



Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 163 (2016) 222–247

JOURNAL OF
**Economic
Theory**

www.elsevier.com/locate/jet

Informational herding with model misspecification [☆]

J. Aislinn Bohren

Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19103, United States

Received 9 December 2014; final version received 25 January 2016; accepted 26 January 2016

Available online 2 February 2016

Abstract

This paper demonstrates that a misspecified model of information processing interferes with long-run learning and allows inefficient choices to persist, despite sufficient information for asymptotic learning. I consider an observational learning environment in which agents observe a private signal about an unknown state and some agents observe the actions of their predecessors. Individuals face an inferential challenge when extracting information from the actions of others, as prior actions aggregate multiple sources of correlated information. A misspecified model allows for the fact that an agent may not be able to distinguish between new and redundant information, and may have an incorrect model of how others process this information. When individuals significantly overestimate the amount of new information, beliefs about the state become entrenched and incorrect learning occurs with positive probability. When individuals sufficiently overestimate the amount of redundant information, beliefs fail to converge and learning is incomplete. Learning is complete when agents have an approximately correct model of inference, establishing that the correctly specified model is robust to perturbation.

© 2016 Elsevier Inc. All rights reserved.

JEL classification: D82; D83

Keywords: Model misspecification; Observational learning; Informational herding

[☆] I thank Nageeb Ali, Vince Crawford, Ernesto Dal Bo, Alex Imas, Frederic Koessler, George Mailath, Craig McKenzie, Matthew Rabin, Joel Sobel and an anonymous referee for useful comments. I also thank participants of the NSF Information Transmission in Networks Conference and the UCSD theory lunch for helpful feedback.

E-mail address: abohren@sas.upenn.edu.

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 biased learning offers an explanation for how inefficient choices can persist, incorrect beliefs can become entrenched, or beliefs can fail to converge, despite sufficient evidence for complete learning.

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 have a correct model of how others process this information, in order 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 \(2000\)](#). In these models, 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 redundant 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 her to overweigh information from the public history, and allows public beliefs about the state to become entrenched, possibly unjustifiably so. 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. Thus, the difference between \hat{p} and p determines the level of model misspecification.

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 through their belief \hat{p} . The main result of the paper ([Theorem 1](#)) specifies thresholds \hat{p}_1 and \hat{p}_2 , such that when $\hat{p} < \hat{p}_1$ both correct and fully incorrect learning occur with positive probability, when $\hat{p} > \hat{p}_2$, beliefs about the state perpetually fluctuate, rendering learning incomplete, while when $\hat{p} \in (\hat{p}_1, \hat{p}_2)$, correct learning occurs with probability one. The first two cases admit the possibility of inefficient learning: with positive probability, informed agents choose the inefficient action infinitely often, despite observing sufficient information to

learn the correct state. In the final case, informed agents will eventually choose the efficient action. This case includes the correctly specified model ($\hat{p} = p$), as demonstrated by the fact that $p \in (\hat{p}_1, \hat{p}_2)$.

Fully incorrect learning or incomplete learning with oscillating beliefs are 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 is general enough that it can be utilized to examine other forms of model misspecification. Thus, the paper develops new techniques to analyze learning in models that are not fully Bayesian.

Model misspecification has important policy implications for interventions aimed at counteracting inefficient social choices. In the presence of information processing errors, the timing, frequency and strength of interventions – such as public information campaigns – are an important determinate of long-run efficiency. 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 in which parents refuse vaccines 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.

The sequential observational learning framework used in this paper was first modeled in Banerjee (1992) and Bikhchandani et al. (1992) with a binary signal space. They conclude that incorrect informational cascades arise with positive probability, but beliefs in these cascades are fragile and easily overturned by the arrival of new information. Moscarini et al. (1998) show that informational cascades are temporary when the state of the world changes frequently enough. Smith and Sorensen (2000) allow for a general signal distribution and crazy types. An unbounded signal space is sufficient to ensure complete learning, eliminating the possibility of inefficient cascades. Acemoglu et al. (2011) examines social learning in a network – the correctly specified model in this paper ($\hat{p} = p$) is a special case of their model.

Recent work examines the implications of information processing biases, particularly correlation neglect and the failure to account for redundant information, in the social learning framework. Eyster and Rabin (2010) study *inferential naivety* – players believe prior agents' actions solely reflect their private information. This confounds learning because the actions of initial agents receive disproportionate weight. Although similar in nature to model misspecification, inferential naivety differs in generality and interpretation. Inferential naivety considers the case in which every repeated action is viewed as being independent with probability one, whereas with model misspecification, informed agents recognize that actions contain some repeated information, but misperceive the exact proportion. The analogue of inferential naivety in my environment corresponds to $\hat{p} = 0$ and $p = 1$.

¹ This example abstracts from the payoff interdependencies of vaccines.

Guarino and Jehiel (2013) apply the analogy based expectation equilibrium solution concept (Jehiel, 2005) to a social learning setting. Agents know the aggregate relationship between the state and distribution of actions, but do not understand the relationship between private information and actions. Learning is complete in a continuous action model – although initial signals are overweighted, the excess weight on a signal increases linearly with time, preventing 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. Levy and Razin (2015) examine the implications of correlation neglect in a network model of learning, and establish that beliefs converge under mild conditions on the network structure. Demarzo et al. (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. Mueller-Frank and Neri (2015) build on Eyster and Rabin (2010)'s concept of inferential naivety to study information aggregation in networks. They establish sufficient conditions on the learning environment to achieve information aggregation in small networks, and show that in any learning environment, information aggregation fails in large enough networks.

Model misspecification is also related to level- k and cognitive hierarchy models.² In the model misspecification framework, uninformed types are level-1 thinkers who follow their private signal while informed types are level-2 thinkers who believe other agents are a mix of level-1 and level-2 thinkers. In a level- k model, informed agents believe that all other agents are level-1 thinkers – this corresponds to $\hat{p} = 0$. Thus, in both frameworks, level-2 agents misperceive the share of other agents who are level-2, but this paper allows level-2 agents to place positive weight on other agents using a level-2 decision rule.

This paper is also related to a broader literature on how information processing biases and model misspecification impact long-run learning. Epstein et al. (2010) show that incorrect learning can arise in a single agent model when an agent overweightes signals, as is the case in this paper, but that complete learning obtains when an agent underweights signals. In this model, agents who underweight information may never learn the state. In earlier work by Eyster and Rabin (2005) on cursed equilibrium, a cursed player does not understand the correlation between a player's type and her action choice, and therefore fails to realize a player's action choice reveals information about her type.³

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.

² Costa-Gomes et al. (2009); Camerer et al. (2004).

³ Other recent work includes Rabin and Schrag (1999), Acemoglu et al. (2016), Gottlieb (2015), Schwartzstein (2014), Wilson (2014) and Esponda and Pouzo (2015). There is also an older statistics literature, including Berk (1966) and DeGroot (1974).

States, actions and payoffs There are two payoff-relevant states of the world, $\omega \in \{L, H\}$ 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 $\mathcal{T} = \{1, 2, \dots\}$ act sequentially and attempt to match the realized state of the world by making a single decision between two actions, $a_t \in \{\ell, h\}$, $t \in \mathcal{T}$. They receive a payoff of 1 if their action matches the realized state, $u(\ell, L) = u(h, H) = 1$, and a payoff of 0 otherwise.

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 $\omega = L$, computed via Bayes' rule 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^H 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. Finally, assume that some signals are informative. This rules out $dF^L/dF^H = 1$ almost surely, and implies $\underline{b} < 1/2 < \bar{b}$. Beliefs are bounded if $0 < \underline{b} < \bar{b} < 1$, and are unbounded if $[\underline{b}, \bar{b}] = [0, 1]$.

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 her 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 her private signal.

Beliefs about types Each informed individual believes that each other individual is informed with probability $\hat{p} \in [0, 1]$, 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 persist because no agent ever learns what the preceding agents actually observed or incorporated into their decision-making processes.⁴

Timing At time t , agent t observes type θ_t and a private signal s_t . If $\theta_t = I$, the agent also observes the public history h_t . Then 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 (misspecified) belief about the type distribution, and maximize payoffs with respect to these beliefs. The decision rule of each type is

⁴ Although it is admittedly restrictive to require that agents hold identical misperceptions about others, and that this misperception takes the form of a 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.

common knowledge, as is the fact that all informed agents compute the same probability of any history h_t .

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

$$\lambda_t = \frac{P(L|h_t; \hat{p})}{P(H|h_t; \hat{p})}, \tag{1}$$

which depends on the history and beliefs about the share of informed agents.⁵ An agent who holds prior belief λ and receives signal s updates to the private posterior belief $q(\lambda, s) = \lambda \times \left(\frac{s}{1-s}\right)$. An uninformed agent has prior belief $\lambda_1 = 1$ and an informed agent has prior belief λ_t . Guided by posterior belief q , the agent maximizes her payoff by choosing $a = \ell$ if $q \geq 1$, and $a = h$ otherwise. An agent’s decision can be represented as a cut-off rule, $s^*(\lambda) = 1/(\lambda + 1)$, such that the agent chooses action ℓ when $s \geq s^*(\lambda)$ and chooses action h otherwise. An informed agent in period t uses cut-off $s^*(\lambda_t)$, while uninformed agents use cut-off $s^*(1) = 1/2$.

The cascade set for action a is the set of prior beliefs such that a is optimal for all realizations of the private signal.

Definition 1 (*Cascade set*). The cascade sets for actions h and ℓ are the sets of beliefs $\mathcal{J}^h = \{\lambda | s < s^*(\lambda) \forall s \in \text{supp}(F)\}$ and $\mathcal{J}^\ell = \{\lambda | s \geq s^*(\lambda) \forall s \in \text{supp}(F)\}$, respectively.

As usual, a cascade occurs when the prior belief outweighs the strongest private belief.

Lemma 1. *When private beliefs are bounded, $\mathcal{J}^h = [0, (1 - \bar{b})/\bar{b}]$ and $\mathcal{J}^\ell = [(1 - \underline{b})/\underline{b}, \infty]$ and when private beliefs are unbounded, $\mathcal{J}^h = \{0\}$ and $\mathcal{J}^\ell = \{\infty\}$.*

Let $\mathcal{J} = \mathcal{J}^\ell \cup \mathcal{J}^h$. An uninformed agent is never in a cascade, since $\lambda_1 \notin \mathcal{J}$. An informed agent is in a cascade if $\lambda_t \in \mathcal{J}$. This agent chooses the same action for all $s \in \text{supp}(F)$ and her action reveals no private information.

When informed agents are in a cascade, information continues to accumulate from the actions of uninformed agents, and the public belief leaves the cascade set with positive probability. Therefore, the formation of a cascade does not necessarily imply belief convergence. If a cascade does not form in finite time, the likelihood ratio may still converge to a point in the cascade set. The following definition introduces the notion of a limit cascade to encompass both of these ideas.

Definition 2 (*Limit cascade*). Suppose there exists a real, nonnegative random variable λ_∞ such that $\lambda_t \rightarrow \lambda_\infty$ almost surely. Then a *limit cascade* occurs if $\text{supp}(\lambda_\infty) \subset \mathcal{J}$.

3. Learning dynamics

3.1. Overview

This section proceeds as follows. After formally defining the stochastic process $\langle \lambda_t \rangle$ governing the evolution of the likelihood ratio, I characterize the set of stationary points; these are candidate limit points for $\langle \lambda_t \rangle$. Next, I determine how the local stability of these stationary points

⁵ I refer to λ_t as the public belief, even though it is not the belief of uninformed agents.

depends on \hat{p} . This establishes the dynamics of the likelihood ratio in the neighborhood of a stationary point. I then use the Law of the Iterated Logarithm (LIL) to show that the likelihood ratio converges to each locally stable point with positive probability from any initial value, which establishes the global stability of locally stable points. Finally, I rule out convergence to unstable stationary points and non-stationary points. The section concludes with a full characterization of the relationship between asymptotic learning outcomes and the degree of model misspecification, as measured by \hat{p} .

3.2. The likelihood ratio

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

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

and

$$\psi(\ell|\omega, \lambda; p) = 1 - \psi(h|\omega, \lambda; p). \quad (3)$$

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 $s^*(\lambda) = 1/(\lambda + 1)$, given likelihood ratio λ .

The likelihood ratio is updated based on the perceived probability of action a , $\psi(a|\omega, \lambda; \hat{p})$. 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 λ_t and action a_t , the likelihood ratio in the next period is $\lambda_{t+1} = \phi(a_t, \lambda_t; \hat{p})$, where

$$\phi(a, \lambda; \hat{p}) = \lambda \left(\frac{\psi(a|L, \lambda; \hat{p})}{\psi(a|H, \lambda; \hat{p})} \right). \quad (4)$$

The joint stochastic process $\langle a_t, \lambda_t \rangle_{t=1}^\infty$ is a discrete-time Markov process defined on state space $\{\ell, h\} \times \mathbb{R}_+$ with $\lambda_1 = 1$. Given state $\{a_t, \lambda_t\}$, the process transitions to state $\{a_{t+1}, \phi(a_t, \lambda_t; \hat{p})\}$ with probability $\psi(a_{t+1}|\omega, \phi(a_t, \lambda_t; \hat{p}); p)$. The stochastic properties of this process determine long-run learning dynamics. The challenge in establishing convergence results for $\langle \lambda_t \rangle$ stems from the dependence of ψ on the current value of the likelihood ratio and the fact that $\langle \lambda_t \rangle$ is not a martingale in a misspecified model.

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 (Stationary). A point λ is stationary if either (i) $\psi(a|\omega, \lambda; p) = 0$ or (ii) $\phi(a, \lambda; \hat{p}) = \lambda$ for $a \in \{\ell, h\}$.

The next Lemma characterizes the set of stationary points.

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

A stationary point λ is locally stable if $\langle \lambda_t \rangle$ converges to λ with positive probability when λ_1 is in the neighborhood of λ .

Definition 4 (Local stability). Let $\lambda \in [0, \infty)$ be a stationary point of $\langle \lambda_t \rangle$. Then λ is locally stable if there exists an open ball N_0 around 0 such that $\lambda_1 - \lambda \in N_0 \Rightarrow P(\lambda_t \rightarrow \lambda) > 0$. A point $\lambda = \infty$ is locally stable if there exists an M such that $\lambda_1 > M \Rightarrow P(\lambda_t \rightarrow \infty) > 0$.

Local stability can be reframed in the context of the expected change in the log likelihood ratio. Suppose $\omega = H$. Given likelihood ratio λ , the probability of action a is $\psi(a|H, \lambda; p)$. Define

$$\gamma(\hat{p}, \lambda) := \sum_{a \in \{l, h\}} \psi(a|H, \lambda; p) \log \left(\frac{\psi(a|L, \lambda; \hat{p})}{\psi(a|H, \lambda; \hat{p})} \right) \tag{5}$$

as the expected change in the log likelihood ratio. Then

$$E_t[\log \lambda_{t+1}] = \log \lambda_t + \gamma(\hat{p}, \lambda_t).$$

Therefore, the sign of $\gamma(\hat{p}, \lambda_t)$ determines whether $E_t[\log \lambda_{t+1}]$ is greater or less than $\log \lambda_t$.

Lemma 3 establishes the relationship between the local stability of $\lambda \in \{0, \infty\}$ and $\gamma(\hat{p}, \lambda)$. Intuitively, 0 is locally stable when the expected change in the log likelihood ratio is negative at 0, and ∞ is locally stable when the expected change in the likelihood ratio is positive at ∞ . Note $\gamma(\hat{p}, 0)$ and $\gamma(\hat{p}, \infty)$ are straightforward to calculate from the primitives of the model.

Lemma 3. Suppose $\omega = H$. Given \hat{p} and $\gamma(\hat{p}, \cdot)$ defined in (5),

1. If $\gamma(\hat{p}, 0) < 0$, then 0 is locally stable, while if $\gamma(\hat{p}, 0) > 0$, then 0 is not locally stable.
2. If $\gamma(\hat{p}, \infty) > 0$, then ∞ is locally stable, while if $\gamma(\hat{p}, \infty) < 0$, then ∞ is not locally stable.
3. If $\gamma(\hat{p}, \lambda) = 0$ for $\lambda \in \{0, \infty\}$ and private beliefs are bounded, then λ is not locally stable.⁶

The condition for the local stability of 0 follows directly from Corollary C.1 of [Smith and Sorensen \(2000\)](#), which derives a criterion for the local stability of a nonlinear stochastic difference equation with state-dependent transitions. The condition for the local stability of ∞ follows from defining Markov process $\langle x_t \rangle$ as $x_t = 1/\lambda_t$ and noting that the analogue of (5), given x , is $-\gamma(\hat{p}, 1/x)$. Thus, 0 is a locally stable point of $\langle x_t \rangle$ when $-\gamma(\hat{p}, \infty) < 0$. If 0 is a locally stable point of $\langle x_t \rangle$, then ∞ is a locally stable point of $\langle \lambda_t \rangle$.

The conditions for when 0 and ∞ are not locally stable follow from the Law of Large Numbers (LLN), which is used to rule out convergence to the relevant stationary point. Consider the case of bounded private beliefs and suppose the likelihood ratio is in the h -cascade set. The probability of each action is constant, $\psi(a|H, \lambda, p) = \psi(a|H, 0, p)$ for all $\lambda \in \mathcal{J}^h$. If the cascade persists, then by the LLN, the share of each action almost surely converges to its expected

⁶ If private beliefs are unbounded and $\gamma(\hat{p}, \lambda) = 0$ for $\lambda \in \{0, \infty\}$, the stability of λ also depends on $\gamma(\hat{p}, \cdot)$ in a neighborhood of λ (for bounded beliefs, $\gamma(\hat{p}, \cdot)$ is constant in a neighborhood of λ). If $\gamma(\hat{p}, 0) = 0$ and there exists an $\epsilon > 0$ such that $\gamma(\hat{p}, \lambda) < 0$ for $\lambda \in (0, \epsilon)$, then 0 is locally stable, while if there exists an $\epsilon > 0$ such that $\gamma(\hat{p}, \lambda) \geq 0$ for $\lambda \in (0, \epsilon)$, then 0 is not locally stable. The condition for ∞ is analogous. These cases are non-generic, since [Lemma 4](#) establishes that there is a unique $\hat{p} \in [0, 1]$ for which $\gamma(\hat{p}, 0) = 0$ and a unique $\hat{p} \in [0, 1]$ for which $\gamma(\hat{p}, \infty) = 0$. I do not consider them, as they significantly complicate the analysis without adding much economic insight. Note that it is straightforward to verify local stability of these cases for specific private belief distributions, and given the local stability properties, all subsequent results carry through.

value, $\psi(a|H, 0, p)$. Therefore, if the cascade persists, the limit of $\log \lambda_t / t$ almost surely converges to a limit determined by the expected share of each action, which is exactly $\gamma(\hat{p}, 0)$. If $\gamma(\hat{p}, 0) > 0$, then when a cascade persists,

$$\lim_{t \rightarrow \infty} \log \lambda_t / t = \gamma(\hat{p}, 0) > 0.$$

But in order to remain inside the cascade set, it must be that

$$\lim_{t \rightarrow \infty} \log \lambda_t / t < \lim_{t \rightarrow \infty} \log (1 - \bar{b}) / \bar{b} t = 0,$$

a contradiction. Therefore, if $\gamma(\hat{p}, 0) > 0$, then the likelihood ratio will almost surely leave the cascade set.

Next I characterize how $\gamma(\cdot, \lambda)$ varies with \hat{p} , which determines how the local stability of $\lambda \in \{0, \infty\}$ depends on \hat{p} . Let

$$\hat{p}_1 := \begin{cases} \{\hat{p} | \gamma(\hat{p}, \infty) = 0\} & \text{if } \{\hat{p} | \gamma(\hat{p}, \infty) = 0\} \neq \emptyset \\ 0 & \text{if } \{\hat{p} | \gamma(\hat{p}, \infty) = 0\} = \emptyset \end{cases} \tag{6}$$

be the set of beliefs \hat{p} such that $\gamma(\cdot, \infty)$ is zero, and let

$$\hat{p}_2 := \{\hat{p} | \gamma(\hat{p}, 0) = 0\} \tag{7}$$

be the set of beliefs \hat{p} such that $\gamma(\cdot, 0)$ is zero. Then (6) and (7) define the cut-offs at which a stationary point switches from being locally stable to not stable and vice versa.

Given $\lambda \in \{0, \infty\}$, Lemma 4 uses the monotonicity of $\gamma(\cdot, \lambda)$ to establish that \hat{p}_1 and \hat{p}_2 are unique.⁷ Below the cutoff, λ is locally stable, and above the cutoff, λ is not locally stable. When $\hat{p} = p$, the likelihood ratio is a martingale, so 0 is locally stable and ∞ is not. This establishes that $\hat{p}_1 < p$ and $\hat{p}_2 > p$. Although there is always a belief \hat{p} such that 0 is not locally stable (i.e. $\hat{p}_2 < 1$), it is possible that there is no \hat{p} such that ∞ is locally stable (i.e. $\hat{p}_1 = 0$ can occur). This latter possibility depends on the actual share of informed agents p and the informativeness of the actions of uninformed agents.

Lemma 4. *Suppose $\omega = H$. There exist unique cutoffs $\hat{p}_1 \in [0, p)$ and $\hat{p}_2 \in (p, 1)$ defined by (6) and (7).*

1. If $\hat{p}_1 > 0$ and $\hat{p} \in [0, \hat{p}_1)$, then the set of locally stable points is $\{0, \infty\}$.
2. If $\hat{p} \in (\hat{p}_1, \hat{p}_2)$, then 0 is the unique locally stable point.
3. If $\hat{p} \in (\hat{p}_2, 1]$, then there are no locally stable points.
4. If $\hat{p} = \hat{p}_1$ and private beliefs are bounded, then 0 is the unique locally stable point, while if $\hat{p} = \hat{p}_2$ and private beliefs are bounded, then there are no locally stable points.

If $p > p^*$, then $\hat{p}_1 > 0$, and otherwise $\hat{p}_1 = 0$, where $p^* \in (0, 1)$ is defined by

$$p^* := 1 - \frac{\log \left(\frac{1 - F^L(1/2)}{1 - F^H(1/2)} \right)}{F^H(1/2) \left[\log \left(\frac{F^H(1/2)}{F^L(1/2)} \right) + \log \left(\frac{1 - F^L(1/2)}{1 - F^H(1/2)} \right) \right]}. \tag{8}$$

⁷ With a slight abuse of notation, I also use \hat{p}_1 and \hat{p}_2 to denote the unique cut-offs.




Fig. 1. Model misspecification leads to local stability at the incorrect state or instability at the correct state.

Intuitively, if informed agents sufficiently overestimate the share of uninformed agents, then both 0 and ∞ are locally stable, whereas if agents sufficiently underestimate the share of uninformed agents, then no points are locally stable. If the belief about the share of informed agents is close to correct, the unique locally stable point is 0.

Beliefs \hat{p} influence the information that accumulates from each action, but not the probability of each action. When λ is close to 0, state H is perceived as very likely. If an informed agent chooses ℓ , this is indicative of a strong signal in favor of state L , whereas if an informed agent chooses h , this is indicative of a weak signal in favor of state H . The informativeness of uninformed agents' actions is independent of λ_t . Fixing λ close to 0, as \hat{p} increases, the perceived informativeness of contrary ℓ actions increases and the perceived informativeness of supporting h actions decreases. The likelihood ratio jumps further away from 0 when an ℓ action is observed, and moves a relatively smaller distance towards 0 when an h action is observed. Eventually, \hat{p} is high enough such that the likelihood ratio moves away from 0 in expectation and 0 is not locally stable. Fig. 1 plots $\gamma(\cdot, 0)$ and $-\gamma(\cdot, \infty)$ for an unbounded private belief distribution.

When p is low enough, then enough new information is generated by uninformed agents such that even in the extreme case in which informed agents believe all other agents are uninformed, and thus do not account for any repeated information, there is still enough new information to prevent the likelihood ratio from converging to the incorrect state. Mathematically, this is captured by the fact that when $p < p^*$, then $\gamma(\cdot, \infty) < 0$ for all $\hat{p} \in [0, 1]$, where p^* is defined in (8) and depends on the relative informativeness of ℓ and h actions from uninformed agents, $\frac{1-F^L(1/2)}{1-F^H(1/2)}$ and $\frac{1-F^L(1/2)}{1-F^H(1/2)}$, respectively.

3.4. Global convergence to limit outcomes

The next Lemma establishes that, from any initial value $\lambda_0 \in (0, 1)$, the likelihood ratio converges to each locally stable point with positive probability and almost surely does not converge to non-stable stationary points or non-stationary points.

Lemma 5. *For any initial value $\lambda_0 \in (0, \infty)$, $P(\lambda_t \rightarrow \lambda) > 0$ iff λ is a locally stable point of $\langle \lambda_t \rangle$.*

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 the LLN and the LIL to establish global convergence to locally stable points.

Consider the case of bounded signals. The probability of each action is constant when the likelihood ratio is in the cascade set. If a cascade persists, then by the LLN, the share of each action almost surely converges to its expected value. Therefore, if the cascade persists, the likelihood ratio almost surely converges to a limit determined by the expected share of each action. When 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. On this set of sample paths, the cascade does indeed persist. In contrast, when this limit lies outside the cascade set, then the likelihood ratio almost surely leaves the cascade set. Precisely the same criterion on γ determines whether the candidate limit lies inside the cascade set and whether a stationary point is locally stable. Therefore, whenever a stationary point is locally 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 *incorrect learning* denote the event where $\lambda_t \rightarrow \infty$, *correct learning* denote the event where $\lambda_t \rightarrow 0$, and *non-stationary incomplete learning* denote the event where λ_t does not converge or diverge.⁸ Learning is *complete* if correct learning occurs almost surely.

When agents attribute too few actions to informed 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 correct and incorrect learning arise. When agents attribute approximately the correct ratio of actions to informed agents, incorrect learning is no longer possible and learning is complete. Finally, when informed agents attribute too many actions to informed agents, they underestimate the informativeness of actions supporting the more likely state and overestimate the informativeness of contrary actions. Beliefs cannot converge, leading to non-stationary incomplete learning and temporary cascades on both actions. [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 4](#).

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

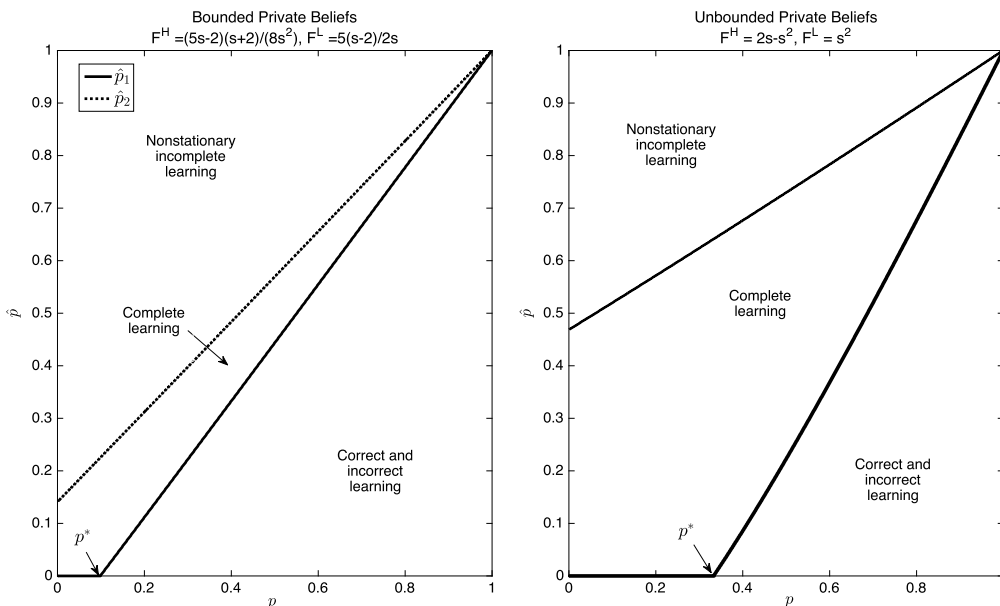


Fig. 2. Long run learning outcomes.

Theorem 1. Suppose $\omega = H$. Given cutoffs $\hat{p}_1 \in [0, p)$ and $\hat{p}_2 \in (p, 1)$ defined by (6) and (7),

1. If $\hat{p} > 0$ and $\hat{p} \in [0, \hat{p}_1)$, then $\lambda_t \rightarrow \lambda_\infty$ almost surely, where λ_∞ is a random variable with $\text{supp}(\lambda_\infty) = \{0, \infty\}$.
2. If $\hat{p} \in (\hat{p}_1, \hat{p}_2)$, then $\lambda_t \rightarrow 0$ almost surely.
3. If $\hat{p} \in (\hat{p}_2, 1]$, then λ_t almost surely does not converge or diverge and $P(\lambda_t \notin \mathcal{J} \text{ i.o.}) = 1$.
4. If private beliefs are bounded and $\hat{p} = \hat{p}_1$, then $\lambda_t \rightarrow 0$ almost surely, while if $\hat{p} = \hat{p}_2$, then λ_t almost surely does not converge or diverge and $P(\lambda_t \notin \mathcal{J} \text{ i.o.}) = 1$.⁹

Lemmas 4 and 5 established that when $\hat{p} < \hat{p}_2$, the likelihood ratio converges to a locally stable point with positive probability and does not converge to a non-stationary or non-locally stable point, and when $\hat{p} > \hat{p}_2$, the likelihood ratio does not converge to any point. Lemma 4 also characterized the stable points when $\hat{p} < \hat{p}_2$. The final step to establish Theorem 1 is 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, a value outside the cascade set occurs infinitely often with probability 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.

Fig. 2 illustrates the three asymptotic learning outcomes outlined in Theorem 1 for a bounded and an unbounded private belief distribution. When \hat{p} lies above the blue line, non-stationary incomplete learning occurs almost surely, whereas when \hat{p} lies below the black line, both incorrect and correct learning occur with positive probability. When \hat{p} lies between the two lines, learning is complete. The 45-degree line along which $\hat{p} = p$ is contained in the complete learning region,

⁹ See Footnote 6 for a discussion of the case of unbounded private beliefs when $\hat{p} \in \{\hat{p}_1, \hat{p}_2\}$.

illustrating the insight that correct beliefs lead to complete learning. Fig. 2 also illustrates p^* . For the bounded private belief distribution, $p^* = 0.10$; for any $p > 0.10$, there exists a belief $\hat{p} > 0$ such that incorrect learning occurs with positive probability.

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 $\omega = H$. Given cutoffs $\hat{p}_1 \in [0, p)$ and $\hat{p}_2 \in (p, 1)$ defined by (6) and (7),

1. If $\hat{p}_1 > 0$ and $\hat{p} \in [0, \hat{p}_1)$, then $a_{t_n} \rightarrow a_\infty$ almost surely, where a_∞ is a random variable with $\text{supp}(a_\infty) = \{\ell, h\}$.
2. If $\hat{p} \in (\hat{p}_1, \hat{p}_2)$, then $a_{t_n} \rightarrow h$ almost surely.
3. If $\hat{p} \in (\hat{p}_2, 1]$, then a_{t_n} almost surely does not converge.
4. If private beliefs are bounded and $\hat{p} = \hat{p}_1$, then $a_{t_n} \rightarrow h$ almost surely, while 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, then learning will be efficient in that informed agents will almost surely 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.

Theorem 1 and Corollary 1 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).

4. Discussion

A misspecified model of information processing interferes with asymptotic learning. 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 $\hat{p} \geq \hat{p}_2$, 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 herd on the correct state. An intervention must be long-term; once an intervention ceases, the herd will break.

Experimental evidence provides support for both the presence of uninformed agents and a misspecified belief about their frequency. In a social learning experiment, Goeree et al. (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. 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. Ziegelmeyer et al. (2010) 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.

Experimental evidence studying how people process correlated information also supports this form of model misspecification. Enke and Zimmermann (2015) show that subjects treat correlated information as independent when updating, and beliefs are too sensitive to correlated information sources.

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.

Appendix A

Proof of Lemma 1. Suppose $\lambda \geq (1 - \underline{b})/\underline{b}$. The strongest signal an agent can receive in favor of state H is \underline{b} . This leads to posterior $q(\lambda, \underline{b}) = \lambda \underline{b}/(1 - \underline{b}) \geq 1$ and an informed agent finds it optimal to choose $a = \ell$. Therefore, for any signal $s \geq \underline{b}$, an informed agent will choose action ℓ . Similarly, if $\lambda < (1 - \bar{b})/\bar{b}$, then an informed agent will choose action h for any signal $s \leq \bar{b}$. \square

Proof of Lemma 2. At a stationary point λ , $\phi(a, \lambda) = \lambda$ for all a such that $\psi(a|\omega, \lambda; p) > 0$. As $p < 1$ and uninformed agents are never in a cascade, $\psi(a|\omega, \lambda; p) > 0$ for all $a \in \{\ell, h\}$ and for all $(\omega, \lambda) \in \{L, H\} \times [0, \infty]$. Also, these actions are informative,

$$\frac{\psi(a|L, \lambda; \hat{p})}{\psi(a|H, \lambda; \hat{p})} \neq 1,$$

for all $a \in \{\ell, h\}$ and $\lambda \in (0, \infty)$. Therefore, $\{0, \infty\}$ are the only two values that satisfy $\phi(a, \lambda) = \lambda$ for all $a \in \{\ell, h\}$. \square

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

Lemma 6 (Condition for locally stable fixed point). *Given a finite set A , and Borel measurable functions $f : A \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ and $\rho : A \times \mathbb{R}_+ \rightarrow [0, 1]$ satisfying $\sum_{a \in A} \rho(a|x) = 1$. Let $x_1 \in \mathbb{R}$. Then the process $\{x_t\}_{t=0}^\infty$ where $x_{t+1} = f(a_t, x_t)$ with probability $\rho(a_t|x_t)$ for $a_t \in A$ is a Markov*

process. Let \tilde{x} be a fixed point of x . Suppose $f(a, \cdot)$ is continuously differentiable and $\rho(a|\cdot)$ is continuous at \tilde{x} for all $a \in A$. If

$$\sum_{a \in A} \rho(a|\tilde{x}) \log |f_x(a, \tilde{x})| < 0 \tag{9}$$

then \tilde{x} is locally stable.

Proof. See Corollary C.1 in [Smith and Sorensen \(2000\)](#). \square

Proof of Lemma 3. Suppose $\omega = H$. Let $(\Upsilon, \mathcal{F}, \mathbb{P})$ denote the underlying probability space for $\langle a_t, \lambda_t \rangle$ and define

$$g(a, \lambda) = \log \frac{\psi(a|L, \lambda; \hat{p})}{\psi(a|H, \lambda; \hat{p})}, \tag{10}$$

and $\rho(a|\lambda) = \psi(a|H, \lambda; p)$. Using this notation, $\log \lambda_{t+1} = \log \lambda_t + g(a_t, \lambda_t)$, $E[g(a, \lambda)] = \gamma(\hat{p}, \lambda)$ and $\gamma(\hat{p}, \lambda) = \rho(\ell|\lambda)g(\ell, \lambda) + \rho(h|\lambda)g(h, \lambda)$. The proof follows from [Claims 1–3](#).

Claim 1. If $\gamma(\hat{p}, 0) < 0$, then 0 is locally stable and if $\gamma(\hat{p}, \infty) > 0$, then ∞ is locally stable.

Proof. Applying [Lemma 6](#) to $\langle \lambda_t \rangle$, $A = \{\ell, h\}$, $\rho(a|\lambda) = \psi(a|H, \lambda; p)$, $f(a, \lambda) = \phi(a, \lambda; \hat{p})$ and

$$\phi_\lambda(a, \lambda; \hat{p}) = \frac{\psi(a|L, \lambda; \hat{p})}{\psi(a|H, \lambda; \hat{p})} + \lambda \frac{d}{d\lambda} \left(\frac{\psi(a|L, \lambda; \hat{p})}{\psi(a|H, \lambda; \hat{p})} \right). \tag{11}$$

Thus, at $\lambda = 0$, (9) is equal to $\gamma(\hat{p}, 0)$. This establishes that 0 is locally stable when $\gamma(\hat{p}, 0) < 0$.

Define Markov process $\langle \Lambda_t \rangle$ with transitions

$$\Psi(h|\omega, \Lambda; p) = pF^\omega \left(\frac{\Lambda}{1 + \Lambda} \right) + (1 - p)F^\omega(1/2) \tag{12}$$

$$\Psi(\ell|\omega, \Lambda; p) = 1 - \Psi(h|\omega, \Lambda; p) \tag{13}$$

$$\Phi(a, \Lambda; \hat{p}) = \Lambda \left(\frac{\Psi(a|H, \Lambda; \hat{p})}{\Psi(a|L, \Lambda; \hat{p})} \right) \tag{14}$$

Note $\Lambda_t = 1/\lambda_t$. The set of stationary points of $\langle \Lambda_t \rangle$ are $\{0, \infty\}$. Define

$$\Gamma(\hat{p}, \Lambda) := \sum_{a \in \{\ell, h\}} \Psi(a|H, \Lambda; p) \log \left(\frac{\Psi(a|H, \Lambda; \hat{p})}{\Psi(a|L, \Lambda; \hat{p})} \right). \tag{15}$$

Analogous to the preceding paragraph, 0 is locally stable when $\Gamma(\hat{p}, 0) < 0$. Note $\Gamma(\hat{p}, 0) = -\gamma(\hat{p}, \infty)$ and $\Lambda = 0$ corresponds to $\lambda = \infty$. Therefore ∞ is a locally stable point of $\langle \lambda_t \rangle$ when $\gamma(\hat{p}, \infty) > 0$. \square

Claim 2. If private beliefs are bounded and $\gamma(\hat{p}, 0) \geq 0$ ($\gamma(\hat{p}, \infty) \leq 0$), then for any $\lambda_0 \in (0, \infty)$, $P(\lambda_t \rightarrow 0) = 0$ ($P(\lambda_t \rightarrow \infty) = 0$). Thus, 0 (∞) is not locally stable.

Proof. Suppose private beliefs are bounded and $\gamma(\hat{p}, 0) \geq 0$. Let τ_1 be the stopping time corresponding to the period in which an h -cascade forms and never breaks,

$$\tau_1 = \inf \left\{ t \geq 1 \mid \lambda_i \in \mathcal{J}^h \ \forall i \geq t \right\},$$

and let $E = \{v \in \Upsilon | \tau_1(v) < \infty\}$ be the event in which an h -cascade forms in finite time and never breaks. If informed agents are in an h -cascade in period t , then $g(a_t, \lambda_t) = g(a_t, 0)$ and $\lambda_t < (1 - \bar{b})/\bar{b}$. Then on any sample path $v \in E$,

$$\log \lambda_t(v) = \log \lambda_{\tau_1(v)} + \sum_{i=\tau_1(v)}^{t-1} g(a_i(v), 0) < \log(1 - \bar{b})/\bar{b} \quad \forall t > \tau_1(v). \tag{16}$$

Suppose $\gamma(\hat{p}, 0) > 0$. In order for (16) to hold for v , it must be the case that

$$\limsup_{t \rightarrow \infty} \frac{1}{t - \tau_1(v)} \sum_{i=\tau_1(v)}^{t-1} g(a_i(v), 0) \leq 0. \tag{17}$$

By the Strong Law of Large Numbers,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^t g(a_i, 0) = \gamma(\hat{p}, 0) > 0 \quad \text{a.s.} \tag{18}$$

Thus, (17) cannot hold for a set of sample paths that occur with positive probability and it must be that $P(E) = 0$.

Suppose $\gamma(\hat{p}, 0) = 0$. Then on set E , $\langle \lambda_t \rangle$ has the same limit properties as a zero mean random walk with increments $g(a_t, 0)$. But

$$\limsup_{t \rightarrow \infty} \sum_{i=1}^t g(a_i, 0) = \infty \quad \text{a.s.} \tag{19}$$

Thus, (16), cannot hold for a set of sample paths that occur with positive probability and it must be that $P(E) = 0$.

Thus, if $\gamma(\hat{p}, 0) \geq 0$, $P(E) = 0$ and every h -cascade breaks with probability 1. Therefore, $P(\lambda_t \rightarrow 0) = 0$ from any $\lambda_0 \in (0, \infty)$. The proof for ∞ is analogous. \square

Claim 3. *If private beliefs are unbounded and $\gamma(\hat{p}, 0) > 0$ ($\gamma(\hat{p}, \infty) < 0$), then for any $\lambda_0 \in (0, \infty)$, $P(\lambda_t \rightarrow 0) = 0$ ($P(\lambda_t \rightarrow \infty) = 0$). Thus, 0 (∞) is not locally stable.*

Proof. Suppose $\gamma(\hat{p}, 0) > 0$ and private beliefs are unbounded. Let τ_1 be the stopping time corresponding to the period in which the likelihood ratio is less than M for all future periods,

$$\tau_1 = \inf \{t \geq 1 | \lambda_i < M \forall i \geq t\},$$

and let $E = \{v \in \Upsilon | \tau_1(v) < \infty\}$ be the event in which this stopping time is finite. Then on any sample path $v \in E$,

$$\log \lambda_t(v) = \log \lambda_{\tau_1(v)} + \sum_{i=\tau_1(v)}^{t-1} g(a_i(v), \lambda_i(v)) < M \quad \forall t > \tau_1(v). \tag{20}$$

In order for (20) to hold, it must be the case that for $v \in E$,

$$\limsup_{t \rightarrow \infty} \frac{1}{t - \tau_1(v)} \sum_{i=\tau_1(v)}^{t-1} g(a_i(v), \lambda_i(v)) \leq 0. \tag{21}$$

Next I construct a process on $(\Upsilon, \mathcal{F}, \mathbb{P})$ that converges to a positive limit almost surely. Define an i.i.d. sequence of random variables $(\alpha_1, \alpha_2, \dots)$ with

$$\alpha_t = \begin{cases} \ell & \text{if } (\theta_t = U \text{ and } s_t \geq 1/2) \\ h & \text{if } (\theta_t = I) \text{ or } (\theta_t = U \text{ and } s_t < 1/2). \end{cases} \tag{22}$$

Then α corresponds to the action that is chosen if $\lambda = 0$, with $P(\alpha) = \rho(\alpha|0)$. Given $\gamma(\hat{p}, 0) > 0$, by continuity of ψ , there exists an $M > 0$ such that

$$\rho(\ell|0)g(\ell, x) + \rho(h|0)g(h, y) > 0, \tag{23}$$

for all $x, y \in [0, M]$. Choose $\lambda_\ell, \lambda_h \in [0, M]$ such that

$$\lambda_\ell = \arg \min_{\lambda \in [0, M]} g(\ell, \lambda) \tag{24}$$

and

$$\lambda_h = \arg \min_{\lambda \in [0, M]} g(h, \lambda). \tag{25}$$

Note $E[g(\alpha_i, \lambda_{\alpha_i})] = \rho(\ell|0)g(\ell, \lambda_\ell) + \rho(h|0)g(h, \lambda_h) > 0$, where the inequality follows from $\lambda_h, \lambda_\ell \in [0, M]$ and (23). By the Strong Law of Large Numbers, for any finite $j \geq 1$,

$$\lim_{t \rightarrow \infty} \frac{1}{t-j} \sum_{i=j}^t g(\alpha_i, \lambda_{\alpha_i}) > 0 \quad \text{a.s.} \tag{26}$$

For $\lambda \in [0, M]$, $g(h, \lambda) \geq g(h, \lambda_h)$, $g(\ell, \lambda) \geq g(\ell, \lambda_\ell)$ and $g(\ell, \lambda) > g(h, \lambda_h)$, where the first two inequalities follow from the definition of λ_h, λ_ℓ , and the third follows from $g(\ell, \lambda) > 1$ and $g(h, \lambda_h) < 1$. Also, $(a_t, \alpha_t) \neq (h, \ell)$ by definition. Therefore, if $\lambda_t \leq M$, then $g(a_t, \lambda_t) \geq g(\alpha_t, \lambda_{\alpha_t})$. Therefore, for $v \in E$,

$$\sum_{i=\tau_1(v)}^t g(a_i(v), \lambda_i(v)) \geq \sum_{i=\tau_1(v)}^t g(\alpha_i(v), \lambda_{\alpha_i(v)}) \tag{27}$$

for all $t > \tau_1(v)$.

Combining (21) and (27), for $v \in E$,

$$0 \geq \limsup_{t \rightarrow \infty} \frac{1}{t - \tau_1(v)} \sum_{i=\tau_1(v)}^{t-1} g(\alpha_i(v), \lambda_{\alpha_i(v)}). \tag{28}$$

But given (26), inequality (28) is satisfied with probability 0. Therefore, $P(E) = 0$. Therefore, almost surely the likelihood ratio exceeds M infinitely often and $P(\lambda_t \rightarrow 0) = 0$ from any $\lambda_0 \in (0, \infty)$. The proof for ∞ is analogous. \square

Proof of Lemma 4. Suppose $\omega = H$. The proof follows from Claims 4–7.

Claim 4. For $a \in \{\ell, h\}$,

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

Proof. Suppose $a = h$. Then

$$\frac{d}{d\hat{p}} \left(\frac{\psi(h|L, \lambda; \hat{p})}{\psi(h|H, \lambda; \hat{p})} \right) = \frac{F^L(1/(\lambda + 1))F^H(1/2) - F^L(1/2)F^H(1/(\lambda + 1))}{[\hat{p}F^H(1/(\lambda + 1)) + (1 - \hat{p})F^H(1/2)]^2} \tag{29}$$

Given F^L/F^H is strictly increasing on $\text{supp}(F)$ (Smith and Sorensen, 2000, 2008), and $1/2 \in \text{supp}(F)$, if $\lambda > 1$, then $\frac{F^L(1/(\lambda+1))}{F^H(1/(\lambda+1))} < \frac{F^L(1/2)}{F^H(1/2)}$, if $\lambda < 1$, then $\frac{F^L(1/(\lambda+1))}{F^H(1/(\lambda+1))} > \frac{F^L(1/2)}{F^H(1/2)}$ and if $\lambda = 1$, then the numerator is 0, which establishes Claim 4 for $a = h$. The proof of $a = \ell$ is analogous. \square

Claim 5 (Local stability of 0). There exists a $\hat{p}_2 \in (p, 1)$ such that 0 is locally stable for $\hat{p} \in [0, \hat{p}_2)$ and 0 is not locally stable for $\hat{p} \in (\hat{p}_2, 1]$. If private beliefs are bounded, 0 is not locally stable for $\hat{p} = \hat{p}_2$.

Proof. By Lemma 3, 0 is locally stable if $\gamma(\hat{p}, 0) < 0$ and 0 is not locally stable if $\gamma(\hat{p}, 0) > 0$. If private beliefs are bounded, 0 is not locally stable if $\gamma(\hat{p}, 0) = 0$. By (5),

$$\begin{aligned} \gamma(\hat{p}, 0) &= \sum_{a \in \{\ell, h\}} \psi(a|H, 0; p) \log \left(\frac{\psi(a|L, 0; \hat{p})}{\psi(a|H, 0; \hat{p})} \right) \\ &= (1 - p)(1 - F^H(1/2)) \log \left(\frac{1 - F^L(1/2)}{1 - F^H(1/2)} \right) \\ &\quad + \left(p + (1 - p)F^H(1/2) \right) \log \left(\frac{\hat{p} + (1 - \hat{p})F^L(1/2)}{\hat{p} + (1 - \hat{p})F^H(1/2)} \right) \end{aligned} \tag{30}$$

Substituting $\hat{p} = 1$ into (30),

$$\gamma(1, 0) = (1 - p)(1 - F^H(1/2)) \log \left(\frac{1 - F^L(1/2)}{1 - F^H(1/2)} \right) > 0 \tag{31}$$

where $\frac{1 - F^L(1/2)}{1 - F^H(1/2)} > 1$ follows from $F^L(1/2) < F^H(1/2)$. Substituting $\hat{p} = p$ into (30),

$$\begin{aligned} \gamma(p, 0) &= \sum_{a \in \{\ell, h\}} \psi(a|H, 0; p) \log \left(\frac{\psi(a|L, 0; p)}{\psi(a|H, 0; p)} \right) \\ &< \log \left(\sum_{a \in \{\ell, h\}} \psi(a|H, 0; p) \left(\frac{\psi(a|L, 0; p)}{\psi(a|H, 0; p)} \right) \right) \\ &= \log \left(\sum_{a \in \{\ell, h\}} \psi(a|L, 0; p) \right) = 0 \end{aligned} \tag{32}$$

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

$$\frac{d\gamma(\hat{p}, 0)}{d\hat{p}} = \left(p + (1 - p)F^H(1/2) \right) \left(\frac{\psi(h|H, 0; \hat{p})}{\psi(h|L, 0; \hat{p})} \right) \frac{d}{d\hat{p}} \left(\frac{\psi(h|L, 0; \hat{p})}{\psi(h|H, 0; \hat{p})} \right) > 0 \quad (33)$$

where the inequality follows from Claim 4.

Therefore, $\gamma(\hat{p}, 0)$ is increasing in \hat{p} , $\gamma(p, 0) < 0$ and $\gamma(1, 0) > 0$. 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 locally stable, while for $\hat{p} > \hat{p}_2$, $\gamma(\hat{p}, 0) > 0$ and 0 is not locally stable. For $\hat{p} = \hat{p}_2$, $\gamma(\hat{p}, 0) = 0$; if private beliefs are bounded, 0 is not locally stable. \square

Claim 6 (Local stability of ∞).

1. If $p > p^*$, where p^* is defined in (8), there exists a $\hat{p}_1 \in (0, p)$ such that ∞ is locally stable for $\hat{p} \in [0, \hat{p}_1)$ and is not locally stable for $\hat{p} \in (\hat{p}_1, 1]$. If private beliefs are bounded, ∞ is not locally stable for $\hat{p} = \hat{p}_1$.
2. If $p < p^*$, then ∞ is not locally stable for all $\hat{p} \in [0, 1]$.
3. If $p = p^*$, then ∞ is not locally stable for all $\hat{p} \in (0, 1]$, and if private beliefs are bounded, ∞ is not locally stable for $\hat{p} = 0$.

Proof. Recall the Markov process $\langle \Lambda_t \rangle$ defined in (12)–(14), where $\Lambda_t = 1/\lambda_t$. At $\Lambda_t = 0$,

$$\begin{aligned} \Gamma(\hat{p}, 0) = & \left(\left(p + (1 - p)(1 - F^H(1/2)) \right) \log \left(\frac{\hat{p} + (1 - \hat{p})(1 - F^H(1/2))}{\hat{p} + (1 - \hat{p})(1 - F^L(1/2))} \right) \right. \\ & \left. + (1 - p)F^H(1/2) \log \left(\frac{F^H(1/2)}{F^L(1/2)} \right) \right) \end{aligned} \quad (34)$$

where Γ is the stability criterion of $\langle \Lambda_t \rangle$ defined in (15). Substituting $\hat{p} = 0$ into (34),

$$\begin{aligned} \Gamma(0, 0) = & \left((1 - (1 - p)F^H(1/2)) \log \left(\frac{1 - F^H(1/2)}{1 - F^L(1/2)} \right) \right. \\ & \left. + (1 - p)F^H(1/2) \log \left(\frac{F^H(1/2)}{F^L(1/2)} \right) \right) \end{aligned} \quad (35)$$

which is less than 0 when

$$p > 1 - \frac{\log \left(\frac{1 - F^L(1/2)}{1 - F^H(1/2)} \right)}{F^H(1/2) \left[\log \left(\frac{F^H(1/2)}{F^L(1/2)} \right) + \log \left(\frac{1 - F^L(1/2)}{1 - F^H(1/2)} \right) \right]} := p^*. \quad (36)$$

Substituting $\hat{p} = p$ into (34),

$$\begin{aligned} \Gamma(p, 0) &= \sum_{a \in \{\ell, h\}} \Psi(a|H, 0; p) \log \left(\frac{\Psi(a|H, 0; p)}{\Psi(a|L, 0; p)} \right) \\ &= - \sum_{a \in \{\ell, h\}} \Psi(a|H, 0; p) \log \left(\frac{\Psi(a|L, 0; p)}{\Psi(a|H, 0; p)} \right) \\ &> - \log \left(\sum_{a \in \{\ell, h\}} \Psi(a|H, 0; p) \frac{\Psi(a|L, 0; p)}{\Psi(a|H, 0; p)} \right) \\ &= 0 \end{aligned}$$

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

$$\frac{d\Gamma(\hat{p}, 0)}{d\hat{p}} = \Psi(\ell|H, 0; p) \frac{\Psi(\ell|L, 0; \hat{p})}{\Psi(\ell|H, 0; \hat{p})} \frac{d}{d\hat{p}} \left(\frac{\Psi(\ell|H, 0; \hat{p})}{\Psi(\ell|L, 0; \hat{p})} \right) > 0 \tag{37}$$

where the inequality follows from Claim 4 and $\Psi(a|\omega, 0; \hat{p}) = \psi(a|\omega, \infty; \hat{p})$.

Therefore, $\Gamma(\hat{p}, 0)$ is increasing in \hat{p} and $\Gamma(p, 0) > 0$.

- Case 1. When $p > p^*$, $\Gamma(0, 0) < 0$. By continuity, when $p > p^*$, 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 locally stable point of $\langle \Lambda_t \rangle$, while for $\hat{p} > \hat{p}_1$, $\Gamma(\hat{p}, 0) > 0$ and 0 is not locally stable. For $\hat{p} = \hat{p}_1$, $\Gamma(\hat{p}, 0) = 0$; if private beliefs are bounded, 0 is not locally stable.
- Case 2. When $p < p^*$, then $\Gamma(\hat{p}, 0) > 0$ for all $\hat{p} \in [0, 1]$ and 0 is not locally stable for any \hat{p} .
- Case 3. When $p = p^*$, then $\Gamma(0, 0) = 0$ and $\hat{p}_1 = 0$. For $\hat{p} > 0$, $\Gamma(\hat{p}, 0) > 0$ and 0 is not a locally stable point of $\langle \Lambda_t \rangle$. For $\hat{p} = 0$, $\Gamma(\hat{p}, 0) = 0$; if private beliefs are bounded, 0 is not locally stable.

For any \hat{p} , if 0 is a locally stable point of $\langle \Lambda_t \rangle$, then ∞ is a locally stable point of $\langle \lambda_t \rangle$. \square

Claim 7. $\hat{p}_1 < \hat{p}_2$.

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

Proof of Lemma 5. Suppose $\omega = H$. Let $(\Upsilon, \mathcal{F}, \mathbb{P})$ denote the underlying probability space for $\langle a_t, \lambda_t \rangle$ and define

$$g(a, \lambda) = \log \frac{\psi(a|L, \lambda; \hat{p})}{\psi(a|H, \lambda; \hat{p})}, \tag{38}$$

and $\rho(a|\lambda) = \psi(a|H, \lambda; p)$. Using this notation, $\log \lambda_{t+1} = \log \lambda_t + g(a_t, \lambda_t)$ and $\gamma(\hat{p}, \lambda) = \rho(\ell|\lambda)g(\ell, \lambda) + \rho(h|\lambda)g(h, \lambda)$.

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

$$\alpha_t = \begin{cases} \ell & \text{if } \{\theta_t = U \text{ and } s_t \geq 1/2\} \\ h & \text{if } \{\theta_t = I\} \text{ or } \{\theta_t = U \text{ and } s_t < 1/2\}. \end{cases} \tag{39}$$

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

$$X_t = \frac{g(\alpha_t, 0) - \gamma(\hat{p}, 0)}{\sigma}. \tag{40}$$

Then (X_1, X_2, \dots) are i.i.d. random variables with mean 0 and variance 1. By the Law of the Iterated Logarithm (LIL) (Hartman and Wintner, 1941),

$$\limsup_{t \rightarrow \infty} \frac{\sum_{i=1}^t X_i}{\sqrt{2t \log \log t}} = 1 \text{ a.s.} \tag{41}$$

Thus, for all $\epsilon > 0$,

$$P \left[\frac{1}{t} \sum_{i=1}^t g(\alpha_i, 0) \geq \gamma(\hat{p}, 0) + (1 + \epsilon)\beta_t \text{ i.o.} \right] = 0 \tag{42}$$

where

$$\beta_t := \sqrt{\frac{2\sigma^2 \log \log t}{t}}. \tag{43}$$

The proof of [Lemma 5](#) follows from [Claims 8–10](#), which establish that if λ is a locally stable point of $\langle \lambda_t \rangle$, then $P(\lambda_t \rightarrow \lambda) > 0$ from any initial value $\lambda_1 \in (0, \infty)$, [Claim 11](#), which rules out convergence to non-stationary points and [Claims 2–3](#), which establish when stationary points are not globally stable.

Claim 8. *Let*

$$E = \left\{ v \in \Upsilon \mid \frac{1}{t} \sum_{i=1}^t g(\alpha_i(v), 0) < \gamma(\hat{p}, 0) + (1 + \epsilon)\beta_t \text{ for all } t \geq 3, \epsilon > 0 \right\} \tag{44}$$

be the event that $\frac{1}{t} \sum_{i=1}^t g(\alpha_i, 0)$ never exceeds $\gamma(\hat{p}, 0) + (1 + \epsilon)\beta_t$ for all $t \geq 3$ and $\epsilon > 0$. Then there exists a $\delta > 0$ such that $P(E) \geq \delta$.

Proof. Let $S_t = \sum_{i=1}^t g(\alpha_i, 0)$. Fixing $\epsilon > 0$, define the number of times that the Law of the Iterated Logarithm bound is exceeded starting at time t ,

$$R_t = \sum_{i=t}^{\infty} I\{S_i > i\gamma(\hat{p}, 0) + (1 + \epsilon)i\beta_i\}$$

where I is the indicator function. From [\(42\)](#), $P(R_3 < \infty) = 1$. Let τ be the stopping time corresponding to the last time that S_t exceeds this boundary,

$$\tau = \inf\{T \geq 3 \mid S_t < t\gamma(\hat{p}, 0) + (1 + \epsilon)t\beta_t \text{ for all } t \geq T\}.$$

Then $R_\tau = 0$ by definition and $P(\tau < \infty) = 1$ by [\(42\)](#). For any $t < \infty$, the probability of no crossings during $i \in \{3, \dots, t\}$ is strictly positive, $P((R_3 - R_t) = 0) > 0$. Thus, for any v with $3 < \tau(v) < \infty$, there is a corresponding sample path v' such that $S_t(v) = S_t(v')$ for $t \geq \tau$ and $R_3(v') = 0$. Therefore, $P(R_3 = 0) > 0$ and there exists a $\delta > 0$ such that $P(E) \geq \delta$. \square

Claim 9. *If private beliefs are bounded and $\lambda \in \{0, \infty\}$ is locally stable, then $P(\lambda_t \rightarrow \lambda) > 0$ from any initial value $\lambda_1 \in (0, \infty)$.*

Proof. Suppose 0 is locally stable and private beliefs are bounded. Fix $\lambda_1 \in (0, \infty)$ and let $\eta - 1$ be the number of consecutive h actions required to start a cascade (η is deterministic and finite).¹⁰ Let E_n be the event that an h -cascade begins in period η and persists at least until period $\eta + n$,

$$E_n = \{\lambda_t \in \mathcal{J}^h \forall t \in \{\eta, \dots, \eta + n\}\} \tag{45}$$

¹⁰ The likelihood ratio in period t is based on actions a_1, \dots, a_{t-1} . Thus, $\eta - 1$ consecutive h actions will start a cascade in period η .

where $E_0 = \{\lambda_\eta \in \mathcal{J}^h\}$ is the event that a cascade begins in period η and $E_\infty = \{\lambda_t \in \mathcal{J}^h \forall t \geq \eta\}$ is the event that a cascade begins in period η and never breaks.

Suppose sample path $v \in E_0$. If $v \in E_n$, the likelihood ratio is equal to

$$\log \lambda_{\eta+n}(v) = \log \lambda_\eta(v) + \sum_{i=\eta}^{\eta+n-1} g(\alpha_i(v), 0), \tag{46}$$

since α_i coincides with a_i in an h -cascade. Thus, a sufficient condition for $v \in E_n$ is

$$\sum_{i=\eta}^{t-1} g(\alpha_i(v), 0) < 0 \quad \forall t \in \{\eta + 1, \dots, \eta + n\}. \tag{47}$$

From Claim 8, we know that

$$P \left(\sum_{i=\eta}^{t-1} g(\alpha_i, 0) < (t - \eta)(\gamma(\hat{p}, 0) + (1 + \epsilon)\beta_{t-\eta}) \forall t > \eta + 2, \epsilon > 0 \right) \geq \delta \tag{48}$$

where β_t is defined in (43). Given 0 is locally stable, by Lemma 3, $\gamma(\hat{p}, 0) < 0$. Given $\beta_t \rightarrow 0$, $\gamma(\hat{p}, 0) + (1 + \epsilon)\beta_t$ is eventually negative for any $\epsilon > 0$. Fix $\bar{\epsilon} > 0$ and let $k + 1$ be the number of periods required for the LIL bound to be negative,

$$k + 1 = \inf \{t \geq 3 | \gamma(\hat{p}, 0) + (1 + \bar{\epsilon})\beta_t < 0\} \tag{49}$$

(k is deterministic and finite). Then $\gamma(\hat{p}, 0) + (1 + \bar{\epsilon})\beta_t < 0$ for all $t \geq k + 1$.

Conditional on an h -cascade beginning in period η , the probability that the likelihood ratio remains in the h -cascade set through period $\eta + k$ is

$$P(E_k | E_0) \geq P \left(\sum_{i=\eta}^{t-1} g(\alpha_i, 0) < 0 \forall t \in \{\eta + 1, \dots, \eta + k\} \right) > \rho(h|0)^k > 0 \tag{50}$$

where the first inequality follows from (47) and the second inequality follows from the probability of k consecutive h actions. The probability of the h -cascade never breaking is

$$\begin{aligned} P(E_\infty | E_0) &\geq P \left(\sum_{i=\eta}^{t-1} g(\alpha_i, 0) < 0 \forall t > \eta \right) \\ &> P \left(\sum_{i=\eta}^{t-1} g(\alpha_i, 0) < \min\{0, (t - \eta)(\gamma(\hat{p}, 0) + (1 - \bar{\epsilon})\beta_{t-\eta})\} \forall t > \eta \right) \\ &\geq \delta \rho(h|0)^k \end{aligned}$$

where the second inequality follows from the LIL bound, which is less than 0 starting in period $\eta + k + 1$, and the third inequality follows from Claim 8 and (50). Finally, the probability an h -cascade forms in period η is bounded below by the probability of $\eta - 1$ consecutive uninformed agents who all choose h ,

$$P(E_0) \geq (1 - p)^{\eta-1} F^H(1/2)^{\eta-1} > 0. \tag{51}$$

Therefore, the probability that an h -cascade forms in period η and never breaks is strictly positive since

$$P(E_\infty) > \delta(1 - p)^{\eta-1} F^H (1/2)^{\eta-1} \rho(h|0)^k > 0. \tag{52}$$

As before, let random variable τ_1 be the stopping time corresponding to the first period in which an h -cascade forms and never breaks,

$$\tau_1 = \inf \left\{ t \geq 1 \mid \lambda_i \in \mathcal{J}^h \forall i \geq t \right\}. \tag{53}$$

Then the probability that an h -cascade forms in finite time and never breaks is strictly positive since

$$P(\tau_1 < \infty) > P(E_\infty) > 0. \tag{54}$$

For all v such that $\tau_1(v) < \infty$,

$$\log \lambda_t(v) < \log \left(\frac{1 - \bar{b}}{\bar{b}} \right) + \sum_{i=\tau_1(v)}^{t-1} g(\alpha_i(v), 0), \tag{55}$$

for all $t > \tau_1(v)$. Also,

$$\lim_{t \rightarrow \infty} \sum_{i=\tau_1(v)}^{t-1} g(\alpha_i(v), 0) = -\infty \quad \text{a.s.}, \tag{56}$$

where the convergence follows from $E[g(\alpha_t, 0)] = \gamma(\hat{p}, 0) < 0$. Therefore, for almost all sample paths v such that $\tau_1(v) < \infty$,

$$\lim_{t \rightarrow \infty} \log \lambda_t(v) = -\infty. \tag{57}$$

Therefore, $P(\lambda_t \rightarrow 0) = P(\tau_1 < \infty) > 0$. The proof for ∞ is analogous. \square

Claim 10. *If private beliefs are unbounded and $\lambda \in \{0, \infty\}$ is locally stable, then $P(\lambda_t \rightarrow \lambda) > 0$ from any initial value $\lambda_1 \in (0, \infty)$.*

Proof. Suppose 0 is locally stable and private beliefs are unbounded. By Lemma 3, $\gamma(\hat{p}, 0) = \rho(\ell|0)g(\ell, 0) + \rho(h|0)g(h, 0) < 0$. First I construct a process on $(\Upsilon, \mathcal{F}, \mathbb{P})$ that converges to a negative limit almost surely. By continuity of ψ , there exists an $M > 0$ such that

$$\rho(\ell|M)g(\ell, x) + \rho(h|M)g(h, y) < 0, \tag{58}$$

for all $x, y \in [0, M]$. Choose $\lambda_\ell, \lambda_h \in [0, M]$ such that

$$\lambda_\ell = \arg \max_{\lambda \in [0, M]} g(\ell, \lambda) \tag{59}$$

and

$$\lambda_h = \arg \max_{\lambda \in [0, M]} g(h, \lambda). \tag{60}$$

Define an i.i.d. sequence of random variables (v_1, v_2, \dots) with

$$v_t = \begin{cases} \ell & \text{if } (\theta_t = I \text{ and } s_t \geq s^*(M)) \text{ or } (\theta_t = U \text{ and } s_t \geq 1/2) \\ h & \text{if } (\theta_t = I \text{ and } s_t < s^*(M)) \text{ or } (\theta_t = U \text{ and } s_t < 1/2). \end{cases} \tag{61}$$

Then v corresponds to the action that is chosen if $\lambda = M$, with $P(v) = \rho(v|M)$.

Note $E[g(v_i, \lambda_{v_i})] = \rho(\ell|M)g(\ell, \lambda_\ell) + \rho(h|M)g(h, \lambda_h) < 0$, where the inequality follows from $\lambda_h, \lambda_\ell \in [0, M]$ and (58). By the Strong Law of Large Numbers, for any finite $j \geq 1$,

$$\lim_{t \rightarrow \infty} \frac{1}{t-j} \sum_{i=j}^t g(v_i, \lambda_{v_i}) < 0 \quad \text{a.s.} \tag{62}$$

Therefore, using similar logic to Claim 9, for any finite $j \geq 1$,

$$P \left(\sum_{i=j}^{t-1} g(v_i, \lambda_{v_i}) < 0 \forall t > j \right) > 0. \tag{63}$$

For $\lambda \in [0, M]$, $g(h, \lambda) \leq g(h, \lambda_h)$, $g(\ell, \lambda) \leq g(\ell, \lambda_\ell)$ and $g(h, \lambda) < g(\ell, \lambda_\ell)$, where the first two inequalities follow from the definition of λ_h, λ_ℓ , and the third follows from $g(\ell, \lambda_\ell) > 1$ and $g(h, \lambda) < 1$. Also, when $\lambda \in [0, M]$, $(a, v) \neq (\ell, h)$ by definition. Therefore, if $\lambda_t \leq M$, then $g(a_t, \lambda_t) \leq g(v_t, \lambda_{v_t})$.

Similar to Claim 9 with \mathcal{J}^h replaced by $[0, M]$, define η as the number of consecutive h actions required for the likelihood ratio to fall below M and let

$$E_n = \{\lambda_t \in [0, M] \forall t \in \{\eta, \dots, \eta + n\}\}. \tag{64}$$

Thus, if $v \in E_0$ and

$$\sum_{i=\eta}^{t-1} g(v_i(v), \lambda_{v_i(v)}) < 0 \quad \forall t \in \{\eta + 1, \dots, \eta + n\}. \tag{65}$$

then $v \in E_n$. By (63), (65) holds on a set of sample paths with positive measure for any n . This establishes that $P(E_\infty|E_0) > 0$. Finally, as in Claim 9, $P(E_0) > 0$. Therefore, the probability that the likelihood ratio falls below M in period η and never again exceeds M is strictly positive, $P(E_\infty) > 0$. By similar logic to Claim 9, on this set of sample paths, $\lambda_t \rightarrow 0$ almost surely. Therefore, $P(\lambda_t \rightarrow 0) > 0$. The proof for ∞ is analogous. \square

Claim 11. If λ is not a stationary point of $\langle \lambda_t \rangle$, then $P(\lambda_t \rightarrow \lambda) = 0$.

Proof. 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 $\langle \lambda_t \rangle$. Therefore, if $P(\lambda_t \rightarrow \lambda) > 0$, then $\lambda \in \{0, \infty\}$. \square

Proof of Theorem 1. Suppose $\omega = H$. The proof follows from Claims 12–14.

Claim 12. If $p > p^*$ and $\hat{p} < \hat{p}_1$, then $\lambda_t \rightarrow \lambda_\infty$ almost surely, where λ_∞ is a random variable with $\text{supp}(\lambda_\infty) = \{0, \infty\}$.

Proof. Suppose $p > p^*$ and $\hat{p} < \hat{p}_1$. By Lemma 4, the set of locally stable points is $\{0, \infty\}$ and by Lemma 5, $P(\lambda_t \rightarrow \lambda) > 0$ iff $\lambda \in \{0, \infty\}$. Thus, it is necessary to rule out incomplete learning to show that there exists an λ_∞ with $\text{supp}(\lambda_\infty) = \{0, \infty\}$ such that $\lambda_t \rightarrow \lambda_\infty$ almost surely.

Suppose private beliefs are bounded. Let $\tau_3 = \inf\{t \geq 1 | \lambda_t \in \mathcal{J}\}$ be the stopping time corresponding to the first time that the likelihood ratio enters the cascade set and let $\tau_4 = \inf\{t > \tau_3 | \lambda_t \notin \mathcal{J}\}$ be the stopping time corresponding to the first time that the likelihood ratio leaves

the cascade set. For any λ_1 , $P(\tau_3 < \infty) = 1$ and by Lemma 5, $P(\tau_4 < \infty) < 1$ since the cascade persists with positive probability. The same holds for subsequent cascades. Therefore, $P(\lambda_t \notin \mathcal{J} \text{ i.o.}) = 0$ and the likelihood ratio eventually remains in the cascade set almost surely. Lemma 5 established belief convergence on any sample path that remains in the cascade set. Thus, there exists a random variable λ_∞ with $\text{supp}(\lambda_\infty) = \{0, \infty\}$ such that $\lambda_t \rightarrow \lambda_\infty$ almost surely.

Suppose private beliefs are unbounded. Similar logic establishes that for any $(\lambda_1, \lambda_2) \subset (0, \infty)$, $P(\lambda_t \in (\lambda_1, \lambda_2) \text{ i.o.}) = 0$. Therefore, there exists a τ such that $P(\lambda_t \in [0, \lambda_1] \cup [\lambda_2, \infty] \forall t > \tau) = 1$. Choosing λ_1 small enough and λ_2 large enough can be used to establish convergence. \square

Claim 13. *If $\hat{p} \in (\hat{p}_1, \hat{p}_2)$, then $\lambda_t \rightarrow 0$ almost surely.*

Proof. Similar logic to Claim 12, substituting \mathcal{J}^h for \mathcal{J} in the case of bounded private beliefs and setting $\lambda_2 = \infty$ in the case of unbounded private beliefs, establishes the claim. \square

Claim 14. *If $\hat{p} > \hat{p}_2$, then λ_t almost surely does not converge or diverge.*

Proof. When $\hat{p} > \hat{p}_2$, the likelihood ratio almost surely doesn't converge to 0 or diverge to ∞ . By Lemma 5, these are the only two candidate limit points. Therefore, learning is incomplete.

$P(\lambda_t \notin \mathcal{J} \text{ i.o.}) = 1$ follows immediately from Claim 2 for bounded private beliefs, which establishes when a cascade breaks with probability 1, and from Claim 14 for unbounded private beliefs, which establishes that the likelihood ratio almost surely does not enter the cascade set. When private beliefs are bounded, Claim 13 also holds for $\hat{p} = \hat{p}_1$ and Claim 14 also holds for $\hat{p} = \hat{p}_2$. \square

References

- Acemoglu, D., Chernozhukov, V., Wildiz, M., 2016. Fragility of asymptotic agreement under Bayesian learning. *Theor. Econ.* 11 (1), 187–225.
- Acemoglu, D., Dahleh, M.A., Lobel, I., Ozdaglar, A., 2011. Bayesian learning in social networks. *Rev. Econ. Stud.* 78 (4), 1201–1236.
- Banerjee, A.V., 1992. A simple model of herd behavior. *Q. J. Econ.* 107, 797–817.
- Berk, R.H., 1966. Limiting behavior of posterior distributions when the model is incorrect. *Ann. Math. Stat.* 37, 51–58.
- Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* 100, 992–1026.
- Bohren, A., 2012. Information and incentives in stochastic games, social learning and crowdsourcing. Ph.D. thesis. University of California, San Diego.
- Camerer, C.F., Ho, T.-H., Chong, J.-K., 2004. A cognitive hierarchy model of games. *Q. J. Econ.* 119, 861–898.
- Costa-Gomes, M.A., Crawford, V.P., Iriberry, N., 2009. Comparing models of strategic thinking in Van Huyck, Battalio, and Beil's coordination games. *J. Eur. Econ. Assoc.* 7, 365–376.
- DeGroot, M.H., 1974. Reaching a consensus. *J. Am. Stat. Assoc.* 69, 118–121.
- Demarzo, P.M., Vayanos, D., Zwiebel, J., 2003. Persuasion bias, social influence, and unidimensional opinions. *Q. J. Econ.* 118, 909–968.
- Enke, B., Zimmermann, F., 2015. Correlation neglect in belief formation. University of Zurich. Mimeo.
- Epstein, L.G., Noor, J., Sandroni, A., 2010. Non-Bayesian learning. *B.E. J. Theor. Econ.* 10.
- Esponda, I., Pouzo, D., 2015. Berk-Nash equilibrium: a framework for modeling agents with misspecified models. Mimeo.
- Eyster, E., Rabin, M., 2005. Cursed equilibrium. *Econometrica* 73, 1623–1672.
- Eyster, E., Rabin, M., 2010. Nave herding in rich-information settings. *Am. Econ. J. Microecon.* 2, 221–243.

- Goeree, J.K., Palfrey, T.R., Rogers, B.W., McKelvey, R.D., 2007. Self-correcting information cascades. *Rev. Econ. Stud.* 74, 733–762.
- Gottlieb, D., 2015. Will you never learn? Self deception and biases in information processing. Mimeo.
- Guarino, A., Jehiel, P., 2013. Social learning with coarse inference. *Am. Econ. J. Microecon.* 5, 147–174.
- Hartman, P., Wintner, A., 1941. On the law of the iterated logarithm. *Am. J. Math.* 63, 169–176.
- Jehiel, P., 2005. Analogy-based expectation equilibrium. *J. Econ. Theory* 123 (2), 81–104.
- Kubler, D., Weizsacker, G., 2004. Limited depth of reasoning and failure of cascade formation in the laboratory. *Rev. Econ. Stud.* 71, 425–441.
- Levy, G., Razin, R., 2015. Perceptions of correlation, communication, and ambiguity, 1–29. Mimeo.
- Moscarini, G., Ottaviani, M., Smith, L., 1998. Social learning in a changing world. *Econ. Theory* 11, 657–665.
- Mueller-Frank, M., Neri, C., 2015. A general model of boundedly rational observational learning: theory and evidence. Mimeo.
- Rabin, M., Schrag, J.L., 1999. First impressions matter: a model of confirmatory bias. *Q. J. Econ.* 114, 37–82.
- Schwartzstein, J., 2014. Selective attention and learning. *J. Eur. Econ. Assoc.* 12, 1423–1452.
- Smith, L., Sorensen, P., 2000. Pathological outcomes of observational learning. *Econometrica* 68, 371–398.
- Smith, L., Sorensen, P.N., 2008. Rational social learning with random sampling. University of Wisconsin. Madison. Mimeo.
- Wilson, A., 2014. Bounded memory and biases in information processing. *Econometrica* 82, 2257–2294.
- Ziegelmeyer, A., Bracht, J., Koessler, F., Winter, E., 2010. Fragility of information cascades: an experimental study using elicited beliefs. *Exp. Econ.*, 121–145.