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 14-007

"Informational Herding with Model Misspecification"

by

J. Aislinn Bohren

Informational Herding with Model Misspecification*

J. Aislinn Bohren†
University of Pennsylvania

June 15, 2013

Abstract

This paper demonstrates that a misspecified model of information processing interferes with long-run learning and offers an explanation for why individuals may continue to choose an inefficient action, despite sufficient public information to learn the true state. I consider a social learning environment where agents draw inference from private signals, public signals and the actions of their predecessors, and sufficient public information exists to achieve asymptotically efficient learning. Prior actions aggregate multiple sources of information; agents face an inferential challenge to distinguish new information from redundant information. I show that when individuals significantly overestimate the amount of new information contained in prior actions, beliefs about the unknown state become entrenched and incorrect learning may occur. On the other hand, when individuals sufficiently overestimate the amount of redundant information, beliefs are fragile and learning is incomplete. When agents have an approximately correct model of inference, learning is complete – the model with no information-processing bias is robust to perturbation.

*I thank Vince Crawford, Ernesto Dal Bo, Frederic Koessler, 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 or cars or 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, and how departures from Bayesian models of learning affect long-run outcomes. 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 even when contradicted by public information. The results have important implications for policies aimed at counteracting inefficient social choices; in the face of information processing errors, the timing and frequency of public information campaigns becomes particularly relevant for long-run efficiency.

Individuals face an inferential challenge when extracting information from the actions of others. An action often aggregates multiple sources of information; full rationality requires agents to parse out the new information and discard redundant information. A critical feature of standard observational learning models in the tradition of Smith and Sorensen (2000) is common knowledge of how individuals process information. 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 to draw 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 have access to public information about the state, including public signals and 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 overweight 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, $\widehat{p}_1$ and $\widehat{p}_2$, such that when $\widehat{p} < \widehat{p}_1$ both correct and fully incorrect learning occur, and when $\widehat{p} > \widehat{p}_2$, beliefs about the state perpetually fluctuate, rendering learning incomplete. Both cases admit the possibility of inefficient learning: informed agents can continue to choose the suboptimal action, despite observing an infinite sequence of new information. When $\widehat{p}$ falls between these two thresholds, $\widehat{p} \in (\widehat{p}_1, \widehat{p}_2)$, correct learning obtains and informed agents will eventually choose the optimal action. Correct beliefs about agent types lead to efficient learning, $p \in (\widehat{p}_1, \widehat{p}_2)$.

Fully incorrect learning or incomplete learning is possible when $\widehat{p} \neq p$ 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 $\hat{p} > \hat{p}_2$. While I describe this approach in the framework of the model misspecification outlined above, it is general in that it could easily be utilized to examine other forms of model misspecification.

Model misspecification has important policy implications. To illustrate the relevance of this result, 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 their 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.

Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) first model social learning in a sequential setting with binary signals. 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 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

[1] This example abstracts from the payoff interdependencies of vaccines.

explanations for the robustness of inefficient learning. Eyster and Rabin (2010) also embed inferentially naive agents in a model with rational agents. When every nth 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.

The recent initial response models, including level-k analysis and cognitive hierarchy models, are similar in spirit to this paper.[2] 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 $\widehat{p} = 0$. 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 Model

The basic set-up of this model mirrors a standard sequential learning environment with binary action and state spaces, and a continuous signal space.

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

[2] Costa-Gomes, Crawford, and Iriberri (2009). Camerer, Ho, and Chong (2004).

of the game. A countably infinite set of agents $T = \{1, 2, ...\}$ 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: $u(a_t, \omega) = 1_{a_t = \omega}$.

Information. Each agent privately observes a signal about the state of the world, s_t. Signals are independent and identically distributed, conditional on the state, with conditional c.d.f. F^ω. Assume that F^L, F^R are differentiable and mutually absolutely continuous, with Radon Nikodym derivative $f = dF^L/dF^R$. This ensures that no private signal perfectly reveals the state, and allows for a normalization such that the signal $s_t = P(\omega = L|s_t)$ corresponds to the private belief that the state is L. Let $[\underline{b}, \bar{b}]$ represent the convex hull of this support. Beliefs are bounded if $0 < \underline{b} < \bar{b} < 1$, and are unbounded if $[\underline{b}, \bar{b}] = [0, 1]$. To ensure that signals are informative, assume that the Radon Nikodym derivative $f(s) = dF^L(s)/dF^R(s)$ is increasing in s. This is equivalent to the monotone likelihood ratio property.

Public signals also reveal information about the state. Each period, a public signal s_t^p is released with probability $\varepsilon > 0$. For simplicity, assume that this public signal is binary, $s_t^p \in \{L, R\}$ with precision $\pi^p = P(s_t^p = \omega) \in (1/2, 1)$. Denote the event in which no public signal is released as $s_t^p = \emptyset$.

Agent Types. There are two types of agents, $\theta \in \{I, U\}$. With probability $p \in (0, 1)$, an agent is a socially informed type I who observes the action choices of her predecessors and the sequence of public signals, $h_t = (a_1, ..., a_{t-1}; s_1^p, ..., s_t^p)$. 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. Individuals face an inferential challenge when extracting information from the actions of others. An action contains information about both an agent's private signal and prior actions and public signals. Full rationality requires informed agents to parse out the private information in an action, and discard the redundant information about other actions and public signals. But what happens if informed agents misunderstand the link between action choices and information? For example, what if informed agents realize that prior actions contain repeated information, but do not fully understand how to disentangle the new and redundant information? Or what if informed agents know how to draw inference from the private information contained in an action, but fail to realize that these actions also contain repeated information?

Motivated by this possibility, I introduce higher-order uncertainty over the level of information possessed by other agents to allow for a distinction between the perceived and actual informational content of actions. Formally, each informed individual believes that any 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.

Remark. Although 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. The working paper Bohren (2010) elaborates on a model in which agents begin with a non-degenerate prior distribution over p, and learn about p

from the action history. Extending the model to allow for heterogenous biases is left for future research.

2.1 The Individual Decision-Problem

Before choosing an action, each agent observes his type θ and a private signal s. Informed agents also observe the public history h. A decision rule specifies an action a 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 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. This outcome departs from the Bayesian equilibrium concept because agents interpret the history through the lens of their potentially misspecified beliefs about the share of informed agents.

Throughout the analysis, it is convenient to express the public belief of informed agents as a likelihood ratio,

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

which depends on the history and beliefs about the share of informed agents. After receiving private signal s, an informed agent updates this public belief to the private belief $l_I = l\left(\frac{s}{1-s}\right)$, while an uninformed agent bases his private posterior belief solely on his private signal, $l_U = \left(\frac{s}{1-s}\right)$.

Guided by posterior belief l_θ, each agent maximizes her payoff by choosing $a = L$ if $l_\theta \geq 1$, and $a = R$ otherwise. An agent's decision can be represented as a cut-off rule. Informed agents have a signal cut-off $s_I(l) = 1/(l+1)$, such that the agent chooses action L when $s \geq s_I(l)$ and chooses action R otherwise. The signal threshold for uninformed agents is independent of the likelihood ratio, $s_U = 1/2$.

2.2 Cascade formation

An *information cascade* occurs when it is optimal for an agent to choose the same action regardless of her private signal realization; therefore, this action reveals no private information. As is standard in the literature, a cascade occurs when public beliefs outweigh the strongest private belief.

Lemma 1. *An informed agent is in an information cascade when $l \geq (1-\underline{b})/\underline{b}$ or $l \leq \left(1-\bar{b}\right)/\bar{b}$. Uninformed agents are never in an information cascade.*

As in Smith and Sorensen (2000), the cascade set of informed agents is an interval for bounded private beliefs, $[0,(1-\underline{b})/\underline{b}] \cup \left[\left(1-\bar{b}\right)/\bar{b}, 1\right]$ and a set of points for unbounded private beliefs, $\{0,1\}$.

With bounded private beliefs, a cascade forms in finite time.[3] Informed agents no longer reveal private information in a cascade, but public information continues to aggregate from uninformed agents and public signals. Therefore, all cascades break with positive probability, and the formation of a cascade does not necessarily imply belief convergence. I say a cascade *persists in the limit* if, when a cascade forms in period t, there exists no $\tau > t$ where the

[3] Public signals and uninformed agents allow beliefs to jump inside the cascade set.

cascade breaks. When private beliefs are unbounded, a cascade does not form in finite time and a weaker notion is necessary. A *limit cascade* occurs when beliefs converge to a point in the cascade set, but do not necessarily enter the cascade set in finite time. For bounded private beliefs, informed agents will *herd* when a cascade persists in the limit, whereas for unbounded private beliefs, informed agents herd when a limit cascade forms.

3 Learning Dynamics

I next characterize the long-run behavior of a misspecified model of inference, and determine how the failure of agents to accurately interpret the action history will interfere with the long-run learning dynamics and action choices of agents. Represent the likelihood ratio $\langle l_t \rangle_{t=1}^{\infty}$ as a stochastic process with $l_0 = 1$ and transitions:

$$\begin{aligned}
\psi(a = L|\omega, l; p) &= p(1 - F^{\omega}(1/(l+1))) + (1-p)(1 - F^{\omega}(1/2)) \\
\psi(a = R|\omega, l; p) &= pF^{\omega}(1/(l+1)) + (1-p)F^{\omega}(1/2) \\
\psi^p(s^p = \omega|\omega) &= \varepsilon\pi^p \\
\psi^p(s^p \neq \omega|\omega) &= \varepsilon(1 - \pi^p) \\
\psi^p(s^p = \emptyset|\omega) &= 1 - \varepsilon \\
\phi(a, s^p, l; \hat{p}) &= l\left(\frac{\psi(a|L, l; \hat{p})}{\psi(a|R, l; \hat{p})}\right)\left(\frac{\psi^p(s^p|L)}{\psi^p(s^p|R)}\right)
\end{aligned}$$

where $\psi(a|\omega, l; p)$ denotes the probability of action a, given likelihood ratio l and state ω, and $\psi^p(s^p|\omega)$ denotes the probability of public signal s^p, given state ω. The probability of action a is a weighted average of the probability that an uninformed type chooses a when using cut-off rule $s_U = 1/2$ and the probability that an informed type chooses a using cut-off rule $s_I(l) = 1/(l+1)$.

The likelihood ratio is updated based on the perceived probability of action a, $\psi(a|\omega, l; \hat{p})$, which depends on agents' beliefs about the share of informed agents. 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, action a_t and public signal s_t^p, the likelihood ratio in the next period is $l_{t+1} = \phi(a_t, s_t^p, l_t; \hat{p})$.

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

3.1 Stationary and Stable Limit Points

If agents accurately interpret the action history, $\hat{p} = p$, then the likelihood ratio is a martingale, as is standard in the literature. However, when agents have an inaccurate model of inference, the likelihood ratio may no longer be a martingale and it is not possible to use standard martingale methods to establish belief convergence. I first characterize the set of stationary points for the likelihood ratio; these are candidate limit points for $\langle l_t \rangle$. Next, I determine how the stability of these stationary points depends on $\hat{p}$ to establish the dynamics of the likelihood ratio near a stationary point. Finally, 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 starting point, and almost surely doesn't converge to unstable points.

At a stationary point, the likelihood ratio remains constant for any action and public signal profile that occur with positive probability.

Definition 1. *A point $\widetilde{l}$ is **stationary** if either (i) $\psi(a|\omega, l; p) = 0$ and $\psi^p(s^p|\omega) = 0$ or (ii) $\phi(a, s^p, l; \hat{p}) = l$ for all $(a, s^p) \in \{L, R\} \times \{L, R, \emptyset\}$.*

Public signals are always informative; therefore, the only stationary points are 0 and ∞. This is also true without public signals, provided at least one of the following is true: (i) signals are unbounded, (ii) there are uninformed agents or (iii) agents believe there are uninformed agents. Lemma 2 formalizes this result.

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

A stationary point $\widetilde{l}$ is stable if the likelihood ratio process $\langle l_t \rangle$ converges to $\widetilde{l}$ with positive probability when l_0 is in the neighborhood of $\widetilde{l}$.

Definition 2. *Let $\widetilde{l}$ be a stationary point of $\langle l_t \rangle$. Then $\widetilde{l} \in (-\infty, \infty)$ is **stable** if there exists an open ball N_0 around 0 such that $l_0 - \widetilde{l} \in N_0 \Rightarrow P(l_t \to \widetilde{l}) > 0$. A point $\widetilde{l} = \infty$ is stable if there exists an M such that $l_0 > M \Rightarrow P(l_t \to \infty) > 0$.*

The challenge in establishing convergence results for $\langle l_t \rangle$ stems from the fact that the transitions ψ and ϕ depend 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 martingale properties of the log likelihood ratio. If $\langle \log l_t \rangle$ is a supermartingale for l_t near zero, then $\langle \log l_t \rangle$ diverges to negative infinity and zero is a stable stationary point of l, whereas if $\langle \log l_t \rangle$ is a submartingale for l_t near infinity, then it diverges and infinity is a stable stationary point of l.

I use this criterion to characterize the relationship between the stability of a stationary point and the beliefs informed agents hold about the informativeness of prior actions. 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. *There exist unique cutoffs $\widehat{p}_1 \in [0, p)$ and $\widehat{p}_2 \in (p, 1]$ such that:*

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

Beliefs $\widehat{p}$ 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 $\widehat{p}$ 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 $\widehat{p}_2$, $\langle \log l_t \rangle$ flips from being a supermartingale to a submartingale near zero. Above $\widehat{p}_2$, zero is no longer a stable stationary point. Similar logic establishes the stability of l_t near infinity for some cut-off $\widehat{p}_1$. When $\widehat{p} = p$, the likelihood ratio is a martingale. Therefore, $\langle \log l_t \rangle$ is a supermartingale and zero is a stable stationary point, establishing that $p < \widehat{p}_2$. The martingale convergence theorem precludes infinity from being a stable point when $\widehat{p} = p$; this establishes that $p > \widehat{p}_1$.

The next Lemma establishes that the likelihood ratio converges to a point $\tilde{l}$ if and only if $\tilde{l}$ is a stable stationary point. Additionally, if $\tilde{l}$ is a stable stationary point, then the likelihood converges to $\tilde{l}$ with positive probability, from any initial value.

Lemma 4. *$P(l_t \to \widetilde{l}) > 0$ iff $\widetilde{l}$ is a stable stationary point of $\langle l_t \rangle$.*

The LIL enables a precise characterization of belief convergence. Consider the case of bounded signals. During a cascade, the probability of each action is independent across time. By the law of large numbers (LLN), the share of each action and public signal converge to their expected values when the cascade persists. These shares determine the limit of the likelihood ratio. If the candidate limit point lies inside the cascade set, then by the LIL Sheu (1974), there is a positive measure of sample paths that converge to this limit without crossing outside the cascade set, and the cascade persists in the limit with positive probability.[4] On the other hand, if the candidate limit point lies outside the cascade set, then the likelihood ratio will eventually leave the cascade set and the cascade breaks almost surely. Precisely the same criterion determines whether a stationary point is stable and whether l_t converges to a value inside or outside the cascade set. Therefore, whenever a stationary point is stable, the likelihood ratio converges to this point with positive probability, from any starting value l_0. The intuition is similar for the case of unbounded signals. I bound the likelihood ratio near zero or infinity with an i.i.d. process, and use the LIL to determine the limiting behavior of this i.i.d. process.

3.2 Asymptotic Learning

This section presents the main result of the paper - a characterization of the learning dynamics in a misspecified model of inference. I use insights from the stability of stationary points to determine whether observing an infinite sequence of new information allows informed agents to eventually learn the true state, and examine what happens when learning fails.

Several possible long-run outcomes may occur. Learning is *complete* if public beliefs converge to a point mass on the true state, whereas learning is *incorrect* if public beliefs converge to a point mass on the incorrect state. If public beliefs about the state remain interior, either by converging to an interior belief or perpetually oscillating, then learning is *incomplete*. Belief convergence implies *action convergence* for informed agents, in that they eventually choose the same action. Note that action convergence never obtains for uninformed agents, as their actions always depend on their private information.

If the likelihood ratio has a limit, then the support of the limit is equal to the set of stable stationary points, by Lemma 4. For complete learning to arise almost surely, zero must be the unique stable stationary point. If both zero and infinity are stable stationary points, then

[4] The LIL bounds the rate at which a sequence converges to its expected value. The probability that the sequence crosses outside this bound infinitely often is zero.

both complete and incorrect learning arise, whereas if there are no stable stationary points, the likelihood ratio perpetually oscillates and learning is incomplete. Theorem 1 characterizes the relationship between asymptotic learning outcomes and informed agents' interpretation of prior actions.

Theorem 1. *There exist unique cutoffs $\widehat{p}_1$ and $\widehat{p}_2$ such that:*

1. *If $\widehat{p} < \widehat{p}_1$, then complete and incorrect learning both occur with positive probability. Incomplete learning almost surely does not occur.*
2. *If $\widehat{p} \in (\widehat{p}_1, \widehat{p}_2)$, then complete learning occurs almost surely.*
3. *If $\widehat{p} > \widehat{p}_2$, then incomplete learning occurs almost surely and beliefs perpetually oscillate.*

Both complete and incorrect learning outcomes arise when agents attribute too many actions to uninformed agents. Agents overestimate the informativeness of actions supporting the more likely state, and underestimate the informativeness of contrary actions, causing beliefs to quickly become entrenched. When beliefs about the share of informed agent are relatively accurate, incorrect learning is no longer possible, while complete learning remains 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. This prevents belief convergence - there are no stable stationary points and learning is incomplete.

It is also necessary to rule out incomplete learning when $\widehat{p} < \widehat{p}_2$. 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.

The asymptotic properties of learning determine whether the action choices of informed agents eventually converge on the optimal action. When $\widehat{p}$ falls between the two thresholds, $\widehat{p} \in (\widehat{p}_1, \widehat{p}_2)$, only correct learning is possible. 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. Correct beliefs about agent types lead to efficient action choices, given that $p \in (\widehat{p}_1, \widehat{p}_2)$.

4 Discussion

This paper demonstrates that a misspecified model of information processing impacts asymptotic learning. Individuals may continue to choose an inefficient action, despite sufficient public information to learn the true state. This has important policy implications. In a correctly specified model, complete learning obtains regardless of the structure of public information releases. However, in the face of model misspecification, the timing and frequency of public information impact the asymptotic properties of the learning process. When $\hat{p} < \hat{p}_1$, immediate release of public information will prevent beliefs from becoming entrenched on the incorrect state. A delayed public response will require stronger or more frequent public signals to overturn an incorrect herd. Policy 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 $\hat{p} > \hat{p}_1$, 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. A policy intervention must be maintained ad infinitum; once an intervention ceases, the herd will eventually break.

Experimental evidence from Goeree, Palfrey, Rogers, and McKelvey (2007) suggests that new information does indeed continue 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 would recognize these off-the-equilibrium-path actions are likely to reveal an agent's private signal, even if they are unsure of the exact prevalence of such actions. Koessler, Ziegelmeyer, Bracht, and Winter (2008) examines 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; 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. This provides support for both the presence of uninformed agents and a misspecified belief about their frequency. 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 model leaves open interesting questions for future research on model misspecification. 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 $l_I = l\underline{b}/(1 - \underline{b}) \geq 1$ 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. The signal threshold of uninformed agents is independent of the likelihood ratio; combined with the assumption that $\underline{b} < \bar{b}$ guarantees that an uninformed agent chooses both actions with positive probability. Q.E.D.

Proof of Lemma 2: At a stationary point $\widetilde{l}$, $\phi(a, s^p, \widetilde{l}) = \widetilde{l}$ for all a and s^p that occur with positive probability when $l = \widetilde{l}$. As $\psi^p(s^p|L)/\psi^p(s^p|R) \neq 1$ for $s^p \in \{L, R\}$ and $P(s^p|\omega) > 0$ for $s^p \in \{L, R\}$, independent of l, the only possible candidates for stationary points are $l \in \{0, \infty\}$.

This is also true without public signals, provided signals are unbounded or $p < 1$ or $\hat{p} < 1$. Q.E.D.

The proof of Lemma 3 makes use of several intermediate lemmas and Corollary C1. 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, probability measure $p(.|x)$ defined on A for any $x \in \mathbb{R}_+$ and Borel measurable function $f : A \times \mathbb{R}_+ \to \mathbb{R}_+$, suppose $\langle x_t \rangle_{t=0}^{\infty}$ is a Markov process with transition rule $x_{t+1} = f(a_t, x_t)$ for $a_t \in A$, governed by the probability measure $P(B|x) = \sum_{\{a: f(a,x) \in B\}} p(a|x)$ for any set B in the Borel $\sigma-$algebra on $\mathbb{R}_+$. Suppose $\widetilde{x}$ is a fixed point of x. Then $\widetilde{x}$ is a stable fixed point if*

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

Proof: See Smith and Sorensen (2000).

Lemma 6. *Assume the ratio of two densities, $f^L(x)/f^R(x)$ satisfies the monotone likelihood ratio property. Then $\frac{d}{dx}\frac{F^L}{F^R}(x) \geq 0$ and $\frac{d}{dx}\frac{1-F^L}{1-F^R}(x) \geq 0$ and $F^L \leq F^R$.*

Proof: Let $x_1 \geq x_0$

$$\begin{aligned}
&\Rightarrow \quad \frac{f^L(x_1)}{f^R(x_1)} \geq \frac{f^L(x_0)}{f^R(x_0)} \\
&\Rightarrow \quad \int_0^{x_1} f^L(x_1) f^R(x_0) dx_0 \geq \int_0^{x_1} f^L(x_0) f^R(x_1) dx_0 \\
&\Rightarrow \quad f^L(x_1) F^R(x_1) \geq F^L(x_1) f^R(x_1) \\
&\Rightarrow \quad \frac{d}{dx}\frac{F^L(x)}{F^R(x)} = \frac{F^R(x) f^L(x) - F^L(x) f^R(x)}{F^R(x)^2} \geq 0
\end{aligned}$$

The proof of the remaining case is analogous. First order stochastic dominance follows directly from the monotone likelihood property. Q.E.D.

Lemma 7. *Assume the ratio of two densities, $f^L(x)/f^R(x)$ satisfies the monotone likelihood ratio property. Then the transition of the likelihood ratio changes with $\hat{p}$ as follows:*

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

Proof:

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

Lemma 6 established that $\frac{d}{dx}\frac{F^L}{F^R}(x) > 0$. Suppose $l > 1$. Then

$$\begin{aligned} \frac{F^L(1/(l+1))}{F^R(1/(l+1))} &< \frac{F^L(1/2)}{F^R(1/2)} \\ \Rightarrow\ F^L(1/(l+1))F^R(1/2) &- F^L(1/2)F^R(1/(l+1)) < 0 \end{aligned}$$

and therefore, $\frac{d}{d\hat{p}}\left(\frac{\psi(R|L,l;\hat{p})}{\psi(R|R,l;\hat{p})}\right) < 0$. The proof of the remaining cases is analogous. Q.E.D.

Proof of Lemma 3: From Lemma 5, when $\omega = R$, the fixed point $\widetilde{l} = 0$ is a stable fixed point of l if

$$\theta(\widehat{p}, \widetilde{l} = 0) := \sum_{(a,s^p)\in\{L,R\}\times\{L,R,\emptyset\}} \psi(a|R,0;p)\psi^p(s^p|R)\log\phi_l(a,s^p,0;\hat{p}) < 0$$

where

$$\phi_l(a,s^p,0;\hat{p}) = \left(\frac{\psi(a|L,0;\hat{p})}{\psi(a|R,0;\hat{p})}\right)\left(\frac{\psi^p(s^p|L)}{\psi^p(s^p|R)}\right)$$

Simplify $\theta(\widehat{p}, 0)$ as

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

where $\delta := \varepsilon\,(2\pi^p - 1)\log\left(\frac{\pi^p}{1-\pi^p}\right) > 0$ is the public signal component of θ.

If $\widehat{p} = 1$, then for small enough ε and $p < 1$,

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

since $F^L \le F^R$ and therefore $\frac{1-F^L(1/2)}{1-F^R(1/2)} > 1$. At $\widehat{p} = p$, l_t is a martingale and $\theta(\widehat{p} = p) < 0$ Also, $\theta(\widehat{p}, 0)$ is increasing in $\widehat{p}$:

$$\frac{d\theta}{d\widehat{p}} = \psi(R|R,0;p)\left(\frac{\psi(R|R,0;\hat{p})}{\psi(R|L,0;\hat{p})}\right)\frac{d}{d\hat{p}}\left(\frac{\psi(R|L,l;\hat{p})}{\psi(R|R,l;\hat{p})}\right) > 0$$

and $\frac{d}{d\hat{p}}\left(\frac{\psi(R|L,l;\hat{p})}{\psi(R|R,l;\hat{p})}\right) > 0$ at $l = 0$. Therefore, there exists a $\widehat{p}_2 \in (p, 1)$ s.t. $\theta(\widehat{p}_2, 0) = 0$. For $\widehat{p} < \widehat{p}_2$, $\theta(\widehat{p}, 0) < 0$ and 0 is a stable limit point, while for $\widehat{p} > \widehat{p}_2$, $\theta(\widehat{p}, 0) > 0$ and 0 is not a stable limit point.

Now consider the limit point $\widetilde{l} = \infty$. Transform the problem to consider the Markov process

$\langle(a_t, x_t = 1/l_t)\rangle$ with transitions

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

When $\omega = R$, the fixed point $\widetilde{x} = 0$ is a stable fixed point if

$$\widetilde{\theta}(\widehat{p}, \widetilde{x} = 0) := \sum_{(a,s^p)\in\{L,R\}\times\{L,R,\emptyset\}} \widetilde{\psi}(a|R, 0; p)\psi^p(s^p|R) \log \widetilde{\phi}_x(a, s^p, 0; \hat{p}) < 0$$

Note $\widetilde{\theta}(\widehat{p}, 0) = -\theta(\widehat{p}, \infty)$. Simplify $\widetilde{\theta}(\widehat{p}, 0)$ as

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

Suppose at $\widehat{p} = p$, $\widetilde{\theta}(p, 0) < 1$, and therefore $x_t = 1/l_t \to 0$ with positive probability in the neighborhood of 0. Then l_t converges to ∞ with positive probability, a contradiction since l_t is a martingale at $\widehat{p} = p$. Therefore, $\widetilde{\theta}(p, 0) > 0$.

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

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

which is less than zero when

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

Also note $\widetilde{\theta}(\widehat{p}, 0)$ is increasing in $\widehat{p}$:

$$\frac{d\widetilde{\theta}}{d\widehat{p}} = \widetilde{\psi}(L|R, 0; p)\frac{\widetilde{\psi}(L|L, 0; \hat{p})}{\widetilde{\psi}(L|R, 0; \hat{p})}\frac{d}{d\hat{p}}\left(\frac{\widetilde{\psi}(L|R, 0; \hat{p})}{\widetilde{\psi}(L|L, 0; \hat{p})}\right) > 0$$

since $\frac{\widetilde{\psi}(L|R,0;\hat{p})}{\widetilde{\psi}(L|L,0;\hat{p})} = \frac{\psi(L|R,\infty;\hat{p})}{\psi(L|L,\infty;\hat{p})}$ and $\frac{d}{d\hat{p}}\left(\frac{\psi(L|L,l;\hat{p})}{\psi(L|R,l;\hat{p})}\right) < 0$ when $l > 1$. Therefore, when $p > p^*$ there exists a $\widehat{p}_1 \in (0, p)$ s.t. $\widetilde{\theta}(\widehat{p}_1, 0) = 0$. For $\widehat{p} < \widehat{p}_1$, $\widetilde{\theta}(\widehat{p}, 0) < 0$ and 0 is a stable limit point of x and for $\widehat{p} > \widehat{p}_1$, $\widetilde{\theta}(\widehat{p}, 0) > 0$ and 0 is not a stable limit point of x. When $p < p^*$, then 0 is not a stable limit point of $x : \widetilde{\theta}(\widehat{p}, 0) > 1$ for all $\widehat{p}$ and incorrect learning is not possible for any beliefs $\widehat{p}$. Q.E.D.

Proof of Lemma 4: Theorem B.1 in Smith and Sorensen (2000) establishes that a martingale cannot converge to a non-stationary point; this is straightforward to extend to the current setting. Therefore, if $P(l_t \to \widetilde{l}) > 0$, then $\widetilde{l} \in \{0, \infty\}$. It remains to establish that $\widetilde{l}$ is a stable stationary point if and only if $P(l_t \to \widetilde{l}) > 0$ from any starting point l_0.

I first use the LIL to bound the convergence of an i.i.d. random variable. Define random variable $\gamma(a, s^p) := \phi_l(a, s^p, 0; \hat{p})$ with $P(\gamma(a, s^p)) := P(a, s^p|\omega = R, l = 0; p)$ and $\sigma^2 := Var(\log \gamma(a, s^p))$. Then $E[\log \gamma(a, s^p)] = \theta(\widehat{p}, 0)$. By the LIL,

$$\limsup_{t\to\infty} \frac{\sum_{i=1}^{t} (\log \gamma(a_i, s_i^p) - \theta(\widehat{p}, 0))}{\sqrt{2\sigma^2 \log\log \sigma^2}} = 1 \text{ a.s.}$$

Thus, for $\delta > 0$,

$$P\left[\frac{1}{t}\sum_{i=1}^{t} \log(a_i, s_i^p) \geq \beta_t + \theta(\widehat{p}, 0) \text{ i.o.}\right] = 0$$

where $\beta_t := (1+\delta)\sqrt{\frac{2\sigma^2 \log\log t\sigma^2}{t}}$. For almost all sample paths, there exist only *finitely many* t such that $\frac{1}{t}\sum_{i=1}^{t} \log \gamma(a_i, s_i^p)$ lies outside $[\theta(\widehat{p}, 0) - \beta_t, \theta(\widehat{p}, 0) + \beta_t]$. Define

$$\zeta := \left\{ \{\widehat{a}_i, \widehat{s}_i\} \mid \frac{1}{t}\sum_{i=1}^{t} \log \gamma(a_i, s_i^p) > \theta(\widehat{p}, 0) + \beta_t \text{ for some } t \right\}$$

as the set of sample paths such that $\frac{1}{t}\sum_{i=1}^{t} \log \gamma(a_i, s_i^p)$ crosses its upper bound at least once. To show that the measure of ζ is strictly less than 1, consider the following. For each $\{\widehat{a}_i, \widehat{s}_i\} \in \zeta$, form a corresponding sample path $\{a_i', s_i'\}$ by changing $\widehat{a}_s$ to $a_s' \neq \widehat{a}_s$ and $\widehat{s}_i$ to $s_s' = \emptyset$ for each s such that $\frac{1}{s}\sum_{i=1}^{s} \log \gamma(\widehat{a}_i, \widehat{s}_i) > B_s$. Then each sample in ζ has a unique corresponding sample path in its complement ζ^c. Therefore, $P(\zeta^c) \geq P(\zeta)$, which implies $P(\zeta^c) \geq 1/2$ and the set of sample paths such that $\frac{1}{t}\sum_{i=1}^{t} \log \gamma(a_i, s_i^p)$ never crosses its upper bound has positive measure.

Case (i): Suppose signals are bounded. Let τ be the first time that the likelihood ratio enters the cascade set for action R :

$$\tau = [t : l_t \in [0, (1-\underline{b})/\underline{b}] \text{ and } l_s \notin [0, (1-\underline{b})/\underline{b}] \;\; \forall s < t]$$

Note that τ is finite with positive probability and $\log l_\tau \leq \log(1-\underline{b})/\underline{b}$. Define a stochastic process $\langle \lambda_t \rangle$:

$$\lambda_t = \begin{cases} \log l_\tau + \sum_{i=\tau+1}^{t} \log \gamma(a_i, s_i^p) \text{ if } t > \tau \\ \log l_t \text{ if } t \leq \tau \end{cases}$$

The processes $\langle \lambda_t \rangle$ and $\langle l_t \rangle$ coincide as long as an R cascade has not formed and broken. By the law of large numbers, $P(\lim_{t\to\infty} \frac{1}{t-\tau}\sum_{i=\tau+1}^{t} \log \gamma(a_i, s_i^p) = \theta(\widehat{p}, 0)) = 1$.

Suppose 0 is a stable point, $\theta(\widehat{p}, 0) < 0$. An R cascade forms in finite time with positive probability, and persists for any finite number of periods with positive probability; combining this with the LIL, there exists a $\delta_1 > 0$ such that $P(\lambda_t < \log(1-\underline{b})/\underline{b} + (t-\delta_1)\theta(\widehat{p}, 0) + \beta_{t-\delta_1}$ for all $t > \delta_1) > 0$ and the cascade persists in the limit with positive probability. On this set of sample paths, λ_t converges to $-\infty$, so $P(\lim_{t\to\infty} \lambda_t = -\infty) > 0$. Given that $\langle \lambda_t \rangle$ and $\langle l_t \rangle$ coincide along the set of sample paths where a cascade forms and persists, the likelihood ratio also diverges to $-\infty$ with positive probability, $P(\lim_{t\to\infty} l_t = -\infty) > 0$.

Now suppose 0 is not a stable point, $\theta(\widehat{p}, 0) > 0$. Given $l_\tau > 0$, there exists a $\varepsilon_1 < \infty$ such that $\lambda_t > -\varepsilon_1 + \sum_{i=\tau+1}^{t} \log \gamma(a_i, s_i^p)$ and for all $\varepsilon_2 > 0$, there exists a δ_2 such that for $t > \delta_2$, $P(\lambda_t > -\varepsilon_1 + (t-\tau)\,\theta(\widehat{p}, 0) - \varepsilon_2) = 1$. Therefore, $P(\exists t > \tau \text{ s.t. } \lambda_t > \log\left(1-\underline{b}\right)/\underline{b}) = 1$, which also implies $P(\exists t > \tau \text{ s.t. } l_t > \log\left(1-\underline{b}\right)/\underline{b}) = 1$ and a cascade breaks almost surely. The proof is analogous for a cascade on action L with ∞ as a stable stationary point.

Case (ii): Suppose signals are unbounded. Let infinity be a stable stationary point, $\theta(\widehat{p}, \infty) > 0$. Then, by continuity, $\exists l^* < \infty$ such that

$$\sum_{a,s^p} \psi(a|R, l^*; p)\psi^p(s^p|R) \log \phi_l(a, s^p, \infty; \hat{p}) = 0.$$

Let τ be the first time that the likelihood ratio crosses l^*, which is finite with positive probability. Define a stochastic process $\langle \lambda_t \rangle$:

$$\lambda_t = \begin{cases} \log l_\tau + \sum_{i=\tau+1}^{t} \log \gamma(a_i, s_i^p, l) \text{ if } t > \tau \\ \log l_t \text{ if } t \leq \tau \end{cases}$$

where $\gamma(a, s^p, l) := \phi_l(a, s^p, \infty; \hat{p})$ is an r.v. with $P(\gamma(a, s^p, l)) := P(a, s^p|\omega = R, l; p)$. Note that $l_t > \lambda_t$, and for $l > l^*$, $E[\log \gamma(a, s^p, l)] > 0$. Similar arguments to the case of bounded signals show that if $\lambda_t > l^*$ at time t, then $P(\lambda_s > l^* \text{ for all } s > t) > 0$ and on the set of sample paths where $\lambda_s > l^*$ for all $s > t$, $\lambda_t \to \infty$. Since $l_t > \lambda_t$, $P(l_t \to \infty) > 0$. A similar argument bounds l_t from above to establish that $P(l_t \to \infty) = 0$ when $\theta(\widehat{p}, \infty) < 0$. An analogous proof establishes that $P(l_t \to 0) > 0$ iff $\theta(\widehat{p}, 0) < 0$. Q.E.D.

Proof of Theorem 1: Combining Lemma 4 and 3 establishes the set of limit points of the likelihood ratio, as a function of $\hat{p}$. When $\widehat{p} < \widehat{p}_1$, 0 and ∞ are both limit points of the likelihood ratio. Therefore, both complete learning and fully incorrect learning occur with positive probability. When $\widehat{p} \in (\widehat{p}_1, \widehat{p}_2)$, 0 is the unique limit point of the likelihood ratio, leading to complete learning with positive probability. When $\widehat{p} > \widehat{p}_2$, the likelihood ratio neither converges to zero nor diverges. These are the only two candidate limit points; therefore, the likelihood ratio does not converge and learning is incomplete.

It is also necessary to rule out incomplete learning when $\widehat{p} < \widehat{p}_2$. First consider bounded signals. When $\widehat{p} < \widehat{p}_2$, once a cascade forms, the probability of returning to no cascade is less than 1 for at least one type of cascade. Therefore, by Theorem 8.2 in Billingsley (1995) $P(l_s \in \left(\log\left(1-\underline{b}\right)/\underline{b}\right), \left(1-\bar{b}\right)/\bar{b}\right) \text{ i.o.}) = 0$. Eventually a cascade forms and persists in the limit, and learning is incomplete with probability 0. Similar logic establishes that when $\widehat{p} < \widehat{p}_2$ and signals are unbounded, $P(l_s = \bar{l} \text{ i.o.}) = 0$ for $\bar{l}$ such that $\theta(\widehat{p}, \bar{l}) = 0$, and beliefs eventually converge. 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.

Billingsley, P. (1995): *Probability and Measure*, Wiley-Interscience, 3 ed.

Bohren, A. (2010): "Information Processing Uncertainty in Social Learning," .

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.

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.

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.

Sheu, S. S. (1974): "Some Iterated Logarithm Results for Sums of Independent Two-dimensional Random Variables," *The Annals of Probability*, 2, 1139–1151.

Smith, L. and P. Sorensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398.