



Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
<http://economics.sas.upenn.edu/pier>

PIER Working Paper 15-022

“Informational Herding with Model Misspecification, Second Version”

by

J. Aislinn Bohren

<http://ssrn.com/abstract=2619757>

Informational Herding with Model Misspecification*

J. Aislinn Bohren[†]

University of Pennsylvania

November 2014

Abstract

This paper demonstrates that a misspecified model of information processing interferes with long-run learning and allows inefficient choices to persist in the face of contradictory public information. I consider an observational learning environment where agents observe a private signal about a hidden state, and some agents observe the actions of their predecessors. Prior actions aggregate multiple sources of correlated information about the state, and agents face an inferential challenge to distinguish between new and redundant information. When individuals significantly overestimate the amount of new information, beliefs about the state become entrenched and incorrect learning may occur. When individuals sufficiently overestimate the amount of redundant information, beliefs are fragile and learning is incomplete. Learning is complete when agents have an approximately correct model of inference, establishing that the correct model is robust to perturbation. These results have important implications for timing, frequency and strength of policy interventions to facilitate learning.

*I thank Vince Crawford, Ernesto Dal Bo, Alex Imas, Frederic Koessler, George Mailath, Craig McKenzie, Matthew Rabin, Joel Sobel, and especially Nageeb Ali for useful comments. I also thank participants of the UCSD theory lunch for helpful feedback.

[†]Email: abohren@sas.upenn.edu; Web: <https://sites.google.com/site/aislinnbohren/>; Address: 3718 Locust Walk, Department of Economics, Philadelphia, PA 19103.

1 Introduction

Observational learning plays an important role in the transmission of information, opinions and behavior. People use bestseller lists to guide their purchases of books, cars and computers. Co-workers' decisions to join a retirement plan influence a person's decision to participate herself. Social learning also influences behavioral choices, such as whether to smoke or exercise regularly, or ideological decisions, such as which side of a moral or political issue to support. Given the gamut of situations influenced by observational learning, it is important to understand how people learn from the actions of their peers. This paper explores how a misspecified model of information processing may interfere with asymptotic learning, and demonstrates that such biases offer an explanation for how inefficient choices can persist in the face of contradictory public information. The results have important implications for policies aimed at counteracting inefficient social choices. In the presence of information processing errors, the timing, frequency and strength of policy interventions – such as public information campaigns – are an important determinate of long-run efficiency.

Individuals face an inferential challenge when extracting information from the actions of others. An action often aggregates multiple sources of correlated information. Full rationality requires an agent to parse out the new information and discard redundant information. This is a critical feature of standard observational learning models in the tradition of [Smith and Sorensen \(2000\)](#). Agents understand exactly how preceding agents incorporate the action history into their decision-making rule, and are aware of the precise informational content of each action. However, what happens if agents are unsure about how to draw inference from the actions of their predecessors? What if they believe the actions of previous agents are more informative than is actually the case, or what if they attribute too many prior actions to repeated information and are not sensitive enough to new information?

Motivated by this possibility, I allow agents to have a misspecified model of the information possessed by other agents. This draws a distinction between the perceived and actual informational content of actions. Consider an observational

learning model where individuals have common-value preferences that depend on an unknown state of the world. They act sequentially, observing a private signal before choosing an action. A fraction p of individuals also observe the actions of previous agents. These socially informed agents understand that prior actions reveal information about private signals, but fail to accurately disentangle this new information from the redundant information also contained in prior actions. Formally, informed agents believe that any other individual is informed with probability \hat{p} , where \hat{p} need not coincide with p . When $\hat{p} < p$, an informed decision maker attributes too many actions to the private signals of uninformed individuals. This leads him to overweigh information from the public history, and allows public beliefs about the state to become entrenched. On the other hand, when $\hat{p} > p$, an informed decision maker underweights the new information contained in prior actions, rendering beliefs more fragile to contrary information.

To understand how model misspecification affects long-run learning requires careful analysis of the rate of information accumulation, and how this rate depends on the way informed agents interpret prior actions. Theorem 1 specifies thresholds on beliefs about the share of informed agents, \hat{p}_1 and \hat{p}_2 , such that when $\hat{p} < \hat{p}_1$ both correct and fully incorrect learning occur, and when $\hat{p} > \hat{p}_2$, beliefs about the state perpetually fluctuate, rendering learning incomplete. Both cases admit the possibility of inefficient learning: with positive probability, informed agents continue to choose the inefficient action infinitely often, despite observing sufficient information to learn the correct state. When \hat{p} falls between these two thresholds, $\hat{p} \in (\hat{p}_1, \hat{p}_2)$, learning is complete and informed agents will eventually choose the efficient action. Efficient learning obtains in the correctly specified model, as demonstrated by the fact that $p \in (\hat{p}_1, \hat{p}_2)$.

Fully incorrect learning or incomplete learning is possible for some values of $\hat{p} \neq p$ because the public belief about the state is no longer a martingale. This also complicates the analysis on a technical level, as it is no longer possible to use the Martingale Convergence Theorem to establish belief convergence. The Law of the Iterated Logarithm (LIL) and Law of Large Numbers (LLN) are jointly used to establish belief convergence when $\hat{p} < \hat{p}_2$, and rule out belief convergence when

$\hat{p} > \hat{p}_2$. This approach could also be utilized to examine other forms of model misspecification.

Model misspecification has important policy implications. Consider a parent deciding whether there is a link between vaccines and autism. The parent observes public signals from the government and other public health agencies, along with the vaccination decisions of peers. If all parents are rational, then a public health campaign to inform parents that there is no link between vaccines and autism should eventually overturn a herd on refusing vaccinations. However, if parents do not accurately disentangle repeated information and attribute too many choices to new information, then observing many other parents refusing to vaccinate their children will lead to strong beliefs that this is the optimal choice, and make it less likely that the public health campaign is effective.¹ When this is the case, the best way to quash a herd on refusing vaccinations is to release public information immediately and frequently. This contrasts with the fully rational case, in which the timing of public information release is irrelevant for long-run learning outcomes.

This paper relates to a rich literature. [Banerjee \(1992\)](#) and [Bikhchandani, Hirshleifer, and Welch \(1992\)](#) first model observational learning in a sequential setting with binary signals. [Moscarini, Ottaviani, and Smith \(1998\)](#) show that in the [Bikhchandani et al. \(1992\)](#) framework, informational cascades are temporary when the state of the world changes frequently enough. [Smith and Sorensen \(2000\)](#) study a social learning framework with a general signal distribution and crazy types. An unbounded signal space is sufficient to ensure complete learning, eliminating the possibility of inefficient cascades. [Acemoglu, Dahleh, Lobel, and Ozdaglar \(2011\)](#) examines social learning in a network - the rational model of sequential learning with uninformed agents is a special case of their model.

This paper is most closely related to concurrent work on social learning by [Eyster and Rabin \(2010\)](#). They extend a sequential learning model with continuous actions and signals to allow for “inferential naivety”: players realize that previous agents’ action choices reflect their signals, but fail to account for the fact that these actions are also based on the actions of agents preceding these players. While continuous

¹This example abstracts from the payoff interdependencies of vaccines.

actions lead to full revelation of players' signals in the absence of inferential naivety, inferential naivety can confound learning by overweighing actions of the first few agents. Although similar in nature, inferential naivety and model misspecification differ in generality and interpretation. Inferential naivety considers the case in which every repeated action is viewed as being independent with probability one, whereas in the current setting, most decision makers are sophisticated and recognize that actions contain some repeated information, but misperceive the exact proportion. Additionally, all agents observe public information in [Eyster and Rabin \(2010\)](#). The analogue of inferential naivety in my environment corresponds to $\hat{p} = 0$ and $p = 1$. As such, both papers provide complementary explanations for the robustness of inefficient learning. [Eyster and Rabin \(2010\)](#) also embed inferentially naive agents in a model with rational agents. When every n th player in the sequence is inferentially naive, rational agents achieve complete learning but inferentially naive agents do not. Augmenting the misspecified and inferentially naive models with fully rational agents who do not know precisely which previous agents are also rational, naive or uninformed is an interesting avenue left open for future research.

Several other papers examine boundedly rational information processing in a social learning framework. [Guarino and Jehiel \(2013\)](#) employ the concept of analogy based expectation equilibrium (ABEE), in which agents best respond to the aggregate distribution of action choices. Learning is complete in a continuous action model - in an ABEE, the excess weight placed on initial signals increases linearly with time, preventing these initial signals from permanently dominating subsequent new information. This contrasts with [Eyster and Rabin \(2010\)](#), in which the excess weight on initial signals doubles each period, allowing a few early signals to dominate all future signals. As in the fully rational model, complete learning no longer obtains in an ABEE when actions are discrete. [Demarzo, Vayanos, and Zwiebel \(2003\)](#) introduce the notion of persuasion bias in a model of opinion formation in networks. Decision makers embedded in a network graph treat correlated information from others as being independent, leading to informational inefficiencies. Although this paper studies a different environment than theirs, it provides a natural analogue for considering persuasion bias in social learning. Earlier work by [Eyster and Rabin](#)

(2005) on cursed equilibrium also examines information processing errors. A cursed player doesn't understand the correlation between a player's type and his action choice, and therefore fails to realize a player's action choice reveals information about his type.²

The recent initial response models, including level-k analysis and cognitive hierarchy models, are similar in spirit to this paper.³ Consider level-k analysis in the context of sequential learning. Anchoring level 0 types to randomize between the two possible actions, level 1 types best respond by following their private signal - this corresponds to uninformed types. Level 2 types believe all other agents follow their private signal, and thus act as informed agents with beliefs $\hat{p} = 0$. Consequently, the main difference between level-k analysis and the model misspecification in this paper is the beliefs informed agents have about other agents' types - in this paper, informed agents can place positive weight on other agents using a level 2 decision rule, whereas in a level k analysis, informed agents believe that all other agents use a level 1 decision rule. In both settings, level 2 agents misperceive the share of other agents who are level 2. The comparison to a cognitive hierarchy model is similar.

The organization of this paper proceeds as follows. Section 2 sets up the model and solves the individual decision-problem. Section 3 characterizes the asymptotic learning dynamics of a misspecified model of inference, while Section 4 discusses the results and concludes. All proofs are in the Appendix.

2 The Common Framework

2.1 The Model

The basic set-up of this model mirrors a standard sequential learning environment.

States, Actions and Payoffs. There are two payoff-relevant states of the world, $\omega \in \{L, R\}$ with common prior belief $P(\omega = L) = 1/2$. Nature selects one of these states at the beginning of the game. A countably infinite set of agents $T = \{1, 2, \dots\}$

²Epstein, Noor, and Sandroni (2010) study non-Bayesian learning in a single-agent framework.

³Camerer, Ho, and Chong (2004); Costa-Gomes, Crawford, and Iriberry (2009).

act sequentially and attempt to match the realized state of the world by making a single decision between two actions, $a_t \in \{L, R\}$. They receive a payoff of 1 if their action matches the realized state, and a payoff of 0 otherwise: $u(a_t, \omega) = 1_{a_t=\omega}$.

Private Beliefs. Before choosing an action, each agent privately observes a signal that is independent and identically distributed, conditional on the state. Following [Smith and Sorensen \(2000\)](#), I work directly with the private belief, $s_t \in (0, 1)$, which is an agent's belief that $\omega = L$ after observing the private signal but not the history. Conditional on the state, the private belief stochastic process $\langle s_t \rangle$ is i.i.d, with conditional c.d.f. F^ω . Assume that no private signal perfectly reveals the state, which implies that F^L, F^R are mutually absolutely continuous and have common support, $\text{supp}(F)$. Let $[\underline{b}, \bar{b}] \subseteq [0, 1]$ denote the convex hull of the support. Beliefs are bounded if $0 < \underline{b} < \bar{b} < 1$, and are unbounded if $[\underline{b}, \bar{b}] = [0, 1]$. Finally, assume that some signals are informative. This rules out $dF^L/dF^H = 1$ almost surely.

Agent Types. There are two types of agents, $\theta_t \in \{I, U\}$. With probability $p \in (0, 1)$, an agent is a socially informed type I who observes the action choices of her predecessors, $h_t = (a_1, \dots, a_{t-1})$. She uses her private signal and this history to guide her action choice. With probability $1 - p$, an agent is a socially uninformed type U who only observes his private signal. An alternative interpretation for this uninformed type is a behavioral type who is not sophisticated enough to draw inference from the history. This type's decision is solely guided by the information contained in his private signal.

Beliefs About Types. Each informed individual believes that each other individual is informed with probability \hat{p} , where \hat{p} need not coincide with p . An informed agent believes that other agents also hold the same beliefs about whether previous agents are informed or uninformed. Incorrect beliefs about p can persist because no agent ever learns what the preceding agents actually observed or incorporated into their decision-making processes.⁴

⁴Although it is admittedly restrictive to require that agents hold identical misperceptions about others, and this misperception takes the form of a potentially incorrect point-mass belief about the distribution of p , it is a good starting point to examine the possible implications of model misspecification. [Bohren \(2012\)](#) also analyzes the model in which agents begin with a non-degenerate prior distribution over p , and learn about p from the action history.

Timing. At time t , agent t observes type θ_t and a private signal s_t ; if $\theta_t = I$, then the agent also observes the public history h_t . Next, she chooses action a_t .

2.2 The Individual Decision-Problem

A decision rule specifies an action for each history and signal realization pair. I look for an outcome that has the nature of a Bayesian equilibrium, in the sense that agents use Bayes rule to formulate beliefs about the state of the world, given their incorrect belief about the type distribution, and seek to maximize payoffs. The decision rule of each type is common knowledge, as is the fact that all informed agents compute the same (possibly inaccurate) probability of any history h_t .

It is standard to express the public belief of informed agents as a likelihood ratio,

$$l_t = \frac{P(L|h_t; \hat{p})}{P(R|h_t; \hat{p})},$$

which depends on the history and beliefs about the share of informed agents.⁵ An agent who holds prior belief l and receives signal s updates to the private posterior belief $q(l, s) = l \times \left(\frac{s}{1-s}\right)$. An uninformed agent has prior belief $l_1 = 1$ and an informed agent has prior belief l_t . Guided by posterior belief q , the agent maximizes her payoff by choosing $a = L$ if $q \geq 1$, and $a = R$ otherwise. An agent's decision can be represented as a cut-off rule, $s^*(l) = 1/(l+1)$, such that the agent chooses action L when $s \geq s^*(l)$ and chooses action R otherwise. An informed agent in period t uses cut-off $s^*(l_t)$, while uninformed agents use cut-off $s^*(1) = 1/2$.

An agent is in an information cascade when it is optimal for the agent to choose the same action regardless of her private signal realization; therefore, this action reveals no private information.

Definition 1 (Cascade Set). *The cascade set is the set of beliefs $\Lambda = \{l|s < s^*(l) \forall s \in \text{supp}(F)\} \cup \{l|s \geq s^*(l) \forall s \in \text{supp}(F)\}$.*

An informed agent is in a cascade if $l_t \in \Lambda$, and an uninformed agent is in a cascade

⁵I describe l_t as the public belief, even though it is not the belief of uninformed agents.

if $l_1 \in \Lambda$. As usual, a cascade occurs when the prior belief outweighs the strongest private belief.

Lemma 1. *The cascade set is $\Lambda = [0, (1 - \bar{b})/\bar{b}] \cup [(1 - \underline{b})/\underline{b}, \infty]$ when signals are bounded and $\Lambda = \{0, \infty\}$ when signals are unbounded.*

By Lemma 1, uninformed agents are never in a cascade, since $l_1 = 1$. When informed agents are in a cascade, information continues to accumulate from the actions of uninformed agents, and the formation of a cascade does not necessarily imply belief convergence. Therefore, if the likelihood ratio enters the cascade set in finite time, I would like to determine whether the likelihood ratio remains in the cascade set. If a cascade doesn't form in finite time, I would like to determine whether the likelihood ratio can converge to a point in the cascade set. The following definition introduces the notion of a limit cascade to encompass both of these ideas.⁶

Definition 2 (Limit Cascade). *Suppose there exists a real, nonnegative random variable l_∞ such that $l_t \rightarrow l_\infty$ almost surely. Then an informed agent is almost surely in a limit cascade if $\text{supp}(l_\infty) \in \Lambda$.*

3 Learning Dynamics

3.1 Overview

This section proceeds as follows. After formally defining the stochastic process $\langle l_t \rangle$ governing the evolution of the likelihood ratio, I characterize the set of stationary points; these are candidate limit points for $\langle l_t \rangle$. Next, I determine how the local stability of these stationary points depends on \hat{p} . This establishes the dynamics of the likelihood ratio in the neighborhood of a stationary point. I use the law of the iterated logarithm (LIL) to show that the likelihood ratio converges to each stable stationary point with positive probability, from any initial value. Finally, I rule out convergence to unstable stationary points and non-stationary points. The

⁶It would also be possible for beliefs to remain in the cascade set, but not converge. I rule this out in Section 3.

section concludes with a full characterization of asymptotic learning outcomes, which depend on \hat{p} .

3.2 The Likelihood Ratio

Let $\psi(a|\omega, l; p)$ denote the probability of action a , given likelihood ratio l , state ω and share of informed agents p . Then

$$\psi(L|\omega, l; p) = p(1 - F^\omega(1/(l+1))) + (1-p)(1 - F^\omega(1/2)) \quad (1)$$

and

$$\psi(R|\omega, l; p) = pF^\omega(1/(l+1)) + (1-p)F^\omega(1/2). \quad (2)$$

This probability is a weighted average of the probability that an uninformed type chooses a when using cut-off rule $s^*(1) = 1/2$ and the probability that an informed type chooses a using cut-off rule $s^*(l) = 1/(l+1)$, given likelihood ratio l .

The likelihood ratio is updated based on the perceived probability of action a , $\psi(a|\omega, l; \hat{p})$. If agents attribute a smaller share of actions to informed agents, $\hat{p} < p$, then they place more weight on the action revealing private information and overestimate the informativeness of prior actions. The opposite holds when agents attribute too large a share to informed agents. Given a likelihood ratio l_t and action a_t , the likelihood ratio in the next period is $l_{t+1} = \phi(a_t, l_t; \hat{p})$, where

$$\phi(a, l; \hat{p}) = l \left(\frac{\psi(a|L, l; \hat{p})}{\psi(a|R, l; \hat{p})} \right). \quad (3)$$

The joint stochastic process $\langle a_t, l_t \rangle_{t=1}^\infty$ is a discrete-time Markov process defined on $A \times \mathbb{R}_+$ with $l_1 = 1$. Given state $\{a_t, l_t\}$, the process transitions to state $\{a_{t+1}, \phi(a_t, l_t; \hat{p})\}$ with probability $\psi(a_{t+1}|\omega, \phi(a_t, l_t; \hat{p}); p)$. The stochastic properties of this process determine long-run learning dynamics.

3.3 Local Stability of Limit Outcomes

At a stationary point, the likelihood ratio remains constant for any action that occurs with positive probability.

Definition 3. A point l is **stationary** if either (i) $\psi(a|\omega, l; p) = 0$ or (ii) $\phi(a, l; \hat{p}) = l$ for $a \in \{L, R\}$.

The next Lemma characterizes the set of stationary points.

Lemma 2. The set of stationary points are $\{0, \infty\}$.

A stationary point l is stable if the likelihood ratio process $\langle l_t \rangle$ converges to l with positive probability when l_1 is in the neighborhood of l .

Definition 4. Let $l \in [0, \infty)$ be a stationary point of $\langle l_t \rangle$. Then l is **stable** if there exists an open ball N_0 around 0 such that $l_1 - l \in N_0 \Rightarrow P(l_t \rightarrow l) > 0$. A point $l = \infty$ is stable if there exists an M such that $l_1 > M \Rightarrow P(l_t \rightarrow \infty) > 0$.

The challenge in establishing convergence results for $\langle l_t \rangle$ stems from the dependence of ψ and ϕ on the current value of the likelihood ratio. Corollary C.1 of [Smith and Sorensen \(2000\)](#) derives a criterion for the local stability of a nonlinear stochastic difference equation with state-dependent transitions. In the current setting, the stability of a stationary point can be reframed in the context of the log likelihood ratio. Suppose $\omega = R$. Then, given likelihood ratio l , the probability of action a is $\psi(a|R, l; p)$. Define

$$\gamma(\hat{p}, l) := \sum_{a \in \{L, R\}} \psi(a|R, l; p) \log |\phi_l(a, l; \hat{p})| \quad (4)$$

where

$$\phi_l(a, l; \hat{p}) = \frac{\psi(a|L, l; \hat{p})}{\psi(a|R, l; \hat{p})} + l \frac{d}{dl} \left(\frac{\psi(a|L, l; \hat{p})}{\psi(a|R, l; \hat{p})} \right) \quad (5)$$

is the derivative of ϕ with respect to l . From Corollary C.1 of [Smith and Sorensen \(2000\)](#), if l is a stationary point and $\gamma(\hat{p}, l) < 0$, then l is a stable stationary point.

A straightforward extension establishes that if l is a stationary point and $\gamma(\hat{p}, l) > 0$, then l is not a stable stationary point.

For intuition, consider the case where $\psi(a|\omega, l; \hat{p})$ is constant with respect to l and l is small.⁷ Then $E[\log l_{t+1}|l_t] \approx \log l_t + \gamma(\hat{p}, l_t)$. If $\gamma(\hat{p}, 0) < 0$ then $E[\log l_{t+1}|l_t] < \log l_t$ for l_t near zero and zero is a stable stationary point. If $\gamma(\hat{p}, 0) > 0$ then $E[\log l_{t+1}|l_t] > \log l_t$ for l_t near zero and zero is not a stable stationary point.

I use this criterion to characterize the relationship between \hat{p} and the stability of a stationary point. Suppose $\omega = R$. If informed agents sufficiently overestimate the share of uninformed agents, both zero and infinity are stable stationary points, whereas if agents sufficiently underestimate the share of uninformed agents, then there are no stable stationary points. When beliefs are close to correct, zero is the only stable stationary point. Lemma 3 formally states this result.

Lemma 3. *Suppose $\omega = R$. There exist unique cutoffs $\hat{p}_1 \in [0, p)$ and $\hat{p}_2 \in (p, 1]$ such that:*

1. *If $\hat{p} < \hat{p}_1$ then the set of stable stationary points are $\{0, \infty\}$.*
2. *If $\hat{p} \in (\hat{p}_1, \hat{p}_2)$ then $\{0\}$ is the unique stable stationary point.*
3. *If $\hat{p} > \hat{p}_2$, then there are no stable stationary points.*

Beliefs \hat{p} influence the information that accumulates from each action, but not the probability of each action. When l_t is near zero, beliefs favor state R . As \hat{p} increases, agents place more weight on the informativeness of contrary L actions and less weight on the informativeness of supporting R actions. This makes the likelihood ratio take a bigger jump away from zero when an L action is observed, and a smaller jump towards zero when an R action is observed. At some cut-off \hat{p}_2 , $E[\log l_{t+1}|l_t]$ changes from decreasing to increasing near zero. Above \hat{p}_2 , zero is no longer a stable stationary point. Similar logic establishes the stability of l_t near infinity for some cut-off \hat{p}_1 . When $\hat{p} = p$, the likelihood ratio is a martingale. Therefore, $\langle \log l_t \rangle$ is a supermartingale and zero is a stable stationary point, establishing that $p < \hat{p}_2$. The

⁷This holds when informed agents are in a cascade.

martingale convergence theorem precludes infinity from being a stable point when $\hat{p} = p$; this establishes that $p > \hat{p}_1$.

3.4 Global Convergence to Limit Outcomes

The next Lemma establishes that if ℓ is a stable stationary point, then the likelihood converges to ℓ with positive probability, from any initial value. The likelihood ratio almost surely does not converge to non-stable stationary points or non-stationary points.

Lemma 4. *For any initial value l_1 , $P(l_t \rightarrow \ell) > 0$ iff ℓ is a stable stationary point of $\langle l_t \rangle$.*

When agents have an inaccurate model of inference, $\hat{p} \neq p$, the likelihood ratio is no longer a martingale and it is not possible to use standard martingale methods to establish belief convergence. I use a two-pronged approach: the LIL establishes global convergence to stable stationary points and the LLN rules out convergence to non-stable stationary points.

Consider the case of bounded signals. The probability of each action is constant when the likelihood ratio is in the cascade set. Suppose a cascade persists. By the law of large numbers (LLN), the share of each action converges to its expected value, which determines the limit of the likelihood ratio. If this limit lies inside the cascade set, then by the LIL, there is a positive measure of sample paths that converge to this limit without leaving the cascade set. On the other hand, if the limit lies outside the cascade set, then the likelihood ratio will almost surely leave the cascade set, a contradiction. Precisely the same criterion determines whether the candidate limit lies inside or outside the cascade set and whether a stationary point is stable. Therefore, whenever a stationary point is stable, the likelihood ratio converges to this point with positive probability, from any initial value. The intuition is similar for the case of unbounded signals. I bound the likelihood ratio with a stochastic process that has state-independent transitions near the stable stationary point, and use the LIL to determine the limiting behavior of this second process.

3.5 Long Run Learning

This section presents the main result of the paper, a characterization of the learning dynamics in a misspecified model of inference. Several possible long-run learning outcomes may occur. Let *complete learning* denote the event where $l_t \rightarrow 0$, and *incorrect learning* denote the event where $l_t \rightarrow \infty$. *Non-stationary incomplete learning* refers to the event where l_t does not converge or diverge.⁸

Suppose there is at least one stable stationary point. Then there exists a random variable l_∞ such that $l_t \rightarrow l_\infty$. By Lemma 4, $\text{supp}(l_\infty)$ is equal to the set of stable stationary points. If 0 and ∞ are stable stationary points, then both complete and incorrect learning arise with positive probability, and incomplete learning almost surely does not occur. For complete learning to arise almost surely, 0 must be the unique stable stationary point. When there are no stable stationary points, then the likelihood ratio almost surely does not converge or diverge and learning is incomplete. If signals are bounded, cascades break almost surely. Theorem 1 formally characterizes the relationship between learning and model misspecification, using the cut-offs \hat{p}_1 and \hat{p}_2 derived in Lemma 3.

Theorem 1. *Suppose $\omega = R$. There exist unique cutoffs \hat{p}_1 and \hat{p}_2 such that*

1. *If $\hat{p} < \hat{p}_1$, then $l_t \rightarrow l_\infty$ almost surely, where l_∞ is a random variable with $\text{supp}(l_\infty) = \{0, \infty\}$.*
2. *If $\hat{p} \in (\hat{p}_1, \hat{p}_2)$, then $l_t \rightarrow 0$ almost surely.*
3. *If $\hat{p} > \hat{p}_2$, then l_t almost surely does not converge or diverge. Additionally, $P(l_t \notin \Lambda \text{ i.o.}) = 1$.*

When agents attribute too many actions to uninformed agents, they overestimate the informativeness of actions supporting the more likely state, and underestimate the informativeness of contrary actions, causing beliefs to quickly become entrenched. Both complete and incorrect learning outcomes arise in this situation. When agents

⁸*Stationary incomplete learning*, or the event where $l_t \rightarrow \ell$ for some $\ell \notin \{0, \infty\}$, is another type of incomplete learning. This does not occur in the current model.

attribute approximately the correct ratio of actions to uninformed agents, incorrect learning is no longer possible. Finally, when informed agents attribute too few actions to uninformed agents, they underestimate the informativeness of actions supporting the more likely state, and overestimate the informativeness of contrary actions. Neither correct or incorrect learning is possible.

It is also necessary to rule out incomplete learning when $\hat{p} < \hat{p}_2$. Consider the case of bounded signals. When a cascade persists with positive probability, the probability that the likelihood ratio returns to any value outside the cascade set is strictly less than one. Therefore, the probability that a value outside the cascade set occurs infinitely often is zero – eventually a cascade forms and persists. When a cascade persists and the likelihood ratio remains inside the cascade set, the LLN guarantees belief convergence.

Action convergence obtains for informed agents, in that they eventually choose the same action, if and only if the likelihood ratio converges or diverges. Action convergence never obtains for uninformed agents, as their actions always depend on their private information. Define a subsequence (a_{t_n}) to represent the actions of informed agents, where $t_n = \inf\{t > t_{n-1} | \theta_t = I\}$ and $t_0 = 0$. Then the following Corollary is an immediate consequence of Theorem 1.

Corollary 1. *Suppose $\omega = R$.*

1. *If $\hat{p} < \hat{p}_1$, then $a_{t_n} \rightarrow a_\infty$ almost surely, where a_∞ is a random variable with $\text{supp}(a_\infty) = \{L, R\}$.*
2. *If $\hat{p} \in (\hat{p}_1, \hat{p}_2)$, then $a_{t_n} \rightarrow R$ almost surely.*
3. *If $\hat{p} > \hat{p}_2$, then a_{t_n} almost surely does not converge.*

The asymptotic properties of learning determine whether the action choices of informed agents eventually converge to the optimal action. If complete learning obtains almost surely, then learning will be efficient in that informed agents will choose the optimal action all but finitely often. Otherwise, there is positive probability that learning will be inefficient and informed agents will choose the suboptimal action infinitely often.

4 Discussion

A misspecified model of information processing impacts asymptotic learning, and these results are robust to the addition of other information sources, such as an infinite stream of public signals or gurus (agents who know the state with probability 1). This insight has important policy implications. Suppose that a social planner can release additional public information. In a correctly specified model, this will affect the speed of learning, but will not impact asymptotic learning. However, in the face of model misspecification, the timing, frequency and strength of public information will play a key role in determining whether asymptotic learning obtains. When $\hat{p} < \hat{p}_1$, immediate release of public information prevents beliefs from becoming entrenched on the incorrect state. A delayed public response requires stronger or more frequent public signals to overturn an incorrect herd. Interventions are required on a short-term basis: once a herd begins on the correct action, it is likely to persist on its own (although another short-term intervention may be necessary in the future). When $\hat{p} > \hat{p}_2$, the important policy dimension is the frequency or strength of public information. As herds become more fragile, more frequent or precise public information is required to maintain a correct herd. An intervention must be long-term; once an intervention ceases, the herd will break.

Experimental evidence studying how people process correlated information supports this form of model misspecification. [Enke and Zimmermann \(2013\)](#) show that subjects treat correlated information as independent when updating, and beliefs are too sensitive to correlated information sources. In a social learning experiment, [Goree, Palfrey, Rogers, and McKelvey \(2007\)](#) find that new information continues to accumulate in cascades. Some agents still follow their private signal, despite the fact that all agents observe the history. In rational models, this off-the-equilibrium-path action would be ignored. However, it seems plausible that subsequent agents recognize these off-the-equilibrium-path actions reveal an agent's private signal, even if they are unsure of the exact prevalence of such actions. [Koessler, Ziegelmeyer, Bracht, and Winter \(2008\)](#) examine the fragility of cascades in an experiment where an expert receives a more precise signal than other participants. The unique Nash

equilibrium is for the expert to follow her signal, and observation of a contrary signal overturns a cascade. However, experts rarely overturn a cascade when equilibrium prescribes that they do so. As the length of the cascade increases, experts become even less likely to follow their signal: they break 65% of cascades when there are two identical actions, but only 15% of cascades when there are five or more identical actions. Elicited beliefs evolve in a manner similar to the belief process that would arise if all agents followed their signals, and each action conveyed private information. In addition, off-the-equilibrium-path play frequently occurs, and these non-equilibrium actions are informative. [Kubler and Weizsacker \(2004\)](#) also find evidence consistent with a misspecified model of social learning. They conclude that subjects do learn from their predecessors, but are uncertain about the share of previous agents who also learned from their predecessors. Particularly, agents underestimate the share of previous agents who herded and overestimate the amount of new information contained in previous actions. This provides support for both the presence of uninformed agents and a misspecified belief about their frequency.

This model leaves open several interesting questions. Individuals may differ in their depth of reasoning and their ability to combine different information sources - introducing heterogeneity into how agents process information would capture this. Allowing for partial observability of histories would also be a natural extension, while introducing payoff interdependencies would make the model applicable to election and financial market settings.

5 Appendix: Proofs

Proof of Lemma 1: Suppose $l \geq (1 - \underline{b}) / \underline{b}$. The strongest signal an agent can receive in favor of state R is \underline{b} . This leads to posterior $q(l, \underline{b}) = l \times \underline{b} / (1 - \underline{b}) \geq 1$ and an informed agent finds it optimal to choose $a = L$. Therefore, for any signal $s \geq \underline{b}$, an informed agent will choose action L. Similarly, if $l \leq (1 - \bar{b}) / \bar{b}$ an informed agent will choose action R for any signal $s \leq \bar{b}$. Q.E.D.

Proof of Lemma 2: At a stationary point l , $\phi(a, l) = l$ for all a such that

$\psi(a|\omega, l; p) > 0$. As $p < 1$ and uninformed agents are never in a cascade, $\psi(a|\omega, l; p) > 0$ for all $a \in \{L, R\}$ and for all $(\omega, l) \in \{L, R\} \times [0, \infty)$. Also, these actions are informative:

$$\left(\frac{\psi(a|L, l; \hat{p})}{\psi(a|R, l; \hat{p})} \right) \neq 1.$$

Therefore, $\{0, \infty\}$ are the only two values that satisfy $\phi(a, l) = l$ for all $a \in \{L, R\}$. Q.E.D.

The proof of Lemma 3 makes use of Corollary C.1 from Smith and Sorensen (2000), which is reproduced below using the notation of this paper.

Lemma 5 (Condition for Stable Fixed Point). *Given a finite set A , and Borel measurable functions $f : A \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ and $p : A \times \mathbb{R}_+ \rightarrow [0, 1]$ satisfying $\sum_{a \in A} p(a|x) = 1$. Let $x_1 \in \mathbb{R}$. Then the process $\langle x_t \rangle_{t=0}^\infty$ where $x_{t+1} = f(a_t, x_t)$ with probability $p(a_t|x_t)$ for $a_t \in A$ is a Markov process. Suppose $f(a, \cdot)$ is continuously differentiable and $p(a|\cdot)$ is continuous for all $a \in A$. Let \tilde{x} be a fixed point of x . If*

$$\sum_{a \in A} p(a|\tilde{x}) \log |f_x(a, \tilde{x})| < 0$$

then \tilde{x} is a stable fixed point.

Proof: See Corollary C.1 in Smith and Sorensen (2000).

Proof of Lemma 3:

Claim 1. For $a \in \{L, R\}$,

1. If $l > 1$, then $\frac{d}{d\hat{p}} \left(\frac{\psi(a|L, l; \hat{p})}{\psi(a|R, l; \hat{p})} \right) \leq 0$.
2. If $l < 1$, then $\frac{d}{d\hat{p}} \left(\frac{\psi(a|L, l; \hat{p})}{\psi(a|R, l; \hat{p})} \right) \geq 0$.
3. If $l = 1$, then $\frac{d}{d\hat{p}} \left(\frac{\psi(a|L, l; \hat{p})}{\psi(a|R, l; \hat{p})} \right) = 0$.

Suppose $a = R$. Then:

$$\frac{d}{d\hat{p}} \left(\frac{\psi(R|L, l; \hat{p})}{\psi(R|R, l; \hat{p})} \right) = \frac{F^L(1/(l+1))F^R(1/2) - F^L(1/2)F^R(1/(l+1))}{[\hat{p}F^R(1/(l+1)) + (1-\hat{p})F^R(1/2)]^2}$$

Given $F^L/F^R \leq dF^L/dF^R \leq (1-F^L)/(1-F^R)$, a standard inequality for the monotone likelihood ratio property (Smith and Sorensen 2013), $dF^L/dF^R \geq 0$. If $l > 1$, then $\frac{F^L(1/(l+1))}{F^R(1/(l+1))} \leq \frac{F^L(1/2)}{F^R(1/2)}$ which establishes (1) for $a = R$. The proof of the remaining cases and $a = L$ is analogous.

Claim 2. *There exists a $\hat{p}_2 \in (p, 1)$ such that $\ell = 0$ is a stable stationary point for $\hat{p} < \hat{p}_2$ and is not a stable stationary point for $\hat{p} > \hat{p}_2$.*

From Lemma 5, 0 is a stable limit point iff $\gamma(\hat{p}, 0) < 0$, where $\gamma(\hat{p}, l)$ is defined in Equation 4. At $l = 0$,

$$\begin{aligned} \gamma(\hat{p}, 0) &= \sum_{a \in \{L, R\}} \psi(a|R, 0; p) \log \left(\frac{\psi(a|L, 0; \hat{p})}{\psi(a|R, 0; \hat{p})} \right) \\ &= (1-p)(1-F^R(1/2)) \log \left(\frac{1-F^L(1/2)}{1-F^R(1/2)} \right) \\ &\quad + (p+(1-p)F^R(1/2)) \log \left(\frac{\hat{p}+(1-\hat{p})F^L(1/2)}{\hat{p}+(1-\hat{p})F^R(1/2)} \right) \end{aligned}$$

If $\hat{p} = 1$, then

$$\gamma(1, 0) = (1-p)(1-F^R(1/2)) \log \left(\frac{1-F^L(1/2)}{1-F^R(1/2)} \right) > 0$$

since $\frac{1-F^L(1/2)}{1-F^R(1/2)} > 1$, given $F^L(1/2) < F^R(1/2)$. If $\hat{p} = p$, then

$$\begin{aligned} \gamma(p, 0) &= \sum_{a \in \{L, R\}} \psi(a|R, 0; p) \log \left(\frac{\psi(a|L, 0; p)}{\psi(a|R, 0; p)} \right) \\ &< \log \left(\sum_{a \in \{L, R\}} \psi(a|R, 0; p) \left(\frac{\psi(a|L, 0; p)}{\psi(a|R, 0; p)} \right) \right) = \log \left(\sum_{a \in \{L, R\}} \psi(a|L, 0; p) \right) = 0 \end{aligned}$$

where the second line follows from the weighted arithmetic mean-geometric mean inequality. Finally,

$$\frac{d\gamma(\hat{p}, 0)}{d\hat{p}} = (p + (1-p)F^R(1/2)) \left(\frac{\psi(R|R, 0; \hat{p})}{\psi(R|L, 0; \hat{p})} \right) \frac{d}{d\hat{p}} \left(\frac{\psi(R|L, l; \hat{p})}{\psi(R|R, l; \hat{p})} \right) \geq 0$$

where the inequality follows from Claim 1. Therefore, $\gamma(\hat{p}, 0)$ is weakly increasing in \hat{p} . By continuity, there exists a unique $\hat{p}_2 \in (p, 1)$ such that $\gamma(\hat{p}_2, 0) = 0$. For $\hat{p} < \hat{p}_2$, $\gamma(\hat{p}, 0) < 0$ and 0 is a stable limit point, while for $\hat{p} > \hat{p}_2$, $\gamma(\hat{p}, 0) > 0$ and 0 is not a stable limit point.

Claim 3. *If*

$$p > p^* := 1 - \frac{\log \left(\frac{1-F^L(1/2)}{1-F^R(1/2)} \right)}{F^R(1/2) \left[\log \left(\frac{F^R(1/2)}{F^L(1/2)} \right) + \log \left(\frac{1-F^L(1/2)}{1-F^R(1/2)} \right) \right]}$$

then there exists a $\hat{p}_1 \in (0, p)$ such that $\ell = \infty$ is a stable stationary point for $\hat{p} < \hat{p}_1$ and is not a stable stationary point for $\hat{p} > \hat{p}_1$. If $p < p^$, then $\ell = \infty$ is not a stable stationary point for all \hat{p} .*

Consider the Markov process $\langle (a_t, x_t) \rangle$ with transitions

$$\begin{aligned}
\Psi(L|\omega, x; p) &= p \left(1 - F^\omega \left(\frac{x}{1+x} \right) \right) + (1-p)(1 - F^\omega(1/2)) \\
\Psi(R|\omega, x; p) &= p F^\omega \left(\frac{x}{1+x} \right) + (1-p) F^\omega(1/2) \\
\Phi(a, x; \hat{p}) &= x \left(\frac{\Psi(a|R, x; \hat{p})}{\Psi(a|L, x; \hat{p})} \right)
\end{aligned}$$

Note $x_t = 1/l_t$. The set of stationary points of $\langle x_t \rangle$ are $\{0, \infty\}$. Suppose $\omega = R$ and define

$$\Gamma(\hat{p}, x) := \sum_{a \in \{L, R\}} \Psi(a|R, x; p) \log |\Phi_x(a, x; \hat{p})|.$$

Then $x^* = 0$ is a stable stationary point iff $\Gamma(\hat{p}, 0) < 0$. Note $\Gamma(\hat{p}, 0) = -\gamma(\hat{p}, \infty)$. At $x = 0$,

$$\begin{aligned}
\Gamma(\hat{p}, 0) &= ((p + (1-p)(1 - F^R(1/2))) \log \left(\frac{\hat{p} + (1-\hat{p})(1 - F^R(1/2))}{\hat{p} + (1-\hat{p})(1 - F^L(1/2))} \right) \\
&\quad + (1-p)F^R(1/2) \log \left(\frac{F^R(1/2)}{F^L(1/2)} \right)).
\end{aligned}$$

If $\hat{p} = 0$, then

$$\Gamma(0, 0) = (1 - (1-p)F^R(1/2)) \log \left(\frac{1 - F^R(1/2)}{1 - F^L(1/2)} \right) + (1-p)F^R(1/2) \log \left(\frac{F^R(1/2)}{F^L(1/2)} \right)$$

which is less than zero when

$$p > 1 - \frac{\log \left(\frac{1-F^L(1/2)}{1-F^R(1/2)} \right)}{F^R(1/2) \left[\log \left(\frac{F^R(1/2)}{F^L(1/2)} \right) + \log \left(\frac{1-F^L(1/2)}{1-F^R(1/2)} \right) \right]} := p^*.$$

Suppose $\Gamma(p, 0) < 0$. Then $x_t \rightarrow 0$ with positive probability in the neighborhood of

0, and therefore $l_t \rightarrow \infty$ with positive probability. This is a contradiction, since l_t is a martingale when $\hat{p} = p$. Therefore, $\Gamma(p, 0) > 0$. Finally,

$$\frac{d\Gamma(\hat{p}, 0)}{d\hat{p}} = \Psi(L|R, 0; p) \frac{\Psi(L|L, 0; \hat{p})}{\Psi(L|R, 0; \hat{p})} \frac{d}{d\hat{p}} \left(\frac{\Psi(L|R, 0; \hat{p})}{\Psi(L|L, 0; \hat{p})} \right) \geq 0$$

where the inequality follows from Claim 1 and $\Psi(a|\omega, 0; \hat{p}) = \psi(a|\omega, \infty; \hat{p})$. Therefore, $\Gamma(\hat{p}, 0)$ is weakly increasing in \hat{p} .

When $p > p^*$, by continuity there exists a unique $\hat{p}_1 \in (0, p)$ such that $\Gamma(\hat{p}_1, 0) = 0$. For $\hat{p} < \hat{p}_1$, $\Gamma(\hat{p}, 0) < 0$ and 0 is a stable limit point of $\langle x_t \rangle$, while for $\hat{p} > \hat{p}_1$, $\Gamma(\hat{p}, 0) > 0$ and 0 is not a stable limit point. When $p < p^*$, then $\Gamma(\hat{p}, 0) > 0$ for all \hat{p} and 0 is not a stable limit point of $\langle x_t \rangle$. Note that ∞ is a stable limit point of $\langle l_t \rangle$ when 0 is a stable limit point of $\langle x_t \rangle$

Claim 4. $\hat{p}_1 < \hat{p}_2$.

This follows immediately from the fact that $\hat{p}_1 < p$ and $\hat{p}_2 > p$.

Proof of Lemma 4:

Claim 5. *If ℓ is not a stationary point of $\langle l_t \rangle$, then $P(l_t \rightarrow \ell) = 0$.*

Theorem B.1 in [Smith and Sorensen \(2000\)](#) establishes that a martingale cannot converge to a non-stationary point; the same result applies to the Markov process $\langle l_t \rangle$. Therefore, if $P(l_t \rightarrow \ell) > 0$, then $\ell \in \{0, \infty\}$.

For the remainder of the proof, suppose $\omega = R$. Define $g(a, l) = \frac{\psi(a|L, l; \hat{p})}{\psi(a|R, l; \hat{p})}$ and $\rho(a, l) = \psi(a|R, l; p)$ given $\psi(\cdot|\cdot)$ from Equation 1. Using this notation, $l_{t+1} = l_t \times g(a_t, l_t)$ and $\gamma(\hat{p}, 0) = \rho(R, 0) \log g(R, 0) + \rho(L, 0) \log g(L, 0)$, where $\gamma(\cdot)$ is defined in Equation 4.

Case I: Bounded Private Beliefs

Let $(\alpha_1, \alpha_2, \dots)$ be an i.i.d. sequence of random variables with

$$\alpha_t = \begin{cases} L & \text{if } (\theta_t = U \text{ and } s_t \geq 1/2) \\ R & \text{if } (\theta_t = I) \text{ or } (\theta_t = U \text{ and } s_t < 1/2). \end{cases}$$

Then α_t corresponds to the action that is chosen if informed agents are in an R-cascade in period t , with $P(\alpha_t = a) = \rho(a, 0)$. Note $E[\log g(\alpha_t, 0)] = \gamma(\hat{p}, 0)$ and let $\sigma^2 := \text{Var}(\log g(\alpha_t, 0))$. Define a sequence of random variables (X_1, X_2, \dots) where

$$X_t := \frac{\log g(\alpha_t, 0) - \gamma(\hat{p}, 0)}{\sigma}.$$

Then (X_1, X_2, \dots) are i.i.d random variables with mean zero and variance one.

Fix $\tau \geq 1$. By the Law of the Iterated Logarithm ([Hartman and Wintner 1941](#)),

$$\limsup_{t \rightarrow \infty} \frac{\sum_{i=\tau}^t X_i}{\sqrt{2(t-\tau+1) \log \log(t-\tau+1)}} = 1 \text{ a.s.}$$

Thus, for all $\varepsilon > 0$,

$$P \left[\frac{1}{t-\tau+1} \sum_{i=\tau}^t \log g(\alpha_i, 0) \geq \gamma(\hat{p}, 0) + \beta_{t-\tau+1} \text{ i.o.} \right] = 0$$

where $\beta_t := (1 + \varepsilon) \sigma \sqrt{\frac{2 \log \log t}{t}}$.

Claim 6. *Let*

$$\zeta_\tau := \left\{ \{\alpha_\tau, \alpha_{\tau+1}, \dots\} \mid \frac{1}{t-\tau} \sum_{i=\tau}^{t-1} \log g(\alpha_i, 0) < \gamma(\hat{p}, 0) + \beta_{t-\tau} \text{ for all } t \geq \tau \right\}$$

be the set of sample paths such that $\frac{1}{t-\tau} \sum_{i=\tau}^{t-1} \log g(\alpha_i, 0)$ never exceeds $\gamma(\hat{p}, 0) + \beta_{t-\tau}$. Then $P(\zeta_\tau) \geq 1/2$.

The complement of ζ_τ , denoted ζ_τ^c , represents the set of sample paths that exceed $\gamma(\hat{p}, 0) + \beta_{t-\tau}$ at least once. For each $\alpha = \{\alpha_\tau, \alpha_{\tau+1}, \dots\} \in \zeta_\tau^c$, form a corresponding sample path α' with $\alpha'_t \neq \alpha_t$ for each t such that $\frac{1}{t-\tau} \sum_{i=\tau}^{t-1} \log g(\alpha_i, 0) > \gamma(\hat{p}, 0) + \beta_{t-\tau}$. Then each $\alpha \in \zeta_\tau^c$ has a unique corresponding sample path $\alpha' \in \zeta_\tau$. Therefore, $P(\zeta_\tau) \geq P(\zeta_\tau^c)$, which implies $P(\zeta_\tau) \geq 1/2$. This establishes that the set of sample paths such that $\frac{1}{t-\tau} \sum_{i=\tau}^{t-1} \log g(\alpha_i, 0)$ never crosses its upper bound has positive measure.

Claim 7. *If $\ell \in \{0, \infty\}$ is a stable point, then $P(l_t \rightarrow \ell) > 0$ from any initial value l_1 .*

Let τ_1 be the smallest t such that an R-cascade forms,

$$\tau_1 = \inf \{t \geq 1 | l_t \in [0, (1 - \bar{b}) / \bar{b}]\}.$$

Fixing l_1 , let $\eta < \infty$ be the number of consecutive R actions required to start a cascade. Then the probability that an R-cascade forms in η periods is $P(\tau_1 = \eta) > ((1 - p)F^R(1/2))^\eta > 0$, where the middle term is the probability of η uninformed agents choosing R .

If $\sum_{i=\tau_1}^j \log g(\alpha_i, 0) < 0$ for all $j \in \{\tau_1, \dots, t-1\}$, then the cascade has not broken in period t and

$$\log l_t = \log l_{\tau_1} + \sum_{i=\tau_1}^{t-1} \log g(\alpha_i, 0).$$

From Claim 6, we know that

$$P \left(\sum_{i=\tau_1}^{t-1} \log g(\alpha_i, 0) < (t - \tau_1)(\gamma(\hat{p}, 0) + \beta_{t-\tau_1}) \forall t > \tau_1 \right) \geq 1/2.$$

Suppose $\ell = 0$ is a stable point, $\gamma(\hat{p}, 0) < 0$. Let $\tau_2 = \inf \{t > \tau_1 | \gamma(\hat{p}, 0) + \beta_{t-\tau_1} < 0\}$ be the smallest value of $t > \tau_1$ such that $\gamma(\hat{p}, 0) + \beta_{t-\tau_1}$ is negative. This is well-defined since $\beta_t \rightarrow 0$. Between periods τ_1 and τ_2 , the probability that the likelihood ratio remains in the R-cascade set is

$$\begin{aligned} & P(l_t \in [0, (1 - \bar{b}) / \bar{b}] \forall t \in \{\tau_1, \dots, \tau_2 - 1\}) \\ & > P \left(\sum_{i=\tau_1}^{t-1} \log g(\alpha_i, 0) < 0 \forall t \in \{\tau_1, \dots, \tau_2 - 1\} \right) \\ & > \rho(R, 0)^{\tau_2 - \tau_1} > 0 \end{aligned}$$

where $\rho(R, 0)^{\tau_2 - \tau_1}$ is the probability of all R actions during this interval. The prob-

ability that an R-cascade doesn't break at any $t > \tau_1$ is

$$\begin{aligned}
& P(l_t \in [0, (1 - \bar{b})/\bar{b}] \ \forall t > \tau_1) \\
& \geq P\left(\sum_{i=\tau_1}^{t-1} \log g(\alpha_i, 0) < 0 \ \forall t > \tau_1\right) \\
& \geq P\left(\sum_{i=\tau_1}^{t-1} \log g(\alpha_i, 0) < \min\{0, (t - \tau_1)(\gamma(\hat{p}, 0) + \beta_{t-\tau_1})\} \ \forall t \geq \tau_1\right) \\
& \geq \frac{1}{2}\rho(R, 0)^{\tau_2 - \tau_1} > 0.
\end{aligned}$$

where the final line follows from the probability of all R actions between periods τ_1 and τ_2 and the probability of not exceeding the LIL bound after τ_2 (Claim 6). Therefore, the probability that an R-cascade forms and never breaks is at least

$$\frac{1}{2}((1 - p)F^R(1/2))^\eta \rho(R, 0)^{\tau_2 - \eta} > 0.$$

Let $\tau_3 = \inf\{t \geq 1 | l_s \in [0, (1 - \bar{b})/\bar{b}] \ \forall s \geq t\}$ be the period in which an R-cascade forms and never breaks. We just established that $P(\tau_3 < \infty) > 0$. Also, when $\tau_3 < \infty$, $l_t = l_{\tau_3} \prod_{i=\tau_3}^{t-1} g(\alpha_i, 0) \rightarrow 0$. Therefore, $P(l_t \rightarrow 0) > 0$. The proof for $\ell = \infty$ is analogous.

Claim 8. *If $\ell \in \{0, \infty\}$ is not a stable point, then $P(l_t \rightarrow \ell) = 0$ from any initial value l_1 .*

1. *Suppose $\ell = 0$ is not stable and there exists a $\tau_1 < \infty$ such that $l_{\tau_1} \in [0, (1 - \bar{b})/\bar{b}]$. Let $\tau_2 = \inf\{t > \tau_1 | l_{\tau_2} \notin [0, (1 - \bar{b})/\bar{b}]\}$. Then $P(\tau_2 < \infty) = 1$.*
2. *Suppose $\ell = \infty$ is not stable and there exists a $\tau_1 < \infty$ such that $l_{\tau_1} \in [(1 - \underline{b})/\underline{b}, \infty]$. Let $\tau_2 = \inf\{t > \tau_1 | l_{\tau_2} \notin [(1 - \underline{b})/\underline{b}, \infty]\}$. Then $P(\tau_2 < \infty) = 1$.*

Suppose $\ell = 0$ is not a stable point, so $\gamma(\hat{p}, 0) > 0$. Let τ_3 be the period in which an R-cascade forms and never breaks, $\tau_3 = \inf\{t \geq 1 | l_s \in [0, (1 - \bar{b})/\bar{b}] \ \forall s \geq t\}$.

Suppose $P(\tau_3 < \infty) > 0$. Then, given $\tau_3 < \infty$,

$$\log l_t = \log l_{\tau_3} + \sum_{i=\tau_3}^{t-1} \log g(\alpha_i, 0) \leq \log(1 - \bar{b}) / \bar{b}$$

for all $t > \tau_3$. Then it must be the case that

$$\lim_{t \rightarrow \infty} \frac{1}{t - \tau_3} \sum_{i=\tau_3}^{t-1} \log g(\alpha_i, 0) \leq 0.$$

But by the Law of Large Numbers,

$$\lim_{t \rightarrow \infty} \frac{1}{t - \tau_3} \sum_{i=\tau_3}^{t-1} \log g(\alpha_i, 0) = \gamma(\hat{p}, 0) > 0 \quad a.s.$$

a contradiction. Therefore the cascade must break almost surely, and $P(l_t \rightarrow 0) = 0$. The proof for $\ell = \infty$ is analogous.

Case II: Unbounded Private Beliefs

Claim 9. *If $\ell \in \{0, \infty\}$ is a stable point, then $P(\lim_{t \rightarrow \infty} l_t = \ell) > 0$ from any initial value l_1 .*

Let 0 be a stable stationary point, $\gamma(\hat{p}, 0) < 0$. I proceed in three steps: (i) construct a stochastic process $\langle \lambda_t \rangle_{t=1}^{\infty}$ with state-independent transitions near 0 and 0 as a stable stationary point, (ii) apply the techniques from Claims 6-7 (which require state-independent transitions near the stable point) to show that $P(\lambda_t \rightarrow 0) > 0$ from any initial value, and (iii) show that $P(l_t \leq \lambda_t \forall t | \lambda_t \rightarrow 0) > 0$, and therefore, $P(l_t \rightarrow 0) > 0$.

By continuity of ψ , it is possible to find an $M \in (0, 1)$ such that

$$\rho(L, M) \log g(L, x) + \rho(R, M) \log g(R, y) < 0.$$

for all $x, y \in [0, M]$. Choose $\ell_L, \ell_R \in [0, M]$ such that

$$\ell_L = \arg \max_{l \in [0, M]} g(L, l)$$

and

$$\ell_R = \arg \max_{l \in [0, M]} g(R, l).$$

Let (ν_1, ν_2, \dots) be an i.i.d. sequence of random variables with

$$\nu_t = \begin{cases} L & \text{if } (\theta_t = I \text{ and } s_t \geq s^*(M)) \text{ or } (\theta_t = U \text{ and } s_t \geq 1/2) \\ R & \text{if } (\theta_t = I \text{ and } s_t < s^*(M)) \text{ or } (\theta_t = U \text{ and } s_t < 1/2). \end{cases}$$

Then ν_t corresponds to the action that is chosen if $l_t = M$, with $P(\nu_t = a) = \rho(a, M)$.

Define a stochastic process $\langle \lambda_t \rangle_{t=1}^{\infty}$ where $\lambda_1 = l_1$ and

$$\lambda_{t+1} = \begin{cases} \lambda_t \times g(L, \ell_L) & \text{if } \nu_t = L \text{ and } \lambda_t \leq M \\ \lambda_t \times g(R, \ell_R) & \text{if } \nu_t = R \text{ and } \lambda_t \leq M \\ \lambda_t \times g(a_t, \lambda_t) & \text{if } \lambda_t > M \end{cases}$$

where a_t is the action chosen in period t . By Lemma 5, 0 is a stable fixed point of $\langle \lambda_t \rangle_{t=1}^{\infty}$ if

$$\rho(L, M) \log g(L, \ell_L) + \rho(R, M) \log g(R, \ell_R) < 0.$$

Given $\ell_R, \ell_L \in [0, M]$ and the definition of M , this condition holds.

Let $\tau_1 = \inf\{t \geq 1 \mid \lambda_t < M\}$ be the first time that $\lambda_t < M$ and $\tau_2 = \inf\{t > \tau_1 \mid \lambda_t > M\}$ be the first time after τ_1 that $\lambda_t > M$. Similar arguments to the case of bounded private beliefs (replacing the R-cascade set with $[0, M]$ and α_t with ν_t) can be used to establish that $P(\tau_2 = \infty) > 0$ and $P(\lambda_t \rightarrow 0) > 0$.

Note $g(R, x) \leq g(R, \ell_R)$, $g(L, x) \leq g(L, \ell_L)$ and $g(R, x) < g(L, \ell_L)$ for $x \in [0, M]$ and it is never the case that $(a_t, \nu_t) = (L, R)$ when $t < \tau_2$. Therefore, $l_t \leq \lambda_t$ when $t < \tau_2$. If $\tau_2 = \infty$, then $l_t \rightarrow 0$. Therefore, $P(l_t \rightarrow 0) > 0$.

Claim 10. *If $\ell \in \{0, \infty\}$ is not a stable point, then $P(l_t \rightarrow \ell) = 0$ from any initial*

value l_1 .

A similar argument can be used to bound l_t from below and establish that $P(l_t \rightarrow \ell) = 0$ when $\gamma(\hat{p}, 0) > 0$. Q.E.D.

Proof of Theorem 1:

Claim 11. *If $\hat{p} < \hat{p}_1$, then $l_t \rightarrow l_\infty$ almost surely, where l_∞ is a random variable with $\text{supp}(l_\infty) = \{0, \infty\}$.*

Suppose $\hat{p} < \hat{p}_1$. Lemma 3 established that the set of stable stationary points is $\{0, \infty\}$, and Lemma 4 established that $P(l_t \rightarrow \ell) > 0$ iff ℓ is a stable stationary point. Thus, we need to rule out non-stationary incomplete learning to show that there exists an l_∞ with $\text{supp}(l_\infty) = \{0, \infty\}$ such that $l_t \rightarrow l_\infty$ almost surely.

Suppose private beliefs are bounded. Let $\tau_1 = \inf\{t \geq 1 | l_t \in \Lambda\}$ be the first time that the likelihood ratio enters the cascade set and let $\tau_2 = \inf\{t > \tau_1 | l_t \notin \Lambda\}$ be the first time that the likelihood ratio leaves the cascade set. For any l_1 , $P(\tau_1 < \infty) = 1$ and by Lemma 4, $P(\tau_2 < \infty) < 1$. Therefore, $P(l_t \notin \Lambda \text{ i.o.}) = 0$. Let $\tau_3 = \inf\{t \geq 1 | l_s \in \Lambda \forall s \geq t\}$. Then $P(\tau_3 < \infty) = 1$ and the likelihood ratio almost surely eventually remains in the cascade set. Lemma 4 established belief convergence on any sample path that remains in the cascade set. Thus, there exists a random variable l_∞ with $\text{supp}(l_\infty) = \{0, \infty\}$ such that $l_t \rightarrow l_\infty$ almost surely.

Similar logic establishes that when private beliefs are unbounded, $P(l_t \in (\ell_1, \ell_2) \text{ i.o.}) = 0$ for any $(\ell_1, \ell_2) \subset (0, 1)$, and therefore, there exists a τ such that $P(l_t \in [0, \ell_1] \cup [\ell_2, 1] \forall t > \tau) = 1$. Choosing ℓ_1 small enough and ℓ_2 large enough can be used to establish convergence.

Claim 12. *If $\hat{p} \in (\hat{p}_1, \hat{p}_2)$, then $l_t \rightarrow 0$ almost surely.*

Similar logic to Claim 11, substituting $[0, (1 - \bar{b})/\bar{b}]$ for Λ in the case of bounded private beliefs and setting $\ell_2 = 1$ in the case of unbounded private beliefs establishes the claim.

Claim 13. *If $\hat{p} > \hat{p}_2$, then l_t almost surely does not converge or diverge.*

When $\hat{p} > \hat{p}_2$, the likelihood ratio neither converges to 0 nor diverges to ∞ . These are the only two candidate limit points; therefore, the likelihood ratio does not converge and learning is incomplete.

$P(l_t \notin \Lambda \text{ i.o.}) = 1$ follows immediately from Claim 8 for bounded private beliefs, and from Claim 13 for unbounded private beliefs. Q.E.D.

References

- ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): “Bayesian Learning in Social Networks,” *The Review of Economic Studies*.
- BANERJEE, A. V. (1992): “A Simple Model of Herd Behavior,” *The Quarterly Journal of Economics*, 107, 797–817.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *The Journal of Political Economy*, 100, 992–1026.
- BOHREN, A. (2012): “Information and Incentives in Stochastic Games, Social Learning and Crowdsourcing,” Ph.D. thesis, University of California, San Diego, chapter 1.
- CAMERER, C. F., T.-H. HO, AND J.-K. CHONG (2004): “A Cognitive Hierarchy Model of Games,” *The Quarterly Journal of Economics*, 119, 861–898.
- COSTA-GOMES, M. A., V. P. CRAWFORD, AND N. IRIBERRI (2009): “Comparing Models of Strategic Thinking in Van Huyck, Battalio, and Beil’s Coordination Games,” *Journal of the European Economic Association*, 7, 365–376.
- DEMARZO, P. M., D. VAYANOS, AND J. ZWIEBEL (2003): “Persuasion Bias, Social Influence, And Unidimensional Opinions,” *The Quarterly Journal of Economics*, 118, 909–968.

- ENKE, B. AND F. ZIMMERMANN (2013): “Correlation Neglect in Belief Formation,”
- EPSTEIN, L. G., J. NOOR, AND A. SANDRONI (2010): “Non-Bayesian Learning,” *The B.E. Journal of Theoretical Economics*, 10.
- EYSTER, E. AND M. RABIN (2005): “Cursed Equilibrium,” *Econometrica*, 73, 1623–1672.
- (2010): “Nave Herding in Rich-Information Settings,” *American Economic Journal: Microeconomics*, 2, 221–43.
- GOEREE, J. K., T. R. PALFREY, B. W. ROGERS, AND R. D. MCKELVEY (2007): “Self-Correcting Information Cascades,” *Review of Economic Studies*, 74, 733–762.
- GUARINO, A. AND P. JEHIEL (2013): “Social Learning with Coarse Inference,” *American Economic Journal: Microeconomics*, 5, 147–74.
- HARTMAN, P. AND A. WINTNER (1941): “On the Law of the Iterated Logarithm,” *American Journal of Mathematics*, 63, 169–176.
- KOESSLER, F., A. ZIEGELMEYER, J. BRACHT, AND E. WINTER (2008): “Fragility of Information Cascades: An Experimental Study using Elicited Beliefs,” Jena Economic Research Papers in Economics 2008-094, Friedrich-Schiller-University Jena, Max-Planck-Institute of Economics.
- KUBLER, D. AND G. WEIZSACKER (2004): “Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory,” *Review of Economic Studies*, 71, 425–441.
- MOSCARINI, G., M. OTTAVIANI, AND L. SMITH (1998): “Social learning in a changing world,” *Economic Theory*, 11, 657–665.
- SMITH, L. AND P. SORENSEN (2000): “Pathological Outcomes of Observational Learning,” *Econometrica*, 68, 371–398.

SMITH, L. AND P. N. SORESENSEN (2013): “Rational Social Learning with Random Sampling,” .