibrium outcome.

Rational social-learning entails that players benefit *individually* from taking into account the information revealed by others' actions. Collectively, however, rational social-learning is *self-defeating* (Chamley, 2004b), because the more a player's decision is influenced by the actions of her predecessors (the more distinct are the probabilities $\Pr(h_t | \theta, \sigma^*)$ for $\theta \in \Theta$), the less new information it conveys to her successors (the more signals $s_t \in S$ yield the same sign of $b_t^*(s_t, h_t) - 1/2$). As a consequence, information aggregation slows down or stops completely which is why the economic outcome is inefficient.

3.3 Adaptive Social-Learning

Consider a (finite) family of social-learning games $\langle T, A, u, \Theta, (G_0^k, G_1^k) \rangle$, $k \in \mathcal{K} = \{1, \dots, K\}$, which differ only in the distribution of private signals, G_0^k and G_1^k , with support $[\underline{b}^k, \bar{b}^k]$. Players are assumed not to know these distributions. They interact repeatedly over the rounds $r = 1, 2, \dots$ where in each round one of the K games is played. Concretely, each round r begins with the random draws of (i) the game k^r where $k^r = k$ with probability $\pi_k > 0$, (ii) the state of nature θ^r where both states are equally likely in each round, (iii) the sequence of private signals $\{s_t^r\}_{t=1}^T$ where signals are independent across periods and drawn according to $G_{\theta^r}^{k^r}$, and (iv) the order of players such that each player eventually occupies any position in the sequence. Random draws are independent across rounds which guarantees that learning about the state and learning about the structure of the game and the strategies of other players are not confounded. At the end of each round, payoffs are realized.

Adaptation through repeated interaction requires players to receive feedback after each round. I suppose that private signals are never publicly revealed. Accordingly, for a given player i (in an abuse of notation) the relevant outcome of the game after round r is given

¹⁰I.e. the set of parameters of the social-learning game for which there exist multiple iteratively undominated outcomes is a null set.

by the tuple $\omega_i^r \equiv (k^r, \theta^r, t_i^r, h_{T+1}^r) \in \Omega \equiv \mathcal{K} \times \Theta \times \{1, \dots, T\} \times H_{T+1}$ where t_i^r denotes the position occupied by the player in round r .¹¹ Following Esponda (2008), I formally capture feedback via the functions $y_{\mathcal{K}} : \Omega \rightarrow 2^{\mathcal{K}}$, $y_{\Theta} : \Omega \rightarrow 2^{\Theta}$, and $y_H : \Omega \rightarrow 2^H$. When the outcome of the game is $\omega^r \in \Omega$, the player observes that the game k^r belongs to the set $y_{\mathcal{K}}(\omega^r)$. Similarly, the player observes that the state θ^r belongs to the set $y_{\Theta}(\omega^r)$ and that histories $h_t^r \in y_H(\omega^r)$ where $t \in \{1, \dots, T\}$ occurred. Let $y(\omega^r) = (y_{\mathcal{K}}(\omega^r)^r, y_{\Theta}(\omega^r)^r, y_H(\omega^r)^r)$. An *adaptation path* for round r is given by $\zeta^r = (\omega^1, \dots, \omega^{r-1}) \in \Omega^{r-1}$.

Since players do not know the information structure, they cannot derive the information contained in a history $h_t \in H_t$, i.e. the (game-specific) probabilities $\Pr(h_t | \theta)$ for each $\theta \in \Theta$ from their knowledge. Instead, players are assumed to *assess* this information based on their feedback from past interactions. Formally, for a given game k an *assessment* for period t is a mapping $\varphi_t^k : \Theta \rightarrow \Delta(H_t)$ assigning to each state $\theta \in \Theta$ a probability distribution $\varphi_t^k(h_t | \theta)$ over histories $h_t \in H_t$.

Definition 1. Let $\eta > 0$. The **adaptive process** is given by a sequence of frequencies $\kappa_t^{k,r} : H_t \times \Theta \times \Omega^{r-1} \rightarrow \mathbb{N}$, a sequence of assessments $\varphi_t^{k,r} : \Theta \times \Omega^{r-1} \rightarrow \Delta(H_t)$, and a sequence of strategic responses $\sigma_t^{k,r} : [\underline{b}_k, \bar{b}_k] \times H_t \times \Omega^{r-1} \rightarrow \Delta(A)$ for each $t = 1, \dots, T$ and each $k \in \mathcal{K}$, such that

(i) for each $r \geq 1$, $\zeta^r \in \Omega^{r-1}$, $k \in \mathcal{K}$, $t = 1, \dots, T$, $h_t \in H_t$, and $\theta \in \Theta$

$$\kappa_t^{k,r}(h_t, \theta | \zeta^r) = |\{1 \leq \rho < r : k \in y_{\mathcal{K}}(\omega^\rho) \wedge h_t \in y_H(\omega^\rho) \wedge \theta \in y_{\Theta}(\omega^\rho)\}|, \quad (2)$$

(ii) for each $r \geq 1$, $\zeta^r \in \Omega^{r-1}$, $k \in \mathcal{K}$, $t = 2, \dots, T$, $h_t \in H_t$, and $\theta \in \Theta$

$$\varphi_t^{k,r}(h_t | \theta; \zeta^r) = \frac{\kappa_t^{k,r}(h_t, \theta | \zeta^r) + \eta}{\sum_{h'_t \in H_t} [\kappa_t^{k,r}(h'_t, \theta | \zeta^r) + \eta]}, \quad (3)$$

(iii) for each $r \geq 1$, $\zeta^r \in \Omega^{r-1}$, $k \in \mathcal{K}$, $t = 1, \dots, T$, $h_t \in H_t$, and $s_t \in [\underline{b}_k, \bar{b}_k]$

$$\sigma_t^{k,r}(s_t, h_t | \zeta^r) = 1(0) \quad \text{if} \quad \frac{s_t}{1 - s_t} > (<) \frac{\varphi_t^{k,r}(h_t | 0; \zeta^r)}{\varphi_t^{k,r}(h_t | 1; \zeta^r)}. \quad (4)$$

The adaptive process relies on two main assumptions:¹² First, in a given round r , period t , and game k , players form beliefs at history h_t by replacing for each state $\theta \in \Theta$ the (unknown) conditional probability $\Pr(h_t | \theta)$ with the *relative frequency* with which history h_t occurred across relevant past rounds in which the state was θ . The past rounds used to calculate this relative frequency are determined by players' feedback. For example,

¹¹I assume that a player always remembers her position and I omit the subscript i henceforth.

¹²The adaptive process is based on the idea of *fictitious play* (Brown, 1951).

players only rely on past rounds in which the same game was played if $y_K(\omega^r) = \{k^r\}$. In contrast, players rely on past rounds regardless of the respective social-learning game if $y_K(\omega^r) = K$. Moreover, some past rounds may not be usable because the realized state was not revealed (i.e. $y_\Theta(\omega^\rho) = \emptyset$ for some $\rho < r$). Since the relative frequency is not well defined if no past rounds are usable for a given couple (history, state), I assume that players have arbitrarily small initial weights $\kappa_t^{k,1}(h_t, \theta | \emptyset) = \eta > 0$ for each k, t, h_t , and θ and I focus on the limit as $\eta \rightarrow 0$. In the limit, players attach probability zero to histories not observed previously and they believe that such histories are uninformative.

The second assumption is that players combine the relative frequencies according to Bayes' rule with their private signal, and that they rationally respond to the resulting belief. This assumption entails that players are myopic, i.e. they do not engage in strategic considerations of repeated play. In section 5, I relax the assumption that players are Bayesian by allowing them to experiment with flexible updating rules.

One feature of the adaptive process deserves special emphasis: While a player assesses the informational content of others' actions using past experience, she is assumed to *know* the informational content of her own private signal. The dichotomy relies on an interpretation of social-learning as a process in which *informed players* (though imperfectly) attempt to learn from others' decisions. Accordingly, private information is the outcome of an active process of information search and selection of the most credible source. In contrast, observed decisions are driven by unknown information sources and garbled through others' strategic thinking.

The adaptive process is very specific about how players form and respond to assessments. More general models of adaptation stay agnostic about how players exactly reach their decisions, allow for active experimentation, and focus on the asymptotic properties of the adaptive process (Fudenberg and Levine, 1998). A generalization of the adaptive process in this direction is discussed in Appendix C.

3.4 Convergence

To formalize convergence of the adaptive process I consider the following two metrics on, respectively, the set of assessments and the strategy space:

Definition 2. Let $\epsilon > 0$. For each $k \in K$ and each $t = 1, \dots, T$

- (i) assessments $\varphi_t^k, \hat{\varphi}_t^k : \Theta \rightarrow \Delta(H_t)$ are ϵ -close if $|\varphi_t^k(h_t | \theta) - \hat{\varphi}_t^k(h_t | \theta)| < \epsilon$ for each $h_t \in H_t$ and each $\theta \in \Theta$,
- (ii) strategy σ_t^k plays ϵ -like strategy $\hat{\sigma}_t^k$ at history $h_t \in H_t$ if there exists $B_\epsilon \subseteq [\underline{b}^k, \bar{b}^k]$ such that $G_\theta^k(B_\epsilon) > 1 - \epsilon$ for each $\theta \in \Theta$, and $|\sigma_t^k(s_t, h_t) - \hat{\sigma}_t^k(s_t, h_t)| < \epsilon$ for each $s_t \in B_\epsilon$ (Jackson and Kalai, 1997).

Lemma 2 establishes that the two distance functions are consistent meaning that ϵ -closeness of assessments implies ϵ -like play of strategies and vice versa. To state the

result, call a strategy profile σ^k and a profile of assessments $\{\varphi_t^k\}_{t=1}^T$ corresponding, if (i) $\varphi_t^k(h_t | \theta) = \prod_{\tau < t} \int_{\underline{b}^k}^{\bar{b}^k} \sigma_\tau^k(a_\tau | s_\tau, h_\tau) dG_\theta^k(s_\tau)$ and (ii) $\sigma_t^k(s_t, h_t) = 1$ (0) if $\frac{s_t}{1-s_t} > (<) \frac{\varphi_t^k(h_t|0)}{\varphi_t^k(h_t|1)}$ for each $t = 1, \dots, T$, each $h_t \in H_t$, and each $s_t \in [\underline{b}^k, \bar{b}^k]$.

Lemma 2. *Fix $k \in \mathcal{K}$ and consider the strategy profiles $\sigma^k, \hat{\sigma}^k$ with corresponding assessments $\{\varphi_t^k\}_{t=1}^T$ and $\{\hat{\varphi}_t^k\}_{t=1}^T$. For each $t = 1, \dots, T$ and each $\epsilon > 0$ there exists $\delta > 0$ such that*

- (i) σ_t^k plays ϵ -like $\hat{\sigma}_t^k$ at each $h_t \in H_t$ satisfying $\hat{\varphi}_t^k(h_t | \theta) > 0$ for each $\theta \in \Theta$, if φ_t^k is δ -close to $\hat{\varphi}_t^k$,
- (ii) φ_t^k is ϵ -close to $\hat{\varphi}_t^k$, if σ_τ^k plays δ -like $\hat{\sigma}_\tau^k$ for each $\tau < t$ at each $h_\tau \in H_\tau$ satisfying $\hat{\varphi}_\tau^k(h_\tau | \theta) > 0$ for each $\theta \in \Theta$.

Based on Lemma 2, I will focus on convergence of strategies henceforth. All convergence results are in probabilistic terms with respect to the distribution \mathbf{P} of (infinite) adaptation paths $(\omega^1, \omega^2, \dots) \in \Omega^\infty$ induced by the objective distributions of the random variables and the rules for the formation of assessments and strategies.

4 Long-run Outcomes of Adaptation

In this section I discuss the long-run outcomes of the adaptive process under different feedback regimes. As a benchmark, I first characterize the long-run outcome if players observe the complete sequence of actions and the state at the end of each round, and are able to distinguish games (4.1). Second, I investigate the consequences of constraints on the observation of actions or the state (4.2). Finally, I study the outcome of the adaptive process under the assumption that players do not distinguish games (4.3). While for the sake of clarity I assume mutual knowledge of the social-learning game, I refrain from assuming any higher-order interactive knowledge.

4.1 Adaptation by Game

I first assume that players observe the outcome of the game after each round.

Proposition 1. *Let $y(\omega^r) = (\{k^r\}, \{\theta^r\}, \{h_t \in h_{T+1}^r\})$ for each $r \geq 1$ and each $\omega^r \in \Omega$. For each $\epsilon > 0$, each $k \in \mathcal{K}$, and each $t = 1, \dots, T$, the limit strategy $\lim_{r \rightarrow \infty} \sigma_t^{k,r}$ almost surely plays ϵ -like a rationalizable strategy $\sigma_t^{k,*}$ at all histories satisfying $\varphi_t^{k,*}(h_t | \theta) > 0$ for each $\theta \in \Theta$ where $\varphi_t^{k,*}(h_t | \theta)$ are the assessments under rational social-learning.*

The proposition establishes that rational social-learning may be justified as the outcome of an adaptive process if players are able to distinguish games, receive ample feedback after each round, and play each game infinitely often. It follows from the dominance-solvability of the game (Milgrom and Roberts, 1991). In a nutshell, since players have

a dominant strategy in period 1, period 2-assessments will eventually be correct by the law of large numbers if players distinguish games. Hence, strategies for period 2 (and by induction strategies for later periods) are eventually rationalizable.¹³

4.2 Feedback Constraints

Proposition 1 relies on strong assumptions regarding players' feedback. For instance, though the state must be revealed through the payoffs for at least one of the available actions, it may not be revealed for all of them. If a player rejects an investment opportunity, she may never know what the outcome could have been. To take this possibility into account, I now investigate the impact of feedback constraints on the state.¹⁴ Formally, I assume that $y_\Theta(\omega^r) = \emptyset$ if $a_{tr}^r = 0$ and $y_\Theta(\omega^r) = \{\theta^r\}$ if $a_{tr}^r = 1$. I maintain the assumption that players get rich feedback on the game and the history.

Proposition 2. *Assume that for each $r \geq 1$ and each $\omega^r \in \Omega$, $y_K(\omega^r) = k^r$, $y_H^r(\omega^r) = \{h_t \subset h_{T+1}^r\}$, $y_\Theta(\omega^r) = \emptyset$ if $a_{tr}^r = 0$, and $y_\Theta(\omega^r) = \{\theta^r\}$ if $a_{tr}^r = 1$. For each $k \in \mathcal{K}$ such that $\frac{G_0^k(1/2)}{G_1^k(1/2)} < \frac{\bar{b}_k}{1-b_k}$ there exists $\bar{\epsilon} > 0$ and $t > 1$ such that with strictly positive probability the limit strategy $\lim_{r \rightarrow \infty} \sigma_t^{k,r}$ does not play ϵ -like the rationalizable strategy $\sigma_t^{k,*}$ at history $h_t = h_t^0 \equiv (0, \dots, 0)$ for each $0 < \epsilon < \bar{\epsilon}$. Concretely, $\lim_{r \rightarrow \infty} \sigma_\tau^{k,r}(s_\tau, h_\tau) = 0$ for each $\tau \geq t$, each $h_\tau \supseteq h_t^0$, and each $s_\tau \in [b_k, \bar{b}_k]$.*

The proposition shows that adaptive social-learning may differ from rational social-learning if feedback on the state is conditional on a player's decision. In particular, players may in the long-run imitate the rejections of their predecessors although they would not do so if they knew the true informational content of previous decisions. More broadly, herds on an action that does not reveal the state through its payoffs are more likely in the long-run of adaptive social-learning than under rational social-learning, and more likely to be wrong.

The rationale behind the result is simple: After any finite number of rounds, players are with strictly positive probability convinced that the evidence conveyed by a sequence of rejections is sufficiently strong to swamp any private signal since private signals are bounded. Players therefore imitate their predecessors and are no longer able to revise their wrong assessments because of the constrained feedback. Formally, the limit outcome is a *self-confirming equilibrium* (see e.g. Dekel et al., 2004). The condition $G_0^k(1/2)/G_1^k(1/2) <$

¹³Mutual knowledge of the social learning game is not necessary for the result. In the long-run any possible state and any possible (open subset of) private signal(s) occurs infinitely often such that frequentists may learn their own private signal distribution as long as its support is well-behaved. Moreover, players are able to explore the (finite) action spaces of other players and their payoff function once a small amount of experimentation is assumed.

¹⁴Another possibility is that a player does not get feedback on subsequent choices. Yet, since a player's position in the sequence is randomly determined in each round, each player acts in period T infinitely often and observes the complete history of actions. Therefore, the result of Proposition 1 is robust with respect to limited feedback on the history.

$\bar{b}_k / (1 - \bar{b}_k)$ guarantees that it is not optimal to imitate the first rejection in game k . The result is constrained to histories containing only rejections since players switch positions in the social-learning game. Thus, players who invest before the herd starts would be able to revise their assessments. The result could be strengthened by assuming, additionally, that players receive no feedback about subsequent choices.

Proposition 2 continues to hold if the first players in the sequence are uninformed (and uninformed players move early for ulterior motives) as long as this is not mutually known. The proposition therefore helps explain why players sometimes follow others based on no apparent reason. For example, in recent years a growing number of parents refuse to vaccinate their children for fear of harmful side effects despite the extensive statistical evidence to the contrary. One reason seems to be a strong reliance on anecdotal evidence (see e.g. Moran et al., 2015). In the model considered here, players are forced to rely on anecdotal reasoning since feedback constraints prevent them from collecting representative statistical information.

4.3 Adaptation Across Multiple Social Learning Games

In the field, individuals rarely encounter exactly the same strategic situation a large number of times. Therefore, they are likely to extrapolate experience across games they deem similar (see e.g. Fudenberg, 2006). In this subsection, I investigate this idea by assuming that players receive a coarse feedback $y_K(\omega^r) = \mathcal{K}$ about the game.

There are several reasons why players adapt across games. First, absent common knowledge of the information structure, players may be unable to distinguish games. Grimm and Mengel (2012) show that experimental subjects extrapolate especially in complex environments where information to distinguish games is scarce. Though a player could assess the informational content of others' actions conditional on her own private signal,¹⁵ she may fail to account for the correlation between her private signal and the actions of others (Esponda, 2008), or ignore the importance of this factor altogether (see e.g. Ross, 1977). Second, players may hold a simplified representation of the world due to limited cognitive abilities (see e.g. Samuelson, 2001; Mengel, 2012). Third, bundling experiences can be optimal when experience with a collection of games is scarce because the larger amount of available data overcompensates the loss in predictive accuracy (Al-Najjar and Pai, 2014; Mohlin, 2014). Appendix B presents simulation results for simple social-learning games which establish this principle for the setup of this paper.

The first result of the subsection shows that adaptation across games converges to an *analogy-based expectations equilibrium* (Jehiel, 2005; Jehiel and Ettinger, 2010, ABEE henceforth). In a general ABEE, each player partitions the decision nodes of other players into *analogy classes* and has a correct understanding of average behavior in each class.

¹⁵This is not possible if games differ with respect to preferences; see Section 6.1.

In the long-run ABEE considered here, players use a separate analogy class for each history-state-pair (h_t, θ) and their expectation about the (average) distribution of actions at history h_t if the state is θ is correct. However, each analogy-class bundles all decision situations with the same history and state regardless of the game k and players' private signals. I therefore refer to the partition as the *information-anonymous analogy partition*. The ABEE additionally assumes that players combine their analogy-based expectations with their private signal using Bayes' rule and best respond to the resulting beliefs.¹⁶

Proposition 3. *Let $y(\omega^r) = (\mathcal{K}, \{\theta^r\}, \{h_t \in h_{T+1}^r\})$ for each $r \geq 1$ and each $\omega^r \in \Omega$. For each $\epsilon > 0$, each $k \in \mathcal{K}$, and each $t = 1, \dots, T$, the limit strategy $\lim_{r \rightarrow \infty} \sigma_t^{k,r}$ almost surely plays ϵ -like an ABEE strategy $\sigma_t^{k,A}$ at all histories satisfying $\bar{\varphi}_t(h_t | \theta) > 0$ for each $\theta \in \Theta$.*

As Proposition 1, Proposition 3 rests upon the dominance-solvability of the social-learning game. The existence of a dominant strategy in period 1 implies that players eventually correctly assess the average choice probabilities conditional on the state, where the average is taken across games. Accordingly, by best responding to assessments, players eventually play the ABEE strategy in period 2, and in later periods by induction.

Comparing analogy-based social-learning with rational social-learning yields two interesting results. First, while the ABEE strategies coincide with the rationalizable strategies in period 1, players will in general not make correct inferences in a given game in later periods. The ABEE strategies may therefore be suboptimal in a given game. Second, ABEE strategies do not discriminate between games. More precisely, $\sigma_t^{k,A}(s_t, h_t) = \sigma_t^{\ell,A}(s_t, h_t)$ for any two games $k, \ell \in \mathcal{K}$, any period t and history $h_t \in H_t$, and any signal $s_t \in [b_k, \bar{b}_k] \cap [b_\ell, \bar{b}_\ell]$. Among all strategies $\sigma_t : \bigcup_k [b^k, \bar{b}^k] \times H_t \rightarrow \Delta(A)$ which do not discriminate between games in this sense, the ABEE strategy need not maximize the *ex ante expected payoff* in the randomly selected social learning game given by $\bar{U}_t(\sigma_t, \{\sigma_{-t}^{k,A}\}_{k \in \mathcal{K}}) = \sum_{k \in \mathcal{K}} \pi_k U_t^k(\sigma_t, \sigma_{-t}^{k,A})$ where $U_t^k(\sigma_t, \sigma_{-t})$ is the expected payoff in game k given in (1). The following proposition formalizes these two results. Following Smith and Sørensen (2000), I call a property *generic* if it holds for an open and dense subset of parameters.

Proposition 4. *Generically,*

1. *there exists $k \in \mathcal{K}$, $t > 1$, $h_t \in H_t$ and $\bar{\epsilon} > 0$ such that the ABEE strategy $\sigma_t^{k,A}$ does not play ϵ -like a rationalizable strategy $\sigma_t^{k,*}$ at history h_t for each $0 < \epsilon < \bar{\epsilon}$.*
2. *there exists $t > 1$ and $\hat{\sigma}_t : \bigcup_k [b^k, \bar{b}^k] \times H_t \rightarrow \Delta(A)$ such that*

$$\bar{U}_t(\hat{\sigma}_t, \{\sigma_{-t}^{k,A}\}_{k \in \mathcal{K}}) > \bar{U}_t(\sigma_t^{k,A}, \{\sigma_{-t}^{k,A}\}_{k \in \mathcal{K}}).$$

¹⁶See appendix A for a formal definition and a characterization of the ABEE. The assumption of Bayesian updating is relaxed in section 5.

Proposition 4 shows that long-run behavior is not optimal when players do not distinguish games while adapting. Accordingly, opportunities to improve upon long-run behavior exist. Section 5 explores one such opportunity by assuming that players use more flexible belief updating rules than Bayes' rule. The proposition also provides an explanation for unjustified contagion or spill-over between games. If a herd forms after history h_t in game k , players may conclude that following others at h_t is also optimal in game $\ell \neq k$ even if this is not warranted by fundamentals. This problem may be exacerbated when allowing for limited feedback on the state as in Proposition 2.

It is finally instructive to compare the approach to related literature. First, the information-anonymous analogy partition is finer than the payoff-relevant analogy partition considered by Guarino and Jehiel (2013) since players distinguish behavior at different histories.¹⁷ On the other hand, it is distinct from the private information analogy partition (Jehiel and Koessler, 2008).¹⁸ Second, players might adapt across games that differ in other aspects such as payoffs.¹⁹ In the present game, the simple structure of payoffs, the absence of payoff externalities, and the realization of payoffs at the end of each round facilitate the identification of a player's own payoff function. An extension of the model where games also differ with respect to the distribution of preferences yields similar results and is investigated in section 6.1.

5 Heterogeneous Belief Updating

Thus far, I have assumed that players form beliefs according to Bayes' rule. A more general adaptive process allows players to also adjust their responses to beliefs. As in Section 2, I consider flexible belief updating rules where player t 's belief in game k given history h_t , signal s_t , and assessments $\varphi_t^k(h_t | \theta)$ for each $\theta \in \Theta$ is given by

$$b_t^k(s_t, h_t) = \left[1 + \left(\frac{1 - s_t}{s_t} \right)^{\beta_t} \frac{\varphi_t^k(h_t | 0)}{\varphi_t^k(h_t | 1)} \right]^{-1}. \quad (5)$$

β_t is player t 's *private information weight*. A player is Bayesian, if $\beta_t = 1$, overweights private information if $\beta_t > 1$, and underweights private information if $\beta_t < 1$ (see e.g. Grether, 1980; Hung and Plott, 2001; Palfrey and Wang, 2012). In the following, I assume that players adjust their private information weight alongside their assessments. Concretely, in the spirit of reinforcement learning, players pick the weight which yields a higher expected payoff. I focus on the long-run. Accordingly, the selected long-run private information weight must maximize the ex ante expected payoff given long-run assessments.

¹⁷Guarino and Jehiel (2013) do not consider multiple games.

¹⁸The coarsest common refinement is given by the analogy partition $\hat{\mathcal{A}}_i = \{\alpha(h_t, \theta, s_i)\}$ for which player i assesses others' behavior separately for each history, state, and realization of his private signal.

¹⁹See e.g. Steiner and Stewart (2008); Mengel (2012); Grimm and Mengel (2012).

The simple example of Section 2 suggests that the long-run belief updating rule depends on the family of social-learning games. I explore this possibility in an extension of the model where the precision of private signals has a player-specific component. Formally, in game k player i draws signals from $\{0, 1\}$ according to probabilities $\Pr(\tilde{s}_i^k = 1 \mid \tilde{\theta} = 1) = \Pr(\tilde{s}_i^k = 0 \mid \tilde{\theta} = 0) = q_i^k$ where

$$q_i^k = \frac{1}{2} + \frac{1}{2} \frac{\exp(\bar{\rho}^k + \nu_i)}{1 + \exp(\bar{\rho}^k + \nu_i)}.$$

$\bar{\rho}^k$ determines the average signal precision in game k and is a measure for the ease with which information can be collected in this game. By contrast, ν_i is a player-specific signal precision component which measures the ease with which player i collects information.²⁰ Player i is better informed the larger is ν_i . Let V denote the cumulative distribution function of player-specific signal precision components.

I study the long-run outcome of adaptation across games when diversely informed players are allowed to combine public and private information differently. The limit outcome is an extended ABEE given by the mapping $\beta^A: \mathbb{R} \rightarrow \mathbb{R}_+$ such that (i) player i with quality component ν_i invests (rejects) in period t of game k given history $h_t \in H_t$ and signal $s_i^k \in [\underline{b}^k, \bar{b}^k]$ if

$$\left(\frac{s_i^k}{1 - s_i^k} \right)^{\beta^A(\nu_i)} > (<) \frac{\bar{\varphi}_t(h_t \mid 0)}{\bar{\varphi}_t(h_t \mid 1)}, \quad (6)$$

and (ii) assessments $\bar{\varphi}_t(h_t \mid \theta) = \sum_k \pi_k \bar{\varphi}_t^k(h_t \mid \theta)$ take into account the distribution of player-specific signal precision components and the associated updating rules via

$$\bar{\varphi}_t^k(h_t \mid \theta) = \prod_{\tau < t} \int_{\nu_\tau} \sum_{s_\tau} \Pr_k(s_\tau \mid \theta, \nu_\tau) \sigma_\tau^{k, \beta^A}(a_\tau \mid s_\tau, h_\tau) dV(\nu_\tau)$$

where $\Pr_k(s_\tau \mid \theta, \nu_\tau)$ denotes the probability that the player in period τ who has signal precision component ν_τ observes signal s_τ in game k when the state is θ , and σ_τ^{k, β^A} are the strategies determined by (6). Proposition 5 characterizes the equilibrium mapping β^A which maximizes the ex ante expected payoff for each $\nu \in \text{supp}(V)$.

Proposition 5. *There exist thresholds $\underline{\nu}, \bar{\nu} \in \mathbb{R}$ such that the extended ABEE β^A that maximizes the ex ante expected payoff for each $\nu \in \text{supp}(V)$ satisfies $\beta^A(\nu_t) \geq 1$ for each $\nu_t > \bar{\nu}_i$ and $\beta^A(\nu_t) \leq 1$ for each $\nu_t < \underline{\nu}$.*

Proposition 5 states that players with a high (low) signal precision benefit from over-weighting (underweighting) private information. Figure 2 illustrates the result for two equally likely social-learning games ($K = 2$, $\pi_1 = \pi_2 = \frac{1}{2}$), and player-specific signal

²⁰I abuse notation and index players with i to emphasize that the second component of signal precision is player-specific rather than period-specific.

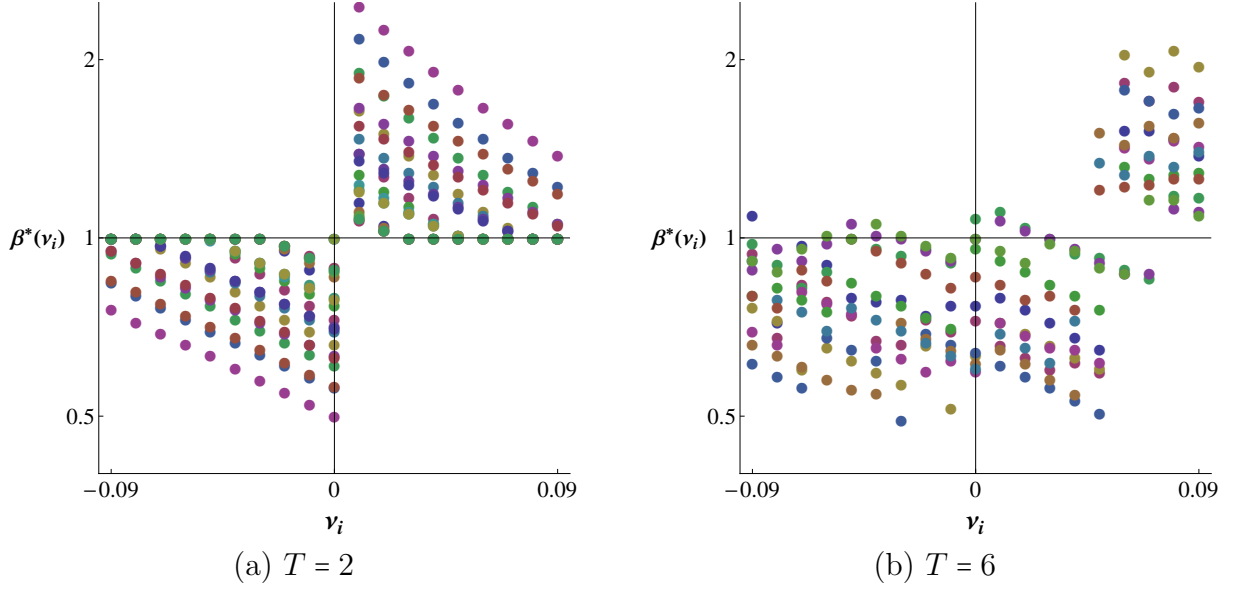


Figure 2: ADAPTATION ACROSS GAMES, HETEROGENOUS QUALITY OF INFORMATION, AND OPTIMAL UPDATING OF BELIEFS.

precisions, q_i^k , $k = 1, 2$, distributed uniformly around average qualities q^k , where $q^k \in \{0.6, 0.65, \dots, 0.9\}$, and $q_i^k \in \{q^k - 0.09, q^k - 0.08, \dots, q^k + 0.09\}$. In the figure, different colors denote equilibrium values $\beta^A(\nu_i)$ for different pairs (q^1, q^2) .²¹ As one can see, especially players whose signal precision is slightly larger (smaller) than the average signal precision deviate the most from Bayesian updating in equilibrium, since they are most likely to suffer from distorted assessments.

It is noteworthy that exactly what constitutes a high or low precision of private information is context-dependent. This may be illustrated with the help of a simple story: University students usually have diverse educational backgrounds. Some students attend specific university-preparatory schools. In such an environment, a student may experience that imitating others is a successful strategy even if she considers herself well informed. This student may therefore develop a tendency to underweight her private information. In contrast, other students come from more rural areas where schools have large catchment areas and a correspondingly diverse student body. A top-of-the class student from such a school may learn that imitation is rarely beneficial even if his own information is weak. He may therefore come to overweight his private information. Though adaptive, these tendencies to apply non-Bayesian updating rules show up as biases at the university.

The proposition suggests that heterogeneous updating rules will evolve. It therefore provides a theoretical background for corresponding evidence (e.g. El-Gamal and Grether, 1995; Delavande, 2008; Palfrey and Wang, 2012). Moreover, it motivates a straightforward, behavioral extension of the standard model of social-learning in which players are assumed to have private information weights drawn from a given distribution. The extension

²¹The equilibria have been calculated numerically. The results and the code are available from the author upon request.

is straightforward, since the equilibrium of the extended game in which the distribution of private information weights is commonly known is equivalent to the Bayesian equilibrium of a standard game with a distorted distribution of private information (March and Ziegelmeyer, 2009, Lemma 3.2). On the other hand, the new, behavioral model of social-learning better captures the established experimental regularities under the assumption that a sufficiently large proportion of the population overweights private information. In this case, compared to the standard model cascades emerge later (since overweighers herd later), beliefs become more extreme (since each action stems from an overweigher with positive probability), and therefore the length and strength of cascades are correlated (since more and more moderate overweighers enter the herd). Moreover, heterogeneous updating rules improve economic efficiency. Indeed, if the support of the distribution of private information weights is unbounded, social-learning is complete even with private signals of bounded strength (March and Ziegelmeyer, 2009, Corollary 3.4). The behavioral model also accommodates new findings such as the coexistence of the ‘overweighting-of-private-information’ bias with a ‘social-confirmation’ bias (March and Ziegelmeyer, 2015). Finally, its adaptive foundations distinguish the model from other suggested explanations such as the level- k model or models based on a limited depth of reasoning.

6 Extensions

6.1 Preferences

I have assumed so far that players share the same utility function. More general social-learning games allow for heterogeneous preferences and assume that the distribution of preference types is commonly known (Smith and Sørensen, 2000). Obviously, this common knowledge assumption is equally questionable, and it is likely that players adapt across games which differ in the underlying distribution of preferences. I illustrate that similar results obtain under this assumption with the help of the simple example of Section 2. Accordingly, let Anna and Bob interact repeatedly in games $k \in \{H, L\}$ which are equally likely. To highlight the role of preferences, I assume that $q_i^H = q_i^L = q_i$ for each $i \in \{A, B\}$.

Assume first that Anna occasionally does not care about the investment payoff and invests or rejects regardless of her private signal. Concretely, each time game k is played Anna picks her action regardless of her private signal with probability γ_A^k where $\gamma_A^L < \gamma_A^H$ without loss of generality. I assume that Anna invests or rejects with equal probability in this case. Conversely, in any given repetition of game k Anna follows her private signal with probability $1 - \gamma_A^k$. The assumptions imply that the relative frequency with which Anna invests (rejects) when the investment payoff is 1 (0) approaches $\hat{q}_A^k = \frac{1}{2} \gamma_A^k + (1 - \gamma_A^k) q_A$. If Bob distinguishes games, he eventually imitates Anna’s decision (follows his private signal) if $q_B < (>) \hat{q}_A^k$. In contrast, if Bob adapts across games he eventually imitates Anna’s decision (follows his private signal) if $q_B < (>) \hat{q}_A \equiv \frac{1}{2} \bar{\gamma}_A + (1 - \bar{\gamma}_A) q_A$

where $\bar{\gamma}_A = \frac{\gamma_A^L + \gamma_A^H}{2}$. This strategy is not optimal for Bob in game L (H) if $\hat{q}_A < q_B < \hat{q}_A^L$ ($\hat{q}_A^H < q_B < \hat{q}_A$) and Bob benefits from underweighting (overweighting) his private signal.

An even stronger result obtains if Anna is assumed to have *opposed* preferences in the two games. Assume for instance that Anna's payoff function satisfies $u_A^L(1, \theta) = \theta - \frac{1}{2}$, $u_A^H(1, \theta) = \frac{1}{2} - \theta$, and $u_A^k(0, \theta) = 0$ for each $k \in \{H, L\}$ and each $\theta \in \Theta$. It follows that Anna invests (rejects) given signal $s_A = 1$ and rejects (invests) given signal $s_A = 0$ in game L (H). Accordingly, the relative frequency with which Anna invests (rejects) across rounds where the investment payoff is 1 (0) approaches q_A in game L , $1 - q_A$ in game H , and $\frac{1}{2}$ across games. Therefore, Bob eventually follows his private signal in each game when adapting across games. In contrast, it would be optimal for Bob to imitate Anna's decision in game L and to anti-imitate Anna's decision in game H if $q_B < q_A$.²²

It is straightforward to extend the results to the general social-learning game. Preference heterogeneity is therefore another hindrance to successful adaptation. Indeed, since players cannot use their private signal realizations to distinguish games, adaptation across games which differ with respect to preferences may be unavoidable. Accordingly, preference heterogeneity strengthens the results.

6.2 Endogenous Timing of Decisions

The analysis has focused on settings where the timing of decisions is given exogeneously. With an endogenous timing of decisions, players must foresee future social-learning opportunities in addition to interpreting the observed history of actions (see e.g. Chamley, 2004a). As shown below, the adaptive process need not converge to a Perfect Bayesian equilibrium (PBE) when feedback constraints prevail.

Following Chamley (2004a), I extend the social-learning game as follows: First, each player $i = 1, \dots, N$ has the option to make one irreversible investment in one of the periods $t = 1, 2, \dots$. Second, the payoff from investing in period t is given by $\delta^{t-1} \cdot (\theta - c)$ where $0 < \delta < 1$ is the discount factor and the cost of investment satisfies $\underline{b} \leq c < \bar{b}$. In any PBE of this game, there exists a threshold s_1^* such that players with private signals $s > s_1^*$ ($s < s_1^*$) invest in period 1 (delay investment for at least one period). Moreover, there exists at least one PBE with $c < s_1^* < \bar{b}$ if $\underline{b} < c < \bar{b}$, there exists a PBE such that $s_1^* = \underline{b}$ if $c = \underline{b}$, and there may be multiple equilibria in both cases (Chamley, 2004a, Theorem 1).

Assume that the distribution of private information and the equilibrium strategies are not commonly known. Accordingly, players must acquire the information necessary to assess the *option value of delay* through repeated play of the social-learning game. I focus on the case of a single game. Let $\omega_i^r \equiv (\theta^r, t_i^r, h^r)$ denote the outcome of the social-learning game in repetition r where $h^r = (x_1^r, \dots, x_N^r)$ is the history of the number of investments x_t^r in each period $t = 1, \dots, N$ and $t_i^r = \infty$ if the player never invested (the game lasts

²²Bob anti-imitates Anna's decision if he chooses $a_B \neq a_A$ regardless of his private signal.

at most N periods). Let $y_H(\omega^r)$ and $y_\Theta(\omega^r)$ denote the feedback on respectively the history and the state of nature given the outcome ω^r . Below, I characterize the long-run outcome of the adaptive process assuming that (i) $y_H(\omega_i^r) = \{h_t \subset h^r : t < t_i^r\}$, i.e. player i who invests in period t in a given round does not receive feedback about the number of investments in period t or later, and (ii) $y_\Theta(\omega_i^r) = \theta^r$ if $t_i^r < \infty$ and $y_\Theta(\omega_i^r) = \emptyset$ if $t_i^r = \infty$, i.e. players only receive feedback about the realized state when they invest.

Proposition 6. *Let $y_H(\omega_i^r) = \{h_t \subset h^r : t < t_i^r\}$, $y_\Theta(\omega_i^r) = \theta^r$ if $t_i^r < \infty$, and $y_\Theta(\omega_i^r) = \emptyset$ if $t_i^r = \infty$. In the long-run of the adaptive process almost surely,*

1. *all players invest in period 1 if $c = \underline{b}$,*
2. *players with $s_i > c$ invest in period 1 and players with $s_i < c$ never invest if $\underline{b} < c < \bar{b}$.*

The proposition entails that players fail to assess the option value of delay. The reason is simple (a formal proof is omitted): A priori, delay has no value and players invest (in period 1) if and only if it is profitable, i.e. if $s_i > c$. Feedback constraints prevent all players from acquiring information since players who invest in period 1 receive no feedback about the number of investments in period 1 and players who never invest are not informed about the state of nature.

7 Conclusion

The paper scrutinizes models of rational social-learning through the lens of adaptation. Adaptation generates rational behavior in the long-run if and only if individuals are able to distinguish social learning settings, receive ample feedback, and their experience with each setting grows without bounds. Limited opportunities for adaptation lead to mistaken inferences from others' actions and render Bayes' rule payoff inferior compared to non-Bayesian belief updating rules.

The paper offers some new directions for experimental and theoretical research on social-learning. First, to the best of my knowledge the existing experimental studies have considered laboratory settings in line with standard economic models of social-learning. These experimental settings were the obvious candidates for testing economic models of social-learning, but they have a limited ecological validity.²³ If subjects perceive the laboratory setting as artificial, deviations from rational behavior might not come as a surprise and do not constitute conclusive evidence against rational social-learning. Novel economic experiments should test the rational view of social-learning in settings resembling those

²³For an experimental study to possess ecological validity, the methods and the setting of the study must approximate the real-life situation that is under investigation. Ecological validity is independent from external validity which relates to the ability of a study's results to generalize.

encountered by human participants in the field with structural uncertainty and sufficient opportunities to adapt.²⁴

Second, the paper shows that a thorough investigation of the assumptions underlying rational social-learning straightforwardly leads to a behavioral model of social-learning with heterogeneous belief updating rules. This behavioral model of social-learning complements recent attempts which have relaxed the assumptions of rational social-learning by arguing for greater psychological realism. Indeed, some of the psychologically more plausible assumptions would disappear when adaptation is considered.²⁵ More generally, the acknowledgment that individuals face limited opportunities for adaptation is likely to deliver fruitful insights in other economic domains.

References

- Acemoglu, D., Munther, D., Lobel, I., Ozdaglar, A., 2011. Bayesian learning in social networks. *Review of Economic Studies* 78, 1201–1236.
- Al-Najjar, N. I., Pai, M. M., 2014. Coarse decision making and overfitting. *Journal of Economic Theory* 150, 467–486.
- Avery, C., Zemsky, P., 1998. Multidimensional uncertainty and herd behavior in financial markets. *The American Economic Review* 88, 724–748.
- Banerjee, A. V., 1992. A simple model of herd behavior. *The Quarterly Journal of Economics* 107, 797–817.
- Bernardo, A., Welch, I., 2001. On the evolution of overconfidence and entrepreneurs. *Journal of Economics & Management Strategy* 10, 301–330.
- Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100, 992–1026.
- Bohren, J. A., 2015. Informational herding with model misspecification. *Journal of Economic Theory* forthcoming.
- Brandenburger, A. M., 1996. Strategic and structural uncertainty in games. In: Zeckhauser, R., Keeney, Sebenius (Eds.), *Wise Choices*. Springer, Berlin, pp. 221–232.
- Brown, G. W., 1951. Iterative solution of games by fictitious play. In: Koopmans, T. (Ed.), *Activity Analysis of Production and Allocation*. Wiley, New York, pp. 374–376.

²⁴For example, Hertwig et al. (2004) find that “decisions from experience and decisions from description can lead to dramatically different choice behavior”.

²⁵For instance individuals which are naïve in the sense of Eyster and Rabin (2010) should be surprised by the clustering of decisions which does not tally with independently distributed private information.

- Cai, H., Chen, Y., Fang, H., 2009. Observational learning: Evidence from a randomized natural field experiment. *American Economic Review* 99, 864–882.
- Chamley, C., 2004a. Delays and equilibria with large and small information in social learning. *European Economic Review* 48, 477–501.
- Chamley, C., 2004b. *Rational herds: Economic models of social learning*. Cambridge, Cambridge University Press.
- Cipriani, M., Guarino, A., 2014. Estimating a structural model of herd behavior in financial markets. *American Economic Review* 104, 224–251.
- Dekel, E., Fudenberg, D., Levine, D. K., 2004. Learning to play bayesian games. *Games and Economic Behavior* 46, 282–303.
- Dekel, E., Gul, F., 1997. Rationality and knowledge in game theory. In: Kreps, D. M., Wallis, K. F. (Eds.), *Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress*. Vol. 1 of *Econometric Society Monographs*, Vol. 26. Cambridge: Cambridge University Press, pp. 87–172.
- Delavande, A., 2008. Measuring revisions to subjective expectations. *Journal of Risk and Uncertainty* 36, 43–82.
- El-Gamal, M., Grether, D., 1995. Are people bayesian? uncovering behavioral strategies. *Journal of the American Statistical Association* 90, 1137–1145.
- Esponda, I., 2008. Behavioral equilibrium in economies with adverse selection. *The American Economic Review* 98, 1269–1291.
- Eyster, E., Rabin, M., 2010. Naive herding in rich-information settings. *American Economic Journal: Microeconomics* 2, 221–243.
- Fudenberg, D., 2006. Advancing beyond “advances in behavioral economics”. *Journal of Economic Literature* 44, 694–711.
- Fudenberg, D., Levine, D. K., 1998. *The Theory of Learning in Games*. Cambridge, MA: MIT Press.
- Goeree, J., Palfrey, T., Rogers, B., McKelvey, R., 2007. Self-correcting information cascades. *Review of Economic Studies* 74, 733–762.
- Grether, D. M., 1980. Bayes rule as a descriptive model: The representativeness heuristic. *Quarterly Journal of Economics* 95, 537–57.
- Grimm, V., Mengel, F., 2012. An experiment on learning in a multiple games environment. *Journal of Economic Theory* 147, 2220–59.

- Guarino, A., Jehiel, P., 2013. Social learning with coarse inference. *American Economic Journal: Microeconomics* 5, 147–174.
- Heal, G., Millner, A., 2014. Uncertainty and decision making in climate change economics. *Review of Environmental Economics and Policy* 8, 120–137.
- Hendricks, K., Kovenock, D., 1989. Asymmetric information, information externalities, and efficiency: The case of oil exploration. *The RAND Journal of Economics* 20, 164–182.
- Hertwig, R., Barron, G., Weber, E. U., Erev, I., 2004. Decisions from experience and the effect of rare events in risky choices. *Psychological Science* 15, 534–539.
- Hung, A. A., Plott, C. R., 2001. Information cascades: Replication and an extension to majority rule and conformity-rewarding institutions. *The American Economic Review* 91, 1508–1520.
- Iyer, R., Khwaja, A. I., Luttmer, E. F. P., Shue, K., 2015. Screening peers softly: Inferring the quality of small borrowers. *Management Science* forthcoming.
- Jackson, M. O., Kalai, E., 1997. Social learning in recurring games. *Games and Economic Behavior* 21, 102–134.
- Jehiel, P., 2005. Analogy-based expectation equilibrium. *Journal of Economic Theory* 123, 81–104.
- Jehiel, P., Ettinger, D., 2010. A theory of deception. *American Economic Journal: Microeconomics* 2, 1–20.
- Jehiel, P., Koessler, F., 2008. Revisiting games of incomplete information with analogy-based expectations. *Games and Economic Behavior* 62, 533–557.
- Kariv, S., 2005. Overconfidence and informational cascades. mimeo, UC Berkeley.
- Kübler, D., Weizsäcker, G., 2004. Limited depth of reasoning and failure of cascade formation in the laboratory. *Review of Economic Studies* 71, 425–441.
- Kübler, D., Weizsäcker, G., 2005. Are longer cascades more stable? *Journal of the European Economic Association* 3, 330–339.
- Kulkarni, G., Ratchford, B. T., Kannan, P., 2012. The impact of online and offline information sources on automobile choice behavior. *Journal of Interactive Marketing* 26, 167–175.
- Lee, I. H., 1993. On the convergence of informational cascades. *Journal of Economic Theory* 61, 395–411.

- Lin, M., Prabhala, N. R., Viswanathan, S., 2013. Judging borrowers by the company they keep: Friendship networks and information asymmetry in online peer-to-peer lending. *Management Science* 59, 17–35.
- March, C., Ziegelmeyer, A., 2009. Behavioral social learning. *Jena Economic Research Papers* 2009-105.
- March, C., Ziegelmeyer, A., 2015. Even experts follow large herds: Naïve herding in the laboratory. mimeo.
- Mengel, F., 2012. Learning across games. *Games and Economic Behavior* 74, 601–19.
- Milgrom, P., Roberts, J., 1991. Adaptive and sophisticated learning in normal form games. *Games and Economic Behavior* 3, 82–100.
- Mohlin, E., 2014. Optimal categorization. *Journal of Economic Theory* 152, 356–381.
- Moran, M., Everhart, K., Lucas, M., Prickett, E., Morgan, A., 2015. Why are anti-vaccine messages so persuasive? a content analysis of anti-vaccine websites to inform the development of vaccine promotion strategies. Paper presented at the annual meeting of the American Public Health Association, Chicago, IL.
- Moretti, E., 2011. Social learning and peer effects in consumption: Evidence from movie sales. *Review of Economic Studies* 78, 356–393.
- Oberlechner, T., Hocking, S., 2004. Information sources, news, and rumors in financial markets: Insights into the foreign exchange market. *Journal of Economic Psychology* 25, 407–424.
- Palfrey, T. R., Wang, S. W., 2012. Speculative overpricing in asset markets with information flows. *Econometrica* 80, 1937–1976.
- Park, A., Sabourian, H., 2011. Herding and contrarian behavior in financial markets. *Econometrica* 79, 973–1026.
- Ross, L. D., 1977. The intuitive psychologist and his shortcomings: Distortions in the attribution process. In: Berkowitz, L. (Ed.), *Advances in Experimental Social Psychology*. Vol. 10. Academic Press, New York, pp. 173–220.
- Samuelson, L., 2001. Analogies, adaptation, and anomalies. *Journal of Economic Theory* 97, 320–366.
- Smith, L., Sørensen, P., 2000. Pathological outcomes of observational learning. *Econometrica* 68(2), 371–398.
- Smith, L., Sørensen, P., 2008. Rational social learning with random sampling. Tech. rep., University of Copenhagen.

- Steiner, J., Stewart, C., 2008. Contagion through learning. *Theoretical Economics* 3, 431–458.
- Tan, T., Werlang, S., 1988. The bayesian foundations of solution concepts of games. *Journal of Economic Theory* 45, 370–391.
- Zhang, J., Liu, P., 2012. Rational herding in microloan markets. *Management Science* 58, 892–912.

Supplementary material for

ADAPTIVE SOCIAL LEARNING

Christoph March

Appendix A contains the proofs and complementary theoretical results. Appendix B reports simulation results for simple social-learning games which establish that adaptation across games can be optimal when experiences are scarce. Finally, Appendix C discusses generalizations of the adaptive process.

Appendix A. Proofs

A.1. Non-Bayesian Updating in the Simple Example

Recall that in the simple example $t \in \{A, B\}$ and $s_t \in \{0, 1\}$ where $\Pr(\tilde{s}_t = 1 \mid \tilde{\theta} = 1) = \Pr(\tilde{s}_t = 0 \mid \tilde{\theta} = 0) = q_t \in (\frac{1}{2}, 1)$. Let $k \in \{L, H\}$ and denote by q_t^k the signal quality in game k . Assume that $q_A^L < q_A^H$ without loss of generality. Finally, $\pi = \pi_L = 1 - \pi_H$ denotes the likelihood to encounter game L .

Proposition A1. *There exist thresholds $\underline{\beta}_B, \bar{\beta}_B$ such that Bob's optimal updating strategy β_B^* satisfies*

- (i) $\beta_B^* > \bar{\beta}_B > 1$ if $q_A^L < q_B^L < \bar{q}_A < \bar{q}_B < q_B^H$,
- (ii) $1 < \underline{\beta}_B < \beta_B^* < \bar{\beta}_B$ if $q_B^H < q_B^L$ and $q_A^L < q_B^L < \bar{q}_A$,
- (iii) $\beta_B^* > \bar{\beta}_B$ where $\bar{\beta}_B \leq 1$ if $q_A^L < \bar{q}_A \leq q_B^L < q_B^H$, or $q_A^L < q_A^H < q_B^H < q_B^L$,
- (iv) $\beta_B^* < \underline{\beta}_B$ where $\underline{\beta}_B < 1$ if $q_B^L < \bar{q}_B < \bar{q}_A < q_B^H < q_A^H$,
- (v) $\underline{\beta}_B < \beta_B^* < \bar{\beta}_B < 1$ if $q_B^H < q_B^L$ and $\bar{q}_A < q_B^H < q_A^H$,
- (vi) $\beta_B^* < \underline{\beta}_B$ where $\underline{\beta}_B \geq 1$ if $q_B^L < \bar{q}_B < q_B^H \leq \bar{q}_A < q_A^H$, or $q_B^H < q_B^L < q_A^L < q_A^H$,
- (vii) $\underline{\beta}_B < \beta_B^* < \bar{\beta}_B$ where $\underline{\beta}_B < 1 < \bar{\beta}_B$ if $q_B^L < q_A^L < q_A^H < q_B^H$ or $q_B^H \leq \bar{q}_A \leq q_B^L$.

Proof. In the long-run Bob's assessments when adapting across games satisfy $\bar{\varphi}_B((1) \mid 1) = \bar{\varphi}_B((0) \mid 0) = \bar{q}_A$. Accordingly, there exist three possibilities:

- (S1) Bob follows private information at each history in each game, if $\left(\frac{q_B^k}{1-q_B^k}\right)^{\beta_B} > \frac{\bar{q}_A}{1-\bar{q}_A}$ or equivalently if

$$\beta_B > \log\left(\frac{\bar{q}_A}{1-\bar{q}_A}\right) / \log\left(\frac{q_B^k}{1-q_B^k}\right)$$

for each $k \in \{L, H\}$. His expected payoff from this strategy is given by $\frac{1}{4}(2\bar{q}_B - 1)$.

- (S2) Bob imitates Anna's decision in each game, if $\left(\frac{q_B^k}{1-q_B^k}\right)^{\beta_B} < \frac{\bar{q}_A}{1-\bar{q}_A}$ or equivalently if

$$\beta_B < \log\left(\frac{\bar{q}_A}{1-\bar{q}_A}\right) / \log\left(\frac{q_B^k}{1-q_B^k}\right)$$

for each $k \in \{L, H\}$. His expected payoff from this strategy is given by $\frac{1}{4}(2\bar{q}_A - 1)$.

- (S3) Bob follows private information at each history in game k_{max} and imitates Anna's decision in game k_{min} where $q_B^{k_{min}} < q_B^{k_{max}}$, if

$$\log\left(\frac{\bar{q}_A}{1-\bar{q}_A}\right) / \log\left(\frac{q_B^{k_{max}}}{1-q_B^{k_{max}}}\right) < \beta_B < \log\left(\frac{\bar{q}_A}{1-\bar{q}_A}\right) / \log\left(\frac{q_B^{k_{min}}}{1-q_B^{k_{min}}}\right).$$

His expected payoff from this strategy is given by $\frac{1}{4} [2 (\pi_{k_{max}} q_B^{k_{max}} + \pi_{k_{min}} q_A^{k_{min}}) - 1]$.

Define

$$\underline{\beta}_B \equiv \log \left(\frac{\bar{q}_A}{1 - \bar{q}_A} \right) / \max_k \log \left(\frac{q_B^k}{1 - q_B^k} \right) \quad \text{and} \quad \bar{\beta}_B \equiv \log \left(\frac{\bar{q}_A}{1 - \bar{q}_A} \right) / \min_k \log \left(\frac{q_B^k}{1 - q_B^k} \right),$$

and note that $\underline{\beta}^B < \bar{\beta}_B$.

Let $q_B^L < q_B^H$.

First, strategy (S1) outperforms strategy (S2) if $\bar{q}_B > \bar{q}_A$ and it outperforms (S3) if $q_B^L > q_A^L$. Hence, $\beta_B^* > \bar{\beta}_B$ under the two conditions. As $\bar{\beta}_B > 1$ if and only if $\bar{q}_A > \min_k q_B^k = q_B^L$, (i) and the first case of (iii) follows.

Second, (S2) outperforms (S1) if $\bar{q}_B < \bar{q}_A$ and it outperforms (S3) if $q_B^H < q_A^H$. Hence, $\beta_B^* < \underline{\beta}_B$ under these two conditions. As $\underline{\beta}_B < 1$ if and only if $\bar{q}_A < q_B^H$, (iv) and the first case of (vi) follows.

Finally, (S3) is optimal if $q_B^L < q_A^L < q_A^H < q_B^H$ which implies $q_B^L < \bar{q}_A < q_B^H$. Therefore, $\underline{\beta}_B < \beta_B^* < \bar{\beta}_B$ where $\underline{\beta}_B < 1 < \bar{\beta}_H$ which yields the first case of (vii).

Let $q_B^L > q_B^H$.

First, (S1) is optimal if $q_B^H > q_A^H$ which implies $\bar{q}_B > \bar{q}_A$ and $q_B^H > \bar{q}_A$. Hence, $\beta_B^* > \bar{\beta}_B$ where $\bar{\beta}_B < 1$ which yields the second case of (iii).

Second, (S2) is optimal if $q_B^L < q_A^L$ which implies $\bar{q}_B < \bar{q}_A$ and $q_B^L < \bar{q}_A$. Hence, $\beta_B^* < \underline{\beta}_B$ where $\underline{\beta}_B > 1$ which yields the second case of (vi).

Finally, (S3) is optimal if $q_B^H < q_A^H$ and $q_B^L > q_A^L$. Moreover, $\underline{\beta}_B > 1$ if $q_B^L < \bar{q}_A$ which yields (ii) and $\bar{\beta}_B < 1$ if $q_B^H > \bar{q}_A$ which yields (v) and conversely the second case of (vii). \square

A.2. Proofs of Lemmas 1 and 2

A.2.1 Proof of Lemma 1

Using $dG_0(s)/dG_1(s) = (1-s)/s$ the ex ante expected payoff of individual t playing strategy σ_t and given strategies σ_{-t} can be written as

$$U_t(\sigma_t, \sigma_{-t}) = \frac{1}{4} \sum_{h_t \in H_t} \int_{\underline{b}}^{\bar{b}} \sigma_t(s_t, h_t) [s_t \Pr(h_t | \tilde{\theta} = 1, \sigma_{-t}) - (1-s_t) \Pr(h_t | \tilde{\theta} = 0, \sigma_{-t})] \frac{1}{s_t} dG_1(s_t). \quad (1)$$

with $\Pr(h_t | \tilde{\theta} = \theta, \sigma_{-t}) = \prod_{\tau < t} \int_{\underline{b}}^{\bar{b}} \sigma_\tau(a_\tau | s_\tau, h_\tau) dG_\theta(s_\tau)$ for each $h_t \in H_t$ and each $\theta \in \Theta$ where $a_\tau = h_t(\tau)$ and $h_\tau \subset h_t$ for each $\tau < t$. Since conditional probabilities of histories in period t only depend upon strategies σ_τ for $\tau < t$, the optimal strategy can be derived inductively.

In period 1, $U_1(\sigma_1, \sigma_{-1}) = \int_{s_1} \sigma_1(s_1, \emptyset) (2s_1 - 1) dG_1(s_1) / (4s_1)$ for each $\sigma_{-1} \in \Sigma_{-1}$ since $H_1 = \{\emptyset\}$ and $\Pr(\emptyset | \theta) = 1$ for each $\theta \in \Theta$. Therefore, any undominated strategy σ_1^* satisfies $\sigma_1^*(s_1, \emptyset) = 1$ (0) if $s_1 > (<) 1/2$. A tie occurs if $s_1 = 1/2$ (the expected payoff at $s_1 = 1/2$ is 0 for each $a \in \{0, 1\}$) which is why any assignment $\sigma_1^*(\frac{1}{2}, \emptyset) \in [0, 1]$ is feasible. This only occurs if the private signal distribution has an atom at $1/2$. The set of private signal distributions with this property is a null set.

Consider period $t \geq 2$ and assume that strategies σ_τ^* for $\tau < t$ are iteratively undominated and obey Bayes' rule and sequential rationality at all histories occurring with strictly positive probability. Assume furthermore that for any $\tau < t$ and any $h_\tau \in H_\tau$ there is at most one private signal $\hat{s}_\tau(h_\tau)$ at which a tie occurs. Under this assumption, ties have measure zero since they can be ruled out by shifting atoms of the private signal distribution at a finite number of points by some $\epsilon > 0$. It is then easily seen that any σ_t^* which is (iteratively) undominated under the two assumptions must also obey Bayes' rule and sequential rationality. Notice first that absent ties probabilities $\Pr(h_t | \tilde{\theta} = \theta, \sigma_{-t}^*)$ are uniquely defined for each $h_t \in H_t$ and each $\theta \in \Theta$. If $\Pr(h_t | \tilde{\theta} = \theta, \sigma_{-t}^*) > 0$ for each $\theta \in \Theta$ it follows from equation (1) that σ_t^* must satisfy

$$\sigma_t^*(s_t, h_t) = \begin{cases} 1 & \text{if } s_t \Pr(h_t | 1, \sigma_{-t}^*) > (1-s_t) \Pr(h_t | 0, \sigma_{-t}^*) \\ 1 & \text{if } s_t \Pr(h_t | 1, \sigma_{-t}^*) < (1-s_t) \Pr(h_t | 0, \sigma_{-t}^*) \end{cases}$$

or equivalently $\sigma_t^*(s_t, h_t) = 1$ (0) if $b_t^*(s_t, h_t) > (<) \frac{1}{2}$ where

$$b_t^*(s_t, h_t) = \frac{s_t \Pr(h_t | \tilde{\theta} = 1, \sigma_{-t}^*)}{s_t \Pr(h_t | \tilde{\theta} = 1, \sigma_{-t}^*) + (1-s_t) \Pr(h_t | \tilde{\theta} = 0, \sigma_{-t}^*)}.$$

Furthermore, for such histories a tie may only occur at the unique private signal $\hat{s}_t(h_t)$ defined via $\frac{\hat{s}_t(h_t)}{1-\hat{s}_t(h_t)} = \frac{\Pr(h_t | 0, \sigma_{-t}^*)}{\Pr(h_t | 1, \sigma_{-t}^*)}$. If $\Pr(h_t | \tilde{\theta} = \theta, \sigma_{-t}^*) = 0$ for each $\theta \in \Theta$, any $\sigma_t^*(s_t, h_t) \in [0, 1]$ is feasible and satisfies Bayes' rule and sequential rationality. However, since

$\Pr(h_\ell \mid \tilde{\theta} = \theta, \sigma_{-\ell}^*) = 0$ for each $\ell > t$, each $h_\ell \supset h_t$, each $\sigma_{-\ell}^* \subset \sigma_{-t}^*$, and each $\theta \in \Theta$, the choice of $\sigma_t^*(s_t, h_t)$ at such histories does not affect the outcome of the game. \square

A.2.2 Proof of Lemma 2

Fix game $k \in \mathcal{K}$. For notational convenience, I omit the superscript k henceforth.

Ad. (i): Fix t and consider $h_t \in H_t$ such that $\hat{\varphi}_t(h_t \mid \theta) > 0$ for each $\theta \in \Theta$. σ_t plays ϵ -like $\hat{\sigma}_t$ at h_t if there exists $B_\epsilon \subseteq [\underline{b}, \bar{b}]$ such that $G_\theta(B_\epsilon) > 1 - \epsilon$ for each $\theta \in \Theta$, and $|\sigma_t(s_t, h_t) - \hat{\sigma}_t(s_t, h_t)| < \epsilon$ for each $s_t \in B_\epsilon$. Since σ_t and $\hat{\sigma}_t$ are cutoff strategies, i.e. $\sigma_t(s_t, h_t) = 1$ (0) if $s_t > c_t(h_t)$ (where

$$c_t(h_t) = \frac{\varphi_t(h_t \mid 0)}{\varphi_t(h_t \mid 1) + \varphi_t(h_t \mid 0)}$$

and similarly for $\hat{\sigma}_t$, there exists $\epsilon_1 > 0$ such that σ_t plays ϵ -like $\hat{\sigma}_t$ at h_t if $|c_t(h_t) - \hat{c}_t(h_t)| < \epsilon_1$ (ties are ruled out). Let $|\varphi_t(h_t \mid \theta) - \hat{\varphi}_t(h_t \mid \theta)| < \delta < \min_\theta \hat{\varphi}_t(h_t \mid \theta)$ for each $\theta \in \Theta$. Then $\varphi_t(h_t \mid \theta) > 0$ for each $\theta \in \Theta$ and

$$\begin{aligned} |c_t(h_t) - \hat{c}_t(h_t)| &= \left| \frac{\varphi_t(h_t \mid 0)}{\varphi_t(h_t \mid 1) + \varphi_t(h_t \mid 0)} - \frac{\hat{\varphi}_t(h_t \mid 0)}{\hat{\varphi}_t(h_t \mid 1) + \hat{\varphi}_t(h_t \mid 0)} \right| \\ &\leq \frac{\varphi_t(h_t \mid 0) |\hat{\varphi}_t(h_t \mid 1) - \varphi_t(h_t \mid 1)| + \varphi_t(h_t \mid 1) |\varphi_t(h_t \mid 0) - \hat{\varphi}_t(h_t \mid 0)|}{[\varphi_t(h_t \mid 1) + \varphi_t(h_t \mid 0)] * [\hat{\varphi}_t(h_t \mid 1) + \hat{\varphi}_t(h_t \mid 0)]} \\ &< \frac{\delta}{\hat{\varphi}_t(h_t \mid 1) + \hat{\varphi}_t(h_t \mid 0)}. \end{aligned}$$

Accordingly, since $\hat{\varphi}_t(h_t \mid \theta) > 0$ for each $\theta \in \Theta$, σ_t plays ϵ -like $\hat{\sigma}_t$ at h_t if φ_t is δ -close to $\hat{\varphi}_t$ for some $\delta < \min\{\hat{\varphi}_t(h_t \mid 1), \hat{\varphi}_t(h_t \mid 0), \epsilon_1 * [\hat{\varphi}_t(h_t \mid 1) + \hat{\varphi}_t(h_t \mid 0)]\}$.

Ad. (ii): Assume first that $\hat{\varphi}_t(h_t \mid \theta) > 0$ for each $\theta \in \Theta$ and therefore $\hat{\varphi}_\tau(h_\tau \mid \theta) > 0$ for each $\theta \in \Theta$ and each $\tau < t$ where $h_\tau \subset h_t$. The proof is by induction. By definition $\varphi_1(h_1 \mid \theta) = \hat{\varphi}_1(h_1 \mid \theta) = 1$ for each $\theta \in \Theta$. Assume therefore that for each $\tau < t$, σ_τ plays δ -like $\hat{\sigma}_\tau$ at each $h_\tau \subset h_t$ and φ_τ is $\frac{\epsilon}{2}$ -close to $\hat{\varphi}_\tau$. Let $x_\tau = \int_{\underline{b}}^{\bar{b}} \sigma_\tau(a_\tau \mid s_\tau, h_\tau) dG_\theta(s_\tau)$ where $a_\tau \in h_t$ and $h_\tau \subset h_t$ for each $\tau < t$ and define \hat{x}_τ accordingly. Hence,

$$|\varphi_t(h_t \mid \theta) - \hat{\varphi}_t(h_t \mid \theta)| = \left| \prod_{\tau < t} x_\tau - \prod_{\tau < t} \hat{x}_\tau \right| \leq x_{t-1} \left| \prod_{\tau < t-1} x_\tau - \prod_{\tau < t-1} \hat{x}_\tau \right| + |x_{t-1} - \hat{x}_{t-1}| \prod_{\tau < t-1} \hat{x}_\tau.$$

By induction assumption, $\left| \prod_{\tau < t-1} x_\tau - \prod_{\tau < t-1} \hat{x}_\tau \right| < \frac{\epsilon}{2}$. Furthermore there exists $0 < \delta' < \delta$ such that

$$|x_{t-1} - \hat{x}_{t-1}| \leq \delta' G_\theta(B_{\delta'}) + \delta' \leq 2\delta' < \frac{\epsilon}{2}$$

if σ_{t-1} plays δ' -close to $\hat{\sigma}_{t-1}$ at $h_{t-1} \subset h_t$. Since $x_\tau, \hat{x}_\tau < 1$, $|\varphi_t(h_t \mid \theta) - \hat{\varphi}_t(h_t \mid \theta)| < \epsilon$.

On the other hand, $\hat{\varphi}_t(h_t | \theta) = 0$ implies existence of a unique $\tau_0 \leq t$ such that $\hat{\varphi}_\tau(h_\tau | \theta) > 0$ for each $\tau < \tau_0$ and $\hat{\varphi}_{\tau_0}(h_{\tau_0} | \theta) = 0$ where $h_\tau, h_{\tau_0} \subset h_t$. Therefore,

$$\hat{x}_{\tau_0} = \int_{\underline{b}}^{\bar{b}} \hat{\sigma}_{\tau_0}(a_{\tau_0} | s_{\tau_0}, h_{\tau_0}) dG_\theta(s_{\tau_0}) = 0$$

which implies that $\hat{\sigma}_{\tau_0}(a_{\tau_0} | s_{\tau_0}, h_{\tau_0}) = 0$ for almost any $s_{\tau_0} \in [\underline{b}, \bar{b}]$. Accordingly, if σ_{τ_0} plays $\frac{\epsilon}{2}$ -like $\hat{\sigma}_{\tau_0}$ at $h_{\tau_0} \subset h_t$, $x_{\tau_0} = \int_{\underline{b}}^{\bar{b}} \sigma_{\tau_0}(a_{\tau_0} | s_{\tau_0}, h_{\tau_0}) dG_\theta(s_{\tau_0}) < \frac{\epsilon}{2} G_\theta(B_{\epsilon/2}) + \frac{\epsilon}{2} < \epsilon$ and therefore $\varphi_t(h_t | \theta) < \epsilon$. □

A.3. Proofs of Propositions 1 and 2

Proof of Proposition 1

Since games can be distinguished, it suffices to focus on a single game. The proof is by induction: In period 1, there is no need to draw inferences from others' actions. Therefore, assessments are equal to rational assessments by definition and play of the dominant strategy follows from the assumption of myopic Bayes-rational strategic responses. By the strong law of large numbers for conditional expectations (e.g. Walk, 2008), assessments in the second period eventually become vanishingly close to rational assessments. Lemma 2 thus implies that strategies eventually play ϵ -like any iteratively undominated strategy at all period 2-histories occurring with strictly positive probability (recall that absent ties iterated dominance uniquely defines probabilities of histories, and behavior at histories occurring with strictly positive probability). This argument can be inductively extended to all subsequent periods $t > 2$. □

Proof of Proposition 2

Fix $k \in \mathcal{K}$ such that $G_0^k(1/2)/G_1^k(1/2) < \bar{b}_k/(1 - \bar{b}_k)$. Assume that during the early rounds in which game k is played the individual deciding in period 1 always rejects and the realized state is 0 sufficiently often such that

$$\frac{\varphi_2^{k,r}(a_1 = 0 | 0; \zeta^r)}{\varphi_2^{k,r}(a_1 = 0 | 1; \zeta^r)} > \frac{\bar{b}^k}{1 - \bar{b}^k}. \quad (2)$$

Obviously, the set of corresponding adaptation paths has strictly positive probability. Equation (2) implies that no individual will invest after $h_2 = (0)$ in any subsequent round. Accordingly, the state is not revealed and individuals can never revise their wrong assessments for this history. Similar adaptation paths may easily be constructed. □

A.4. Definition of ABEE and Proofs of Propositions 3–5

A.4.1 Analogy-based Expectations Equilibrium

The long-run outcome of adaptation across games is captured by an *analogy-based expectations equilibrium* (Jehiel, 2005; Jehiel and Ettinger, 2010, ABEE henceforth). In an ABEE, each player partitions the decision nodes of others into *analogy classes* and has a correct understanding of average behavior in each class. For the family of social-learning games, a decision node is a tuple $(k, \theta, \{s_t\}_{t=1}^T, h_t) \in \mathcal{K} \times \Theta \times [0, 1]^T \times H$. I focus on the specific analogy partition $\mathcal{A} = \{\alpha(h_t, \theta) : h_t \in H \wedge \theta \in \Theta\}$ where $\alpha(h_t, \theta)$ bundles all decision nodes at which (h_t, θ) appears and I refer to this as the *information-anonymous analogy partition*. A player with information-anonymous analogy partition ignores the dependence of behavior upon the distribution of private signals and the private signal realizations when assessing the informational content of a history. Let \mathcal{A}_t denote the subset of analogy classes for period t . An *analogy based expectation* $\bar{\sigma}_t : \mathcal{A}_t \rightarrow \Delta(A)$ assigns to each analogy class $\alpha_t \in \mathcal{A}_t$ a probability distribution over actions. In equilibrium, players sequentially best respond to analogy-based expectations and expectations are consistent with behavior.

Definition A1. *The strategies $\sigma_t^{k,A}$ for $k \in \mathcal{K}$ and $t = 1, \dots, T$ constitute an ABEE with information-anonymous analogy partition if and only if there exist analogy-based expectations $\bar{\sigma} = \{\bar{\sigma}_t\}_{t=1}^T$ such that for each $t = 1, \dots, T$*

(A1) $\sigma_t^{k,A}$ is a best response to $\bar{\sigma}$ for each $k \in \mathcal{K}$, i.e.

$$\sigma_t^{k,A}(s_t^k, h_t) = 1 \text{ (0)} \quad \text{if} \quad \frac{s_t^k}{1 - s_t^k} > (<) \prod_{\tau < t} \frac{\bar{\sigma}_\tau(a_\tau | \alpha(0, h_\tau))}{\bar{\sigma}_\tau(a_\tau | \alpha(1, h_\tau))}$$

for each $s_t^k \in [\underline{b}^k, \bar{b}^k]$ and each $h_t \in H_t$ where $a_\tau = h_t(\tau)$ and $h_\tau \subset h_t$ for each $\tau < t$,

(A2) $\bar{\sigma}$ is consistent with $\{\sigma_t^{k,A}\}$, i.e. for each $\theta \in \Theta$ and each $h_t \in H_t$,

$$\bar{\sigma}_t(h_t, \theta) \equiv \bar{\sigma}_t(\alpha(h_t, \theta)) = \sum_{k=1}^K \int_{\text{supp}(G_\theta^k)} \nu^A(k, ds_t | h_t, \theta) \sigma_t^{k,A}(s_t, h_t)$$

where ν^A is the distribution on tuples (k, s_t, h_t, θ) induced by the strategies $\sigma_t^{k,A}$ and the fundamentals.

Lemma A1 characterizes the *analogy-based assessments* $\bar{\varphi}_t(h_t | \theta)$ and shows that the ABEE derives from iterated elimination of dominated strategies when expected payoffs are defined with respect to analogy-based expectations. Except for a null set of parameters, the ABEE strategies $\sigma_t^{k,A}$ are uniquely defined at histories satisfying $\bar{\varphi}_t(h_t | \theta) > 0$ for each $\theta \in \Theta$.

Lemma A1. *In any ABEE with information-anonymous analogy partition the analogy-based assessments $\bar{\varphi}_t(h_t | \theta) \equiv \prod_{\tau < t} \bar{\sigma}_\tau(a_\tau | \alpha(\theta, h_\tau))$ satisfy for each $t = 1, \dots, T$, each $h_t \in H_t$, and each $\theta \in \Theta$*

$$\bar{\varphi}_t(h_t | \theta) = \sum_{k=1}^K \pi_k \varphi_t^{k,A}(h_t | \theta) = \sum_{k=1}^K \pi_k \prod_{\tau < t} \int_{\underline{b}^k}^{\bar{b}^k} \sigma_\tau^{k,A}(a_\tau | s_\tau, h_\tau) dG_\theta^k(s_\tau) \quad (3)$$

where $a_\tau = h_t(\tau)$ and $h_\tau \subset h_t$ for each $\tau < t$. Furthermore, the set of ABEE is the set of strategy profiles which are iteratively undominated with respect to the analogy-based expectations and the ABEE outcome is almost always unique.

Proof. Let $\sigma_t^{k,A}$ denote the ABEE strategies. I first establish that the analogy-based assessments satisfy for each $t = 1, \dots, T$, each $h_t \in H_t$, and each $\theta \in \Theta$

$$\bar{\varphi}_t(h_t | \theta) \equiv \prod_{\tau < t} \bar{\sigma}_\tau(a_\tau | \alpha(\theta, h_\tau)) = \sum_{k=1}^K \pi_k \varphi_t^{k,A}(h_t | \theta) \quad (4)$$

where for each $k \in \mathcal{K}$, each $t = 1, \dots, T$, and each $h_t \in H_t$

$$\varphi_t^{k,A}(h_t | \theta) = \prod_{\tau < t} \int_{\underline{b}^k}^{\bar{b}^k} \sigma_\tau^{k,A}(a_\tau | s_\tau^k, h_\tau) dG_\theta^k(s_\tau^k)$$

with $a_\tau = h_t(\tau)$ and $h_\tau \subset h_t$ for each $\tau < t$. For each $k \in \mathcal{K}$, each $t = 1, \dots, T$, each $h_t \in H_t$, and each $\theta \in \Theta$, $\varphi_t^{k,A}(h_t | \theta)$ is the equilibrium probability that history h_t arises in game k if the state is θ . Accordingly, $\nu^A(k, \theta, s_t^k, h_t) = \frac{1}{2} \pi_k \varphi_t^{k,A}(h_t | \theta) dG_\theta^k(s_t^k)$ and

$$\nu^A(k, s_t^k | h_t, \theta) = \frac{\pi_k dG_\theta^k(s_t^k) \bar{\varphi}_t^k(h_t | \theta)}{\sum_{k'=1}^K \pi_{k'} \bar{\varphi}_t^{k'}(h_t | \theta) \int_{\underline{b}^k}^{\bar{b}^k} dG_\theta^k(s)} = \frac{\pi_k \bar{\varphi}_t^k(h_t | \theta)}{\sum_{k'=1}^K \pi_{k'} \bar{\varphi}_t^{k'}(h_t | \theta)} dG_\theta^k(s_t^k).$$

Therefore for each $t = 1, \dots, T$

$$\bar{\sigma}_t(h_t, \theta) = \sum_{k=1}^K \int_{\underline{b}^k}^{\bar{b}^k} \frac{\pi_k \varphi_t^{k,A}(h_t | \theta)}{\sum_{\ell=1}^K \pi_\ell \varphi_t^{\ell,A}(h_t | \theta)} dG_\theta^k(s_t^k) \sigma_t^{k,A}(s_t^k, h_t).$$

I proceed by induction. First, in period 1,

$$1 = \bar{\varphi}_1(h_1 | \theta) = \sum_{k=1}^K \pi_k \varphi_1^{k,A}(h_1 | \theta) = \sum_{k=1}^K \pi_k = 1.$$

Assume (4) holds for all $\tau \leq t$. It follows for each $h_t \in H_t$, each $a_t \in A$, and each $\theta \in \Theta$

$$\begin{aligned}\bar{\varphi}_{t+1}((h_t, a_t) | \theta) &= \bar{\varphi}_t(h_t | \theta) * \bar{\sigma}_t(a_t | h_t, \theta) \\ &= \left[\sum_{k=1}^K \pi_k \varphi_t^{k,A}(h_t | \theta) \right] \cdot \left[\sum_{j=1}^K \int_{\underline{b}^j}^{\bar{b}^j} \frac{\pi_j \varphi_t^{j,A}(h_t | \theta)}{\sum_{\ell=1}^K \pi_\ell \varphi_t^{\ell,A}(h_t | \theta)} dG_\theta^j(s_t^j) \sigma_t^{j,A}(a_t | s_t^j, h_t) \right] \\ &= \sum_{j=1}^K \pi_j \varphi_t^{j,A}(h_t | \theta) \int_{\underline{b}^j}^{\bar{b}^j} \sigma_t^{j,A}(a_t | s_t^j, h_t) dG_\theta^j(s_t^j) = \sum_{j=1}^K \pi_j \bar{\varphi}_{t+1}^j((h_t, a_t) | \theta)\end{aligned}$$

where the second equality uses the induction assumption.

Second, I show that any ABEE is the outcome of iterated elimination of strategies which are strictly dominated with respect to the analogy-based expectations. Given the profile of analogy-based expectations $\bar{\sigma}$, the ex-ante expected payoff of strategy σ_t^k in game $k \in \mathcal{K}$ is given by

$$U_t^k(\sigma_t^k, \bar{\sigma}) = \frac{1}{2} \sum_{h_t \in H_t} \int_{\underline{b}^k}^{\bar{b}^k} \sigma_t^k(s_t^k, h_t) \left[s_t^k \bar{\varphi}_t(h_t | 1) - (1 - s_t^k) \bar{\varphi}_t(h_t | 0) \right] \frac{dG_1^k(s_t^k)}{s_t^k}.$$

Therefore, σ_t^k best responds to $\bar{\sigma}$, if and only if

$$\sigma_t^k(s_t^k, h_t) = 1(0) \quad \text{if} \quad \frac{s_t^k}{1 - s_t^k} > (<), \frac{\bar{\varphi}_t(h_t | 0)}{\bar{\varphi}_t(h_t | 1)}.$$

As the analogy-based assessments in period t only depend upon strategies σ_τ^k for $\tau < t$, the ABEE can be derived inductively: There is a dominant strategy in period 1 (as $\bar{\varphi}_1(h_1 | \theta) = 1$ for each $\theta \in \Theta$). Therefore, analogy-based assessments in period 2 are uniquely determined, and there is a unique best response in period 2 etc. Moreover, for each $t = 1, \dots, T$, each $h_t \in H_t$ and each $k \in \mathcal{K}$ there exists a unique private signal at which a tie occurs. Therefore, the families of social-learning games in which a tie arises constitute a null set. \square

A.4.2 Proof of Proposition 3

The proof is by induction, similar to the proof of Proposition 1. In period 1, for each $k \in \mathcal{K}$ and each $r \geq 1$, $\hat{\varphi}_1^{k,r}(h_1 | \theta) = 1$ holds for each $\theta \in \Theta$ and implies that $\sigma_t^{k,r}$ coincides with the ABEE strategy.

Fix a period t and assume that for each $\delta > 0$ there exists r^* such that strategic responses $\sigma_\tau^{k,r}$ play δ -like the ABEE strategies $\sigma_\tau^{k,A}$ at each history $h_\tau \in H_\tau$ for each period $\tau < t$, each game $k \in \mathcal{K}$, and each round $r > r^*$. For r sufficiently large, assessments for history

$h_t \in H_t$ and each game $k \in \mathcal{K}$ satisfy

$$\begin{aligned}\varphi_t^{k,r}(h_t | \theta; \zeta_r) &= \frac{\sum_{k=1}^K \kappa_t^{k,r}(h_t, \theta | \zeta^r)}{\sum_{h'_t \in H_t} \sum_{\ell \in \mathcal{K}} \kappa_t^{\ell,r}(h'_t, \theta | \zeta^r)} \\ &= \sum_{k=1}^K \frac{\kappa_t^{k,r}(h_t, \theta | \zeta^r)}{\sum_{\ell \in \mathcal{K}} \sum_{h'_t \in H_t} \kappa_t^{\ell,r}(h'_t, \theta | \zeta^r)} \frac{\sum_{h'_t \in H_t} \kappa_t^{k,r}(h'_t, \theta | \zeta^r)}{\sum_{h'_t \in H_t} \kappa_t^{k,r}(h'_t, \theta | \zeta^r)} \\ &= \sum_{k=1}^K \frac{\kappa_t^{k,r}(h_t, \theta | \zeta^r)}{\sum_{h'_t \in H_t} \kappa_t^{k,r}(h'_t, \theta | \zeta^r)} \frac{\sum_{h'_t \in H_t} \kappa_t^{k,r}(h'_t, \theta | \zeta^r)}{\sum_{\ell \in \mathcal{K}} \sum_{h'_t \in H_t} \kappa_t^{\ell,r}(h'_t, \theta | \zeta^r)}.\end{aligned}$$

Using the induction assumption and applying the second part of Lemma 2 implies that the assessment for game k

$$\frac{\kappa_t^{k,r}(h_t, \theta | \zeta^r)}{\sum_{h'_t} \kappa_t^{k,r}(h'_t, \theta | \zeta^r)}$$

is ϵ -close to $\varphi_t^{k,A}(h_t | \theta)$ provided r is sufficiently large. On the other hand,

$$\frac{\sum_{h'_t} \kappa_t^{k,r}(h'_t, \theta | \zeta^r)}{\sum_{\ell \in \mathcal{K}} \sum_{h'_t} \kappa_t^{\ell,r}(h'_t, \theta | \zeta^r)}$$

converges to π_k since the game and the state are drawn independently in each round. Thus, by the strong law of large numbers for conditional expectations $\hat{\varphi}_t^{k,r}(h_t | \theta; \zeta_r)$ eventually becomes ϵ -close to $\bar{\varphi}_t(h_t | \theta) = \sum_{k \in \mathcal{K}} \pi_k \varphi_t^{k,A}(h_t | \theta)$, the analogy-based assessments. By the first part of Lemma 2, $\sigma_t^{k,r}$ eventually plays ϵ -like the ABEE strategy $\sigma_t^{k,A}$ for each game k and at each history h_t that satisfies $\bar{\varphi}_t(h_t | \theta) > 0$ for each $\theta \in \Theta$. \square

A.4.3 Proof of Proposition 4

I give two generic examples for $K = 2$. Let $\pi = \pi_1 = 1 - \pi_2$ where $0 < \pi < 1$.

Assume first that the distributions of private signals are continuous for each $k = 1, 2$ and satisfy

$$\frac{G_0^1(\frac{1}{2})}{G_1^1(\frac{1}{2})} < \frac{G_0^2(\frac{1}{2})}{G_1^2(\frac{1}{2})}. \quad (5)$$

Consider history $h_2 = (0)$. ABEE assessments at h_2 are given by $\bar{\varphi}_2(0 | \theta) = \pi G_\theta^1(\frac{1}{2}) + (1 - \pi) G_\theta^2(\frac{1}{2})$ and the ABEE strategies satisfy

$$\sigma_2^{k,A}(s_2, h_2) = 1(0) \quad \text{if} \quad s_2 > (<) c_2^A(h_2) \equiv \frac{\pi G_0^1(\frac{1}{2}) + (1 - \pi) G_0^2(\frac{1}{2})}{\pi [G_1^1(\frac{1}{2}) + G_0^1(\frac{1}{2})] + (1 - \pi) [G_1^2(\frac{1}{2}) + G_0^2(\frac{1}{2})]}$$

for each game $k \in \{1, 2\}$. On the other hand, game-specific assessments satisfy $\varphi_2^k(0 | \theta) =$

$G_\theta^k(\frac{1}{2})$ for $k = 1, 2$ and the rational strategy for game k is given by

$$\sigma_2^{k,*}(s_2, h_2) = 1(0) \quad \text{if} \quad s_2 > (<) c_2^k(h_2) \equiv \frac{G_0^k(\frac{1}{2})}{G_1^k(\frac{1}{2}) + G_0^k(\frac{1}{2})}.$$

Assumption (5) implies $c_2^2(h_2) < c_2^A(h_2) < c_2^1(h_2)$. Therefore, since private signal distributions are continuous, individuals imitate the first player's rejection too often in game 1 and too seldom in game 2.

Second, the ex ante expected payoff from the family of social learning games for strategy σ_t that does not discriminate between games is given by

$$\begin{aligned} U_t(\sigma_t | \varphi_t^1, \dots, \varphi_t^K) &= \frac{1}{4} \sum_{k=1}^K \pi_k \sum_{h_t \in H_t} \int_{\underline{b}^k}^{\bar{k}} \sigma_t(s_t, h_t) [\varphi_t^k(h_t | 1) dG_1^k(s_t) - \varphi_t^k(h_t | 0) dG_0^k(s_t)] \\ &= \frac{1}{4} \sum_{h_t \in H_t} \int_{\min_k \underline{b}^k}^{\max_k \bar{k}} \sigma_t(s_t, h_t) \sum_{k=1}^K \pi_k [\varphi_t^k(h_t | 1) dG_1^k(s_t) - \varphi_t^k(h_t | 0) dG_0^k(s_t)]. \end{aligned}$$

Hence, the optimal strategy satisfies

$$\sigma_t^{opt}(s_t, h_t) = 1(0) \quad \text{if} \quad \sum_{k=1}^K \pi_k \varphi_t^k(h_t | 1) dG_1^k(s_t) > (<) \sum_{k=1}^K \pi_k \varphi_t^k(h_t | 0) dG_0^k(s_t).$$

Assume that private signals are continuously distributed on $[1 - a_k, a_k] \subset [0, 1]$ according to densities $g_0^k(s) = 2(1 - s)/(2a_k - 1)$ and $g_1^k(s) = 2s/(2a_k - 1)$ for each $k = 1, 2$ where without loss of generality $\frac{1}{2} < a_1 < a_2 \leq 1$. In this case, the optimal strategy can be written as

$$\sigma^{opt}(s_t, h_t) = 1(0) \quad \text{if} \quad s_t > (<) c_t^*(h_t) \equiv \frac{\sum_{k=1}^2 \frac{\pi_k}{2a_k - 1} \varphi_t^k(h_t | 0)}{\sum_{k=1}^2 \frac{\pi_k}{2a_k - 1} [\varphi_t^k(h_t | 1) + \varphi_t^k(h_t | 0)]}.$$

while the ABEE strategy is given by

$$\sigma_t^A(s_t, h_t) = 1(0) \quad \text{if} \quad s_t > (<) c_t^A(h_t) \equiv \frac{\sum_{k=1}^2 \pi_k \varphi_t^k(h_t | 0)}{\sum_{k=1}^2 \pi_k [\varphi_t^k(h_t | 1) + \varphi_t^k(h_t | 0)]}.$$

Since the distribution of private signals is continuous, those strategies agree at h_t if either $c_t^*(h_t), c_t^A(h_t) < 1 - a_2 < 1 - a_1$, or $c_t^*(h_t), c_t^A(h_t) > a_2 > a_1$, or $1 - a_2 < c_t^*(h_t) = c_t^A(h_t) < a_2$ where the latter is equivalent to

$$\frac{\varphi_t^2(h_t | 1) \varphi_t^1(h_t | 0) - \varphi_t^1(h_t | 1) \varphi_t^2(h_t | 0)}{2a_2 - 1} = \frac{\varphi_t^2(h_t | 1) \varphi_t^1(h_t | 0) - \varphi_t^1(h_t | 1) \varphi_t^2(h_t | 0)}{2a_1 - 1}$$

or $\varphi_t^2(h_t | 1) \cdot \varphi_t^1(h_t | 0) = \varphi_t^1(h_t | 1) \cdot \varphi_t^2(h_t | 0)$ since $a_1 \neq a_2$. Generically, neither of those conditions is satisfied. For instance at $h_2 = (1)$ (similarly for $h_2 = (0)$)

$$\varphi_2^k((2) | \theta) = 1 - G_\theta^k(1/2) = \begin{cases} \frac{1}{4} + \frac{a_k}{2} & \text{if } \theta = 1 \\ \frac{1}{4} + \frac{1-a_k}{2} & \text{if } \theta = 0 \end{cases}$$

and therefore $c_t^*(h_t) \neq c_t^A(h_t)$ if $a_2 \neq a_1$. Furthermore $1 - a_2 < c_t^A(h_t) < \frac{1}{2}$ since $a_2 > a_1 > \frac{1}{2}$. \square

A.4.4 Proof of Proposition 5

Since T is finite, there exists for each game $k = 1, \dots, K$ a maximal amount of private information that can be inferred from any history. Formally,

$$\lambda_{max}^k = \max_{t, h_t, \theta} \frac{\varphi_t^k(h_t | \theta)}{\varphi_t^k(h_t | 1 - \theta)}$$

is finite. Any individual with private signal quality $q_i / (1 - q_i) > \lambda_{max}^k$ or equivalently individual quality component

$$\nu_i > \max_{k \in \mathcal{K}} \bar{\nu}_i \equiv \log \left(\frac{\lambda_{max}^k - 1}{2} \right) - \bar{\rho}^k$$

finds it optimal to follow her private information at each history in each game. Since adaptation across games may prevent such an individual from following her optimal strategy, it is weakly optimal for her to overweight private information, i.e. $\beta^A(\nu_i) \geq 1$ for each $\nu_i > \bar{\nu}$.

On the other hand, there also exists a minimal amount of public information given by

$$\lambda_{min}^k = \min_{t, \theta} \min_{h_t \in \hat{H}_t} \frac{\varphi_t^k(h_t | \theta)}{\varphi_t^k(h_t | 1 - \theta)}$$

where for each t the minimum is taken over histories $h_t \in \hat{H}_t$ such that $\varphi_t^k(h_t | 1) \neq \varphi_t^k(h_t | 0)$. Therefore, $|\lambda_{min}^k|$ is also finite. An individual with quality component

$$\nu_i < \underline{\nu} \equiv \min_{k \in \mathcal{K}} \log \left((\lambda_{min}^k - 1) / 2 \right) - \bar{\rho}^k$$

finds it optimal to follow the public information at each history. Since adaptation across games may lead this individual to follow private information at some histories, it is weakly optimal for her to underweight private information, i.e. $\beta^A(\nu_i) \leq 1$ for each $\nu_i < \underline{\nu}$. \square

Appendix B. Adaptation Across Games and Scarce Experiences

In this appendix, I discuss the results of a simulation study. The adaptive process is simulated multiple times for simple social-learning games and a finite number of rounds. The simulation is run twice assuming (i) that games are distinguished and (ii) that players adapt across games and the resulting average payoffs across rounds are compared. The results illustrate that adaptation across games can be optimal when the number of repetitions of any of the social-learning games is small. Section B.1. presents the simulation framework. Section B.2. reports results for the simple example presented in Section 2 of the main text. Finally, Section B.3. presents results for social-learning games with symmetric, binary private signals and more than two players.

B.1. Simulation Framework

The setup of the simulation closely follows the definition of the adaptive process (Definition 1 in the main text). Denote by \mathcal{K} the (finite) set of social-learning games with (common) number of players T , and denote the (finite) number of rounds by R . Furthermore, let $\omega^r = (k^r, \theta^r, h_{T+1}^r)$ denote the outcome of the game in round $r \in \{1, \dots, R\}$ and let $\zeta^r = (\omega^1, \dots, \omega^{r-1})$ where $\zeta^1 = \emptyset$ by definition. I assume that $y_\Theta(\omega^r) = \{\theta^r\}$ and $y_H(\omega^r) = \{h_t \in h_{T+1}^r\}$ for each $r = 1, \dots, R$ and each $\omega^r \in \Omega$, i.e. feedback on the state of nature and the history of actions is complete in each round. Furthermore, $y_K(\omega^r t) \in \{\{k^r\}, \mathcal{K}\}$, i.e. players either fully distinguish games or they adapt across all games.

A single simulation run proceeds as follows: First, history-state frequencies are initialized as $\kappa_t^{k,1}(h_t, \theta \mid \zeta^1) = \eta > 0$ for each $t = 1, \dots, T$, each $h_t \in H_t$, and each $\theta \in \Theta \equiv \{0, 1\}$. Second, for each round $r = 1, \dots, R$

- (i) the game k^r , the state of nature θ^r , and the sequence of private signals $\{s_t^r\}_{t=1}^T$ are drawn according to the vector $\pi_K = (\pi_1, \dots, \pi_K)$, the uniform distribution on Θ , and the cumulative distribution function $G_{\theta^r}^{k^r}$, respectively;

- (ii) for each period $t = 1, \dots, T$

- (a) assessments are calculated for each $h_t \in H_t$ and each $\theta \in \Theta$ according to

$$\varphi_t^{k,r}(h_t \mid \theta; \zeta^r) = \frac{\kappa_t^{k,r}(h_t, \theta \mid \zeta^r)}{\sum_{h'_t \in H_t} \kappa_t^{k,r}(h'_t, \theta \mid \zeta^r)},$$

- (b) the action a_t^r is determined as $a_t^r = 1$ (0) if $\frac{s_t^r}{1-s_t^r} > (<) \frac{\varphi_t^r(h_t^r \mid 0; \zeta^r)}{\varphi_t^r(h_t^r \mid 1; \zeta^r)}$ where $h_1^r = \emptyset$ and $h_t^r = (a_1^r, \dots, a_{t-1}^r)$ for $t > 1$;

- (iii) frequencies are updated by setting $\kappa_t^{k,r+1}(h_t^r, \theta^r \mid \zeta^{r+1}) = \kappa_t^{k,r}(h_t^r, \theta^r \mid \zeta^r) + 1$ for each $t = 1, \dots, T$ and each $k \in y_{\mathcal{K}}(k^r, \theta^r, h_{T+1}^r)$.

Given the complete adaptation path $\zeta^{R+1} = (\omega^1, \dots, \omega^R) = (k^r, \theta^r, h_{T+1}^r)_{r=1}^R$, average payoffs for period $t \in \{1, \dots, T\}$ are given by

$$U_t^{y_{\mathcal{K}}}(R) = \frac{1}{R} \sum_{r=1}^R \hat{u}_t(a_t^r, \theta^r)$$

where I use the normalized payoff function

$$\hat{u}_t(a_t, \theta) = |a_t - \theta| = \frac{u(a_t, \theta) - \min_{a \in A} u(a, \theta)}{\max_{a \in A} u(a, \theta) - \min_{a \in A} u(a, \theta)}$$

to facilitate the interpretation of the results. Indeed, $U_t^{y_{\mathcal{K}}}(R)$ denotes the average frequency with which players select the ex post optimal action in period t across the R repetitions of the game. The superscript makes clear that the average payoffs depend on the feedback about the game $y_{\mathcal{K}}$. Similarly, the average payoff across periods is given by $U^{y_{\mathcal{K}}}(R) = \sum_t U_t^{y_{\mathcal{K}}}(R) / T$.

For a given family of social learning games, a simulation consists of N simulation runs for each $y_{\mathcal{K}} \in \{\{k^r\}, \mathcal{K}\}$, and all results rely on the mean average payoff $\bar{U}_t^{y_{\mathcal{K}}}(R)$ where the mean is taken across simulation runs. For all results reported below, $N = 50,000$, $\eta = 0.01$, $K = 2$, $\pi_{\mathcal{K}} = (\frac{1}{2}, \frac{1}{2})$, and $R \in \{1, \dots, 400\}$.¹

B.2. Results for the Simple Example

I first report results for the simple example discussed in section 2 of the main text. Hence, $T = 2$, $\mathcal{K} = \{L, H\}$, and signals in game $k \in \mathcal{K}$ are drawn from $S = \{0, 1\}$ according to probabilities $\Pr(\tilde{s}_t = 1 \mid \tilde{\theta} = 1, k) = \Pr(\tilde{s}_t = 0 \mid \tilde{\theta} = 0, k) = q_t^k$.²

Recall that Anna always follow her private signal. Accordingly, mean average payoffs for period 1 simply reflect the mean of a binomial distribution with success probability $(q_t^1 + q_t^2)/2$ regardless of the underlying feedback structure. I therefore focus on the mean average payoffs for period 2, $\bar{U}_2^{y_{\mathcal{K}}}(R)$.

Figure 1 illustrates the fundamental structure of the results for $q_A^L = 0.6$, $q_A^H = 0.8$, $q_B^L = 0.55$, and $q_B^H = 0.75$. The solid black (gray) line plots the mean average payoff $\bar{U}_2^{\mathcal{K}}(R)$ ($\bar{U}_2^{\{k^r\}}(R)$) for Bob when he adapts across games (by game). In addition, the dotted lines indicate the long-run payoffs $\lim_{R \rightarrow \infty} \bar{U}_2^{\mathcal{K}}(R)$ and $\lim_{R \rightarrow \infty} \bar{U}_2^{\{k^r\}}(R)$, respectively. The selected parameters imply that Bob will suboptimally follow his private signal in game

¹The simulation has been programmed in Mathematica and is available from the author upon request.

²Apparently, this signal structure does not match the signal structure of the general social learning game. This can be adjusted by identifying $s_t = 1$ with $\hat{s}_t = q_t^k$ and $s_t = 0$ with $\hat{s}_t = 1 - q_t^k$ in game k .

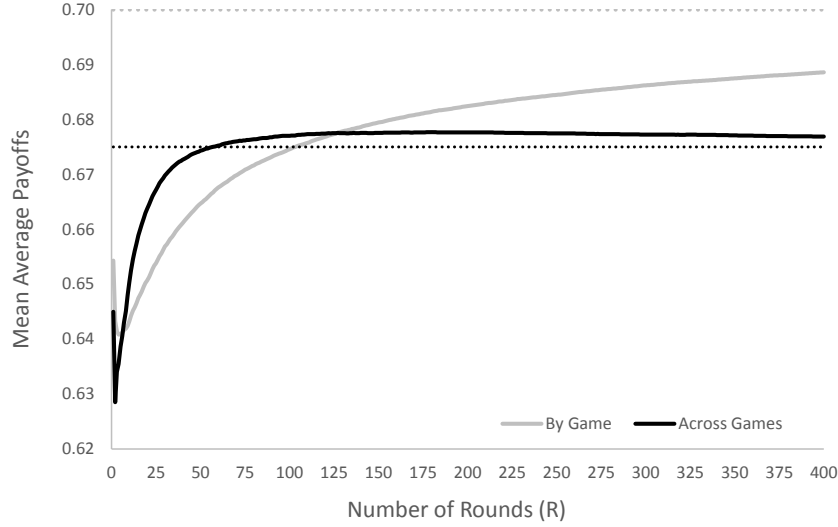


Figure 1: Mean Average Payoffs of Bob for $q_A^L = 0.6$, $q_A^H = 0.8$, $q_B^L = 0.55$, $q_B^H = 0.75$

H in the long-run if he adapts across games since $\bar{q}_A = 0.7 < q_B^H$. Still, adaptation across games is optimal for Bob if $7 \leq R \leq 126$.

The figure isolates three distinct intervals for the number of repetitions. First, adaptation across games is a worse strategy than distinguishing games if experience with any of the games is almost non-existent ($R < 7$). In this case, distinguishing games usually implies that Bob has no relevant feedback for one of the games which is why he resorts to the assumed default strategy of following his private signal. In contrast, Bob is for sure able to rely on feedback in each game by adapting across games, but such feedback can be grossly misleading. For instance, after as little as 5 or 6 repetitions there is a non-negligible probability that Bob anti-imitates Anna's decision, i.e. that he invests when Anna rejects and vice versa regardless of his private signal. Of course, this strategy yields a much lower expected payoff than the default strategy. Second, adaptation across games is optimal if players have some experience with each game, but their experience is restricted ($7 \leq R \leq 126$). By adapting across games, players can rely on a larger database which strongly reduces the likelihood of being grossly misled. These benefits outweigh the costs of adaptation across games which result from players' inability to draw perfect inferences. The larger database is also reflected in the faster convergence of the mean average payoff to the long-run payoff for adaptation across games. Finally, adaptation across games becomes suboptimal as players' experience with each game grows sufficiently large such that errors in inferences diminish even when players distinguish games.

Obviously, the size of the three intervals strongly depends on the parameters of the social-learning games. In particular, some intervals may be empty. Figure 2 plots the mean average payoffs for $q_A^L = 0.6$, $q_A^H = 0.8$, $q_B^L = 0.55$, and $q_B^H = 0.6$. In this case, adaptation across games does not result in suboptimal behavior in the long-run which is why long-run payoffs coincide. Therefore, the costs of adapting across games are strongly

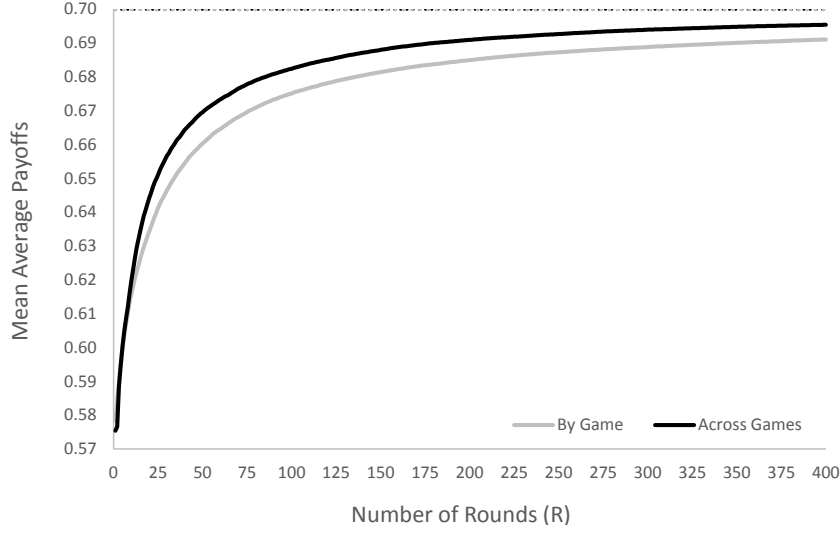


Figure 2: Mean Average Payoffs of Bob for $q_A^L = 0.6$, $q_A^H = 0.8$, $q_B^L = 0.55$, $q_B^H = 0.6$

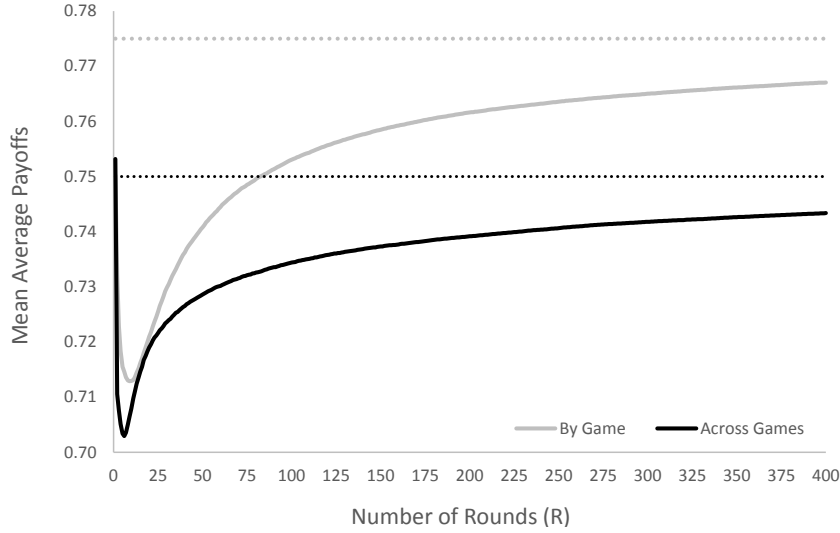


Figure 3: Mean Average Payoffs of Bob for $q_A^L = 0.6$, $q_A^H = 0.8$, $q_B^L = 0.75$, $q_B^H = 0.75$

reduced,³ and adaptation across games is optimal for any $R > 4$. On the other hand, Figure 3 illustrates that adaptation across games is not necessarily optimal for some R . In this case, $q_A^L = 0.6$, $q_A^H = 0.8$, and $q_B^L = q_B^H = 0.75$, and Bob suboptimally follows his private signal in game H in the long-run. Though long-run payoffs are close, the fact that \bar{q}_A is closer to q_B^L than q_A^L generates additional costs of adapting across games. After any finite number of rounds, Bob is more likely to erroneously imitate Anna's decision in game L when he adapts across games than when he distinguishes games.

To give an impression of the prevalence of adaptation across games and the conditions which facilitate it, Tables 1 to 4 report the intervals $[\underline{R}, \bar{R}] = \{\underline{R}, \underline{R} + 1, \dots, \bar{R}\}$ of the

³The costs are not zero. Bob (implicitly) assesses Anna's signal precision by $\bar{q}_A = 0.7$ in the long-run when adapting across games. Since \bar{q}_A is closer to q_B^H than q_A^H , Bob is more likely to pick a wrong strategy in game H after a finite number of repetitions when adapting across games.

number of repetitions for which adaptation across games results in a larger mean average payoff than adaptation by game. I consider signal precisions $q_i^k \in \{0.55, 0.60, \dots, 0.85\}$ for $i \in \{A, B\}$ and $k \in \{L, H\}$ such that $q_A^H - q_A^L \geq 0.15$. Stars indicate parameter constellations for which adaptation across games induces suboptimal long-run behavior.

The results demonstrate that adaptation across games is often optimal when experiences are scarce even if it results in suboptimal long-run behavior. This holds in particular when games are sufficiently similar. In the simple example, similarity of games is captured by the (absolute) difference $q_A^H - q_A^L$ between Anna's signal precisions in the two games. Indeed, Tables 1 to 4 are sorted by this difference. As the difference increases, intervals $[\underline{R}, \overline{R}]$ shrink and parameter constellations for which $[\underline{R}, \overline{R}] = \emptyset$ become more frequent.

q_B^L	q_B^H						
	0.55	0.60	0.65	0.70	0.75	0.80	0.85
0.55	[15, 400]	[14, 73]	[11, 102]*	[11, 400]	[10, 400]	[9, 400]	[9, 400]
0.60	\emptyset^*	[17, 28]*	[14, 42]*	[13, 76]*	[10, 171]*	[9, 171]*	[9, 159]*
0.65	[15, 32]	[13, 30]	[12, 38]*	[11, 66]	[10, 143]	[9, 156]	[9, 142]
0.70	\emptyset	[17, 30]	[15, 33]*	[14, 82]	[12, 400]	[11, 400]	[10, 400]
0.75	[16, 34]	[15, 32]	[13, 47]*	[12, 400]	[11, 400]	[10, 400]	[10, 400]
0.80	[13, 60]	[11, 53]	[12, 62]*	[11, 400]	[10, 400]	[10, 400]	[10, 400]
0.85	[11, 82]	[12, 61]	[11, 80]*	[11, 400]	[10, 400]	[9, 400]	[9, 400]

(a) $(q_A^L, q_A^H) = (0.55, 0.70)$

q_B^L	q_B^H						
	0.55	0.60	0.65	0.70	0.75	0.80	0.85
0.55	[3, 400]	[4, 400]	[7, 154]	[5, 160]*	[7, 400]	[7, 400]	[7, 400]
0.60	[4, 400]	[4, 95]	[8, 58]	[7, 71]*	[7, 400]	[7, 400]	[7, 400]
0.65	[6, 53]*	[9, 39]*	[9, 34]*	[8, 41]*	[8, 69]*	[8, 143]*	[8, 150]*
0.70	[12, 35]	[14, 25]	[16, 19]	[13, 30]*	[13, 43]	[11, 109]	[11, 123]
0.75	[12, 33]	[13, 25]	[13, 23]	[12, 30]*	[12, 59]	[11, 400]	[10, 400]
0.80	[10, 147]	[11, 48]	[11, 35]	[11, 45]*	[10, 400]	[10, 400]	[10, 400]
0.85	[11, 400]	[12, 72]	[11, 54]	[10, 69]*	[10, 400]	[10, 400]	[10, 400]

(b) $(q_A^L, q_A^H) = (0.60, 0.75)$

q_B^L	q_B^H						
	0.55	0.60	0.65	0.70	0.75	0.80	0.85
0.55	[1, 400]	[1, 400]	[3, 400]	[3, 133]	[3, 145]*	[3, 400]	[4, 400]
0.60	[1, 400]	[3, 400]	[3, 400]	[3, 117]	[4, 132]*	[4, 400]	[5, 400]
0.65	[3, 400]	[3, 400]	[4, 60]	[4, 47]	[5, 64]*	[5, 400]	[6, 400]
0.70	[3, 44]*	[4, 35]*	[11, 23]*	[9, 26]*	[10, 30]*	[11, 42]*	[10, 116]*
0.75	[7, 38]	[9, 27]	[10, 20]	[10, 22]	[11, 24]*	[11, 37]	[9, 95]
0.80	[8, 66]	[10, 38]	[10, 25]	[9, 26]	[10, 31]*	[10, 60]	[10, 400]
0.85	[8, 400]	[10, 400]	[10, 49]	[9, 42]	[10, 49]*	[10, 400]	[10, 400]

(c) $(q_A^L, q_A^H) = (0.65, 0.80)$

q_B^L	q_B^H						
	0.55	0.60	0.65	0.70	0.75	0.80	0.85
0.55	[2, 400]	[1, 400]	[2, 400]	[1, 400]	[2, 106]	[1, 112]*	[3, 400]
0.60	[1, 400]	[1, 400]	[2, 400]	[3, 400]	[3, 118]	[3, 122]*	[3, 400]
0.65	[2, 400]	[2, 400]	[3, 400]	[3, 400]	[3, 92]	[3, 111]*	[4, 400]
0.70	[2, 400]	[1, 400]	[3, 75]	[3, 46]	[3, 35]	[5, 35]*	[7, 400]
0.75	[2, 55]*	[3, 45]*	[3, 30]*	[4, 23]*	[7, 21]*	[8, 24]*	[9, 37]*
0.80	[2, 44]	[3, 38]	[7, 23]	[5, 24]	[10, 20]	[11, 22]*	[10, 33]
0.85	[3, 77]	[7, 68]	[10, 35]	[7, 31]	[10, 26]	[11, 29]*	[11, 57]

(d) $(q_A^L, q_A^H) = (0.70, 0.85)$

Table 1: Repetition Numbers Favoring Adaptation Across Games for Bob if $q_A^H - q_A^L = 0.15$.

q_B^L	q_B^H						
	0.55	0.60	0.65	0.70	0.75	0.80	0.85
0.55	[33, 400]	\emptyset	\emptyset^*	[13, 71]*	[11, 400]	[10, 400]	[9, 400]
0.60	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	[12, 61]*	[10, 128]*	[10, 136]*
0.65	\emptyset^*	\emptyset^*	\emptyset^*	[17, 21]*	[12, 40]*	[12, 76]*	[10, 92]*
0.70	\emptyset	\emptyset	\emptyset^*	\emptyset^*	[16, 32]	[13, 96]	[12, 111]
0.75	\emptyset	\emptyset	\emptyset^*	\emptyset^*	[14, 54]	[12, 400]	[11, 400]
0.80	[17, 24]	[16, 22]	[19, 20]*	[13, 36]*	[13, 400]	[11, 400]	[10, 400]
0.85	[14, 45]	[14, 34]	[14, 35]*	[11, 52]*	[11, 400]	[11, 400]	[10, 400]

(a) $(q_A^L, q_A^H) = (0.55, 0.75)$

q_B^L	q_B^H						
	0.55	0.60	0.65	0.70	0.75	0.80	0.85
0.55	[1, 400]	[5, 400]	[3, 134]	[4, 79]*	[7, 126]*	[6, 400]	[6, 400]
0.60	[3, 400]	[3, 85]	[8, 34]	[6, 37]*	[5, 56]*	[8, 400]	[7, 400]
0.65	[3, 39]*	[9, 25]*	[12, 21]*	[9, 25]*	[9, 29]*	[7, 55]*	[8, 108]*
0.70	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	[14, 26]*	[12, 53]*
0.75	\emptyset	\emptyset	\emptyset	\emptyset^*	\emptyset^*	[14, 23]	[11, 65]
0.80	\emptyset	\emptyset	\emptyset	\emptyset^*	[17, 20]*	[12, 41]	[11, 400]
0.85	[13, 400]	[13, 33]	[15, 26]	[12, 29]*	[11, 38]*	[11, 400]	[11, 400]

(b) $(q_A^L, q_A^H) = (0.60, 0.80)$

q_B^L	q_B^H						
	0.55	0.60	0.65	0.70	0.75	0.80	0.85
0.55	[2, 400]	[1, 400]	[3, 400]	[1, 126]	[3, 79]*	[3, 115]*	[4, 400]
0.60	[1, 400]	[1, 400]	[3, 400]	[3, 86]	[4, 58]*	[4, 93]*	[4, 400]
0.65	[3, 400]	[3, 400]	[3, 39]	[3, 34]	[4, 33]*	[5, 45]*	[6, 400]
0.70	[2, 35]*	[12, 18]*	\emptyset^*	[9, 18]*	\emptyset^*	[12, 19]*	[12, 33]*
0.75	[10, 20]*	\emptyset	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	[13, 20]*
0.80	[11, 22]	\emptyset	\emptyset	[10, 15]	\emptyset^*	\emptyset^*	[13, 24]
0.85	[11, 46]	[14, 28]	\emptyset	[12, 21]	[13, 20]*	[14, 23]*	[12, 50]

(c) $(q_A^L, q_A^H) = (0.65, 0.85)$

Table 2: Repetition Numbers Favoring Adaptation Across Games for Bob if $q_A^H - q_A^L = 0.20$.

q_B^L	q_B^H						
	0.55	0.60	0.65	0.70	0.75	0.80	0.85
0.55	[44, 400]	\emptyset	\emptyset	\emptyset^*	[13, 56]*	[12, 400]	[11, 400]
0.60	\emptyset	\emptyset	\emptyset	\emptyset^*	\emptyset^*	[15, 47]	[11, 104]
0.65	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	[16, 27]*	[12, 53]*
0.70	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	[16, 40]*
0.75	\emptyset	\emptyset	\emptyset	\emptyset^*	\emptyset^*	\emptyset	[14, 73]
0.80	\emptyset	\emptyset	\emptyset	\emptyset^*	\emptyset^*	[15, 46]	[12, 400]
0.85	\emptyset	\emptyset	\emptyset	[18, 21]*	[15, 32]*	[13, 400]	[12, 400]

(a) $(q_A^L, q_A^H) = (0.55, 0.80)$

q_B^L	q_B^H						
	0.55	0.60	0.65	0.70	0.75	0.80	0.85
0.55	[1, 400]	[8, 400]	[3, 400]	[4, 47]	[4, 53]*	[5, 102]*	[7, 400]
0.60	[11, 400]	[14, 36]	\emptyset	[5, 18]	[7, 26]*	[8, 43]*	[8, 400]
0.65	[14, 20]*	\emptyset^*	\emptyset^*	\emptyset^*	[11, 13]*	[12, 21]*	[9, 42]*
0.70	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*
0.75	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset^*	\emptyset^*	\emptyset
0.80	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset^*	\emptyset^*	\emptyset
0.85	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset^*	\emptyset^*	[14, 49]

(b) $(q_A^L, q_A^H) = (0.60, 0.85)$

Table 3: Repetition Numbers Favoring Adaptation Across Games for Bob if $q_A^H - q_A^L = 0.25$.

q_B^L	q_B^H						
	0.55	0.60	0.65	0.70	0.75	0.80	0.85
0.55	[41, 400]	[316, 382]	\emptyset	\emptyset^*	\emptyset^*	[22, 35]*	[11, 400]
0.60	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	[16, 36]*
0.65	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*
0.70	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset^*
0.75	\emptyset	\emptyset	\emptyset	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset
0.80	\emptyset	\emptyset	\emptyset	\emptyset^*	\emptyset^*	\emptyset^*	\emptyset
0.85	\emptyset	\emptyset	\emptyset	\emptyset^*	\emptyset^*	\emptyset^*	[216, 226]

Table 4: Repetition Numbers Favoring Adaptation Across Games for Bob if $q_A^L = 0.55$ and $q_A^H = 0.85$.

B.3. Results with Symmetric, Binary Private Signals

There is a legitimate concern that advantages of adapting across games vanish in later periods of the social-learning game. Indeed, with a finite number of private signals none of which perfectly reveals the state, an information cascade quickly emerges in which players ignore their private information and imitate the majority of previous decisions. If the evidence in favor of imitation is strong meaning that beliefs are close to zero or one, a small number of repetitions might suffice to induce optimal behavior with high probability. On the other hand, information cascades are fragile which implies that the evidence in favor of imitation is usually NOT very strong. In addition, predictions of rational social-learning are likely to be similar across different games in late periods which is why adaptation across games may not result in large deviations from optimal behavior in the long-run.

To address this concern, I report simulation results for simple social-learning games with more than two players and symmetric, binary private signals. Concretely, in each game $k \in \mathcal{K}$ signals for each player $t \in \{1, \dots, T\}$ are drawn from $S = \{0, 1\}$ according to probabilities $\Pr(\tilde{s}_t = 1 \mid \tilde{\theta} = 1, k) = \Pr(\tilde{s}_t = 0 \mid \tilde{\theta} = 0, k) = q^k \in (\frac{1}{2}, 1)$. I assume that $\mathcal{K} = \{L, H\}$ where $\pi_L = \pi_H = \frac{1}{2}$ and without loss of generality that $q^L < q^H$. Furthermore, for computational reasons $T \leq 6$.

I first characterize the long-run outcome of adaptation for the specific family of social-learning games. If games are distinguished, the long-run outcome coincides in each game with the Perfect Bayesian equilibrium derived in Bikhchandani, Hirshleifer, and Welch (1992): A player follows her private signal (i) in period 1, (ii) if each action has been chosen by the same number of predecessors, and (iii) if her private signal and the majority of previous decision are concordant. In contrast, a player follows the majority of previous actions regardless of her signal as soon as one action has been chosen by her predecessors at least twice more than the other action. Finally, a player is indifferent between the two actions if one of the actions has been chosen by her predecessors exactly once more than the other action and her private signal conflicts with the majority action. In this case, the player chooses each action with equal probability, as her long-run belief at such histories is distributed symmetrically around the true belief. In the following, let $\sigma_t^{k,*}$ denote the long-run strategy of adaptation when games are distinguished.

Corollary B.1. characterizes the long-run outcome of adaptation across games, i.e. the ABEE with information-anonymous analogy partition. Adaptation across games leads players to follow any majority in game L . In contrast, players mimic the rational strategy in game H (except in period 2, where players are not indifferent and follow their private signal), if games are sufficiently similar, i.e. if $q^L \approx q^H$. Finally, players deviate from the rational strategy if differences between games are large by following their private signal whenever all predecessors agree.

Corollary B1. For any pair $\mathcal{K} = \{L, H\}$ of social-learning games with symmetric, binary private signals of precision q^L and q^H , respectively, the long-run outcome of adaptation across games, $\{\sigma_t^{L,A}, \sigma_t^{H,A}\}_{t=1}^T$, satisfies

- (i) Period 1: $\sigma_1^{k,A}(s_1, \emptyset) = 1(0)$ if $s_1 = 1(0)$ for each $k \in \mathcal{K}$;
- (ii) Game L : $\sigma_t^{L,A}(s_t, h_t) = 1(0)$ if $a_1 = h_t(1) = 1(0)$ for each $t = 2, \dots, T$;
- (iii) Game H – Impure Histories: $\sigma_2^{H,A}(s_2, h_2) = 1(0)$ if $s_2 = 1(0)$ for each $h_2 \in H_2$, and $\sigma_t^{H,A}(s_t, h_t) = \sigma_t^{H,*}(s_t, h_t)$ for each $t = 3, \dots, T$ and each $h_t \in H_t \setminus \{(0, \dots, 0), (1, \dots, 1)\}$;
- (iv) Game H – Pure Histories: for each $t = 3, \dots, T$ and each $a \in A$,
 - $\sigma_t^{H,A}(s_t, (a, \dots, a)) = 1(0)$ if $s_t = 1(0)$, if $q^H - q^L > \delta^H \equiv q^H(1 - q^H)(2q^H - 1)$,
 - $\sigma_t^{H,A}(s_t, (a, \dots, a)) = 1(0)$ if $a = 1(0)$, if $q^H - q^L < \delta^H$.

Proof. By assumption, players follow their private signal in period 1 in each round and (i) holds. Hence, period-2 assessments satisfy $\lim_{R \rightarrow \infty} \bar{\varphi}_2(h_2 = (1) \mid \tilde{\theta} = 1) = \bar{\varphi}_2(h_2 = (0) \mid 0) = \frac{q^L + q^H}{2}$. Since $q^L < \frac{q^L + q^H}{2} < q^H$, players imitate the first action in period 2 of game L , and they follow their private signal in period 2 of game H . We therefore obtain for period 3:

h_3	$\varphi_3^{L,A}(h_3 \mid \tilde{\theta} = 1)$	$\varphi_3^{H,A}(h_3 \mid \tilde{\theta} = 1)$	$\bar{\varphi}_3^A(h_3 \mid \tilde{\theta} = 1)$
(1, 1)	q^L	$(q^H)^2$	$\frac{1}{2} [q^L + (q^H)^2]$
(1, 0)	0	$q^H(1 - q^H)$	$\frac{1}{2} q^H(1 - q^H)$
(0, 1)	0	$q^H(1 - q^H)$	$\frac{1}{2} q^H(1 - q^H)$
(0, 0)	$1 - q^L$	$(1 - q^H)^2$	$\frac{1}{2} [1 - q^L + (1 - q^H)^2]$

and $\varphi_t^{k,A}(h_3 \mid \tilde{\theta} = 0) = \varphi_t^{k,A}(\bar{h}_3 \mid \tilde{\theta} = 1)$ for each $k \in \mathcal{K}$ and each $h_3 = (a_1, a_2) \in H_3$ where $\bar{h}_3 = (1 - a_1, 1 - a_2)$. It follows that players also imitate the first action in period 3 of game L . In contrast, players follow their private signal in period 3 of game H at a pure history $h_3 \in \{(1, 1), (0, 0)\}$ if and only if

$$\frac{q^H}{1 - q^H} > \frac{q^L + (q^H)^2}{1 - q^L + (1 - q^H)^2} \Leftrightarrow q^H - q^L > \delta^H \equiv q^H(1 - q^H)(2q^H - 1).$$

Finally, players follow their private signal in period 3 of game H at any mixed history $h_3 \in \{(1, 0), (0, 1)\}$ which corresponds to the rational strategy.

Ad (iv): Let $t \geq 3$ and $a \in A$, and consider the pure history $h_t^a = (a, \dots, a) \in H_t$. Assume first that $q^H - q^L < \delta^H$ such that players who observe history h_3^a in period 3 of game H choose a for each private signal $s_3 \in S$. Since players also choose a at history h_3^a for each private signal in game L , assessments in period 4 satisfy $\bar{\varphi}_4^A(h_4^a \mid \theta) = \bar{\varphi}_3^A(h_3^a \mid \theta)$ for each $\theta \in \Theta$. Accordingly, players choose a for each private signal at history h_t^a in each period $t > 3$ of each game $k \in \mathcal{K}$, if $q^H - q^L < \delta^H$. Assume second that players follow their

private signal at history h_τ^a in each period $\tau < t$ of game H . Long-run assessments at the pure history h_t^a then satisfy

$$\frac{\bar{\varphi}_t^A(h_t^a \mid \tilde{\theta} = 1)}{\bar{\varphi}_t^A(h_t^a \mid \tilde{\theta} = 0)} = \frac{q^L + (q^H)^{t-1}}{1 - q^L + (1 - q^H)^{t-1}}$$

and players follow their private signal at this history in game H if and only if

$$\frac{q^H}{1 - q^H} > \frac{q^L + (q^H)^{t-1}}{1 - q^L + (1 - q^H)^{t-1}} \quad \Leftrightarrow \quad q^H - q^L > q^H(1 - q^H) \left[(q^H)^{t-2} - (1 - q^H)^{t-2} \right].$$

Since $(q^H)^{t-2} - (1 - q^H)^{t-2} \leq (q^H)^{t-3} - (1 - q^H)^{t-3} \leq \delta^H$ for each $t \geq 4$, players follow their private signal at history h_t^a in each period $t = 4, \dots, T$ of game H if they do so in period 3, i.e. iff $q^H - q^L > \delta^H$.

Ad (ii): Property (iv) implies that for each period $t \geq 3$ long-run assessments at the pure history h_t^a satisfy

$$\frac{\bar{\varphi}_t^A(h_t \mid \tilde{\theta} = 1)}{\bar{\varphi}_t^A(h_t \mid \tilde{\theta} = 0)} = \frac{q^L + (q^H)^j}{1 - q^L + (1 - q^H)^j} > \frac{q^L}{1 - q^L}$$

where $j \in \{2, t-1\}$. Therefore, players choose action $a \in A$ for each private signal at history h_t^a in each period $t = 2, \dots, T$ of game L .

Ad (iii): From (ii) it follows that only pure histories arise with strictly positive probability in game L . Hence,

$$\frac{\bar{\varphi}_t^A(h_t \mid \tilde{\theta} = 1)}{\bar{\varphi}_t^A(h_t \mid \tilde{\theta} = 0)} = \frac{\varphi_t^{H,A}(h_t \mid \tilde{\theta} = 1)}{\varphi_t^{H,A}(h_t \mid \tilde{\theta} = 0)}$$

at any impure history $h_t \in H_t \setminus \{h_t^0, h_t^1\}$. Therefore, the ABEE strategy and the rational strategy coincide in game H at such histories. \square

Tables 5 collects the expected payoffs in the long-run of the two adaptive processes.

	Period		
	$t \in \{1, 2\}$	$t \in \{3, 4\}$	$t \in \{5, 6\}$
Game L			
By Game	q^L	$q^L + \frac{1}{2} \delta^L$	$q^L + \frac{1}{2} \delta^L [1 + q^L (1 - q^L)]$
Across	q^L	q^L	q^L
Game H			
By Game	q^H	$q^H + \frac{1}{2} \delta^H$	$q^H + \frac{1}{2} \delta^H [1 + q^H (1 - q^H)]$
Across if $q^H - q^L < \delta^H$	q^H	$q^H + \delta^H$	$q^H + \delta^H [1 + q^H (1 - q^H)]$
Across if $q^H - q^L > \delta^H$	q^H	q^H	$q^H + \frac{2t-5}{2} \delta^H q^H (1 - q^H)$

Table 5: Expected Payoffs in the Long-Run of the Adaptive Process

While players have a smaller expected payoff in game L when adapting across games, they benefit from the improved information aggregation in game H which stems from the absence of a tie in period 2. Overall, players may benefit in the long-run from adapting across games. Figure 4 indicates the range of signal precisions for which this happens. Players benefit from adapting across games (i) in periods 3 and 4 for signal precision pairs (q^L, q^H) in the light-gray area indicated by A , (ii) in period 5 for signal precisions in the (light-gray) areas A and B , and (iii) in period 6 for signal precisions in the areas A , B , and C . The dark-gray area highlights the restriction $q^L < q^H$. The figure also shows the grid of signal precisions considered in the simulation below.

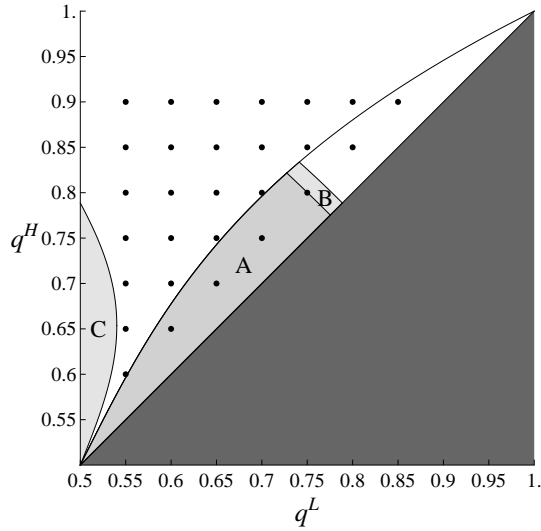


Figure 4: Signal Precisions of the Symmetric, Binary Signals for which Players Benefit from Adapting Across Games

Figure 5 illustrates the simulation results for $q^L = 0.6$ and $q^H = 0.8$. For these signal precisions, adaptation across games leads to a lower long-run expected payoff for each period $t > 2$. Still, for each period $t > 1$ there exists a range of the number of repetitions R for which adaptation across games is optimal. In addition, payoff gains of adaptation across games in the range where it is optimal are comparable across periods.

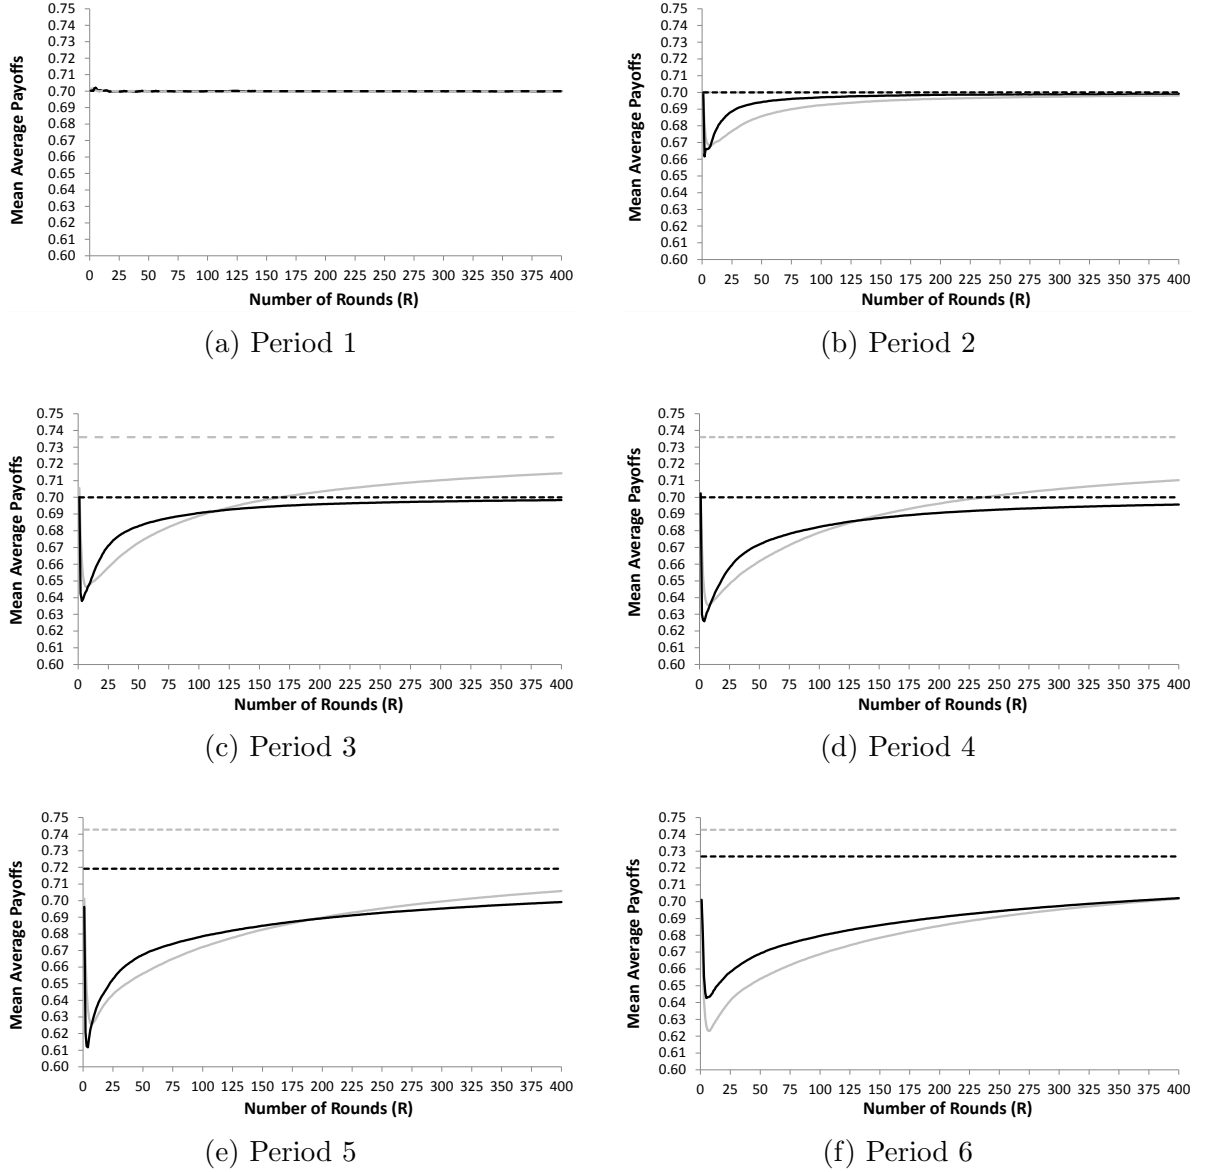


Figure 5: Mean Average Payoffs for Symmetric, Binary Private Signals with Signal Precisions $q^L = 0.6$ and $q^H = 0.8$.

The figure also establishes how convergence of the adaptive processes depends on the period. While mean average payoffs quickly approach long-run expected payoffs for the early periods of the social-learning game, convergence is much slower in later periods. This offers players the opportunity to benefit even in later periods from adapting across games when experiences are scarce.

To verify the robustness of the example, Table 6 reports the intervals of the number

of rounds R for which adaptation across games yields a larger mean average payoff than adaptation with distinction of games. I focus on periods $t \geq 3$. For each period $t \geq 3$, the table contains a distinct panel which states for each pair of signal precisions from the finite set $\{(q^L, q^H) \in \{0.55, 0.60, \dots, 0.90\} : q^L < q^H\}$ the interval $[\underline{R}, \overline{R}]$. A star indicates that adaptation across games yields in the long-run a higher expected payoff than adaptation with distinction of games.

The results clearly establish that adaptation across games can be optimal even for later periods of the social-learning game. Indeed, later periods may favor adaptation across games since convergence of the adaptive process is slower. As for the simple example, the number of repetitions for which adaptation across games is optimal largely depends on the similarity of the games. It is measured here by the difference between the signal precisions $q^H - q^L$. The larger this difference is, the smaller the interval $[\underline{R}, \overline{R}]$.

q^L	q^H						
	0.60	0.65	0.70	0.75	0.80	0.85	0.90
0.55	[14, 400]	[14, 400]	[18, 240]	[18, 137]	[18, 108]	[26, 98]	[24, 123]
0.60		[9, 400]*	[11, 310]	[10, 155]	[9, 113]	[9, 95]	[8, 99]
0.65			[7, 400]*	[6, 226]	[6, 116]	[6, 88]	[7, 81]
0.70				[6, 400]*	[5, 172]	[5, 90]	[5, 67]
0.75					[5, 400]*	[5, 135]	[5, 74]
0.80						[4, 400]	[5, 103]
0.85							[4, 300]

(a) Period 3

q^L	q^H						
	0.60	0.65	0.70	0.75	0.80	0.85	0.90
0.55	[19, 400]	[16, 400]	[24, 281]	[33, 126]	[46, 63]	\emptyset	[2, 136]
0.60		[11, 400]*	[12, 398]	[10, 193]	[9, 132]	[10, 110]	[1, 141]
0.65			[8, 400]*	[7, 289]	[7, 137]	[6, 103]	[1, 105]
0.70				[6, 400]*	[5, 200]	[5, 103]	[1, 89]
0.75					[5, 400]*	[5, 146]	[2, 92]
0.80						[5, 400]	[1, 122]
0.85							[1, 334]

(b) Period 4

q^L	q^H						
	0.60	0.65	0.70	0.75	0.80	0.85	0.90
0.55	[24, 400]	[16, 400]	[30, 400]	[43, 335]	[50, 66]	[2, 166]	[31, 328]
0.60		[8, 400]*	[11, 400]	[9, 298]	[8, 189]	[1, 182]	[7, 219]
0.65			[7, 400]*	[6, 392]	[6, 177]	[2, 147]	[5, 138]
0.70				[6, 400]*	[6, 240]	[2, 134]	[3, 104]
0.75					[5, 400]*	[1, 177]	[3, 104]
0.80						[1, 400]	[4, 131]
0.85							[2, 360]

(c) Period 5

q^L	q^H						
	0.60	0.65	0.70	0.75	0.80	0.85	0.90
0.55	[28, 400]	[18, 400]	[15, 400]	[2, 400]	[1, 400]	[45, 400]	[36, 400]
0.60		[10, 400]*	[8, 400]	[1, 400]	[1, 400]	[9, 325]	[8, 357]
0.65			[6, 400]*	[1, 400]	[1, 241]	[7, 188]	[6, 181]
0.70				[2, 400]*	[1, 287]	[5, 149]	[4, 120]
0.75					[2, 400]*	[1, 198]	[2, 117]
0.80						[2, 400]	[2, 147]
0.85							[1, 359]

(d) Period 6

Table 6: Repetition Numbers Favoring Adaptation Across Games For Symmetric, Binary Private Signals

Appendix C. Generalization of the Adaptive Process

This appendix discusses generalizations of the adaptive process. First, I establish the connection with Milgrom and Roberts (1991) who focus on the asymptotic properties of an adaptive process. Second, I extend the adaptive process by allowing for more sophisticated individuals and active experimentation following Fudenberg and Kreps (1995, *FK* henceforth). Throughout, I focus on a single social learning game ($K = 1$) and I assume that feedback is not restricted (i.e. the realized state of nature and the entire history of actions are revealed at the end of each round).

C.1. Relation to Milgrom and Roberts (1991)

Following Milgrom and Roberts (1991), an adaptive process is *consistent with adaptive learning* if each individual eventually (ϵ)-best responds to strategies which are played infinitely often. The Lemma below establishes that this property is shared by the adaptive process defined in the main text. Any adaptive process consistent with adaptive learning converges to the unique iteratively undominated strategy profile in dominance-solvable (normal-form) games (Milgrom and Roberts, 1991, Theorem 7). This result is reflected in Proposition 1.

Lemma C1. *If $K = 1$, $y_H(\omega^r) = \{h_t \subset h_{T+1}^r\}$, and $y_\Theta(\omega^r) = \{\theta^r\}$, the adaptive process is almost surely consistent with adaptive learning.*

Proof. Define for each $t = 1, \dots, T$ the set of ϵ -best responses to a subset $\hat{\Sigma}_{-t} \subseteq \Sigma_{-t}$ via

$$\Sigma_t^\epsilon(\hat{\Sigma}_{-t}) = \left\{ \sigma_t \in \Sigma_t : \forall \sigma'_t \in \Sigma_t \exists \sigma_{-t} \in \hat{\Sigma}_{-t} \text{ s.t. } U_t(\sigma_t, \sigma_{-t}) + \epsilon > U_t(\sigma'_t, \sigma_{-t}) \right\}$$

where

$$U_t(\sigma_t, \sigma_{-t}) = \frac{1}{4} \sum_{h_t \in H_t} \int_{\underline{b}}^{\bar{b}} \sigma_t(s_t, h_t) \left[\Pr(h_t | \tilde{\theta} = 1, \sigma_{-t}) - \frac{1-s_t}{s_t} \Pr(h_t | \tilde{\theta} = 0, \sigma_{-t}) \right] dG_1(s_t).$$

An adaptive process is *consistent with adaptive learning* if for each $\epsilon > 0$ and each $R > 0$ there exists $\hat{R} > R$ such that for each $r > \hat{R}$ and each $t = 1, \dots, T$, $\sigma_t^r \in \Sigma_t^\epsilon(\Sigma_{-t}^{>R})$ where $\Sigma_{-t}^{>R} = \{\times_{\tau \neq t} \sigma_\tau^\rho : \rho \geq R\}$ is the set of profiles of strategies σ_τ^ρ , $\tau \neq t$, chosen in rounds later than R . The proof will establish that this property holds along almost any adaptation path $\zeta^\infty = (\omega^1, \omega^2, \dots)$.

Consider an adaptation path ζ^r , $r = 1, 2, \dots$. Fix $\epsilon > 0$ and $0 < R < \infty$ and suppose by way of contradiction that there exists $t \in \{1, \dots, T\}$, a sub-sequence of rounds $r_1, r_2, \dots > R$, and a sequence of (alternative) strategies $\zeta_t^{r_1}, \zeta_t^{r_2}, \dots$ such that for each $r = r_1, r_2, \dots$ and each $\sigma_{-t} \in \Sigma_{-t}^{>R}$

$$U_t(\sigma_t^r, \sigma_{-t}) + \epsilon < U_t(\zeta_t^r, \sigma_{-t}).$$

I focus on rounds along the sub-sequence henceforth. Let $\kappa(\sigma_{-t} \mid \zeta^r) = |\{1 \leq \rho < r : \sigma_{-t}^r = \sigma_{-t}\}|$ denote the number of rounds before t in which strategy profile σ_{-t} was chosen. Similarly, let $\kappa(h_t, \theta, \sigma_{-t} \mid \zeta^r) = |\{1 \leq \rho < r : \theta^r = \theta \wedge h_c h_{T+1}^r \wedge \sigma_{-t}^r = \sigma_{-t}\}|$ denote the number of rounds before r in which σ_{-t} was chosen, the realized state was θ , and history h_t occurred in period t . Clearly, for each $\sigma_{-t} \notin \Sigma_{-t}^{>R}$, $\kappa(\sigma_{-t} \mid \zeta^r)/r \rightarrow 0$. Hence, there exists r_1^* such that for $r > r_1^*$

$$\sum_{\sigma_{-t} \in \Sigma_{-t}} \frac{\kappa(\sigma_{-t} \mid \zeta^r)}{r} U_t(\zeta_t^r, \sigma_{-t}) > \sum_{\sigma_{-t} \in \Sigma_{-t}} \frac{\kappa(\sigma_{-t} \mid \zeta^r)}{r} U_t(\sigma_t^r(\zeta^r), \sigma_{-t}) \quad (6)$$

In addition, for each $h_t \in H_t$ and each $\theta \in \Theta$,

$$\frac{\kappa(h_t, \theta, \sigma_{-t} \mid \zeta^r)}{\sum_{h'_t \in H_t} \kappa(h'_t, \theta, \sigma_{-t} \mid \zeta^r)} \rightarrow \Pr(h_t \mid \theta, \sigma_{-t})$$

as $r \rightarrow \infty$ by the *strong law of large numbers for conditional expectation (SSLNCE)* (see e.g. Walk, 2008).⁴ Therefore, for each $\delta > 0$ there exists r_1^* such that for each $r > r_1^*$

$$\left| \frac{\kappa(h_t, \theta, \sigma_{-t} \mid \zeta^r)}{\sum_{h'_t \in H_t} \kappa(h'_t, \theta, \sigma_{-t} \mid \zeta^r)} - \Pr(h_t \mid \theta, \sigma_{-t}) \right| < \frac{\delta}{3}.$$

Hence, there exists $r_2^* > r_1^*$ such that for $r > r_2^*$

$$\sum_{\sigma_{-t} \in \Sigma_{-t}} \frac{\kappa(\sigma_{-t} \mid \zeta^r)}{r} U_t(\sigma_t^r(\zeta^r), \sigma_{-t}) > \sum_{h_t \in H_t} \int_{\underline{b}}^{\bar{b}} \frac{dG_1(s_t)}{4s_t} \sigma_t^r(s_t, h_t \mid \zeta^r) \left[U_t(s_t, h_t \mid \zeta^r) - \frac{\delta}{3} \right] + \epsilon \quad (7)$$

where

$$\begin{aligned} U_t(s_t, h_t \mid \zeta^r) = & s_t \sum_{\sigma_{-t} \in \Sigma_{-t}} \frac{\kappa(\sigma_{-t} \mid \zeta^r)}{r} \frac{\kappa(h_t, 1, \sigma_{-t} \mid \zeta^r)}{\sum_{h'_t \in H_t} \kappa(h'_t, 1, \sigma_{-t} \mid \zeta^r)} \\ & - (1 - s_t) \sum_{\sigma_{-t} \in \Sigma_{-t}} \frac{\kappa(\sigma_{-t} \mid \zeta^r)}{r} \frac{\kappa(h_t, 0, \sigma_{-t} \mid \zeta^r)}{\sum_{h'_t \in H_t} \kappa(h'_t, 0, \sigma_{-t} \mid \zeta^r)}. \end{aligned}$$

For sufficiently large r , (i) $\kappa(\sigma_{-t} \mid \zeta^r)/r \approx \kappa(\theta, \sigma_{-t} \mid \zeta^r)/\kappa(\theta \mid \zeta^r)$ for each $\theta \in \Theta$ and each $\sigma_{-t} \in \Sigma_{-t}^{>R}$ since the strategies and the realized private signals are independent in each round, and (ii) $\sum_{h'_t \in H_t} \kappa(h'_t, \theta, \sigma_{-t} \mid \zeta^r) \approx \kappa(\theta, \sigma_{-t} \mid \zeta^r)$. Therefore, there exists $r_3^* > r_2^*$ such

⁴This crucially relies on non-correlation of *strategies* with the state of Nature or the private belief in the current repetition.

that for each $r > r_3^*$

$$U_t(s_t, h_t | \zeta^r) > s_t \sum_{\sigma_{-t} \in \Sigma_{-t}} \frac{\kappa(1, \sigma_{-t} | \zeta^r)}{\kappa(1 | \zeta^r)} \frac{\kappa(h_t, 1, \sigma_{-t} | \zeta^r)}{\kappa(1, \sigma_{-t} | \zeta^r)} \\ - (1 - s_t) \sum_{\sigma_{-t} \in \Sigma_{-t}} \frac{\kappa(0, \sigma_{-t} | \zeta^r)}{\kappa(0 | \zeta^r)} \frac{\kappa(h_t, 0, \sigma_{-t} | \zeta^r)}{\kappa(0, \sigma_{-t} | \zeta^r)} - \frac{\delta}{3}.$$

In addition, for sufficiently large r (iii) $\sum_{\sigma_{-t} \in \Sigma_{-t}} \kappa(h_t, \theta, \sigma_{-t} | \zeta^r) \approx \kappa(h_t, \theta | \zeta^r)$, and (iv) $\varphi_t^r(h_t | \theta; \zeta^r) \approx \kappa(h_t, \theta | \zeta^r) / \sum_{h'_t \in H_t} \kappa(h'_t, \theta | \zeta^r)$. Hence, there exists $r_4^* > r_3^*$ such that for each $r > r_4^*$

$$U_t(s_t, h_t | \zeta^r) > s_t \varphi_t^r(h_t | 1; \zeta^r) - (1 - s_t) \varphi_t^r(h_t | 0; \zeta^r) - \frac{2}{3} \delta.$$

Accordingly, for each $r > r_4^*$ the RHS of (7) is strictly larger than

$$\sum_{h_t \in H_t} \int_{\underline{b}}^{\bar{b}} \frac{dG_1(s_t)}{4 s_t} \sigma_t^r(s_t, h_t | \zeta^r) [s_t \varphi_t^r(h_t | 1; \zeta^r) - (1 - s_t) \varphi_t^r(h_t | 0; \zeta^r) - \delta] + \epsilon \quad (8)$$

Finally, since σ_t^r best responds to φ_t^r for each (s_t, h_t) occurring a non-vanishing fraction of the time,

$$\frac{1}{4} \sum_{h_t \in H_t} \int_{\underline{b}}^{\bar{b}} \frac{dG_1(s_t)}{s_t} \sigma_t^r(s_t, h_t | \zeta^r) [s_t \varphi_t^r(h_t | 1; \zeta^r) - (1 - s_t) \varphi_t^r(h_t | 0; \zeta^r) - \delta] + \epsilon \\ > \frac{1}{4} \sum_{h_t \in H_t} \int_{\underline{b}}^{\bar{b}} \frac{dG_1(s_t)}{s_t} \varsigma_t^r(s_t, h_t) [s_t \varphi_t^r(h_t | 1; \zeta^r) - (1 - s_t) \varphi_t^r(h_t | 0; \zeta^r) - \delta] + \epsilon \quad (9)$$

for $r > r_4^*$. Combining equations (6), (7), (8), and (9) and noting that $\delta > 0$ has been chosen arbitrarily yields the desired contradiction. \square

C.2. Extension Following Fudenberg and Kreps (1995)

In the simple adaptive process, each individual updates a single collection of assessments $\varphi_t^r(\zeta^r)$ over time. The general adaptive process allows individuals to assign positive probability to various alternative interpretations of the observed history. A *general assessment rule* $\gamma^r : \Omega^{r-1} \rightarrow \Delta(\Phi)$ assigns to each adaptation path ζ^r in round r a probability distribution on the set of possible assessments $\Phi = \{\{\varphi_t\}_{t=1}^T \mid \varphi_t : \Theta \rightarrow \Delta(H_t) \text{ for each } t\}$. The updating of assessments over time is not modeled explicitly, instead restrictions are placed upon the long-run relationship between assessments and feedback. The general process therefore encompasses various models such as Bayesian learning (FK, Section 3.2).

Definition C1. *The general assessment rules γ^r , $r = 1, 2, \dots$ are **asymptotically empirical** if for each $\epsilon > 0$, each ζ^r , each $\theta \in \Theta$, each $t = 1, \dots, T$, and each $h_t \in H_t$ such that $\liminf_{r \rightarrow \infty} \kappa_t^r(h_t, \theta \mid \zeta^r) / r > 0$,*

$$\lim_{r \rightarrow \infty} \gamma^r(\zeta^r) \left(\left\{ \{\varphi_t\}_{t=1}^T : \left\| \varphi_t(h_t \mid \theta) - \frac{\kappa_t^r(h_t, \theta \mid \zeta^r)}{\sum_{h'_t \in H_t} \kappa_t^r(h'_t, \theta \mid \zeta^r)} \right\| < \epsilon \right\} \right) = 1$$

The general adaptive process is also less restrictive with regard to the strategic responses to assessments. In particular, individuals are allowed to experiment with suboptimal strategies, which is important to rule out unstable self-confirming equilibria. Given general assessment γ^r let

$$U_t(\sigma_t \mid \gamma^r) = \sum_{h_t \in H_t} \int_{\underline{b}}^{\bar{b}} \sigma_t(s_t, h_t) U_t(s_t, h_t \mid \gamma^r)$$

for each $t = 1, \dots, T$ where

$$U_t(s_t, h_t \mid \gamma^r) = \int_{\varphi_t \in \Phi_t} [\varphi_t(h_t \mid 1) dG_1(s_t) - \varphi_t(h_t \mid 0) dG_0(s_t)] \gamma(d\varphi).$$

The strategic responses σ_t^r are **asymptotically myopic** with regard to γ^r , if there exists a sequence $\{\epsilon_r\}_{r=1}^\infty$ such that $\epsilon_r > 0$ for each r , $\lim_{r \rightarrow \infty} \epsilon_r = 0$, and

$$U_t(\sigma_t^r(\zeta^r) \mid \gamma^r(\zeta^r)) + \epsilon_r \geq \max_{\sigma_t \in \Sigma_t} U_t(\sigma_t \mid \gamma^r(\zeta^r)).$$

for each $r = 1, 2, \dots$, each ζ^r , and each $t = 1, \dots, T$. Asymptotic myopia permits individuals to choose suboptimal strategies where the suboptimality vanishes over time. Yet, while players may consciously experiment with suboptimal strategies in early rounds, they must eventually confine themselves to *random experimentation* with decreasing overall probability. FK therefore introduce the following more general idea.

Definition C2. The strategic responses σ_t^r are **asymptotically myopic with calendar-time limitations on experimentation** with respect to γ^r if

$$\sigma_t^r(s_t, h_t | \zeta^r) = \alpha_t^r(\zeta_r, h_t) \cdot \hat{\sigma}_t^r(s_t, h_t | \zeta^r) + [1 - \alpha_t^r(\zeta_r, h_t)] \cdot \zeta_t^r(s_t, h_t | \zeta^r).$$

where (i) $\hat{\sigma}_t^r$ is asymptotically myopic for each t and r , and (ii) there exists a non-decreasing positive sequence δ_r with $\lim_{r \rightarrow \infty} \delta_r/r = 0$ such that for each r , t , ζ^r , and h_t , $\alpha_t^r(\zeta_r, h_t) < 1$ only if $\kappa_t^r(a, h_t | \zeta^r) < \delta_r$ for some $a \in A$, and for each s_t and a , $\zeta_t^r(a | s_t, h_t; \zeta^r) > 0$ if and only if $\kappa_t^r(a, h_t | \zeta^r) < \delta_r$ where $\kappa_t^r(a, h_t | \zeta^r)$ is the number of times action a has been chosen at history h_t along ζ_r by the same individual.

Hence, individuals are allowed to actively experiment as long as they confine their experimentation to actions that have been taken infrequently relative to calendar time. An *individual learning model* is an array of general assessment rules and strategic responses, one each for each individual. It is *conforming*, if assessment rules are asymptotically empirical and strategic responses are asymptotically myopic with calendar-time limitations on experimentation with respect to the assessment rules.

Definition C3. A strategy profile σ^{**} is **locally stable**, if there exists some individual learning model with asymptotically empirical general assessment rules and strategic responses that are asymptotically myopic with calendar-time limitations on experimentation, such that $\mathbf{P}(\lim_{r \rightarrow \infty} \hat{\sigma}_t^r(\zeta^r) = \sigma_t^{**}) > 0$ for each $t = 1, \dots, T$.

Proposition C1. For a (single) social learning game, a strategy profile is locally stable if and only if it is iteratively undominated.

Proof. As for the simple adaptive process, the proof is by induction: In period 1, assessments coincide with rational assessments by definition, and the strategy σ_1^r eventually plays arbitrarily close to the iteratively undominated strategy σ_1^* , once experimentation has vanished sufficiently. Since assessments are asymptotically empirical, eventually only assessments which are arbitrarily close to rational assessments in period 2 can be assigned strictly positive probability. This argumentation can be extended to all periods $t > 2$.

Finally, it may be shown that any iteratively undominated strategy profile can be locally stable (notice that besides in non-generic settings these strategies differ only at histories reached with probability zero) by explicitly constructing for a given profile an asymptotically empirical assessment rule and an associated myopic strategic response (see Fudenberg and Kreps, 1995, Proposition 6.3). \square

According to the Proposition, *any* iteratively undominated strategy profile might arise as the long-run outcome of the general adaptive process. Indeed, selection of a unique strategy profile in the long-run seems to be challenging since even refined equilibrium concepts (perfect, sequential) do not restrict the set of strategy profiles significantly.

However, under mild additional conditions a unique strategy profile is selected in the limit (as payoff disturbances approach zero) of a sequence of regular quantal response equilibria (Goeree, Holt, and Palfrey, 2005). I conjecture that this strategy profile may be selected by an adaptive process satisfying some mild additional conditions. Notice that any selection must uniquely define behavior at histories reached with probability zero. Such histories may occur either because players best respond to mistaken beliefs or because players experiment. Yet, a suboptimal decision is least costly if it is favored by an individual's private information. Hence, choices which occur with probability zero in the limit may reveal a player's private signal in the medium run. This is the property of the limit of quantal response equilibria.

References

- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A theory of fads, fashion, custom, and cultural change as informational cascades,” *Journal of Political Economy*, 100, 992–1026.
- FUDENBERG, D., AND D. M. KREPS (1995): “Learning in Extensive-Form Games, I: Self-Confirming Equilibria,” *Games and Economic Behavior*, 8, 20–55.
- GOEREE, J. K., C. A. HOLT, AND T. R. PALFREY (2005): “Regular Quantal Response Equilibrium,” *Experimental Economics*, 8, 347–367.
- JEHIEL, P. (2005): “Analogy-Based Expectation Equilibrium,” *Journal of Economic Theory*, 123, 81–104.
- JEHIEL, P., AND D. ETTINGER (2010): “A Theory of Deception,” *American Economic Journal: Microeconomics*, 2, 1–20.
- MILGROM, P., AND J. ROBERTS (1991): “Adaptive and Sophisticated Learning in Normal Form Games,” *Games and Economic Behavior*, 3, 82–100.
- WALK, H. (2008): “A universal strong law of large numbers for conditional expectations via nearest neighbors,” *Journal of Multivariate Analysis*, 99, 1035–1050.