Information-Processing Bias in Social Learning*

J. Aislinn Bohren†
University of California, San Diego

March 22, 2010
First Draft: April, 2009

Abstract

This paper explores how individuals learn from their predecessors when they are subject to biased beliefs about the information processing capabilities of others. I consider a social learning environment in which individuals observe private signals, and learning is asymptotically efficient in the absence of information processing biases. When individuals underestimate others' information processing capabilities, herds are less fragile and incorrect herds may persist forever. On the other hand, when individuals overestimate others' information processing, correct herds may break infinitely often, causing beliefs to perpetually fluctuate. When beliefs about others are approximately correct, learning is complete, so the model with no information-processing bias is robust to slight perturbations of beliefs.

*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@ucsd.edu; Web: http://econ.ucsd.edu/∼abohren; Address: 9500 Gilman Drive, Department of Economics, La Jolla, CA 92093-0508.

1 Introduction

Observational learning plays an important role in the transmission of information, opinions and behavior. People may use a bestseller list to guide their purchase of a book or a car. Observing high participation rates amongst co-workers may increase the likelihood that a person contributes to their retirement plan. Social learning can also influence 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 wide range of situations influenced by observational learning, it is important to understand how biases in information processing affect learning. This paper explores how an information processing bias may interfere with the efficiency of social learning, and demonstrates that such biases can partially explain how inefficient choices can persist even when contradicted by public information.

This paper extends standard herding models in the tradition of Banerjee [1992] and Bikhchandani et al. [1992] to incorporate the idea of information-processing bias (BIP). In the standard model with binary actions and signals, individuals have common-value preferences that depend on an unknown state of the world. Agents act sequentially, observing a private signal and the actions of previous agents before choosing an action. An information cascade occurs when it is optimal for an agent to ignore his private signal and act only on the basis of the information contained in the actions of previous agents. When this occurs, all subsequent agents follow suit and new information ceases to aggregate. With positive probability, agents herd on the suboptimal action and thus the equilibrium is inefficient.

A critical feature of this model is common knowledge of how individuals process information. Agents understand exactly how preceding agents incorporated the action history into their decision-making process, and are therefore aware of which actions contain no information. Since the herd is based on only a few initial signals, public beliefs about the state remain fragile and are easily reversed by the arrival of new information. Thus, if some agents don't observe prior actions and follow their private signal, or if public information is released periodically, social learning is asymptotically efficient.

However, what happens if agents are unsure about the information-processing capabilities of other agents? What if they believe the actions of previous agents reveal more information about private signals than is actually the case during a cascade, or what if they attribute too many actions to herding and are not sensitive enough to new information? This paper examines how a behavioral bias in information processing, which I refer to as information-processing bias (BIP), can interfere with optimal information aggregation even in settings where new information continues to arrive frequently during a cascade. Individuals subject to BIP are biased in their perception of the information-processing capabilities of others, and consequently fail to accurately disentangle repeated and new information.

In particular, suppose that a fraction of individuals do not observe preceding actions and select an action solely based on their private information. These uninformed agents simply do not have access to all available information, or are boundedly rational and unable to process multiple sources of information. Regardless of the justification for their presence, these uninformed agents always reveal their private signal. Individuals who incorporate the action history into their decision observe the full sequence of preceding actions but are uncertain about the information-processing capabilities of others. Consequently, these informed decision makers face an inferential challenge when extracting information from the actions of others, and their behavior will hinge on their beliefs about the population.

To fix ideas, suppose that each individual observes the history and is fully informed with probability p, and with probability $(1-p)$ is an uninformed type that only observes his private signal. Each informed individual believes that any other individual is informed with probability $\hat{p}$, where $\hat{p}$ need not coincide with p. The difference between p and $\hat{p}$ may arise because even very sophisticated individuals may underestimate or overestimate the information possessed by others, and so it is natural to allow for the distinction.

When $\hat{p} < p$ then an informed decision maker underestimates the fraction of preceding informed individuals. Accordingly, when this decision maker observes a series of identical actions, he incorrectly attributes too many of these actions to the private signals of uninformed individuals. This effect leads him to overweight information from the public history, and may allow public beliefs about the state to become entrenched. On the other hand, when $\hat{p} > p$, then an informed decision maker underweights the new information contained in correlated actions, rendering herds more fragile to contrary information.

To understand how BIP affects eventual efficiency and learning requires careful analysis of the rate of information accumulation. I characterize conditions that allow a herd to persist with positive probability, and conditions that ensure a herd breaks. Using these conditions and fixing the share of informed agents, I establish thresholds on beliefs about the share of informed agents, $\widehat{p}_1$ and $\widehat{p}_2$, such that when $\widehat{p} < \widehat{p}_1$ an incorrect herd can persist with positive probability and when $\widehat{p} > \widehat{p}_2$ a correct herd will always break. When beliefs fall between these two thresholds, $\widehat{p} \in (\widehat{p}_1, \widehat{p}_2)$, incorrect herds always break but correct herds can persist, so eventually a correct herd will persist. Herding will be efficient in that informed agents will choose the optimal action all but finitely often. Otherwise, there is positive probability that herding will be inefficient and informed agents will choose the suboptimal action infinitely often. Correct beliefs about agent types lead to efficient herding since $p \in (\widehat{p}_1, \widehat{p}_2)$.

During a herd, beliefs continue to strengthen. When a herd persists in the long run, public beliefs will converge to a point mass on the state matching the action agents are herding on. Thus, if a correct herd persists, then learning is complete, while if an incorrect herd persists, learning is fully incorrect. If no herd persists, then beliefs remain interior and fluctuate, and learning remains incomplete. Fully incorrect learning or the perpetual fluctuation of beliefs are possible because the conditional public likelihood ratio is no longer a martingale when BIP is present. In fact, when $\widehat{p} < p$, agents overweight herding actions and the conditional public likelihood ratio in an incorrect herd is a submartingale. This explains why fully incorrect learning is possible. On the other hand, when $\widehat{p} > p$, agents underweight herding actions. The conditional public likelihood ratio is a supermartingale in incorrect herds and a submartingale in correct herds. When beliefs are sufficiently incorrect, the submartingale diverges and both types of herds to eventually break.

BIP in the context of a herding model has important implications. A failure to recognize repeated information can confound learning by allowing an incorrect herd to persist even when new information counteracts the incorrect herd, whereas a failure to recognize new information can cause correct herds to continually break (as well as incorrect herds). Every time a correct herd is broken, there is a chance that an incorrect herd may form in its place. These results are robust to the inclusion of other sources of new information, such as public signals, gurus (perfectly informed agents) or continuous signals. Whenever BIP is severe enough that repeated information accumulates at a faster rate than new information, then incorrect herds persist with positive probability.

To illustrate the relevance of this result, consider a public health campaign to increase awareness about the risks of HIV. Agents need to decide if HIV is a threat to them, and whether to take appropriate precautions. They observe public signals from the government and other public health agencies, along with the actions of previous agents. If all agents are herding on the actions of a few initial agents who didn't believe that HIV was a significant threat, then the public health campaign should eventually overturn this herd. However, if agents are subject to BIP, then observing many preceding agents who didn't believe that HIV was a threat will lead to strong beliefs that this is the case, making it less likely that the public health campaign will successfully overturn the herd.[1] In this scenario, the best way to quash the false cascade is to release public signals immediately and frequently. This contrasts with the case of no BIP, in which the timing of public signal release is irrelevant.

Individuals may also use the history to learn about the information processing capabilities of other agents. In Section 4, I examine what happens when agents can also learn about p. Although fully incorrect learning is generally precluded in this setting, incorrect herds may still persist with positive probability. Herding is more likely to be efficient if several herds form and break before a herd persists or if a low share of agents observe the history.

BIP relates to the notion of persuasion bias first introduced by DeMarzo et al. [2001] in a model of opinion formation in networks. In their paper, decision makers embedded in a network graph treat correlated information from others as being independent, leading to informational inefficiencies. Although my paper studies a very different environment than theirs, BIP provides a natural analogue for considering persuasion bias in social learning.

Banerjee [1992] and Bikhchandani et al. [1992] first modelled social learning in a sequential

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

setting, as discussed above. Smith and Sorensen [2000] extend their models to include continuous signals. An unbounded signal space is sufficient to ensure complete learning, eliminating the possibility of inefficient cascades. Acemoglu et al. [2010] examines social learning in a general social network, which includes the sequential learning and uninformed agents networks as special cases. As such, BIP can also be viewed as a boundedly rational extension of the uninformed agents network topology in Acemoglu et al. [2010].

This paper is most closely related to concurrent work on social learning by Eyster and Rabin [2009]. 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 information-processing bias differ in generality and interpretation. Inferential naivety considers the case in which every repeated action is viewed as being independent with probability 1, whereas in the BIP model, most decision makers are sophisticated and recognize that some repeated actions may stem from herding behavior, but misperceive the exact proportion of repeated information. The analogue of inferential naivety in my environment corresponds to $\hat{p} = 0$ and $p = 1$. As such, both papers provide complementary explanations for the robustness of inefficient learning. Eyster and Rabin [2009] also embed inferential 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 BIP and inferentially naive models with rational agents who do not know precisely which previous agents are also rational, naive or uninformed, and perhaps are even uncertain about the share of each type of agent is an interesting avenue left open for future research.

Guarino and Jehiel [2009] explore boundedly rational information processing in a sequential learning environment using 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 degree to which agents overweigh initial signals increases in a linear fashion, preventing these initial signals from permanently dominating subsequent new information. This contrasts with Eyster and Rabin [2009], the degree to which agents overweigh initial signals doubles each period, allowing a few early signals to overwhelm all future signals. As in the fully rational model, complete learning no longer obtains in an ABEE when actions are discrete.

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. A player with BIP understands the correlation between a player's type and their action choice, but incorrectly predicts the distribution of action choices in equilibrium.

BIP also relates to the recent literature on initial response models, including level-k analysis and cognitive hierarchy models.[2] The premise of these models is that agents best respond to their beliefs about how others act, but unlike equilibrium analysis, these beliefs are not required to be correct. 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 in the BIP model. Level 2 types believe all other agents follow their private signal, and thus act as BIP informed agents with beliefs $\widehat{p} = 0$. Consequently, the main difference between the two models stems from the beliefs informed agents have about other agents' types - BIP informed agents can place positive weight on other agents using a level 2 decision rule, whereas "informed agents" in a level k analysis believe that all other agents use a level 1 decision rule. BIP itself stems from level 2 agents misperceiving the share of other agents who are level 2. There is no such misperception in level k models, as all level k agents place probability 1 on other agents being level k-1. The comparison to a cognitive hierarchy (CH) model is similar - level 1 agents correspond to BIP uninformed agents, while level 2 agents act like BIP informed agents with beliefs $\widehat{p} = 0$, who also believe some previous actions convey only noise (i.e. CH level 2 agents place positive probability on level 0 and level 1 types, but probability 0 on other level 2 types).

The organization of this paper proceeds as follows. Section 2 sets up the model and explores the conditions under which learning is confounded in the presence of BIP. Section 3 explores the robustness of the model to several extensions, including public signals, continuous signals and private values. Section 4 allows agents to also learn about p, while Section 5 discusses experimental evidence in support of BIP and concludes. All proofs are in the Appendix.

2 Model

The basic set-up of this model mirrors the standard sequential learning model with binary action and signal spaces. There are two payoff-relevant states of the world, $\omega \in \{L, R\}$ with common prior belief $P(\omega = L) = \pi^L \in (1/2, 1)$.[3] Nature selects one of these states at the beginning 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: $u(a_t, \omega) = 1_{a_t = \omega}$.

Before selecting an action, each agent privately observes a binary signal about the state of the world, $s_t \in \{l, r\}$, which is i.i.d. conditional on the state with precision $\pi^s \in (\pi^L, 1)$.[4]

[2] Costa-Gomes et al. [2009]Camerer et al. [2004]

[3] An asymmetric prior obviates the need for breaking indifference.

[4] π^s is defined such that $P(s_t = l|\omega = L) = P(s_t = r|\omega = R) = \pi^s$

There are two types of agents. With probability $p > 0$, an agent is a socially informed type who observes the prior action choices of other agents. This agent uses her private signal and the action history to guide her action choice. The public history observed by informed agents is represented as $h_t = (a_1, ..., a_{t-1})$. With probability $1 - p$, an agent is a socially uninformed type 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 action history. This type's decision is solely guided by the information contained in his private signal.

Informed agents may misperceive the information-processing capabilities of others. Each informed individual believes that any other individual is informed with probability $\hat{p}$, where $\hat{p}$ need not coincide with p. This captures the fact that there is higher-order uncertainty over the level of information possessed by other agents, which we will refer to as information-processing bias (BIP). The difference between p and $\hat{p}$ may arise because even very sophisticated individuals may underestimate or overestimate the information possessed by others, and so it is natural to allow for the distinction. Incorrect beliefs about p can persist because no agent ever learns what the preceding agents actually observed or incorporated into their decision-making processes. Consequently, these informed decision makers face an inferential challenge when extracting information from the actions of others, and their behavior will hinge on their beliefs about the population. This bias interferes with optimal information aggregation if agents fail to accurately disentangle repeated and new information. An informed agent believes that other agents also hold the same beliefs about whether previous agents are informed or uninformed. Although requiring agents to hold identical misperceptions about others is admittedly restrictive, it is a good starting point to examine the possible implications of BIP. Extending the model to allow for heterogenous biases is left for future research.

Agents use Bayes rule to formulate beliefs about the state of the world. Denote public beliefs of informed agents at the beginning of period t by $\mu_t = P(\omega = L|h_t)$. Public beliefs depend on the history and beliefs about the share of informed agents. Denote private beliefs by μ_t^r if agent t observes a private r signal and μ_t^l if agent t observes a private l signal. Private beliefs for informed agents depend on public beliefs and their private signal realization, while private beliefs for agents who don't observe the history depend on only their private signal realization. Each agent maximizes expected payoffs with respect to their private beliefs about ω. For the uninformed type, this implies an agent chooses the action that corresponds to his private signal, while for an informed type, an agent chooses the action that corresponds to the state he believes is more likely, given his beliefs about p.

An *information cascade* occurs when it is optimal for an agent who observes the history to choose the same action regardless of his private signal realization. Throughout the paper, such an agent's action choice is described as *herding*. When herding arises, the agent's action reveals nothing about his private information, and social learning is impeded. An information cascade breaks when it becomes optimal for an informed agent to follow his private signal (i.e.

in an L-herd, it is optimal for an informed agent to choose R if he receives an r signal). For a given sample path, we say the information cascade *persists in the limit* if it persists in every period for $t = 1, 2, ...$ and the information cascade *breaks* if $\exists \tau < \infty$ such that the information cascade breaks at period τ. The probability that such sample paths occur will determine the probability that a given information cascade persists or breaks.

This paper examines how the efficiency of information cascades depends on the relationship informed agents perceive between prior actions and signals - that is, their beliefs about the share of informed agents, $\hat{p}$, which determines the accuracy of the inference drawn from the history during a herd.

2.1 Cascades in the Benchmark Model

In the benchmark model with common knowledge of no informed agents, inefficient herding arises but the herds are not robust. To see this, consider the following. Define Δ_t as the difference between the number of L and R actions at the beginning of time t. The unique Nash Equilibrium when $p = \widehat{p} = 1$ is to herd on L whenever Δ_t reaches 1, and to herd on R whenever Δ_t reaches -2.[5] An information cascade begins with probability 1, and occurs on the suboptimal action with positive probability. However, these cascades are very fragile - because new information ceases to aggregate once a cascade begins, the cascade is easily reversed if additional information becomes available. For example, a bounded public signal could overturn the herd.

Augmenting the benchmark model with uninformed agents allows information from action choices to continue accumulating in a cascade.[6] We will see in section 2.2.3 that in the absence of BIP, the addition of uninformed agents ($p = \widehat{p} < 1$) leads to complete learning. However, what if players are uncertain about the share of agents who are informed? What if they believe the actions of previous agents reveal more information about private signals than is actually the case during a herd, or what if they attribute too many actions to herding and are not sensitive enough to the new information? The remainder of this section explores the conditions that allow an information cascade to persist and the conditions that ensure an information cascade breaks, and uses these conditions to examine the impact of BIP on learning.

[5]Note that the asymmetry in Δ_t required for the formation of an L-herd versus R-herd stems from the specification of the prior $\pi^L > 1/2$.

[6]The conditions for a herd to begin in the presence of uninformed agents are identical to the conditions in the benchmark case of $p = 1$, regardless of beliefs $\widehat{p}$. Before the formation of a herd, all agents are following their signal regardless of whether they observed the history, and the public likelihood ratio evolves in the same manner as in the benchmark case, leading to the same conditions for herd formation.

2.2 How Does BIP Affect Learning?

Before a herd forms, all agents follow their private signal. Since the decisions of informed and uninformed agents coincide, informed agents correctly infer that previous actions perfectly reveal private information and BIP doesn't affect behavior prior to the onset of a herd. However, BIP interferes with information aggregation during a herd.

Suppose a herd has begun on action L. Each subsequent L action during the herd is attributed to (i) an uninformed agent who followed his private signal with probability $(1-\widehat{p})$; and (ii) an agent who observed the history and followed the herd with probability $\widehat{p}$. The public likelihood ratio following an L action is updated as follows:

$$\frac{\mu_t}{1-\mu_t} = \left(\frac{\widehat{p}+(1-\widehat{p})\pi^s}{\widehat{p}+(1-\widehat{p})(1-\pi^s)}\right)\left(\frac{\mu_{t-1}}{1-\mu_{t-1}}\right)$$

Thus, an action that follows the herd still reveals some information. When $\widehat{p} < p$, informed agents overweight the informativeness of this action, leading to an upward bias in the likelihood ratio relative to correct beliefs. The opposite occurs when $\widehat{p} > p$: agents attribute too many L actions to herding rather than private signals, resulting in a downward bias in the likelihood ratio. Let $\phi^h = \left(\frac{\widehat{p}+(1-\widehat{p})\pi^s}{\widehat{p}+(1-\widehat{p})(1-\pi^s)}\right)$ represent the information accumulating from a supporting action, or an action that follows the herd.

When it is still optimal for informed agents to herd based on $\widehat{p}$ and the history, a decision-maker will attribute a contrary action to an uninformed agent. In an L-herd, each contrary action R is attributed to an agent who did not observe the history and received a private r signal. The public likelihood ratio following an R action is updated as follows:

$$\frac{\mu_t}{1-\mu_t} = \left(\frac{1-\pi^s}{\pi^s}\right)\left(\frac{\mu_{t-1}}{1-\mu_{t-1}}\right)$$

Let $\phi^c = \left(\frac{1-\pi^s}{\pi^s}\right)$ represent the information accumulating from a contrary action, or an action that doesn't follow the herd. Note that beliefs $\widehat{p}$ do not bias the informativeness of an R action.

The public likelihood ratio increases with each supporting action and decreases with each contrary action. Supporting actions are belived to reveal new information with probability $(1-\widehat{p})$, whereas contrary actions reveal new information with probability 1. Therefore, a contrary action is more informative than a supporting action.

Without loss of generality, normalize to zero the period in which the action that begins the herd is chosen, so the length of the herd at the beginning of period t is equal to t. Let $\Delta_t \in [0,1]$ be the fraction of contrary actions chosen after the onset of a herd.[7] In an L-herd, $\Delta_t = \frac{1}{t-1}\sum_{s=1}^{t-1} 1_{a_s=R}$ represents the share of R actions and $1-\Delta_t$ represents the share of L actions. In a herd, Δ_t is a sufficient statistic for the history when examining the evolution of

[7] If $p = 1$ then no contrary actions are observed and $\Delta_t^a = 0 \ \forall t$.

the public likelihood ratio.

2.2.1 When Does a Herd Break?

A herd breaks when sufficient information accrues in favor of the alternative state such that an agent who observes the history finds it optimal to follow her signal.. When contrary actions are possible ($p < 1$), this happens with positive probability in any herd. A finite number of contrary actions can overturn a herd of any length. When $p = 1$, no contrary actions occur and the herd will never break. Theorem 1 demonstrates these results.

Theorem 1. *Suppose a herd is occurring in period t with contrary action share Δ_t. If $p < 1$ then there exists a set of sample paths that occur with positive probability along which the herd breaks within a finite number of periods after t. Otherwise, the herd will never break.*

This Theorem demonstrates that every herd breaks with positive probability when some agents don't observe the history, and the result holds for any belief $\widehat{p}$.

2.2.2 When Does a Herd Persist?

Given that a herd breaks with positive probability when $p < 1$, I now examine when such a herd can also persist with positive probability, and when the herd breaks with probability 1. Define the *herd breaking threshold* $\Delta_t^*(\widehat{p})$ as the minimum share of contrary actions that will result in an agent following her private signal in period t. Whenever the actual share of contrary actions is greater than this threshold, the cascade will break. This threshold depends on informed agents' beliefs about uninformed agents.

The herd breaking threshold is calculated by finding the value of Δ_t such that the private likelihood ratio is equal to one when a contrary private signal is realized. Consider an L-herd. If an informed agent's private likelihood ratio remains greater than 1 when a private r signal is realized, then the agent will continue the herd. However, if the realization of a private r signal will flip her private likelihood ratio below 1, then the agent will choose action R when she receives an r signal and L when she receives an l signal. Thus, her action choice reveals her signal and the herd is broken. Lemma 1 formally characterizes the herd breaking threshold.

Lemma 1. *For each $t > 1$, the herd breaking threshold $\Delta_t^*(\widehat{p})$ can be represented as:*

$$\Delta_t^*(\widehat{p}) = \frac{\ln \Lambda_0 + (t-1)\ln \phi^h}{(t-1)\ln \phi^h - (t-1)\ln \phi^c}$$

where Λ_0 depends on beliefs at the beginning of the herd, μ_0. If Δ_t crosses above the threshold, the herd breaks, whereas if Δ_t remains below the threshold, the herd persists in period t.[8]

[8] When $t = 1$, the herd persists by definition. Recall that the definition of a herd beginning in period 0 is that the subsequent agent (i.e. agent 1) chooses the same action regardless of his signal.

The herd breaking threshold depends on the relative rate of information accumulation from supporting and contrary actions. This threshold rises with an increase in the relative informativeness of supporting actions, and falls when the relative informativeness of contrary increases. If the sample path Δ_t lies above the herd breaking threshold, then the herd will break, whereas if Δ_t lies below, then the herd will persist.

To examine how the behavior of the herd evolves across periods, we need to also characterize the limit behavior of the herd breaking threshold. Lemma 2 shows that the herd breaking threshold monotonically converges to a finite limit.

Lemma 2. *The herd breaking threshold $\Delta_t^*(\widehat{p})$ monotonically converges to a finite limit, which can be represented as:*

$$\Delta^*(\widehat{p}) = \frac{\ln \phi^h}{\ln \phi^h - \ln \phi^c}$$

The actual share of contrary actions in a herd, Δ_t, converges a.s. to its expected value, which is finite and depends on the state, by the strong Law of Large Numbers. Let $\Delta_\infty = \lim_{t\to\infty} \Delta_t$ represent this limit. Comparing the expected share of contrary actions to the limit of the herd breaking threshold allows us to determine whether the herd breaking threshold is crossed with probability one.

Recall that during a herd, the share of contrary actions lies below the herd breaking threshold. If the limit share of contrary actions lies above the limit of the herd breaking threshold, and thus lies in the region where a herd is broken, then almost surely every sample path $\{\Delta_t\}_{t=0}^{\infty}$ crosses the herd breaking threshold at some point as it converges to its limit. When this is the case, a herd breaks with probability 1.

On the other hand, if the expected share of contrary actions lies below the limit of the herd breaking threshold, then the information accumulating *on average* from supporting actions outweighs the information accumulating *on average* from contrary actions and a herd persists in the limit with positive probability. This result is due to the Law of the Iterated Logarithm Sheu [1974], which bounds the rate at which the sequence $\{\Delta_t\}_{t=0}^{\infty}$ converges to its expected value. The probability that $\{\Delta_t\}_{t=0}^{\infty}$ crosses outside this bound infinitely often is zero. This is used to show that there exists a set of sample paths of positive measure such that the actual share of contrary actions never crosses the herd breaking threshold as it converges to its limit, and on such sample paths the herd will never break.

Theorem 2 outlines the conditions under which a herd can persist with positive probability in the long run, and the conditions under which a herd almost always breaks.

Theorem 2. *Given state ω, suppose a herd has formed on action a. Let Δ_∞ represent the limit of the share of contrary actions.*

(i) If beliefs $\widehat{p}$ are such that Δ_∞ lies below the limit of the herd breaking threshold,

$$\Delta_\infty < \Delta^*(\widehat{p})$$

Figure 1: The Herd Breaking Threshold

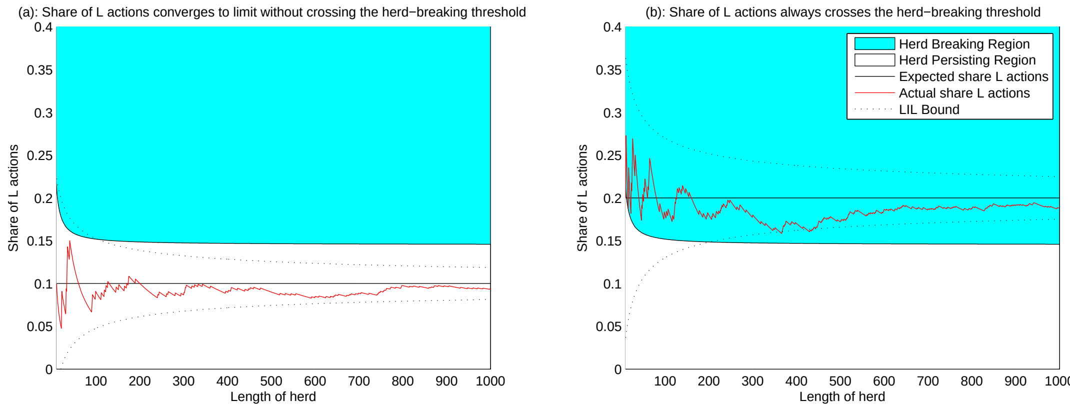

where $\Delta^*(\widehat{p})$ *is as defined above, then there exists a set of sample paths* $\widetilde{\mathbf{\Delta}}$ *that occur with positive probability such that the herd breaking threshold is never crossed: for* $\{\Delta_t\}_{t=0}^{\infty} \in \widetilde{\mathbf{\Delta}}$*,* $\Delta_t < \Delta_t^*(\widehat{p})\ \forall t$*. On such sample paths, the herd is never broken.*

(ii) If beliefs $\widehat{p}$ *are such that* Δ_∞ *lies above the limit of the herd breaking threshold,*

$$\Delta_\infty > \Delta^*(\widehat{p})$$

then for almost every sample path $\{\Delta_t\}_{t=0}^{\infty}$*, there exists a period* τ *such that the herd breaking threshold is crossed at* τ *and the herd breaks:* $\Delta_\tau > \Delta_\tau^*(\widehat{p})$*. Thus, the herd is broken with probability* 1*.*

This result is illustrated in Figure 1 for an R-herd, using arbitrary parameter values. The conditions outlining when a herd can persist and when a herd breaks with probability one will be used in the next section to determine how BIP affects the efficiency of herding.

2.2.3 Efficiency of Herding

The optimal action choice is the action that matches the state. Herding is *efficient* when informed agents choose the optimal action for all but finitely many periods. Thus far, the analysis has proceeded without specifying whether the herd is on the optimal action. In order to determine whether herding is efficient, it is now necessary to consider incorrect and correct herds separately.

Suppose agents are herding on action a. If the herd is correct, then the limit of the share of contrary actions is $\Delta_\infty^{\omega=a} = (1-p)(1-\pi^s)$ and if the herd is incorrect, the limit of the share of contrary actions is $\Delta_\infty^{\omega\neq a} = (1-p)\pi^s$. These limits combined with Theorem 2 can be used to determine which herds persist. The expected share of contrary actions is higher when the herd is incorrect, so if an incorrect herd persists with positive probability then so does a correct herd, and if a correct herd almost always breaks then so does an incorrect herd.[9] It is precisely when incorrect herds break but correct herds can persist that herding will be efficient.

The position of the herd breaking threshold depends on beliefs about the share of informed agents, $\widehat{p}$. An increase in $\widehat{p}$ means less information is accumulating from supporting actions. Beliefs that the herd is correct do not strengthen as quickly and a lower share of contrary actions is required to overturn the herd. Therefore, the herd breaking threshold shifts down as $\widehat{p}$ increases.

Define $\widehat{p}_1$ as the cutoff point such that when $\widehat{p} > \widehat{p}_1$, the expected share of contrary actions in an incorrect herd lies in the herd breaking region, causing incorrect herds to break with probability 1, and when $\widehat{p} < \widehat{p}_1$, the expected share of contrary actions in an incorrect herd lies below the herd breaking region, allowing incorrect herds persist with positive probability. At $\widehat{p}_1$, $\Delta_\infty^{\omega\neq a}$ lies on the limit of the herd breaking threshold, so $\widehat{p}_1$ solves:

$$(1-p)\pi^s = \Delta^*(\widehat{p}_1)$$

Likewise, let $\widehat{p}_2$ be the cutoff point such that when $\widehat{p} > \widehat{p}_2$, the expected share of contrary actions in a correct herd list in the herd breaking region, breaking a correct herd with probability 1, and when $\widehat{p} < \widehat{p}_2$, the expected share of contrary actions lies below the herd breaking region, allowing correct herds persist with positive probability. Similarly, $\Delta_\infty^{\omega=a}$ lies on the limit of the herd breaking threshold at $\widehat{p}_2$, so $\widehat{p}_2$ solves:

$$(1-p)(1-\pi^s) = \Delta^*(\widehat{p}_2)$$

Since $(1-p)(1-\pi^s) < (1-p)\pi^s$ for $p \in [0,1)$ and the limit of the herd breaking threshold decreases with $\widehat{p}$, the cutoff point for correct herds to break is higher than the cutoff point for incorrect herds to break ($\widehat{p}_2 > \widehat{p}_1$). Theorem 3 uses these cutoff points to characterize the efficiency of herding for any beliefs $\widehat{p}$.

Theorem 3. *Let $\widehat{p}_1$ and $\widehat{p}_2$ represent the cutoff points such that incorrect herds break with probability 1 when $\widehat{p} > \widehat{p}_1$ and correct herds break with probability 1 when $\widehat{p} > \widehat{p}_2$. Consider beliefs $\widehat{p} \in [0,1)$*[10]*:*

(i) If $\widehat{p} < \widehat{p}_1$ then incorrect and correct herds both persist with positive probability. There is positive probability that informed agents' action choices converge on the suboptimal action and herding is inefficient.

[9] Suppose $\Delta_\infty^{\omega=a} < \Delta_\infty^{\omega\neq a}$. Then $\Delta_\infty^{\omega\neq a} < \Delta_\infty^* \Rightarrow \Delta_\infty^{\omega=a} < \Delta_\infty^*$ and if $\Delta_\infty^{\omega=a} > \Delta_\infty^* \Rightarrow \Delta_\infty^{\omega\neq a} > \Delta_\infty^*$

[10] In the case where $\widehat{p} = 1$ and $p < 1$, observing contrary actions would be inconsistent with informed agents' beliefs. Therefore, interpreting this case necessitates assumptions on how informed agents interpret contrary actions. For example, if agents attributed contrary actions to a crazy type, then they would ignore these actions. A herd would always persist since no new information accumulates, but learning would be incomplete.

(ii) If $\widehat{p} \in (\widehat{p}_1, \widehat{p}_2)$ then incorrect herds break but correct herds persist with positive probability. A correct herd forms and persists in the limit with probability 1, so informed agents' action choices converge to the optimal action with probability 1 and herding is efficient.

(iii) If $\widehat{p} > \widehat{p}_2$ then incorrect herds and correct herds both break with probability 1. No herd persists in the limit, so informed agents choose the suboptimal action infinitely often and action choices are inefficient.

(iv) Suppose $p < 1$. Then $p \in (\widehat{p}_1, \widehat{p}_2)$ and herding is efficient when beliefs are correct ($\widehat{p} = p$).

Part (iv) of Theorem 3 demonstrates that when beliefs about the share of informed agents are correct ($\widehat{p} = p$), then herding is efficient. Informed agents attribute the correct share of supporting actions to herding, and the correct share to new information from uninformed agents. In an incorrect herd, agents don't overestimate the informativeness of supporting actions, and the public likelihood ratio doesn't become too extreme. This allows the incorrect herd to break. During a correct herd, agents attribute enough supporting actions to private signals and the correct herd is not oversensitive to contrary actions, allowing a correct herd to form and persist in the long run. This result is a special case of Theorem 4 of Acemoglu et al. [2010], which establishes complete learning for the network topology in which all agents are rational, and some agents only observe their own signal.

In fact, as long as $\widehat{p}$ is approximately correct, herding will be efficient. Although agents slightly overestimate or underestimate the informativeness of herding actions, when $\widehat{p}$ lies in the window of efficient herding then this bias is not significant enough to outweigh the accurate information accumulating from uninformed agents. An incorrect herd may persist for longer or a correct herd may break more often than would have been the case if beliefs were correct, but ultimately a correct herd will form and persist. The herding model is robust to perturbations of beliefs about the share of informed agents, and efficient information aggregation is still achieved in the long run.

When $\widehat{p}$ is too extreme in either direction, then efficient herding will no longer obtain. If agents significantly underestimate the share of informed agents (part (i) of Theorem 3), they overestimate the informativeness of supporting actions. The repeated information from actions of agents who herded accumulates at a faster rate than the new information from uninformed agents. If an incorrect herd forms, then it may persist in the long run and agents will perpetually choose the suboptimal action.

On the other hand, if agents significantly overestimate the share of agents who are informed (part (iii) of Theorem 3), they attribute new information from uninformed agents to herding. Too little information accumulates from supporting actions, preventing the herd from persisting. Both incorrect and correct herds form with positive probability, so when a herd breaks, the next herd to form may be correct or incorrect. Beliefs about the true state remain fragile, and agents oscillate between correct and incorrect herds. Both types of cascades occur

Figure 2: When Can a Herd Persist?

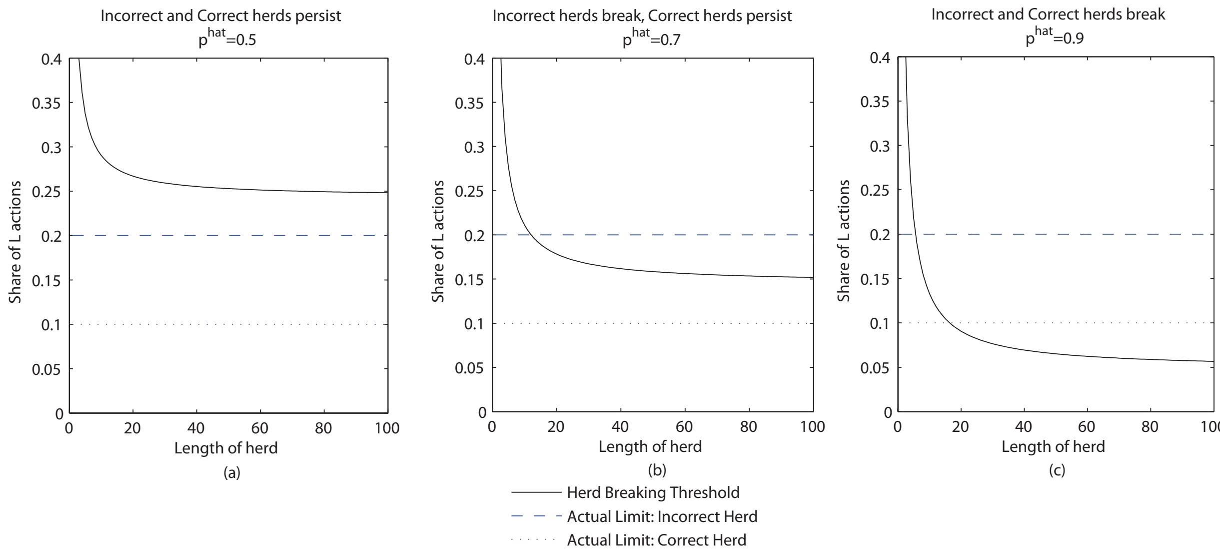

infinitely often, which results in agents choosing both the optimal and suboptimal actions infinitely often. Figure 2 illustrates the areas corresponding to parts (i) - (iii) of Theorem 3 for an R-herd.

The robustness of efficient herding depends on the actual share of informed agents. When p is low, many agents reveal their private signal so accurate information accumulates at a faster rate. Misinterpreting the informativeness of supporting actions has a small impact on the public likelihood ratio, and the interval of efficient herding is large. On the other hand, when the actual share of informed agents is large, most agents are herding so new information accumulates slowly. The public likelihood ratio is very sensitive to any information, and inaccurate information can have a significant impact. Thus, efficient herding is robust to larger perturbations over beliefs when p is small. Figure 3 shows how the regions of beliefs over agent types depend on the share of informed agents.

Whenever a herd persists in the long run, public beliefs about ω will converge to a point mass on the state matching the cascade action. Learning is *complete* if public beliefs converge to a point mass on the true state, whereas learning is *fully incorrect* if public beliefs converge to a point mass on the incorrect state. If public beliefs about the state remain interior ($\mu_t \in (0,1)$), then learning is *incomplete*. Corollary 1 demonstrates that learning is complete when a correct herd persists in the limit, but learning is fully incorrect when an incorrect herd persists in the limit. If no herd persists, then the public likelihood ratio does not converge and learning is incomplete.

Corollary 1. *Learning is as follows:*

(i) If a correct herd persists, then learning is complete.

Figure 3: Learning is complete when beliefs are approximately correct

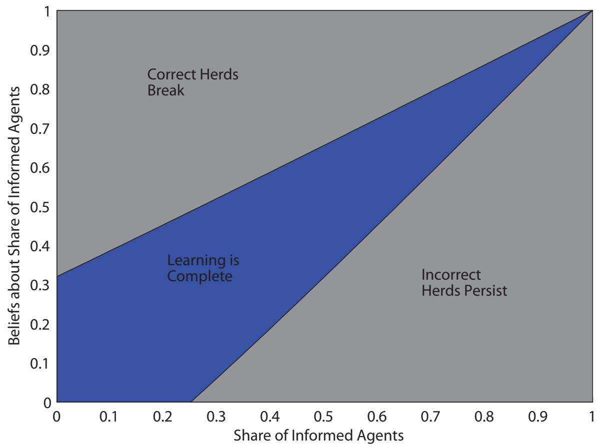

(ii) If an incorrect herd persists, then learning is fully incorrect.

(iii) If no herd persists, then the public likelihood ratio perpetually oscillates, and learning is incomplete.

Usually, the conditional public likelihood ratio is a martingale and therefore converges, preventing fully incorrect learning or perpetually fluctuating beliefs. However, this result hinges on correct beliefs about the share of informed agents. When $\widehat{p} \neq p$, the conditional public likelihood ratio is no longer a martingale. Suppose $\omega = R$ and $\widehat{p} < p$. In an incorrect L-herd, Λ_t^R is a submartingale because agents overweight L actions, and L actions increase the likelihood ratio. Therefore, when the L-herd persists, Λ_t^R diverges to infinity and learning is fully incorrect. In contrast, in a correct R-herd, Λ_t^R is a supermartingale because agents overweight R actions, which decrease the likelihood ratio. Non-negative supermartingales converge, so Λ_t^R does converge in an R-herd. In fact, when an R-herd persists, Λ_t^R converges to 0 and learning is complete.

Now suppose that $\widehat{p} > p$. In an incorrect L-herd, Λ_t^R is now a supermartingale because agents underweight L actions relative to R actions, so Λ_t^R decreases in expectation. Because $\Lambda_t^R > 1$ in an L-herd, eventually Λ_t^R crosses below 1 and the herd breaks. The opposite happens in a correct R-herd. Agents underweight R actions, so Λ_t^R is a submartingale and increases in expectation. Provided beliefs are far enough away from the truth, the Λ_t^R eventually crosses above 1, breaking the herd. Thus, both types of herds break before the conditional public likelihood ratio can converge or diverge. Once a herd breaks, Λ_t^R oscillates until another herd forms, at which point the process repeats.

2.3 Numerical Example

The following example illustrates the potential for BIP to confound learning. Consider a model where nature selects state L with probability $\pi^L = 0.51$. With probability $p = 0.7$, agents are "socially informed", and with probability $1 - p = 0.3$, agents are "socially uninformed" and do not observe the history. Both types of agents observe a private binary signal which reveals the true state with probability $\pi^s = 2/3$. "Socially informed" agents hold a common belief $\widehat{p}$ that previous agents are "socially informed".

Suppose that an L herd has formed. In the case of an incorrect herd, the probability of a contrary R action is equal to the probability that the agent is uninformed times the probability that this agent observes a correct signal: $\Delta_\infty^{\omega=R} = 0.2$. If the herd is correct, the probability of an R action is equal to the probability that the agent is uninformed times the probability that this agent observes an incorrect signal: $\Delta_\infty^{\omega=L} = 0.2$

We can use these parameters and Theorem 3 to characterize the efficiency of herding. If $\widehat{p} < 0.59$ then incorrect and correct herds both persist with positive probability. Learning may or may not be complete, depending on whether an incorrect or correct herd persists. If $\widehat{p} \in (0.59, 0.79)$ then incorrect herds break, but correct herds persist with positive probability. Eventually a correct herd will form and persist, leading to complete learning. If $\widehat{p} > 0.79$ then beliefs about the true state are too fragile - both incorrect and correct herds break, so learning is incomplete. Note that correct beliefs about the share of informed agents ($\widehat{p} = 0.7$) fall in the interval that leads to complete learning, as established in Theorem 3.

2.4 Comparative Statics

The precision of the private signal and the probability that agents are informed affects the relative positions of the herd breaking threshold and the expected share of contrary actions. Therefore, a change in any of these parameters will affect whether an incorrect herd breaks or a correct herd persists.

An increase in the probability that agents are informed, p, reduces the frequency of contrary actions since more agents observe the history and follow the herd. This can move the limit of the sample path for a correct or incorrect herd into the herd persisting region. In the former case, the increase is beneficial because it allows correct herds to persist, while in the latter case, the increase introduces inefficiency by allowing incorrect herds to persist.

An increase in the precision of the private signal, π^s, has an ambiguous effect. This change affects information accumulation through two channels: the information accumulating from each individual action, and the frequency of each type of action. More informative contrary actions decreases the herd breaking threshold, which makes it more likely that both types of herds break. The frequency of contrary actions decreases in a correct herd, and increases in an incorrect herd. Incorrect herds are less likely to persist as there are more informative and more frequent contrary actions. The overall impact on correct herds is ambiguous, as there are fewer contrary actions but each of these actions have a larger impact on beleifs.

This comparative static presents an interesting insight: more precise information may not always improve welfare. If more precise information increases the likelihood that a correct herd breaks, then herding is more likely to be inefficient. However, more precise information may also increase the probability that a correct herd forms in the first place, which would increase the efficiency of herding. The tradeoff between the efficiency gains and losses from more precise information leaves open an interesting question for future research.

3 Extensions

Thus far, we have established how the efficiency of information cascades depends on the relationship informed agents perceive between prior actions and signals. These results are robust to different modifications of the benchmark model, several of which are discussed informally below.

3.1 Public Signals

Suppose that in addition to learning from their own private information and the actions of others, informed agents also observe a sequence of public signals. We allow a public signal $\sigma_t \in \{l, r\}$ of precision $\pi^\sigma \in (1/2, 1)$ to be released with probability $\varepsilon > 0$ each period, and examine whether inefficient herding can still persist in the presence of this infinite sequence of new information.

Now the public likelihood ratio evolves to incorporate new information from action choices and public signal realizations. An r public signal multiplies the public likelihood ratio by $\phi^\sigma = \left(\frac{1-\pi^\sigma}{\pi^\sigma}\right)$, while an l public signal multiplies the likelihood ratio by $1/\phi^\sigma$. Define the contrary public signal lead $\Delta_t^\sigma \in [-1, 1]$ as the difference between the share of contrary and supporting public signals. In an L-herd, $\Delta_t^\sigma = \frac{1}{t}\sum_{s=1}^{t}\left(1_{\sigma s=r} - 1_{\sigma_s=l}\right)$ represents the difference between the share of r and l public signals.[11] In a herd, Δ_t and Δ_t^σ are sufficient statistics for the history when examining the evolution of the public likelihood ratio.

Theorem 1 is still valid with the addition of public signals, so all cascades break with positive probability.[12] In a similar fashion to Section 2, we characterize when cascades can

[11] The additive structure of information accumulating from public signals is due to equal precision of signals across states. An l signal exactly cancels an r signal, so the difference between the number of r and l signals is a sufficient statistic for the public signal history. In the case of unequal signal precision across states, two variables would be necessary to keep track of the public signal history.

[12] With public signals, it is now possible for a herd to break even when $p = 1$, provided that the public signal is more informative than a supporting action $\left(\pi^\sigma > \frac{\widehat{p}+(1-\widehat{p})\pi^s}{1+\widehat{p}}\right)$. When agents correctly believe all previous agents observed the history ($\widehat{p} = 1$), this condition corresponds to $\pi^\sigma > 1/2$ and a herd can break as long as the public signal isn't pure noise. At the other extreme, if agents believe no preceding agents observe the

Figure 4: When Does a Herd Persist?

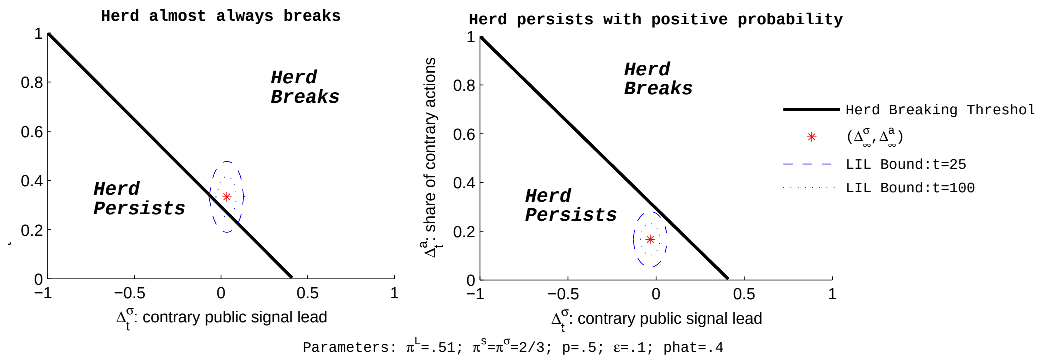

also persist with positive probability. When information accumulates from two sources, the herd breaking threshold at time t is represented as a line in (Δ, Δ^σ) space, such that an agent will follow her private signal when $(\Delta_t, \Delta_t^\sigma)$ lies in the half-plane above this threshold. The slope of the herd breaking threshold is negative and independent of t, capturing the tradeoff between contrary public signals and contrary actions: as the contrary public signal lead increases, fewer contrary actions are necessary to reach the herd breaking threshold.

The actual contrary public signal lead converges to its finite expected value, conditional on the state. Comparing $(\Delta_\infty, \Delta_\infty^\sigma,)$ to the limit of the herd breaking threshold allows us to determine whether the herd breaking threshold is crossed with probability one. If $(\Delta_\infty, \Delta_\infty^\sigma)$ lies in the half-plane above the limit of the herd breaking threshold, and thus lies in the region where a herd is broken, then almost surely every sample path $\{(\Delta_t, \Delta_t^\sigma)\}_{t=0}^\infty$ crosses the herd breaking threshold at some point as it converges to its limit and a herd breaks with probability 1. On the other hand, if the limit $(\Delta_\infty, \Delta_\infty^\sigma)$ lies in the half-plane below the limit of the herd breaking threshold, then a herd persists in the limit with positive probability. This result is due to the Law of the Iterated Logarithm for two-dimensional processes, which bounds the sequence $\{(\Delta_t, \Delta_t^\sigma)\}_{t=0}^\infty$ by a sequence of disks of decreasing radius, centered around the limit $(\Delta_\infty, \Delta_\infty^\sigma)$.[13] Figure 4 illustrates when a herd can persist.

As in the previous section, these conditions can be used to characterize the efficiency of herding by characterizing cut-off points. The results of Theorem 3 extend directly to the case of public signals. Thus, although the addition of public signals may reduce the scope for inefficient herding, it is not eliminated entirely. A formal characterization of the results from this section is available in a Supplementary Appendix.

These results demonstrate that BIP in the context of a sequential learning model with public signals has important implications. If agents overestimate the amount of new information contained in the history, BIP can confound learning by allowing an incorrect herd to

history ($\widehat{p} = 0$), then the public signal needs to be more informative than the private signal, $\pi^\sigma > \pi^s$, for a herd to break with positive probability.

[13] Sheu [1974]

persist even when public signals are released to counteract the incorrect herd. In this scenario, the best way to quash a false herd is to release public signals immediately and frequently, and a rumor may be near impossible to break once it becomes entrenched. This contrasts with the case of no BIP, in which the timing of public signal release is irrelevant - public information that breaks a cascade at time t will also break the herd at time $t+\tau$ for any τ.

3.2 Continuous Signals

In another variation on the model of Section 2, suppose that rather than receiving a binary private signal, agents receive a signal drawn from a continuous support. Smith and Sorensen [2000] show that allowing for unbounded continuous signals eliminates the possibility of incomplete learning. I examine whether this result remains true in the presence of BIP.

Let $s_t = P(\omega = L|\sigma_t)$ represent an agent's private belief that $\omega = L$ after receiving signal σ_t, computed using Bayes rule. Conditional on the state, s_t is i.i.d. with conditional distribution $F^\omega(s)$.and support $(0,1)$, so private signals are unbounded but no signal perfectly reveals the state. There exists a cutoff $s^*(\mu) = 1-\mu$ such that an informed agent chooses L for $s \geq s^*(\mu)$ and R for $s < s^*(\mu)$. Note $s^*(\mu)$ is decreasing in μ - an agent chooses action L for a broader range of private signals when public beliefs are more in favor of state L. An uninformed agent, who only observes her private signal, uses the cutoff $s^* = 1/2$ to determine his action choice, independent of current public beliefs.

An information cascade forms when informed agents choose the same action for every signal in the support of F. With unbounded signals, there is always a signal that will overturn an interior public belief and cascades only occur in the limit. However, with continuous signals, informed and uninformed agents act differently even when no herd is occurring. Repeated actions become less and less informative as beliefs strengthen, because informed agents choose this action for a wider range of private signals.

In order to establish how BIP influences the limiting properties of the public likelihood ratio, I first consider how beliefs about the share of informed agents affects the rate at which information accumulates from actions. When state L is more likely, attributing more actions to uninformed agents dampers the impact of R actions and raises the impact of L actions. An R action from an uninformed type indicates a private signal $s_t < 1/2$ while an R action from an informed type indicates a stronger private signal $s_t < 1-\mu_t < 1/2$. On the other hand, an L action from an uninformed type is more informative than an L action from an informed type as the former indicates a private signal stronger that falls in the interval $[1/2, 1)$, whereas the latter indicates a private signal in the interval $[1-\mu_t, 1)$, which is wider interval. The opposite is true when state R is more likely.

For the remainder of the analysis, consider the case where $\omega = L$ and define the likelihood ratio as $\Lambda_t = \left(\frac{1-\mu_t}{\mu_t}\right)$. Given that the signal space is unbounded, the only stationary limit beliefs about the state are placing probability 1 on either state L or state R. When $\widehat{p} \neq p$, Λ_t is no longer a martingale and convergence may not obtain. The following theorem characterizes which stationary limit points of Λ_t are reached with positive probability, as a function of $\widehat{p}$.

Theorem 4. *Let $\overline{\mathcal{L}}$ represent the set of stationary limit points that Λ_t converges to with positive probability. There exists cutoff points $\widehat{p}_1$ and $\widehat{p}_2$ such that*

1. *If $\widehat{p} < \widehat{p}_1$, then $\overline{\mathcal{L}} = \{0, \infty\}$*
2. *If $\widehat{p} \in (\widehat{p}_1, \widehat{p}_2)$, then $\overline{\mathcal{L}} = \{0\}$*
3. *If $\widehat{p} > \widehat{p}_2$ then $\overline{\mathcal{L}} = \{\emptyset\}$*

   where $\widehat{p}_2 \in (0,1)$ for all p, and $\widehat{p}_1 \in (0,1)$ for $p > \frac{1-2F^L(1/2)}{2(1-F^L(1/2))}$

First consider the scenario when agents believe actions reveal more private information than is actually the case. When beliefs favor L over R ($\Lambda < 1$), the conditional likelihood ratio decreases in expectation and converges to 0 with positive probability 0, so complete learning is possible. In the case where beliefs favor R over L ($\Lambda > 1$), the conditional likelihood ratio increases in expectation. When $\widehat{p}$ is far enough away from the truth, the conditional likelihood ratio also converges to infinity with positive probability, and fully incorrect learning is also possible.

When agents believe actions contain less private information than is actually the case, the results are flipped. If beliefs favor L over R, then the conditional likelihood ratio increases in expectation, and if beliefs favor R over L, the conditional likelihood ratio decreases in expectation. When $\widehat{p}$ is far enough away from the truth, the conditional likelihood ratio converges to infinity when it is less than 1, and converges to 0 when it is greater than 1. Therefore, neither fixed point is stable and the likelihood ratio perpetually fluctuates. Provided agents' beliefs about the information content of actions are approximately correct, 0 is the only stable fixed point of the likelihood ratio and learning is correct. These results are presented formally in the Supplemental Appendix.

3.3 Private Values Types

Suppose there are two private value types, θ_L and θ_R, who choose actions L and R, respectively, regardless of the history, and let both types occur with positive probability. The result: less information accumulates from both supporting and contrary actions, but the conclusions of Section 2 are still valid. In fact, allowing for uncertainty over the share of private value types may lead to similar conclusions as information-processing bias: if agents underestimate the share of private value types, they will overestimate the informational content of actions; and if agents overestimate the share of private value types, they will underestimate the informational content of actions.

4 Learning About BIP

In the previous section, agents begin with exogenous and possibly incorrect beliefs about the information processing capabilities of other agents. Informed agents use these beliefs and the history to update their beliefs about the state of the world. This section will use a simple example to examine what happens when agents can also learn about the information processing capabilities of other agents.

Suppose p is a random variable distributed according to a common prior. Upon observing the history, informed agents use Bayes rule to update their beliefs about p. Let $g(p, \omega)$ represent the common prior beliefs held by informed agents (after observing their own type). For this example, suppose there are two possible shares of informed agents, $p \in \{0.4, 0.8\}$, and (p, ω) are independent of each other with marginal distributions $P(\omega = L) = \pi^L$ and $P(p = 0.4) = \pi^p$. Let the precision of the private signal be $\pi^s = 0.75$.

I will examine whether incomplete learning is possible in the case where $\omega = L$ and $p = 0.8$. Now informed agents will use the history to learn about both p and ω. Let $g(p, \omega | h_t)$ represent an agent's joint probability distribution over p and ω after observing history h_t. Then the conditional likelihood ratio between $(0.4, L)$ and any other pair (ω, p) is a martingale, represented as:

$$\Lambda_t(p, \omega) = \frac{g(p, \omega | \Delta_t^R, \Delta_t^L)}{g(.8, L | \Delta_t^R, \Delta_t^L)} = \Lambda_{t-1}(p, R) \left(\frac{P(a_t | p, \omega)}{P(a_t | .4, L)} \right)$$

where Δ_t^R is the share of contrary R actions observed in L-herds and Δ_t^L is the share of contrary L actions observed in R-herds as of time t. The probability of a given action depends on which type of herd is occurring.

By the Martingale Representation Theorem, this likelihood ratio converges. Whether complete learning obtains depends on several factors, including whether there are pairs (p, ω) that are indistinguishable from each other, the relative weight that the prior places on these indistinguishable pairs, and the duration of previous herds on the opposite action.

If a herd persists, then there are two realizations of (p, ω) that may be indistinguishable. Consider an R-herd: if it persists, the share of L actions converges to its expected value, $\Delta_\infty^L(0.8, L) = 0.15$ (recall $\Delta_\infty^L = (1-p)\pi^s$ in an incorrect herd). There are two possible pairs (p, ω) that would give rise to this share of L actions, as $\Delta_\infty^L(0.4, R)$ is also equal to 0.15 (recall $\Delta_\infty^L = (1-p)(1-\pi^s)$ if the state is R). It is impossible to distinguish between $(0.8, L)$ and $(0.4, R)$ in an R-herd since both of these pairs result in the same expected share of contrary actions. Therefore $\frac{P(a_t | 0.4, R)}{P(a_t | .8, L)} = 1$ in an R-herd and no information is gained about the relative likelihood of these two events. The probability of all other pairs (p, ω) converges to 0 when the R-herd persists, since $\left(\frac{P(a_t | p, \omega)}{P(a_t | .4, L)} \right) \neq 1$ and zero is the only finite stationary point of $\Lambda_t(p, \omega)$ for such pairs.

Next consider information from previous L-herds. The share of contrary R actions in previous L-herds also yields information about p, and these contrary R actions help distinguish between $(0.8, L)$ and $(0.4, R)$. The only pair that is indistinguishable from $(0.8, L)$ in an L-herd is $(0.93, R)$, which differs from the pair that is indistinguishable in an R-herd (and in this example, is not in the support of (p, ω)). Thus, the true value of p would be identified if both expected shares are observed. Even a finite number of observations from a previous L-herd helps distinguish between $(0.8, L)$ and $(0.4, R)$. Suppose previous L-herds occurred for τ_L periods and yielded a share Δ^R contrary actions. Then the information gleaned from these L-herds multiplies the relative likelihood of $(0.8, L)$ and $(0.4, R)$ by

$$\phi^{\tau_L} = \left[\left(\frac{P(a = L|0.4, R)}{P(a = L|0.8, L)} \right)^{(1-\Delta^R)} \left(\frac{P(a = R|0.4, R)}{P(a = R|0.8, L)} \right)^{\Delta^R} \right]^{\tau_L}$$

When Δ^R is close to its expected value, $\Delta^R_\infty = .05$, this expression is less than 1, and therefore increases the relative likelihood of $(0.8, L)$ compared to $(0.4, R)$

We can now characterize the limit of $\Lambda_t(0.4, R)$ in an R-herd:

$$\Lambda_\infty(0.4, R) = 3 \left(\frac{\pi^p}{1 - \pi^p} \right) \left(\frac{1 - \pi^L}{\pi^L} \right) \phi^{\tau_L}$$

If $\Lambda_\infty(0.4, R) > 3$ then an R-herd can persist with positive probability. In the limit, agents believe that state R is more likely even when they receive a private l signal. Because Λ_t is a martingale, fully incorrect learning is not possible as in the previous section.[14] However, learning is incomplete and beliefs remain interior at $\Lambda_\infty(0.4, R)$. If $\Lambda_\infty(0.4, R) < 3$ then the R-herd breaks with probability 1 as agents will eventually come to believe state L more likely and end the herd. Note that the duration of previous L-herds affects whether a given herd can persist, as longer previous L-herds decrease $\Lambda_\infty(0.4, R)$.

In this simple example, there are no indistinguishable pairs in L-herds. Thus, an L-herd always persists with positive probability, and learning will be complete when this occurs. However, a more general prior over p allows the possibility of indistinguishable pairs in both correct and incorrect herds, precluding complete learning for both cases. This section illustrates that, when agents can learn about others information processing capabilities, the scope for inefficient herding is reduced and the possibility of fully incorrect learning is generally eliminated.

[14] Fully incorrect learning about ω is not generally possible when agents can learn about p. This will only occur if agents have a common posterior that puts no weight on the correct value of p.

5 Discussion and Conclusion

This paper demonstrates that a bias about how others' process information can significantly affect the efficiency of learning. Particularly, it is possible for agents to continue to choose the suboptimal action despite the release of new information contradicting the herd. In the benchmark model, this would be impossible. Inefficient herding occurs because information ceases to aggregate; when even the smallest amount of information continues to accumulate, inefficient herding no longer occurs. Experimental evidence from Goeree et al. [2007] suggests that new information does indeed continue to accumulate in a herd: regardless of how many previous agents chose the same action, some agents still follow their private signal. In the benchmark model, this off-the-equilibrium-path action would be ignored since it is not rational. However, it seems plausible that subsequent agents would recognize these off-the-equilibrium-path actions are likely to reveal an agent's private signal, and therefore contain information. Thus, BIP allows new information to continue to enter the model, and provides an explanation for inefficient herding when this is the case. Inefficient herding occurs because the rate of information accumulation from repeated information outweighs the rate of information accumulation from new information, and these herds can persist even when contradicted by public information. Additionally, it explains how convergence may fail to obtain even when public information is released that supports the correct herd.

Experimental evidence from Koessler et al. [2008] supports the possibility of BIP in an observational learning model. They examine the fragility of cascades in a model where one agent receives a more precise signal than others. The unique Nash equilibrium of such a model is for the high informed agent to follow her signal. Thus, receipt of a contrary signal overturns a cascade. Koessler et al. [2008] find that highly informed agents rarely overturn a cascade when equilibrium prescribes that they do so. As the length of the cascade increases, highly informed agents become even less likely to follow their signal: highly informed participants break 65% of cascades when there are two identical actions, but only 15% of cascades when there are 5 or more identical actions. This phenomenon is likely explained by the evolution of participants' beliefs. The evolution of elicited beliefs is similar to the belief process that would arise if all agents followed their signal, and thus conveyed their private information. In addition, Koessler et al. [2008] find that off-the-equilibrium-path play frequently occurs, and these non-equilibrium actions are informative, providing support for the actual presence of some uninformed agents, in addition to a strong belief about their presence.

Kubler and Weizsacker [2004] also find evidence consistent with BIP. 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 therefore overestimate the amount of new information contained in previous actions.

Another interesting consequence of BIP is that agents may actually be worse off if more information accumulates than was expected. As p decreases, more private information accumulates since fewer agents observe the history. However, if p is far enough below $\widehat{p}$, correct herds will become too fragile and herding will be inefficient.

Conformity preferences is another bias that could make agents more likely to herd as the length of the herd increases, but equilibrium play in such a model differs significantly from a model with BIP and public signals. With BIP, if the contrary public signal lead is high enough to break a herd of a given length, then subsequent agents do not continue to herd. This model has a unique equilibrium where agents choose whichever state that they believe is more likely based on the history and the degree of BIP. However, with conformity preferences, if the preference to conform is large enough, then it is irrelevant whether or not the contrary public signal lead is high enough to make agents believe the alternative state more likely. Agents simply want to choose the action that the majority of other agents will choose. So there are multiple equilibria where agents choose the state that they believe the majority of other agents will choose, independent of the likelihood that this state is the true state.

This model leaves open interesting questions for future research on information processing capabilities. Individuals may differ in their depth of reasoning and their ability to combine different information sources. Such biases may have important implications for the way information is aggregated. Examining the implications of BIP in a more general model may yield interesting insights into this issue. While the assumption of common beliefs over the informational content of the history is a good starting point, a valid criticism is that this model requires implausible levels of belief coordination. Thus, examining how the model fares with heterogenous beliefs about the information processing capabilities of other agents is another avenue for future research. Allowing partial observability of histories would be natural extensions to generalize the model, while introducing payoff interdependencies would make the model applicable to election and financial market settings.

References

Acemoglu, D., M. Dahleh, I. Lobel, and A. E. Ozdaglar (2010): "Bayesian Learning in Social Networks," *SSRN eLibrary*.

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.

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. H. Zwiebel (2001): "Persuasion Bias, Social Influence, and Uni-Dimensional Opinions," *SSRN eLibrary*.

Eyster, E. and M. Rabin (2005): "Cursed Equilibrium," *Econometrica*, 73, 1623–1672.

——— (2009): "Rational and Naive Herding," .

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 (2009): "Social Leanring with Course Inference," WEF Working Papers 0050, ESRC World Economy and Finance Research Programme, Birkbeck, University of London.

Hartman, P. and A. Wintner (1941): "On the Law of the Iterated Logarithm," *American Journal of Mathematics*, 63, 169–176.

Koessler, F., A. Ziegelmeyer, J. Bracht, and E. Winter (2008): "Fragility of Information Cascades: An Experimental Study using Elicited Beliefs," Jena Economic Research Papers in Economics 2008-094, Friedrich-Schiller-University Jena, Max-Planck-Institute of Economics.

Kubler, D. and G. Weizsacker (2004): "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory," *Review of Economic Studies*, 71, 425–441.

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.

6 Appendix: Proofs

Proof of Theorem 1 on pp. 10

Let $\left(\frac{\mu_0}{1-\mu_0}\right)$ represent public beliefs before the action that begins the herd.

Suppose $p < 1$. Let there be an L-herd in period t with $y = (t-1)(1-2\Delta_t)$ net L actions (i.e. the number of L actions minus the number of R actions). Each R action decreases the likelihood ratio by a factor ϕ^c and each L action increases the likelihood ratio by a factor ϕ^h. Observe $\phi^h * \phi^c \leq 1$, so the net effect of an R action and an L action decreases the likelihood ratio. Then $y > 0$ R actions will outweigh the y net L actions. Let κ be the greatest k such that $\left(\frac{\mu_0}{1-\mu_0}\right)\left(\frac{\pi^s}{1-\pi^s}\right)(\phi^c)^k < 1$. Then κ R actions outweigh initial public beliefs. Note κ is finite since $\left(\frac{\mu_0}{1-\mu_0}\right) < \infty$. Then $y+\kappa$ R actions will break this herd. The probability that the next $y+\kappa$ actions are R is:

$$
\begin{aligned}
[(1-p)(1-\pi^s)]^{y+\kappa} &> 0 \text{ if the herd is correct} \\
[(1-p)\pi^s]^{y+\kappa} &> 0 \text{ if the herd is incorrect}
\end{aligned}
$$

This is a lower bound on the probability that the herd breaks. Similar analysis yields the same results for R-herds. Q.E.D.

Proof of Lemma 1 on pp. 10

Let μ_0 represent public beliefs before the action that begins the herd and suppose agents are herding on action a^h. Let $\Lambda_t = \frac{P(\omega=a^h|h_t)}{P(\omega\neq a^h|h_t)}$ be the public likelihood ratio of the probability that the herd is correct to the probability that the herd is incorrect, let $\Lambda_t^{s^c}$ be the updated private likelihood ratio after a contrary private signal, and let $\Lambda_t^{s^h}$ be the updated private likelihood ratio after a supporting private signal. In an L-herd, $\Lambda_0 = \left(\frac{\mu_0}{1-\mu_0}\right)$ while in an R-herd, $\Lambda_0 = \left(\frac{1-\mu_0}{\mu_0}\right)$. Note $\Lambda_0 > 1$ by definition. Provided it is still optimal for agents who observe the history to herd, the likelihood ratio in a herd evolves as follows:

$$
\begin{aligned}
\Lambda_t &= \Lambda_0\left(\frac{\pi^s}{1-\pi^s}\right)\left(\phi^h\right)^{(t-1)(1-\Delta_t)}\left(\phi^c\right)^{(t-1)\Delta_t} \\
&= \Lambda_0\left(\phi^h\right)^{(t-1)(1-\Delta_t)}\left(\phi^c\right)^{(t-1)\Delta_t-1} \\
s_t = s^c \Rightarrow \Lambda_t^{s^c} &= \Lambda_0\left(\phi^h\right)^{(t-1)(1-\Delta_t)}\left(\phi^c\right)^{(t-1)\Delta_t} \\
s_t = s^h \Rightarrow \Lambda_t^{s^h} &= \Lambda_0\left(\phi^h\right)^{(t-1)(1-\Delta_t)}\left(\phi^c\right)^{(t-1)\Delta_t-2}
\end{aligned}
$$

The herd breaking threshold Δ_t^* is calculated by finding the value of Δ_t that satisfies $\Lambda_t^{s^c} = 1$, so the private likelihood is equal to one when a contrary private signal is realized:

$$
\begin{aligned}
& \Lambda_t^{s^c} = 1 \\
\Rightarrow\ & \Lambda_0 \left(\phi^h\right)^{(t-1)(1-\Delta_t^*)} \left(\phi^c\right)^{(t-1)\Delta_t^*} = 1 \\
\Rightarrow\ & \ln \Lambda_0 + (t-1)(1-\Delta_t^*) \ln \phi^h + (t-1)\Delta_t^* \ln \phi^c = 0 \\
\Rightarrow\ & \Delta_t^* = \frac{\ln \Lambda_0 + (t-1)\ln \phi^h}{(t-1)\ln \phi^h - (t-1)\ln \phi^c}
\end{aligned}
$$

Q.E.D.

Proof of Lemma 2 on pp. 11

$$
\Delta_\infty^* = \lim_{t\to\infty} \Delta_t^* = \lim_{t\to\infty} \frac{\ln \Lambda_0 + (t-1)\ln \phi^h}{(t-1)\ln \phi^h - (t-1)\ln \phi^c} = \frac{\ln \phi^h}{\ln \phi^h - \ln \phi^c}
$$

Also note that Δ_t^* monotonically decreases to its limit, a fact which will be used in the proof of Theorem 2.

$$
\begin{aligned}
\frac{d}{dt}\Delta_t^* &= \frac{d}{dt}\left[\left(\frac{1}{t-1}\right) * \frac{\ln \Lambda_0 + (t-1)\ln \phi^h}{\ln \phi^h - \ln \phi^c}\right] \\
&= \left(\frac{1}{t-1}\right)^2 * \frac{-\ln \Lambda_0}{\ln \phi^h - \ln \phi^c} < 0
\end{aligned}
$$

Q.E.D.

Proof of Theorem 2 on pp. 11

(a) Let $\Delta_\infty < \Delta^*$.

Define a numeric representation of the action choice space as an infinite sequence of spaces $\{Y_i\}, i = 1, 2, ...$ over which a probability measure is defined, where $Y_i = \{0, 1\}\ \forall i$. Let $Y = Y_1 \times Y_2 \times ... \times Y_i \times ...$ be the product of such spaces. Let $\{y_i\}$ be an infinite sequence of mutually independent and identically distributed random variables drawn from this space, with distribution function $g(y)$ defined as follows, depending on whether the herd is correct or incorrect:

Correct Herd	Incorrect Herd
$g(0) = p + (1-p)\pi^s$	$g(0) = p + (1-p)(1-\pi^s)$
$g(1) = (1-p)(1-\pi^s)$	$g(1) = (1-p)\pi^s$

A supporting action corresponds to $y_i = 0$, and a contrary action corresponds to $y_i = 1$. Let μ_y represent the expected value of y_i and σ_y^2 represent the variance of y_i.The expected value and variance of y_i are as follows:

Correct Herd	Incorrect Herd
$\mu_y = (1-p)(1-\pi^s)$	$\mu_y = (1-p)\pi^s$
$\sigma_y^2 = (1-p)(1-\pi^s) - (1-p)^2(1-\pi^s)^2$	$\sigma_y^2 = (1-p)\pi^s - (1-p)^2(\pi^s)^2$

Note the equivalence of $\sum_{i=1}^t y_i$ and $t\Delta_{t+1}$, and therefore the equivalence of μ_y and $\Delta_\infty = E[\Delta_t]$. Transform y_i as follows:

$$\overline{y}_i = y_i - \mu_y$$

and calculate $E[\overline{y}_i] = 0$ and $Var(\overline{y}_i) = \sigma_y^2$ and $\sum_{i=1}^t Var(\overline{y}_i) = t\sigma_y^2$. We have $|\overline{y}_i|$ bounded above by 1, which is independent of t. So trivially,

$$\sup_{i \le t} \text{l.u.b. } |\overline{y}_i| = o\left(\frac{t\sigma_y^2}{\log\log t\sigma_y^2}\right)^{\frac{1}{2}}$$

for each t as $t \to \infty$. The necessary assumptions for the law of the iterated logarithm (LIL) applied to a one-dimensional independent random variable are satisfied.[15] The LIL can be used to bound $\sum_{i=1}^t \overline{y}_i$:

$$\lim\sup_{t\to\infty} \frac{\sum_{i=1}^t \overline{y}_i}{\sqrt{2t\sigma_y^2 \log\log t\sigma_y^2}} = 1 \text{ a.s.}$$

Thus, for $\delta > 0$.

$$P\left[\sum_{i=1}^t \overline{y}_i \ge (1+\delta)\sqrt{2t\sigma_y^2 \log\log t\sigma_y^2} \text{ i.o.}\right] = 0$$

Define

$$B_t = (1+\delta)\sqrt{\frac{2\sigma_y^2 \log\log t\sigma_y^2}{t}}$$

This means that for almost all realizations of $\{y_i\}$, there exist only *finitely many* t such that $\frac{1}{t}\sum_{i=1}^t y_i$ lies outside $\Theta_t = [\mu_y + B_t, \mu_y - B_t]$ Define

$$\zeta = \left\{ \{\widehat{y}_i\} \mid \frac{1}{t}\sum_{i=1}^t \widehat{y}_i > \mu_y + B_t \text{ for some } t \right\}$$

[15] The necessary assumptions are:

(i) $\overline{y}_1, \overline{y}_2, ...$ is an an infinite sequence of real-valued independent random variables of class L^2

(ii) $E[\overline{y}_i] = 0$

(iii) $\sum_{i=1}^n Var(\overline{y}_i) \to \infty$ as $n \to \infty$

(v) $\sup_{i\le t}(\text{l.u.b. } |\overline{y}_i||) = o\left(\frac{\sum_{i=1}^n Var(\overline{y}_i)}{\log\log \sum_{i=1}^n Var(\overline{y}_i)}\right)^{\frac{1}{2}}$ for each n as $n \to \infty$

as the set of realizations of $\{y_i\}$ such that $\frac{1}{t}\sum_{i=1}^{t} y_i$ crosses its upper bound at least once. To show that the measure of ζ is strictly less than 1, consider the following. For each $\{\widehat{y}_i\} \in \zeta$, form a corresponding sample path $\{y_i'\}$ by changing $\widehat{y}_\tau$ to $y_\tau' = 0$ for each τ such that $\frac{1}{\tau}\sum_{i=1}^{\tau} \widehat{y}_\tau > \mu_y + B_\tau$ (for any $\{\widehat{y}_i\} \in \zeta$ there are only finitely many such τ). Then each element in ζ has a unique corresponding element in $Z \setminus \zeta$. So the measure of the set $Z \setminus \zeta$ is at least as large as the measure of ζ and this implies that the measure of ζ is strictly less than 1. Therefore, the measure of $Z \setminus \zeta$ is strictly positive. So there exists a set of realizations of $\{y_i\}$ that occur with positive probability such that $\frac{1}{t}\sum_{i=1}^{t} y_i$ never crosses outside $\mu_y + B_\tau$.

Given that $\{r_t\}$ and $\{\Delta_t^*\}$ are monotonic with respect to t,

$$\begin{aligned}\{B_t\} &\rightarrow 0\\ \{\Delta_t^*\} &\rightarrow \Delta^*\\ \Delta_\infty &= \mu_y < \Delta^*\end{aligned}$$

there are at most a finite number of periods k such that Δ_t^* lies inside $\Theta_t = [\mu_y + B_t, \mu_y - B_t]$. The probability that Δ_t doesn't cross Δ_t^* during these k periods is bounded below by $g(0)^k > 0$ (the probability of k supporting actions, which will never break a herd), and thus is strictly positive. Once Δ_t^* lies above Θ_t, all realizations in the set $Z \setminus\zeta$ never cross outside Θ_t, and therefore never cross Δ_t^*. So there is a set of sample path realizations that occur with positive probability such that the actual share of contrary actions never crosses the threshold required to break the herd, allowing the herd to persist with positive probability in the limit. QED. Hartman and Wintner [1941]

(b): Suppose $\Delta_\infty > \Delta^*$. Then Δ_∞ lies in the region that breaks a herd. By the law of large numbers, almost all sample paths of $\{\Delta_t\}_{t=0}^{\infty}$ converge to Δ_∞, so the threshold required to break the herd is crossed with probability 1. Thus the herd is broken with probability 1. Q.E.D.

Proof of Theorem 3 on pp. 13

Recall $\phi^h = \left(\frac{\widehat{p}+(1-\widehat{p})\pi^s}{\widehat{p}+(1-\widehat{p})(1-\pi^s)}\right)$. Note $\frac{d\phi^h}{d\widehat{p}} = \frac{1-2\pi^s}{(1-\pi^s+\widehat{p}\pi^s)^2} < 0$ and $\ln(\phi^c) = \ln\left(\frac{1-\pi^s}{\pi^s}\right) < 0$

$$\frac{d\Delta^*(\widehat{p})}{d\widehat{p}} = \frac{-\ln(\phi^c)\frac{d\phi^h}{d\widehat{p}}}{\phi^h\left(\ln(\phi^h) - \ln(\phi^c)\right)^2} < 0$$

(i) By definition, $\Delta_\infty^{\omega\neq a} = \Delta^*(\widehat{p}_1)$ so the limit of the sample path for an incorrect herd lies on the herd breaking threshold at $\widehat{p}_1$. Since $\frac{d\Delta^*(\widehat{p}_1)}{d\widehat{p}} < 0$, and the limit of the sample path doesn't depend on $\widehat{p}$, for $\widehat{p} < \widehat{p}_1$, $\Delta_\infty^{\omega\neq a} = \Delta^*(\widehat{p}_1) < \Delta^*(\widehat{p})$ so $\Delta_\infty^{\omega\neq a}$ lies below the herd breaking threshold and incorrect herds persist with positive probability, by theorem 2. Thus, there is positive probability that an incorrect herd persists, and agents choose only the suboptimal action infinitely often. Correct herds also persist with positive probability since $\Delta_\infty^{\omega=a} < \Delta_\infty^{\omega\neq a} < \Delta^*(\widehat{p}) \Rightarrow \Delta_\infty^{\omega=a}$ also lies below the herd breaking threshold. Thus, agents' action choices also converge on the optimal action with positive probability.

(iii) By definition, $\Delta_\infty^{\omega\neq a} = \Delta^*(\widehat{p}_2)$ so the limit of the sample path for a correct herd lies on the herd breaking threshold at $\widehat{p}_2$. Since $\frac{dk(\widehat{p})}{d\widehat{p}} < 0$, and the limit of the sample path doesn't depend on $\widehat{p}$, for $\widehat{p} > \widehat{p}_2$, $\Delta_\infty^{\omega=a} = \Delta^*(\widehat{p}_2) > \Delta^*(\widehat{p})$ so $\Delta_\infty^{\omega=a}$ lies above the herd breaking threshold and correct herds break with probability 1, by theorem 2. Since $\Delta_\infty^{\omega\neq a} > \Delta_\infty^{\omega=a} > \Delta^*(\widehat{p})$, incorrect herds also break with probability 1. Thus, no herd persists in the limit. Each time a herd breaks, correct and incorrect herds both form with positive probability. A new herd will form if the same action is played in the two periods subsequent to the herd breaking. The probability of two correct actions is $(\pi^s)^2 > 0$ and the probability of two incorrect actions is $(1-\pi^s)^2 > 0$. Since neither type of herd persists in the limit, both correct and incorrect herds form infinitely often. Therefore both the optimal and suboptimal action are chosen infinitely often, and herding is inefficient.

(ii) If $\widehat{p} \in (\widehat{p}_1, \widehat{p}_2)$ then $\Delta_\infty^{\omega\neq a} = \Delta^*(\widehat{p}_1) > \Delta^*(\widehat{p})$ and $\Delta_\infty^{\omega=a} = \Delta^*(\widehat{p}_2) < \Delta^*(\widehat{p})$ so incorrect herds break with probability 1 but correct herds persist with positive probability. Let A represent the event where no herd is occurring. Let $q_C(\mu)$ represent the probability that a correct herd forms and $q_I(\mu)$ represent the probability that an incorrect herd forms, given that no herd is occurring. Let $r_C(\mu)$ represent the probability that a correct herd persists in the limit. These probabilities depend on current public beliefs, but all are positive. Suppose event A is occurring in period t; that is, in period t, no herd is occurring. Then the probability that event A occurs again at some future period is $1 - q_C(\mu_t)r_C(\mu_t) < 1$. Event A occurs again if (a) no herd forms in period t (then A occurs in $t+1$) (b) an incorrect herd forms in period t (since this herd breaks with probability 1) (c) a correct herd forms in period t and breaks. The periods that A occur in form an increasing sequence $\tau_1 < \tau_2 < \ldots$ and for each τ_k, the probability that A occurs again is $1 - q_C(\mu_{\tau_k})r_C(\mu_{\tau_k}) < 1$. Thus, the probability that event A occurs infinitely often is $\lim_{n\to\infty}\prod_{k=1}^{n} 1 - q_C(\mu_{\tau_k})r_C(\mu_{\tau_k}) = 0$. So with probability 1, A occurs only a finite number of times. Thus, a correct herd forms and persists with probability 1, and agents will choose only the optimal action infinitely often.

(iv) (a) First show that $p > \widehat{p}_1$ by showing that the expected share of contrary actions in an incorrect herd lies above the herd breaking threshold when $\widehat{p} = p$ i.e. show $\Delta_\infty^{\omega\neq a} > \Delta^*(p)$

Consider the following equations:

$$\begin{aligned}
f(p) &= \frac{(1-p)\pi^s}{(1-(1-p)\pi^s)} \ln\left(\frac{\pi^s}{1-\pi^s}\right) \\
g(p) &= \ln(\phi^h) \\
&= \ln\left(\frac{p+(1-p)\pi^s}{p+(1-p)(1-\pi^s)}\right) \\
&= \ln\left(\frac{(1-\pi^s)p+\pi^s}{1-\pi^s+\pi^s p}\right)
\end{aligned}$$

At $p = 1$

$$f(1) = g(1) = 0$$

At $p = 0$

$$f(0) = \frac{\pi^s}{1-\pi^s} \ln\left(\frac{\pi^s}{1-\pi^s}\right) > g(0) = \ln\left(\frac{\pi^s}{1-\pi^s}\right)$$

since $\frac{\pi^s}{1-\pi^s} > 1$. Take the derivative of each expression with respect to p:

$$\begin{aligned}
\frac{d}{dp} f(p) &= -\frac{\pi^s}{(1-\pi^s+\pi^s p)^2} \ln\left(\frac{\pi^s}{1-\pi^s}\right) < 0 \\
\frac{d}{dp} g(p) &= \frac{-(2\pi^s-1)}{((1-\pi^s)p+\pi^s)(1-\pi^s+\pi^s p)} < 0
\end{aligned}$$

Take the second derivative of each expressions with respect to p:

$$\begin{aligned}
\frac{d^2}{dp^2} f(p) &= \frac{2(\pi^s)^2}{(1-\pi^s+\pi^s p)^3} \ln\left(\frac{\pi^s}{1-\pi^s}\right) > 0 \\
\frac{d^2}{dp^2} g(p) &= \frac{((1-\pi^s)^2+(\pi^s)^2+2p\pi^s(1-\pi^s))(2\pi^s-1)}{[((1-\pi^s)p+\pi^s)(1-\pi^s+\pi^s p)]^2} > 0
\end{aligned}$$

Given $f(0) > g(0)$, $f(1) = g(1)$ and both functions monotonically decrease at a decreasing rate, we can conclude that $f(p) > g(p)$ over the interval $[0, 1)$.

$$
\begin{aligned}
& f(p) > g(p) \\
\Rightarrow \quad & \frac{-(1-p)\pi^s}{(1-(1-p)\pi^s)} \ln\left(\phi^c\right) > \ln\left(\phi^h\right) \\
\Rightarrow \quad & (1-p)\pi^s > \frac{\ln\left(\phi^h\right)}{\ln\left(\phi^h\right) - \ln\left(\phi^c\right)} \\
\Rightarrow \quad & \Delta_\infty^{\omega \neq a} > \Delta^*(p) \\
\Rightarrow \quad & p > \widehat{p}_1
\end{aligned}
$$

Thus, $p > \widehat{p}_1$ for all $p \in [0, 1)$.

(b) Next show that $p < \widehat{p}_2$ by showing that the limit of the realized sample path for a correct herd lies below the herd breaking threshold limit when $\widehat{p} = p$ i.e. show $\Delta_\infty^{\omega=a} < \Delta^*(p)$. Consider the following equations:

$$
\begin{aligned}
f(p) &= \frac{(1-p)(1-\pi^s)}{(1-(1-p)(1-\pi^s))} \ln\left(\frac{\pi^s}{1-\pi^s}\right) \\
&= \frac{1-p-\pi^s+\pi^s p}{(p+\pi^s-\pi^s p)} \ln\left(\frac{\pi^s}{1-\pi^s}\right) \\
g(p) &= \ln(\phi^h) \\
&= \ln\left(\frac{(1-\pi^s)p+\pi^s}{1-\pi^s+\pi^s p}\right)
\end{aligned}
$$

At $p = 1$

$$
f(1) = g(1) = 0
$$

At $p = 0$

$$
f(0) = \frac{1-\pi^s}{\pi^s} \ln\left(\frac{\pi^s}{1-\pi^s}\right) < g(0) = \ln\left(\frac{\pi^s}{1-\pi^s}\right)
$$

since $\frac{1-\pi^s}{\pi^s} < 1$. Take the derivative of each expression with respect to p:

$$
\begin{aligned}
\frac{d}{dp} f(p) &= \frac{-(1-\pi^s)}{(p+\pi^s-\pi^s p)^2} \ln\left(\frac{\pi^s}{1-\pi^s}\right) < 0 \\
\frac{d}{dp} g(p) &= \frac{-(2\pi^s-1)}{((1-\pi^s)p+\pi^s)(1-\pi^s+\pi^s p)} < 0
\end{aligned}
$$

Take the second derivative of each expressions with respect to p:

$$
\begin{aligned}
\frac{d^2}{dp^2} f(p) &= \frac{2(1-\pi^s)^2}{(p+\pi^s-\pi^s p)^3} \ln\left(\frac{\pi^s}{1-\pi^s}\right) > 0 \\
\frac{d^2}{dp^2} g(p) &= \frac{((1-\pi^s)^2+(\pi^s)^2+2p\pi^s(1-\pi^s))(2\pi^s-1)}{[((1-\pi^s)p+\pi^s)(1-\pi^s+\pi^s p)]^2} > 0
\end{aligned}
$$

Given $f(0) < g(0)$, $f(1) = g(1)$ and both functions monotonically decrease at a decreasing rate, we can conclude that $f(p) < g(p)$ over the interval $[0, 1)$.

$$
\begin{aligned}
& f(p) < g(p) \\
\Rightarrow \quad & \frac{(1-p)(1-\pi^s)}{(1-(1-p)(1-\pi^s))} \ln\left(\frac{\pi^s}{1-\pi^s}\right) < \ln\left(\phi^h\right) \\
\Rightarrow \quad & (1-p)(1-\pi^s) < \frac{\ln\left(\phi^h\right)}{\ln\left(\phi^h\right)-\ln(\phi^c)} \\
\Rightarrow \quad & \Delta_\infty^{\omega=a} < \Delta^*(p) \\
\Rightarrow \quad & p < \widehat{p}_2
\end{aligned}
$$

Thus, $p < \widehat{p}_2$ for all $p \in [0, 1)$. This illustrates that correct herds persist and incorrect herds break when $\widehat{p} = p$. Q.E.D.

Proof of Corollary 1 on pp. 15

Suppose an L-herd forms and persists. In order for a herd to persist,

$$
\Delta_t < \Delta_t^*(\widehat{p}) \ \forall t
$$

Let $\Delta_t = \Delta_t^* - c_t$ for some $c_t > 0$. Note $\Lambda_0 \left(\phi^h\right)^{(t-1)(1-\Delta_t^*)} (\phi^c)^{(t-1)\Delta_t^*} = 1 \ \forall t$ since this represents beliefs at the herd breaking threshold. We can rewrite the public likelihood ratio in an L-herd as:

$$
\begin{aligned}
\frac{\mu_t}{1-\mu_t} &= \Lambda_0 \left(\frac{\pi^s}{1-\pi^s}\right) \left(\phi^h\right)^{(t-1)(1-\Delta_t)} (\phi^c)^{(t-1)\Delta_t} \\
&= \Lambda_0 \left(\frac{\pi^s}{1-\pi^s}\right) \left(\phi^h\right)^{(t-1)(1-\Delta_t^*+c_t)} (\phi^c)^{(t-1)(\Delta_t^*-c_t)} \\
&= \Lambda_0 \left(\frac{\pi^s}{1-\pi^s}\right) \left(\phi^h\right)^{(t-1)(1-\Delta_t^*)} (\phi^c)^{(t-1)(\Delta_t^*)} \left(\phi^h\right)^{(t-1)c_t} (\phi^c)^{-(t-1)c_t} \\
&= \left(\frac{\pi^s}{1-\pi^s}\right) \left(\frac{\phi^h}{\phi^c}\right)^{(t-1)c_t}
\end{aligned}
$$

Note $\left(\frac{\phi^h}{\phi^c}\right) > 1$.

(i) If the L-herd is correct, then learning is complete when the inverse public likelihood ratio converges to zero. The limit of the public likelihood ratio is as follows:

$$\lim_{t\to\infty} \frac{1-\mu_t}{\mu_t} = \lim_{t\to\infty} \left(\frac{1-\pi^s}{\pi^s}\right)\left(\frac{\phi^c}{\phi^h}\right)^{(t-1)c_t} = 0$$

Similar analysis shows $\lim_{t\to\infty} \frac{\mu_t}{1-\mu_t} = 0$ in a correct R-herd. Thus, when a correct herd persists, learning is complete.

(ii) If the L-herd is incorrect, then learning is fully incorrect when the public likelihood ratio converges to infinity:

$$\lim_{t\to\infty} \frac{\mu_t}{1-\mu_t} = \lim_{t\to\infty} \left(\frac{\pi^s}{1-\pi^s}\right)\left(\frac{\phi^h}{\phi^c}\right)^{(t-1)c_t} = \infty$$

Similar analysis shows $\lim_{t\to\infty} \frac{1-\mu_t}{\mu_t} = \infty$ in an incorrect R-herd. Thus, learning is fully incorrect when an incorrect herd persists.

(iii) Suppose incorrect and correct L-herds both break. Then $\exists \tau$ s.t. $\Delta_\tau > \Delta^*_\tau(\widehat{p})$ and the public likelihood ratio falls below 1. If an L-herd forms again, this will repeat whereas if an R-herd forms, then the public likelihood ratio eventually rises above 1, at which point the R-herd breaks. Thus, the public likelihood ratio never converges in either an L-herd or an R-herd and public beliefs remain interior, resulting in incomplete learning. Q.E.D.