UTMD Working Paper
The University of Tokyo
Market Design Center

UTMD-113

Social Learning with Correlated Information

Yu Awaya
The University of Tokyo

Vijay Krishna
Penn State University

January 20, 2026

Social Learning with Correlated Information*

Yu Awaya[†] and Vijay Krishna[‡]

January 20, 2026

Abstract

We study a standard binary social learning model where agents' information is serially correlated—it is generated by a Markov process. There is a unique equilibrium in which a herd, sometimes incorrect, always forms. In the long run, does greater persistence increase the likelihood that an incorrect herd forms? In the medium run (prior to the formation of a herd), does a greater similarity information—higher persistence—lead to a greater similarity of actions? The answer to both questions is no.

1 Introduction

It is well-known that when economic agents seek to learn relevant information from other agents, the possibility of herd behavior emerges—agents take the same action, perhaps even ignoring their own information. This possibility was first recognized by Banerjee (1992) and Bikhchandani, Hirschleifer and Welch (1992).[1] These papers have then spawned a vast literature on "social learning" that explores, in various contexts, the specific conditions under which herd behavior emerges as an equilibrium phenomenon. In most of this work, there is some unknown payoff-relevant state of nature and, conditional on the realized state of nature, agents' information is *independently* distributed.

*The research reported here was supported by a grant from the National Science Foundation (SES-2048806). We are grateful to Stepan Aleksenko, Yuhta Ishii, Min-Feng Lu and Eduard Osipov for helpful comments.

[†]Department of Economics, University of Tokyo, E-mail: yuawaya@gmail.com.

[‡]Department of Economics, Penn State University, E-mail: vkrishna@psu.edu.

[1]Bikhchandani et al. (1992) use the term "information cascade" to describe herd behavior.

In many economic situations, however, it seems natural that, conditional on the state, agents' information is serially correlated (see, for instance, Graham, 1999). In financial markets, investors get information from advisors, newsletters, and acquaintances. But these sources themselves have a strong common component—Fed announcements, political developments, market rumors. Thus, the information provided to investors today is correlated with that provided to investors yesterday.

In this paper, we study a situation in which the information of agents is serially (positively) correlated—that is, *persistent*. Specifically, it is generated by a simple Markov process—given the fundamentals, if one agent gets a high signal today, then this makes it more likely that the next agent will also get a high signal tomorrow. This structure results in a tractable model that allows for serial correlation of information.

In a simple binary-state, binary-action setting we first choose a particular candidate equilibrium and show that in each state, the resulting public beliefs—those of an observer who sees only the actions chosen by the agents—form a finite Markov chain with absorbing states. This is true even though the beliefs of an observer who sees all the signals are not Markovian. Given this, we then establish that the candidate equilibrium is indeed an equilibrium. This is the (essentially) unique Nash equilibrium (Theorem 1).

In equilibrium, a herd forms for sure. Sometimes the herd is the correct one—where all but a finite number of agents choose the action appropriate for the underlying state. But in other instances, an incorrect herd forms—all but a finite number of agents choose the wrong action.

We then address two questions regarding the effects of greater persistence on welfare and behavior.

1. In the long run, does greater persistence increase the likelihood that an incorrect herd forms?

2. In the medium run (prior to the formation of a herd), does a greater similarity information—higher persistence—lead to a greater similarity of actions?

As a first step, suppose that agents' information was perfectly persistent—that is, they all receive the same information. This, of course, is the worst case since a wrong initial signal will immediately trigger the wrong herd. Also, everyone would take the same action. One may then reasonably conjecture that higher persistence would lead

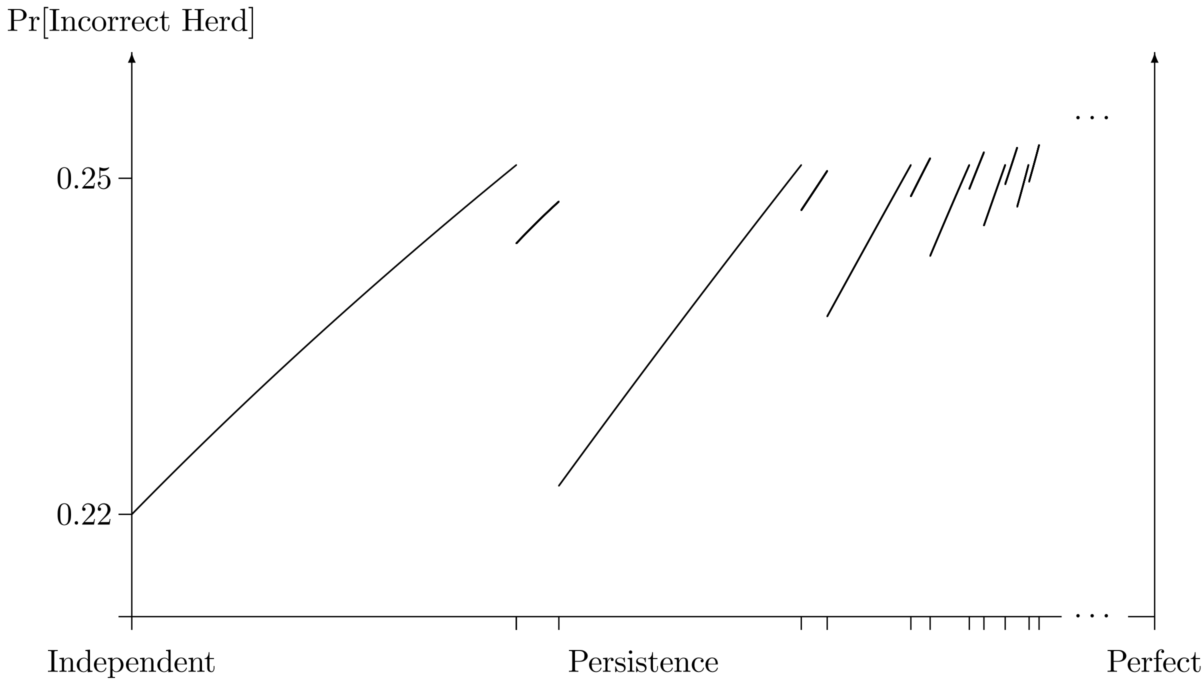

Figure 1: The figure depicts how the probability of an incorrect herd changes as signals become more correlated.

to a higher probability of a wrong herd and also lead to a greater similarity of actions. But we find that this is not the case.

First, it turns out that the relationship between persistence and the probability of a wrong herd is "highly" non-monotonic. Figure 1 depicts an example showing how the probability of an incorrect herd changes non-monotonically with an increase in the serial correlation of signals, ranging from the case of (conditionally) independent signals to those that are perfectly persistent[2]. As the persistence increases, the probability of an incorrect herd increases smoothly and then jumps downward. This pattern repeats a countably infinite number of times. Immediately to the right of a discontinuity, the probability of an incorrect herd is smaller—and the welfare is higher—than at lower levels of correlation to the left of a discontinuity. The total variation of the function—a measure of non-monotonicity—is unbounded. We will show that these features of the example are general (Theorem 2).

Why is this? How can an increase in the persistence of signals lead to a decrease in the probability of an incorrect herd—and therefore an increase in welfare? As usual in social learning models, if an agent observes that sufficiently many agents

[2]The parameter values for this example are reported in Section 5.1.

who precede him take the same action, say L, then this agent infers that others' information is favorable for L and so is convinced to choose L even if his private information indicates otherwise—a herd to L occurs. When signals become more persistent, it takes a larger number of L's to convince an agent to ignore his own information. In other words, now this agent chooses a decision that reflects his own information. This provides all succeeding agents with better information which in turn decreases the chances of an incorrect herd. In this way, more persistent signals can *delay* the formation of a herd.

Even though higher persistence can lead to a lower probability of incorrect herds, one may still surmise that independent information is the best (as in Figure 1). If my private information is independent of that of agents' whose actions I observe, then this should be the best case for paying attention to my own information. We demonstrate that this intuition is also incorrect and it may be that some persistence is better than none. Although this is not a feature of the example in Figure 1, we identify circumstances in which some persistence is better than independence (Proposition 5.3).

Second, we also find that prior to the formation of a herd, a greater similarity of signals—greater persistence—implies a greater *dissimilarity* of actions. As argued above, more persistent signals can lead to a delay in herds and so lead agents to rely on their private information for longer. And since signals are noisy, this leads their choices to be more dissimilar than if the signals were less persistent (Proposition 6.1).

Allowing for correlated signals also leads to an "anything goes" result. Suppose an outsider observes a finite history of action choices. We show that any such history can be rationalized as resulting from equilibrium behavior for some degree of correlation.

In this sense, correlated signals can provide an explanation of seemingly *contrarian* behavior. A contrarian is an agent who does not follow the herd and makes a non-conforming choice. Suppose we observe even though agent 11 sees 10 agents before him choose action L, he chooses H. This behavior would be deemed contrarian (non-equilibrium) under the hypothesis that agents' signals are serially independent. But since persistent signals delay the formation of a herd, it is entirely possible that agent 11's choice of H is an equilibrium response. Seeing 10 people before him choose L is not enough to convince him that he should neglect his own signal. His choice of H after seeing 10 choices of L may well be rational once signals are persistent.

Related literature Starting from Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992, henceforth BHW), a vast literature—too extensive to discuss here—has developed exploring various aspects of the theory as well as applying it to varied environments. Fortunately, there are excellent surveys available, including a recent comprehensive one by Bikhchandani, Hirshleifer, Tamuz and Welch (2024), and we defer to these.[3] Almost all of this work assumes that, conditional on the fundamental state, agents' information is independently distributed over time, that is, serially uncorrelated. In this paper, we work with the very simple binary state, binary signal, binary action model of Banerjee (1992) and BHW (1992) and explore the implications of serially correlated (persistent) signals.

Smith and Sorensen (2000) introduce the possibility of *continuously* distributed signals (as well as noise and payoff shocks) while retaining the assumption that these are serially independent. In this richer environment, asymptotic behavior is determined by whether or not the posterior likelihood ratio of private beliefs (based only on private signals) is bounded or not. When it is bounded, the possibility of incorrect herd emerges as in the finite-signal models Banerjee (1992) and BHW (1992); if it is unbounded, incorrect herds cannot arise. In the case of unbounded likelihood ratios, actions converge stochastically so that even following a long sequence of "right" actions, an agent chooses the "wrong" action. In other words, with positive probability equilibrium behavior is "contrarian." In recent work, Kartik, Lee, Liu and Rappoport (2024) observe that unbounded beliefs cannot occur with more than two states. These authors go well beyond the binary-signal model to allow for general sets of states, signals and utilities and find sufficient and (almost) necessary conditions on the information structure to guarantee that social learning maximizes welfare in the limit. Needless to say, these conditions are different from unbounded beliefs condition of Smith and Sorensen (2000) for the case of binary states. While the Kartik et al. (2024) model is very general in many respects, it retains the standard assumption of serially independent signals.

In work that predates the Banerjee and BHW papers, Scharfstein and Stein (1990) study a model in which two investors may receive signals and then sequentially decide whether to invest or not. But it is not known if the signals are informative or pure noise. When signals are informative, they are perfectly correlated; otherwise, they are independently distributed. There are reputational concerns as well—the second

[3] Other surveys include Chamley (2004) and Smith and Sorensen (2011).

mover is interested in making outside observers believe that he is smart, that is, his signals is informative. In equilbrium, the second agent follows the investment choice of the first regardless of his information. In contrast to the Banerjee (1992) and BHW (1992) models, correlated signals are necessary for a herd to emerge—with independent signals there is no herding. In this sense, Sharfstein and Stein (1990) compare correlated signals with independent ones, albeit in a model with reputational concerns.

Signals are also correlated in the model of Schaal and Taschereau-Dumouchel (2023) where in each period a continuum of agents decide simultaneously whether or not to undertake a risky investment. In addition to his or her private signal, each agent observes only the fraction of investors who invested in previous periods. Signals have a (noisy) common component and when this is large, a boom ensues which is then followed by a bust. Again, the model is very different from the one we study but does contain the essential elements of correlated signals in a herding environment.

Wu (2021) studies a model of herding in which there are *two* separate queues of agents. The states of nature in the two queues are different but correlated and an agent in either queue can observe the history of actions in both queues. Signals are still independent conditional on the underlying state in a particular queue. She finds that although agents are now better informed than in the standard one-queue model, the probability of an incorrect herd may increase as a result of this additional information.

Varying the extent of serial correlation varies the underlying information structure. While retaining the assumption of independent signals, Sato and Shimizu (2025) exhibit an example with binary states and many signals in which a more informative information structure can result in decrease in welfare.

In all of work discussed above, the state of nature is drawn once and for all and stays fixed. Moscarini, Ottaviani and Smith (1998) study a model in which the state of nature changes over time via a simple Markov process defined over two possible states of nature. Conditional on the current state, signals are still independent. Depending on the persistence of the state of nature, a herd may or may not occur. Even if one does occur, it is necessarily temporary. The reason is that once sufficient time has elapsed, there is a good chance that the state has changed and that one's private signal is relevant. Thus, in the Moscarini et al. model, the state is serially correlated while the signals are not. In our model, the state is fixed while the signals are correlated. In the same model, Wang (2024) shows that when the state is quite persistent, actions change more rapidly than the state.

Organization of the paper The rest of this paper is organized as follows. The model of social learning with persistent (serially correlated) signals is described in the next section. Section 3 constructs an equilibrium and shows that it is unique (modulo tie-breaking). The equilibrium generates public beliefs—those of an outside observer—that form a finite Markov chain with three absorbing states. These absorbing states result in herds. Section 4 then studies how changes in persistence affect the Markov chain. In Section 5 we derive the probabilities of the two different herds. We then study how an increase in persistence affects the probability that an "incorrect" herd forms. and show that, as in Figure 1, this probability behaves highly non-monotonically with a change in persistence. Finally, in Section 6 we examine the implications of our results on agents' behavior in the medium run.

2 Model

There are two states of nature, θ^H and θ^L, with prior probabilities ρ and $1-\rho$, respectively. We assume that $\frac{1}{2}<\rho<1$.

A countable infinity of agents arrive sequentially one-at-a-time at $t=1,2,...$ and each chooses an action $a_t \in A=\{H,L\}$. Each agent receives a payoff of 1 if the chosen action matches the unknown state and a payoff of 0 if it does not. In other words, each agent has the same payoff function $u\left(H,\theta^H\right)=1$ and $u\left(L,\theta^H\right)=0$. Similarly, $u\left(L,\theta^L\right)=1$ and $u\left(H,\theta^L\right)=0$.

The agent arriving at time t observes (i) the choices of all preceding agents—the vector $\boldsymbol{a}^{t-1}=(a_1,a_2,...,a_{t-1})$; and (ii) a private signal $s_t\in\{\ell,h\}$.

The signals $s_t\in\{\ell,h\}$ are generated as follows: in state θ^H, the period 1 signals are such that

$$\Pr\left[s_1=h\mid\theta^H\right]=q \text{ and } \Pr\left[s_1=\ell\mid\theta^H\right]=1-q$$

where $\frac{1}{2}<q<1$. In state θ^L, the corresponding probabilities of h and ℓ are $1-q$ and

q, respectively. To economize on notation from now on we will write

$$\mathbb{P}^H[\cdot] \equiv \Pr\left[\cdot \mid \theta^H\right] \text{ and } \mathbb{P}^L[\cdot] \equiv \Pr\left[\cdot \mid \theta^L\right]$$

To avoid trivialities we will assume that the first period signal is sufficiently informative so that agent 1's action a_1 responds to his signal. In other words, $a_1 = H$ if and only if $s_1 = h$. Thus, we suppose that

$$q > \rho \tag{1}$$

In other periods, the signals evolve according to a *Markov* process. In state θ^H, the transition probabilities $\mathbb{P}^H[s_{t+1} \mid s_t]$ are given by

$S^H =$

θ^H	$s_{t+1} = h$	$s_{t+1} = \ell$
$s_t = h$	α	$1-\alpha$
$s_t = \ell$	$1-\beta$	β

where α and β are given parameters. Thus, $\mathbb{P}^H[s_{t+1} = h \mid s_t = h] = \alpha$ etc. In state θ^L, α and β are exchanged so that $\mathbb{P}^L[s_{t+1} = h \mid s_t = h] = \beta$ etc.

In order to isolate the effects of changes in the serial correlation of signals, we will also assume that the signal process is such that in each state, the *marginal* distribution of signals is preserved over time.

Condition 1 (Preservation of Marginals) *For all t, $\mathbb{P}^H[s_t = h] = q$ and $\mathbb{P}^L[s_t = h] = 1-q$. Equivalently,*

$$\frac{1-\beta}{1-\alpha} = \frac{q}{1-q}$$

To see why the condition ensures that the marginals are preserved, note that

$$\begin{aligned} q &= \mathbb{P}^H[s_{t+1} = h] \\ &= \mathbb{P}^H[s_{t+1} = h \mid s_t = h] \times \mathbb{P}^H[s_t = h] + \mathbb{P}^H[s_{t+1} = h \mid s_t = \ell] \times \mathbb{P}^H[s_t = \ell] \\ &= \alpha q + (1-\beta)(1-q) \end{aligned}$$

which yields the condition above. The requirement that $\mathbb{P}^L[s_t = h] = 1-q$ also yields the same condition.

Condition 1 is equivalent to: $(q, 1-q)$ is the stationary distribution of the Markov matrix

$$S^H = \begin{bmatrix} \alpha & 1-\alpha \\ 1-\beta & \beta \end{bmatrix}$$

and also equivalent to: $(1-q, q)$ is the stationary distribution of the Markov matrix S^L obtained from S^H by exchanging α and β. Condition 1 also implies that

$$\alpha > \beta$$

since by (1), $q > \rho > \frac{1}{2}$.

It will be useful to define

$$\gamma \equiv \frac{\beta}{\alpha} \tag{2}$$

and the condition that the marginal distribution of signals is preserved implies that α, β and γ are *co-monotonic*—if any one of the three increases, the other two also increase. We will then refer to γ as the *persistence* parameter since a higher γ implies that the correlation between tomorrow's signal and today's signal is higher.

Notice that the signals are independently distributed over time if and only if $\gamma = \frac{1-q}{q}$ since, together with Condition 1, this is equivalent to $\alpha = q$ and $\beta = 1-q$. We will thus suppose throughout that

$$\gamma = \frac{\beta}{\alpha} \geq \frac{1-q}{q}$$

Interpretation as a renewal process The Markovian signal process described above is equivalent to the following (discrete) *renewal process.* In period 1, the signal s_1 is drawn with the marginal probabilities appropriate for the state. In every subsequent period $t > 1$, with constant probability p, the signal s_t is the same as s_{t-1}; and with probability $1-p$, s_t is an independent draw—that is, the signal is renewed—with the relevant marginal probabilities (that is, the distribution of s_1). Note that p is the same in both states θ^H and θ^L.

Suppose we have the Markovian signal structure outlined above with parameters α and β. If we set $p = \alpha + \beta - 1$, then the resulting renewal process is the same as defined by the Markov matrices S^H and S^L. Conversely, suppose we have a renewal process in which with probability p, a signal is repeated and with probability $1-p$, the next period's signal is renewed. Then it is easy to verify that if we set $\alpha = p + q(1-p)$

and $\beta = 1 - q(1-p)$, then the resulting Markovian signal process is the same as the renewal process.[4]

The renewal process can also be interpreted as follows. With probability p, agents t and $t-1$ have a "common" source of information (and so get identical signals) and with probability $1-p$, t's source of information is independent of $t-1$'s source. Similarly, with probability p^2, t, $t-1$ and $t-2$ have a common source, etc.

3 Equilibrium

In this section, we establish that the game described above has a unique equilibrium (modulo tie-breaking). In this equilibrium, the agents' actions do not depend on the entire history of the actions of preceding agents; rather they depend only on a "state variable" that takes only a finite number of values.

A (pure) *strategy* for agent t is a function $\sigma_t : A^{t-1} \times \{\ell, h\} \rightarrow \{L, H\}$, where $A^{t-1} = \{H, L\}^{t-1}$ is the set of histories of past actions up to $t-1$. Thus, $\sigma_t(\boldsymbol{a}^{t-1}, s_t) \in \{H, L\}$ is the action chosen by t given the history of actions $\boldsymbol{a}^{t-1} \in A^{t-1}$ and his private signal $s_t \in \{\ell, h\}$. We will say that an agent *follows his signal* if $\sigma_t(\boldsymbol{a}^{t-1}, h) = H$ and $\sigma_t(\boldsymbol{a}^{t-1}, \ell) = L$. An agent *ignores his signal* if $\sigma_t(\boldsymbol{a}^{t-1}, h) = \sigma_t(\boldsymbol{a}^{t-1}, \ell)$.

We will say that a history of actions $\boldsymbol{a}^{t-1} \in A^{t-1}$ is *consistent* with $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, ...)$ if there is a signal history $\boldsymbol{s}^{t-1} \in \{\ell, h\}^{t-1}$ such that for all $\tau < t$, $a_\tau = \sigma_\tau(\boldsymbol{a}^{\tau-1}, s_\tau)$.

Definition 1 *Given strategies $\boldsymbol{\sigma}$ and a history of actions $\boldsymbol{a}^{t-1}$ consistent with $\boldsymbol{\sigma}$, the* ***public belief*** *at time $t > 1$, denoted by $B_t(\boldsymbol{\sigma}, \boldsymbol{a}^{t-1})$, is the likelihood ratio of the posteriors of an outside observer who sees only the history of actions and knows that these were chosen according to $\boldsymbol{\sigma}$.*[5]

Our first result is that in the social learning model with Markovian information

[4] It may be verified that in both states, the corresponding φ *coefficient*, the standard measure of serial correlation for binary processes, is also $\alpha + \beta - 1$.

[5] Precisely,

$$B_t(\boldsymbol{\sigma}, \boldsymbol{a}^{t-1}) = \frac{\Pr\left[\theta^H\right]}{\Pr\left[\theta^L\right]} \times \frac{\mathbb{P}^H\left[\left\{\boldsymbol{s}^{t-1} : \forall\tau < t,\ a_\tau = \sigma_\tau\left(\boldsymbol{a}^{\tau-1}, s_\tau\right)\right\}\right]}{\mathbb{P}^L\left[\left\{\boldsymbol{s}^{t-1} : \forall\tau < t,\ a_\tau = \sigma_\tau\left(\boldsymbol{a}^{\tau-1}, s_\tau\right)\right\}\right]}$$

Theorem 1 *There is a unique equilibrium* $\boldsymbol{\sigma}^*$ *(modulo tie-breaking). It is such that:*
(i) The public beliefs $B_t^*(\cdot) \equiv B_t(\boldsymbol{\sigma}^*, \cdot)$ *form a finite Markov chain with three absorbing states.*
(ii) Agents follow their private signals when the public beliefs are transient and ignore these when they are absrorbing.

The theorem will be established as follows. First, we will identify a finite set $\Lambda(m,n)$ of public beliefs (defined below) and use this to define a candidate equilibrium $\boldsymbol{\sigma}^*$. We will then show that

- The public beliefs B_t^* generated by $\boldsymbol{\sigma}^*$ lie in $\Lambda(m,n)$. (Proposition 3.1 below.)
- The public beliefs B_t^* constitute a finite Markov chain on $\Lambda(m,n)$ with three absorbing states. (Proposition 3.2 below.)
- Given the public beliefs B_t^*, the strategies $\boldsymbol{\sigma}^*$ constitute an equilibrium. This is the unique equilibrium (modulo tie-breaking). (Proposition 3.3 below.)

3.1 Likelihood ratios

For any t-period history of signals $\boldsymbol{s}^t = (s_1, s_2, ..., s_t) \in \{\ell, h\}^t$, denote the resulting (likelihood ratio of) *posterior beliefs* about the states of nature by

$$\lambda\left(\boldsymbol{s}^t\right) = \frac{\Pr\left[\theta^H \mid \boldsymbol{s}^t\right]}{\Pr\left[\theta^L \mid \boldsymbol{s}^t\right]} = \frac{\Pr\left[\theta^H\right]}{\Pr\left[\theta^L\right]} \frac{\mathbb{P}^H\left[s_1\right]}{\mathbb{P}^L\left[s_1\right]} \prod_{\tau=2}^{t} \frac{\mathbb{P}^H\left[s_\tau \mid s_{\tau-1}\right]}{\mathbb{P}^L\left[s_\tau \mid s_{\tau-1}\right]}$$

We will argue that in equilibrium, even though an outside observer at time t does not observe the signal history $\boldsymbol{s}^{t-1}$ directly, the public beliefs $B_t^*(\cdot)$ must equal $\lambda(\boldsymbol{s}^\tau)$ for some $\tau \le t-1$.

Denote by ℓ^k a string of k consecutive signals of ℓ and similarly, by h^j a string of j consecutive signals of h. Thus, $\ell^{k_1} h^{j_1} \ldots \ell^{k_p} h^{j_p}$ denotes a signal history which consists of k_1 consecutive ℓ's followed by j_1 consecutive h's followed by k_2 consecutive ℓ's etc.

Definition 2 *Given* γ, *let* m *be the smallest positive integer* j *such that*

$$\lambda\left(\ell h^{j+1} \ell\right) \ge 1$$

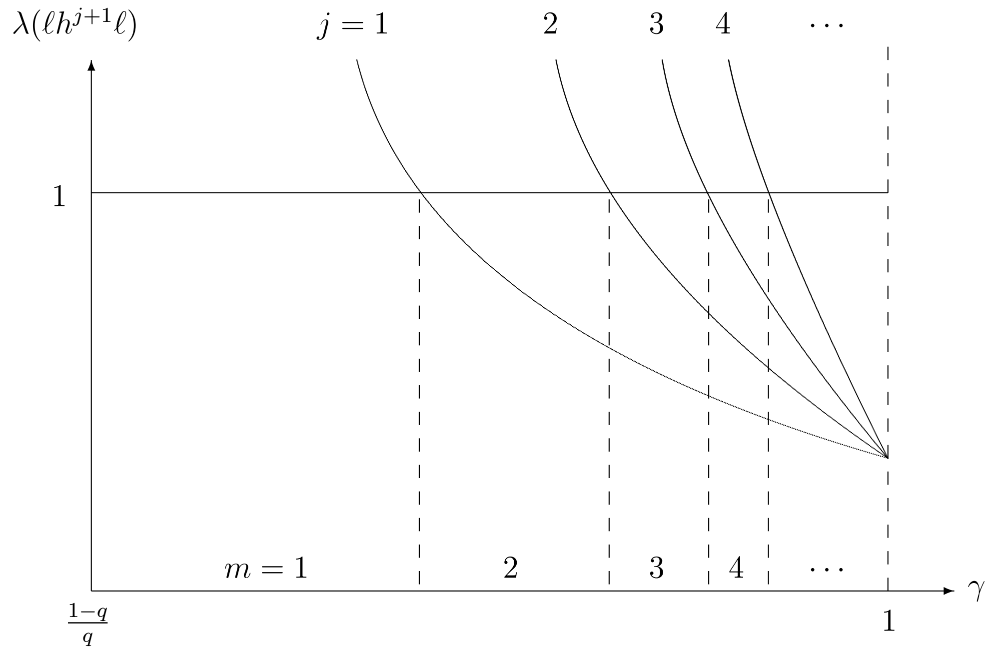

Figure 2: In the first interval, $m = 1$, In the second, $m = 2$ and so on.

Such an m exists and is unique since (i) $\lambda(\ell h \ell) < 1$; (ii) $\lambda(\ell h^{j+1}\ell)$ is increasing in j and unbounded as j goes to infinity. Figure 2 depicts how m is determined.

Analogously,

Definition 3 *Given γ, let n be the smallest positive integer k such that*

$$\lambda\left(\ell^{k+1}h\right) < 1$$

Again such an n exists and is unique since (i) $\lambda(\ell h) > 1$; (ii) $\lambda\left(\ell^{k+1}h\right)$ is decreasing in k and goes to zero as k goes infinity.

Next, given a persistence parameter γ, suppose m and n are determined as above and define the set

$$\begin{aligned}\Lambda(m,n) &= \{\lambda(h)\} \\ &\cup \left\{\lambda\left(\ell h^k\right) : k = 1,2,...,m+1\right\} \cup \left\{\lambda\left(\ell h^k \ell\right) : k = 1,2,...,m\right\} \\ &\cup \left\{\lambda\left(\ell^k h\right) : k = 1,2,...,n\right\} \cup \left\{\lambda\left(\ell^k\right) : k = 1,2,...,n+1\right\}\end{aligned}$$

The set $\Lambda(m,n)$ has $2(m+n)+1$ elements.[6] We will refer to the elements of $\Lambda(m,n)$ as *belief states.*

Notation 1 *Let*

$$\mathcal{A} = \left\{\lambda(h), \lambda\left(\ell h^{m+1}\right), \lambda\left(\ell^{n+1}\right)\right\}$$

Also, let

$$\mathcal{T} = \Lambda(m,n) \diagdown \mathcal{A}$$

We will show below that in the (essentially) unique equilibrium, the public beliefs form a Markov chain on $\Lambda(m,n)$ such that the belief states in $\mathcal{A}$ are *absorbing* while those in $\mathcal{T}$ are *transient.*

3.2 Equilibrium strategy

Given a γ, let m and n be determined as in Definitions 2 and 3.

Consider the following strategy $\boldsymbol{\sigma}^* = (\sigma_1^*, \sigma_2^*, ...)$.

- agent 1 follows her signal;
- agent $t > 1$,
  - follows her signal if the public belief $B_t^*\left(\boldsymbol{a}^{t-1}\right)$ is transient, that is, in $\mathcal{T}$;
  - chooses $a_t = a_{t-1}$ if the public belief $B_t^*\left(\boldsymbol{a}^{t-1}\right)$ is absorbing, that is, in $\mathcal{A}$.

For any $t > 1$ and any $\boldsymbol{a}^{t-1}$ not consistent with $\left(\sigma_1^*, \sigma_2^*, ..., \sigma_{t-1}^*\right)$, agent t's posterior beliefs are arbitrary and the agent is assumed to optimize given these beliefs.

We will show below that if the action history $\boldsymbol{a}^{t-1}$ is consistent with $\boldsymbol{\sigma}^*$, then the public beliefs $B_t^*\left(\boldsymbol{a}^{t-1}\right) \in \Lambda(m,n)$ and so the strategy $\boldsymbol{\sigma}^*$ is well-defined.

3.3 Public beliefs lie in $\Lambda(m,n)$

We now show

Proposition 3.1 *For all $t > 1$, the public beliefs generated by $\boldsymbol{\sigma}^*$, $B_t^* \in \Lambda(m,n)$.*

[6] In the description above, some elements are listed twice. For example, if $k = 1$, $\lambda\left(\ell h^k\right) = \lambda\left(\ell^k h\right)$ and $\lambda\left(\ell h^k \ell\right) = \lambda\left(\ell^k\right)$.

First, consider an arbitrary signal history $\boldsymbol{s}^t$ with $s_1 = h$. Then $\boldsymbol{\sigma}^*$ prescribes that agent 1 choose $a_1 = H$ and so the public belief after period 1, $B_2^* = \lambda(h) \in \Lambda(m, n)$. Now $\boldsymbol{\sigma}^*$ prescribes that $a_2 = a_1 = H$ regardless of s_2. This means that public beliefs cannot be updated and so $B_3^* = \lambda(h)$ as well. Proceeding in this manner shows that then in all subsequent periods, $B_t^* = \lambda(h) \in \Lambda(m, n)$. Thus, the conclusion of the proposition holds if $s_1 = h$.

So it remains to consider signal histories $\boldsymbol{s}^t$ with $s_1 = \ell$. The number of such signal histories is 2^{t-1} and so grows exponentially with t. Nevetheless, it will turn out that the resulting posterior beliefs $\lambda(\boldsymbol{s}^t)$ can be classified into four simple types.

To see how, note that in any history of signals $\boldsymbol{s}^t$, the likelihood ratio of an $\ell \to \ell$ transition is the inverse of the likelihood ratio of an $h \to h$ transition and so these ratios "cancel" each other. Similarly, the likelihood ratio of an $\ell \to h$ transition is the inverse of the likelihood ratio of an $h \to \ell$ transition and so these also "cancel" out. As an example,

$$\begin{aligned}
\lambda(\ell\ell\ell hh) &= \underbrace{\frac{\rho}{1-\rho}}_{\text{Prior}} \underbrace{\frac{1-q}{q}}_{s_1=\ell} \underbrace{\frac{\beta}{\alpha}}_{\ell\to\ell} \underbrace{\frac{\beta}{\alpha}}_{\ell\to\ell} \underbrace{\frac{1-\beta}{1-\alpha}}_{\ell\to h} \underbrace{\frac{\alpha}{\beta}}_{h\to h} \\
&= \frac{\rho}{1-\rho} \frac{1-q}{q} \frac{\beta}{\alpha} \frac{1-\beta}{1-\alpha} \\
&= \lambda(\ell\ell h)
\end{aligned}$$

Using such reductions, in Appendix A we show that if $s_1 = \ell$, then any $\lambda(\boldsymbol{s}^t)$ is equal to one of the following canonical forms: (i) $\lambda\left(\ell^{d+1} h\right)$; (ii) $\lambda\left(\ell h^{d+1}\right)$; (iii) $\lambda\left(\ell^{d+1}\right)$; or (iv) $\lambda\left(\ell h^{d+1} \ell\right)$ where, in each case d is some non-negative integer. We will refer to this as the *reduction property*.

The proof of Proposition 3.1 is by induction and can be found in Appendix B. The reduction property is crucial to the proof. Here we indicate the basic reasoning.

First, it is clear that if the public belief B_t^* is in the set of absorbing belief states, then in subsequent periods the public belief will remain in $\mathcal{A}$. The reason is that $\boldsymbol{\sigma}^*$ prescribes that in absorbing states, agents choose actions that ignore their own private signals. Since no new information is added, the public beliefs are not revised after t.

This implies that if the public belief B_t^* is in the set of transient belief states, then it must be that the public beliefs in all previous periods were also transient.

Now $\boldsymbol{\sigma}^*$ prescribes that in transient states, agents follow their private signals. This means that in all previous periods the actions reveal the signals and so $B_t^* = \lambda\left(\boldsymbol{s}^{t-1}\right)$. For later use this may be summarized as: *if the public belief* B_t^* *is transient, then* $B_t^* = \lambda\left(\boldsymbol{s}^{t-1}\right)$ *and in all previous periods* $\tau < t$, $B_\tau^* = \lambda\left(\boldsymbol{s}^{\tau-1}\right)$ *as well.*

Now the reduction property (Lemma A.1) derived above implies that B_t^* must equal one of the four canonical forms (i) to (iv). As an example, suppose that $B_t^* = \lambda\left(\ell h^d \ell\right)$ where $d < m$ since B_t^* is transient. If the action chosen at time t is H, then the signal s_t must have been h. So by definition of λ, $B_{t+1}^* = \lambda\left(\ell h^d \ell h\right) = \lambda\left(\ell h^d\right)$ which is in $\Lambda(m,n)$. Similarly, if the action chosen at time t is L, then s_t must have been ℓ. Now $B_{t+1}^* = \lambda\left(\ell h^d \ell \ell\right) = \lambda\left(\ell h^{d-1}\ell\right)$ which is also in $\Lambda(m,n)$. It is easy to see that the same argument applies to all canonical forms.

3.4 Public beliefs are Markovian

We have argued above that in transient states, the public belief $B_t^* = \lambda\left(\boldsymbol{s}^{t-1}\right)$, which is as if the outside observer had seen the history of signals and not just the history of actions. Is the process $\lambda\left(\boldsymbol{s}^{t-1}\right) \to \lambda\left(\boldsymbol{s}^{t}\right)$ Markovian? To see that this is not true, notice that even though $\lambda(\ell h) = \lambda(h\ell)$, it is the case that $\lambda(\ell h h) = \frac{\rho}{1-\rho}\frac{\alpha}{\beta}$ whereas $\lambda(h\ell h) = \frac{\rho}{1-\rho}\frac{q}{1-q}$. But public beliefs are not always transient and so it may be that $B_t^* \neq \lambda\left(\boldsymbol{s}^{t-1}\right)$. Our next result is

Proposition 3.2 *Suppose agents follow the strategy* $\boldsymbol{\sigma}^*$*. Then in each state of nature, the public beliefs* B_t^* *form a Markov chain on* $\Lambda(m,n)$*.*

The proof proceeds as follows. We know from Proposition 3.1 above that the public beliefs lie in $\Lambda(m,n)$.

If B_t^* is in an absorbing state, then $\boldsymbol{\sigma}^*$ implies $B_{t+1}^* = B_t^*$ and so the Markov property is obvious.

If B_t^* is in a transient state, then as above $B_t^* = \lambda\left(\boldsymbol{s}^{t-1}\right)$. Suppose there is another signal history such that $\overline{\boldsymbol{s}}^{\tau-1}$ (with τ possibly different from t) such that the corresponding public belief $\overline{B}_\tau^* = B_t^*$. Since $\overline{B}_\tau^*$ must also be transient, this implies that $\lambda\left(\boldsymbol{s}^{t-1}\right) = \lambda\left(\overline{\boldsymbol{s}}^{\tau-1}\right)$. Now Lemma A.2 shows that if two signal histories have the same canonical form, then the last signal in each must be the same, that is, $s_{t-1} = \overline{s}_{\tau-1}$. Note that the probability of a transition from $\boldsymbol{s}^{t-1}$ to $\boldsymbol{s}^t$ depends only on the signal s_{t-1} (because signals are Markovian). Similarly, the probability of a

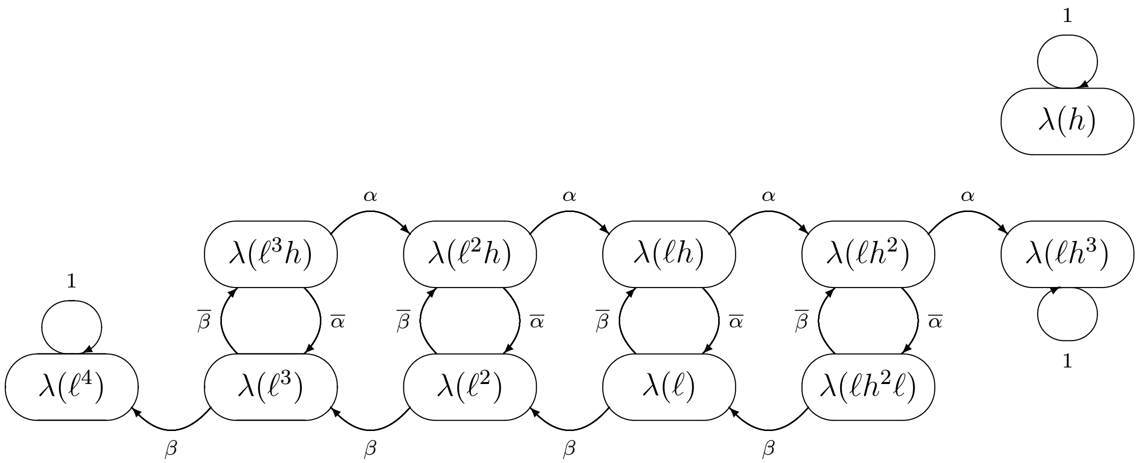

Figure 3: Markov Chain of Public Beliefs in State θ^H

transition from $\overline{\boldsymbol{s}}^{\tau-1}$ to $\overline{\boldsymbol{s}}^{\tau}$ depends only on the signal $\overline{s}_{\tau-1}$. But since, by Lemma A.2, $s_{t-1} = \overline{s}_{\tau-1}$, the probabilities of these transitions must be same. In other words, the history of how B_t^* was reached is irrelevant—the transitions between belief states are Markovian. This completes the proof of Proposition 3.2.

As an example, when $m = 2$ and $n = 3$, the resulting Markov chain of public beliefs in state θ^H is depicted in Figure 3. The belief states $\lambda(h)$, $\lambda(\ell h^3)$ and $\lambda(\ell^4)$ are absorbing. For the remaining, transient, belief states, the transition probabilities are given in the figure using $\overline{\alpha} = 1 - \alpha$ and $\overline{\beta} = 1 - \beta$.

3.5 Optimality of σ^*

Here we establish

Proposition 3.3 *The strategies $\boldsymbol{\sigma}^*$ constitute an equilibrium. This is the unique equilibrium (satisfying the tie-breaking rule that if indifferent, choose H).*

First, note that σ_1^* specifies that agent 1 should follow his signal and this is optimal since $q > \rho$.

Now suppose that agents $1, 2, ..., t-1$ follow $\boldsymbol{\sigma}^*$. We will show that σ_t^* is then a best-response to $\sigma_1^*, \sigma_2^*, ..., \sigma_{t-1}^*$.

Definition 4 *Let $B_t^{**}\left(\boldsymbol{a}^{t-1}, s_t\right)$ denote the belief of agent t after seeing the history of actions $\boldsymbol{a}^{t-1}$ and his private signal $s_t \in \{\ell, h\}$. We will refer to B_t^{**} as the* ***private belief*** *of agent t.*

Note that the private belief of agent t differs from the public belief $B_t^*\left(\boldsymbol{a}^{t-1}\right)$ only in that agent t also knows his private signal s_t. An agent will choose action $a_t = H$ if and only if his private belief $B_t^{**}\left(\boldsymbol{a}^{t-1}, s_t\right) \geq 1$ (ties are broken in favor of H). We begin by deriving these private beliefs and then show that the actions dictated by $\boldsymbol{\sigma}^*$ are optimal.

If the public belief B_t^* is in a *transient* state, then $B_t^*\left(\boldsymbol{a}^{t-1}\right) = \lambda\left(\boldsymbol{s}^{t-1}\right)$. Moreover, since public beliefs must have been transient in all previous periods, agent t can infer $\boldsymbol{s}^{t-1}$ from observing $\boldsymbol{a}^{t-1}$. Thus, the private belief at time t, $B_t^{**}\left(\boldsymbol{a}^{t-1}, s_t\right) = \lambda\left(\boldsymbol{s}^{t-1}, s_t\right)$ and this is the same as B_{t+1}^*, the public belief at time $t+1$.

If the public belief B_t^* is in an *absorbing* state, then let $\tau < t$ be the last period in which B_τ^* was transient. Thus, $B_{\tau+1}^* = \lambda\left(\boldsymbol{s}^\tau\right)$ and since this is absorbing, the public belief stays the same in all subsequent periods. Moreover,

$$\frac{\mathbb{P}^H\left[s_t = \ell \mid s_\tau = h\right]}{\mathbb{P}^L\left[s_t = \ell \mid s_\tau = h\right]} = \frac{\mathbb{P}^H\left[s_t = \ell\right]}{\mathbb{P}^L\left[s_t = \ell\right]} \tag{3}$$

and the same holds if we exchange ℓ and h. When $\tau = t-1$, this is the same as Condition 1 that guarantees that the marginal distribution of signals is preserved over time. To see why this is true even when $\tau < t-1$, it is helpful to recast the signal generation as a renewal process as in Section 2. Recall that in the renewal process interpretation, if $s_t \neq s_\tau$, then it must be that at some point after τ the process restarted with a new independent draw. Let $\tau' \leq t$ be the last period in which the signal was the result of a new draw. Since τ' was the last time there was a new draw, $s_t = s_{\tau'} = \ell$. Now (3) follows.

We (3) in hand, we will derive the private belief of t when the public beliefs are in absorbing states. To be concrete, suppose in fact that the public belief B_t^* is in the absorbing state $\lambda\left(\ell h^{m+1}\right)$. As above, let $\tau < t$ be the last period in which B_τ^* was transient and $s_\tau = h$ (Lemma A.2). Then the private belief of agent t who gets signal $s_t = \ell$ is

$$\begin{aligned} B_t^{**}\left(\boldsymbol{a}^{t-1}, \ell\right) &= B_t^{*}\left(\boldsymbol{a}^{t-1}\right) \times \frac{\mathbb{P}^H\left[s_t = \ell \mid s_\tau = h\right]}{\mathbb{P}^L\left[s_t = \ell \mid s_\tau = h\right]} \\ &= B_t^{*}\left(\boldsymbol{a}^{t-1}\right) \times \frac{\mathbb{P}^H\left[s_t = \ell\right]}{\mathbb{P}^L\left[s_t = \ell\right]} \\ &= \lambda\left(\ell h^{m+1}\right) \times \frac{\mathbb{P}^H\left[s_t = \ell \mid s_{t-1} = h\right]}{\mathbb{P}^L\left[s_t = \ell \mid s_{t-1} = h\right]} \\ &= \lambda\left(\ell h^{m+1} \ell\right) \end{aligned}$$

where the second and third equalities follow because of (3).

In the same manner, if the public belief B_t^* is in the absorbing state $\lambda(h)$, then $B_t^{**}(\boldsymbol{a}^{t-1}, \ell) = \lambda(h\ell) = \rho/(1-\rho)$. Finally, if the public belief B_t^* is in the absorbing state $\lambda(\ell^{n+1})$, then $B_t^{**}(\boldsymbol{a}^{t-1}, s_t = h) = \lambda(\ell^{n+1}h)$.

We now argue that σ_t^* is indeed a best response to $\left(\sigma_1^*, \sigma_2^*, ..., \sigma_{t-1}^*\right)$. We know that public beliefs can only equal one of the four canonical forms. Suppose, for example, $B_t^* = \lambda\left(\ell h^d\right)$ where $d < m+1$ and so transient. If $s_t = h$, then the resulting private belief B_t^{**} is $\lambda\left(\ell h^{d+1}\right) > 1$ and so it is optimal for agent t to choose H. If $s_t = \ell$, then the resulting private belief B_t^{**} is $\lambda\left(\ell h^d \ell\right) < 1$ from the definition of m. Thus, it is optimal for agent t to choose L. Finally, if $d = n+1$ and $s_t = \ell$, then from the argument above, $B_t^{**}(\boldsymbol{a}^{t-1}, \ell) = \lambda(\ell h^{m+1}\ell) > 1$ (again from the definition of m) and so it is optimal for t to choose H. But if it is optimal for t to choose H when $s_t = \ell$, it is also optimal when $s_t = h$. Thus, when the public belief is $\lambda(\ell h^{m+1})$, the optimal response is H regardless of s_t.

The arguments for the optimality of σ_t^* in all remaining belief states are analogous and omitted.

To see that the equilibrium is unique, first notice that for agent 1, σ_1^* is strictly dominant since $q > \rho$. Now for any $t > 1$, given $\left(\sigma_1^*, \sigma_2^*, ..., \sigma_{t-1}^*\right)$, σ_t^* is the unique best response under the assumption that ties are broken in favor of action H.[7]

This completes the proof of Proposition 3.3 and hence also of Theorem 1.

[7] In fact, ties can occur only for a countable values of γ. For all other values, the incentives are (iteratively) strict.

4 Comparative Statics

In what follows we will study how the equilibrium outcomes are affected by a change in the persistence γ of signals. The following simple result (apparent from Figure 2) is key.

Lemma 4.1 *Suppose m and n are determined as in Definitions 2 and 3. Both m and n are non-decreasing functions of γ and unbounded as $\gamma \to 1$.*

Proof. It is easy to verify that

$$\lambda\left(\ell h^{j+1}\ell\right) = \frac{\rho}{1-\rho}\frac{1-q}{q}\frac{1}{\gamma^j}$$

Consider $\gamma < \gamma'$ and let m and m' be determined for γ and γ', as in Definition 2). Now $\lambda\left(\ell h^{m+1}\ell\right) \geq 1$ and $\lambda'(\ell h^{m'+1}\ell) \geq 1$. Since an increase from γ to γ' causes $\lambda\left(\ell h^{m+1}\ell\right)$ to decrease, m' cannot be smaller than m.

A similar argument applies to n.

Finally, since $q > \rho$, for fixed m as $\gamma \to 1$, $\lambda\left(\ell h^{m+1}\ell\right) < 1$. Thus, to compensate $m \to \infty$. ■

Lemma 4.1 implies that if an increase in γ to γ' such that the corresponding $m' > m$ or $n' > n$, then the number of transient states in the Markov chain of public beliefs increase. In other words, when the persistence is $\gamma' > \gamma$, the strategy $\boldsymbol{\sigma}^*$ prescribes that agents wait longer before neglecting their own signals than when the persistence is γ.

To see this, suppose the history of actions consists of $n+1$ choices of L. When the persistence is γ, the resulting public belief is $\lambda\left(\ell^{n+1}\right)$ which is absorbing. Thus, it is optimal even for agent $n+2$ with signal $s_{n+2} = h$ to choose L as well since

$$\lambda\left(\ell^{n+1}h\right) = \frac{\rho}{1-\rho}\gamma^n < 1$$

Thus, a herd to L forms after $n+1$ periods.

But when the persistence is $\gamma' > \gamma$, it is optimal for agent $n+2$ with signal $s_{n+2} = h$ to choose H since now the corresponding public belief is

$$\lambda'\left(\ell^{n+1}h\right) = \frac{\rho}{1-\rho}\left(\gamma'\right)^n \geq 1$$

This is because the same event $\boldsymbol{s}^{n+1} = \ell^{n+1}$ is more indicative of the state θ^L when the persistence is γ than when the persistence is $\gamma' > \gamma$.

The same reasoning applies to the action history consisting of one choice of L followed by $m+1$ choices of H.

5 Welfare Implications

Proposition 3.2 established that in both states of nature, the public beliefs—those of an outside observer who sees only the actions chosen by the agents—form a Markov chain on the set $\Lambda(m,n)$. Since the public beliefs form a Markov chain with the property that from every transient state it is possible to reach an absorbing state in a finite number steps, the process is absorbed with probability one (see for instance, Theorem 11.3 in Grinstead and Snell, 1997). In our context, this means that a herd forms with probability one.

What is the ex ante probability Ψ that an incorrect herd forms—that is, the probability that the absorbing states $\lambda(h)$ and $\lambda(\ell h^{m+1})$ are reached in state θ^L or the absorbing state $\lambda(\ell^{n+1})$ is reached in state θ^H?

One may reasonably conjecture that an increase in persistence would increase the probability of an incorrect herd (decrease welfare). But we show

Theorem 2 *The ex ante probability Ψ that an incorrect herd forms is locally increasing in γ but has a countably infinite number of discontinuous jumps downwards.*

In what follows, we find an explicit expression for Ψ in terms of the basic parameters ρ, q and γ.

To begin, note that the absorbing state, $\lambda(h)$, is isolated—it can only be reached if the signal in the first period $s_1 = h$. On the other hand, if the first-period signal $s_1 = \ell$, then in both states, θ^H and θ^L, the process may reach either of the other two absorbing states—$\lambda(\ell^{n+1})$ or $\lambda(\ell h^{m+1})$—and we wish to determine the relative likelihood of each.

Let H^∞ denote the event (set of signal sequences $\boldsymbol{s}^\infty$) that a herd to H eventually forms—this is the same as the event that $\lim B_t^* = \lambda(\ell h^{m+1})$ or $\lim B_t^* = \lambda(h)$. Similarly, let L^∞ denote the event that a herd to L eventually forms—this is the same as the event $\lim B_t^* = \lambda(\ell^{n+1})$. In the first case, all but a finite number of agents choose H; in the second, all but a finite number choose L.

First, let us determine the probability of an incorrect herd in state θ^L, that is, $\mathbb{P}^L[H^\infty]$. If the first signal $s_1 = h$, then a herd to H occurs with probability one. If $s_1 = \ell$, then both herds occur with positive probability. Thus,

$$\begin{aligned} \mathbb{P}^L[H^\infty] &= \mathbb{P}^L[s_1 = h]\,\mathbb{P}^L[H^\infty \mid s_1 = h] + \mathbb{P}^L[s_1 = \ell]\,\mathbb{P}^L[H^\infty \mid s_1 = \ell] \\ &= (1-q) + q\mathbb{P}^L[H^\infty \mid s_1 = \ell] \end{aligned} \tag{4}$$

To determine $\mathbb{P}^L[H^\infty \mid s_1 = \ell]$ we exploit the *martingale property of likelihood ratios* (public beliefs).[8] This is just the well-known result in statistics that under the null hypothesis, the likelihood ratio is a martingale (see, for instance, Feller (1966, Vol. II), p. 211). If we postulate the null hypothesis that the true state is θ^L, then the martingale property implies that the limiting public beliefs after $s_1 = \ell$ must be such that

$$\mathbb{P}^L[H^\infty \mid s_1 = \ell] \times \lambda\left(\ell h^{m+1}\right) + \mathbb{P}^L[L^\infty \mid s_1 = \ell] \times \lambda\left(\ell^{n+1}\right) = \lambda(\ell) \tag{5}$$

Using the reduction formulae $\lambda(\ell h^{m+1}) = \frac{\rho}{1-\rho}\gamma^{-m}$ and $\lambda(\ell^{n+1}) = \frac{\rho}{1-\rho}\frac{1-q}{q}\gamma^n$ in (5), we obtain that

$$\mathbb{P}^L[H^\infty \mid s_1 = \ell] = \frac{\gamma^m(1-\gamma^n)}{\chi - \gamma^{m+n}} \tag{6}$$

where $\chi = \frac{q}{1-q}$. Substituting (6) into (4) yields

$$\mathbb{P}^L[H^\infty] = (1-q) + q\frac{\gamma^m(1-\gamma^n)}{\chi - \gamma^{m+n}}$$

Finally, by interchanging the roles of q and $(1-q)$ and α and β, we obtain $\mathbb{P}^H[H^\infty]$ and since $\mathbb{P}^H[L^\infty] = 1 - \mathbb{P}^H[H^\infty]$,

$$\mathbb{P}^H[L^\infty] = (1-q)\frac{\gamma^n(\chi - \gamma^m)}{\chi - \gamma^{m+n}}$$

Thus,

Proposition 5.1 *The ex ante probability that an incorrect herd forms is*

$$\Psi \equiv \rho(1-q)\left(r + \frac{r\chi\gamma^m(1-\gamma^n) + \gamma^n(\chi - \gamma^m)}{\chi - \gamma^{m+n}}\right) \tag{7}$$

[8]Its use in our context was suggested by a referee and greatly simplified our earlier direct proof.

where $r = \frac{1-\rho}{\rho}$, $\chi = \frac{q}{1-q}$ *and* m *and* n *are determined according to Definitions 2 and 3.*

Proof. The ex ante probability that an incorrect herd forms is just

$$(1-\rho)\,\mathbb{P}^L\left[H^\infty\right] + \rho\mathbb{P}^H\left[L^\infty\right]$$

and now substituting from the expressions above yields (7). ■

The probability Ψ that an incorrect herd forms is a measure of the asymptotic inefficiency of social learning. Following Rosenberg and Vielle (2022), a utilitarian measure of inefficiency in state θ^L, say, is

$$W^L = \lim_{\delta \to 1} (1-\delta) \sum_{t=1}^{\infty} \delta^t \mathbf{1}_{a_t=H}$$

where $\mathbf{1}_{a_t=H}$ is the indicator of the event $a_t = H$ (the incorrect action for state θ^L) and $\delta < 1$ is a discount factor. By Abel's Theorem, we have that

$$W^L = \lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \mathbf{1}_{a_t=H}$$

and since in our model, a herd to either H or L forms for sure,

$$E_{\theta^L}\left[W^L\right] = \mathbb{P}^L\left[H^\infty\right]$$

Thus, as $\delta \to 1$, our measure of inefficiency is the same as measure of inefficiency used by Rosenberg and Vielle (2022).[9]

5.1 Implications of persistence

To see how Ψ is affected by a change in γ, note that an increase in γ affects Ψ in two ways. First, γ affects Ψ directly. Second, γ affects Ψ indirectly via the induced changes m and n. The first effect increases Ψ and so decreases welfare. The second effect, as will see, goes in the opposite direction and causes a discontinuous increase in welfare.

[9] We note that the model of Rosenberg and Vielle (2022) again assumes that signals are conditionally independent.

To see the first effect, one may differentiate the expression for Ψ in (7) while holding m and n fixed. This yields

$$\frac{d\Psi}{d\gamma} = \rho(1-q)\frac{\chi\left(n\gamma^{n-1}(1-r\gamma^m)(\chi-\gamma^m)+m\gamma^{m-1}(1-\gamma^n)(r\chi-\gamma^n)\right)}{(\chi-\gamma^{m+n})^2}$$

and note that the numerator is positive since $\gamma < 1$, $r = \frac{1-\rho}{\rho} < 1 < r\chi < \chi = \frac{q}{1-q}$.

But while the probability that an incorrect herd forms is locally increasing, it is *not* monotone—Ψ has a countable number of downward discontinuities. In particular, it jumps downwards whenever an increase in γ causes either m or n to increase.

Note that

$$\lambda\left(\ell h^k \ell\right) = \frac{\rho}{1-\rho}\frac{1-q}{q}\frac{1}{\gamma^k} = \frac{1}{r\chi}\frac{1}{\gamma^k}$$

and Definition 2 implies that the change from m to $m+1$ occurs at $\gamma_m = (r\chi)^{-\frac{1}{m}}$. Analogously,

$$\lambda\left(\ell^k h\right) = \frac{\rho}{1-\rho}\frac{1-q}{q}\gamma^k = \frac{1}{r\chi}\gamma^k$$

and Definition 3 implies that the change from n to $n+1$ occurs at $\gamma_n = (r\chi)^{\frac{1}{n}}$.

The next result shows that the ex ante probability of an incorrect herd, Ψ, jumps downward at any γ where there is an increase in either m or n.

Proposition 5.2 *Suppose that γ_n is a point of discontinuity of Ψ such that n increases by 1 at γ_n while m stays fixed. Then there exists an ε such that for all $\gamma' \in (\gamma_n - \varepsilon, \gamma_n)$ and $\gamma'' \in (\gamma_n, \gamma_n + \varepsilon)$,*

$$\Psi(\gamma') > \Psi(\gamma'').$$

The same is true if n is replaced with m and vice-versa.

Proof. It is sufficient to show that for any fixed γ, Ψ is a decreasing function of n. Using (7), this is the same as

$$\frac{r\chi\gamma^m(1-\gamma^n)+\gamma^n(\chi-\gamma^m)}{\chi-\gamma^{m+n}} > \frac{r\chi\gamma^m(1-\gamma^{n+1})+\gamma^{n+1}(\chi-\gamma^m)}{\chi-\gamma^{m+n+1}}$$

and the latter in turn is equivalent to

$$\gamma^n\chi(1-r\gamma^m)(1-\gamma)(\chi-\gamma^m) > 0$$

which holds since $\chi = \frac{q}{1-q} > 1$, $r = \frac{1-\rho}{\rho} < 1$ and $\gamma < 1$.

The proof that for any fixed γ, Ψ is a decreasing function of m is similar. ■

This completes the proof of Theorem 2.

Example 1 Suppose that the prior probability of state θ^H, $\rho = 0.6$ and the quality of the signal $q = 0.7$. Figure 1 in the introduction depicts the ex ante probability of an incorrect herd, that is, the function Ψ. In the first interval of γ's, $m = 1$ and $n = 1$. In the second interval, $m = 2$ and $n = 1$; in the third, $m = 2$ and $n = 2$; in the fourth $m = 3$ and $n = 2$ and so on.

What about the two extremes of independent signals ($\gamma = (1-q)/q$) and perfect persistence ($\gamma = 1$)?

Is the probability of an incorrect minimized when signals are independent? Not always.

Proposition 5.3 *If q is close to ρ, then there exists a $\gamma > \frac{1-q}{q}$ such that the probability of an incorrect herd when the persistence is γ is* smaller *than when signals are independently distributed. Formally,*

$$\Psi\left(\frac{1-q}{q}\right) > \Psi\left(\frac{1-\rho}{\rho}\right)$$

Proof. For all $\gamma \in [\frac{1}{\chi}, r) \cap [\frac{1}{\chi}, \frac{1}{r\chi})$, both $m = 1$ and $n = 1$ and so from (7)

$$\Psi(\gamma) = \rho(1-q)\left(r + \frac{r\chi\gamma(1-\gamma) + \gamma(\chi-\gamma)}{\chi-\gamma^2}\right)$$

Thus, in the independent case, the ex ante probability of an incorrect herd is

$$\Psi\left(\frac{1}{\chi}\right) = \rho(1-q)\left(r + \frac{\chi + r\chi + 1}{\chi^2+\chi+1}\right)$$

Since χ is close to $\frac{1}{r}$, the first discontinuity in Ψ occurs when $m = 1$ and $n = 2$ and so for $\gamma \in [r, r^{\frac{1}{2}})$

$$\Psi(\gamma) = \rho(1-q)\left(r + \frac{r\chi\gamma(1-\gamma^2) + \gamma^2(\chi-\gamma)}{\chi-\gamma^3}\right)$$

and so

$$\Psi(r) = \rho(1-q)\left(r + \frac{2\chi r^2 - r^3 - \chi r^4}{\chi - r^3}\right)$$

It is easy to verify that when $r\chi \approx 1$, $\Psi\left(\frac{1}{\chi}\right) > \Psi(r)$. ■

At the other extreme, Ψ is discontinuous at $\gamma = 1$ (perfect persistence) since

$$\lim_{\gamma \to 1} \Psi(\gamma) = (1-q)(1-q(2\rho-1)) < 1-q = \Psi(1)$$

The probability of an incorrect herd varies non-monotonically with changes in persistence and, as depicted in Figure 1, the function relating the two is highly non-monotonic. One measure of how non-monotonic a function is its the total variation (defined in Appendix C).

Proposition 5.4 *The total variation of the function* $\Psi : \left[\frac{1-q}{q}, 1\right] \to [0,1]$ *is unbounded.*

Proof. See Appendix C. ■

6 Behavioral Implications

In the last section we showed that an increase in the persistence of signals can actually improve welfare. Here we study the behavioral implications of such a change.

6.1 Similarity of actions

Social learning leads agents to take actions that are more similar than dictated by their private information. Herds are an extreme manifestation of this idea—after a while, all agents take the same action. In some contexts, the similarity of actions is harmless. If everyone in society decides to dress in similar fashion, there is no apparent social cost. But in other contexts, the similarity of actions may be detrimental. If all banks hold the same portfolio—heavily weighted toward mortgage backed securities, say—then this increases the systemic risk that the banking sector as whole faces. Even if the same portfolio is optimal for an individual bank, from society's perspective, it would be better if banks' portfolios were dissimilar. These considerations form the basis of what are known as *macroprudential* policies.[10]

In this section we ask how an increase in the similarity (persistence) of *signals* affects the similarity of equilibrium *actions*. We will argue that a small increase in the serial correlation of signals—that causes either m or n to increase by 1—actually causes a *decrease* in the similarity of actions. The reason is simple. Suppose for instance that we are at a point of discontinuity of Ψ. Now suppose a small increase in γ results in a change from n to $n+1$. This delays the formation of a herd—equivalently, this causes public beliefs to remain longer in transient states. This is because there exist signal histories of length greater than $n+1$ such that the public belief reaches the absorbing state $\lambda\left(\ell^{n+1}\right)$ when γ is to the left of the discontinuity but remains transient to the right. In transient states, actions respond to signals and so the probability that tomorrow's action is different from today's is the same as the probability that tomorrow's signal is different from today's.

But an increase in γ also has a countervailing effect—the probability that tomorrow's signal is different from today's decreases. This second effect is small relative to the first when the increase in γ is small. The change from n to $n+1$ is a discontinuous change whereas the change in probabilities is continuous.

Precisely, we have

Proposition 6.1 *Suppose that γ_n is a point of discontinuity of Ψ such that n increases by 1 at γ_n while m stays fixed. Then there exists an ε such that for all $\gamma' \in (\gamma_n - \varepsilon, \gamma_n)$ and $\gamma'' \in (\gamma_n, \gamma_n + \varepsilon)$ and $t > n+1$,*

$$\mathbb{P}_{\gamma'}\left[a'_{t+1} \neq a'_t\right] < \mathbb{P}_{\gamma''}\left[a''_{t+1} \neq a''_t\right]$$

where $\mathbb{P}_{\gamma'}$ and $\mathbb{P}_{\gamma''}$ denote the probability measures on equilibrium actions a'_t and a''_t induced by γ' and γ'', respectively.

[10] A paper from the International Monetary Fund states: "... what constitutes prudent behavior from the point of view of one institution may create broad problems when all institutions engage in similar behavior—whether by selling questionable assets, tightening credit standards, or holding onto cash." (Jacome and Nier, 2012)

6.2 Anything goes

Fix T. Then there exists a γ_T such that for all $\gamma > \gamma_T$ and all T-period action sequences $\boldsymbol{a}^T = (a_1, a_2, ..., a_T)$ such that $a_1 \neq H$ there is a positive probability that $\boldsymbol{a}^T$ is an equilibrium outcome.

The proof is simple. Let γ_T be such that the corresponding m and n both exceed T. This means that as long as $a_1 \neq H$ there is no possibility that in the first T periods, the public beliefs reach either $\lambda(\ell h^{m+1})$ or $\lambda(\ell^{n+1})$. This means that for all $\gamma > \gamma_T$, and all $t \leq T$, the public beliefs remain transient. Let the signal profile $\boldsymbol{s}^T = (s_1, s_2, ..., s_T)$ be such that $s_t = \ell$ if and only if $a_t = L$. Since, in equilibrium, every agent follows his or her signal when the public beliefs are transient, the signal profile $\boldsymbol{s}^T$ results in the action profile $\boldsymbol{a}^T$. Thus with persistent signals *any* observed behavior $\boldsymbol{a}^T$ such that $a_1 \neq H$ can be rationalized (for a large enough γ).

This is not true if information were serially independent ($\gamma = \frac{1-q}{q}$). Now every rationalizable sequence $\boldsymbol{a}^T$ must be one of the following forms: (1) $(H, H, ..., H)$; (2) $(L, L, ..., L)$; (3) $(L, H, H, ..., H)$; (4) $(L, H, L, H, L, ..., H, L, L, L, L, ..., L)$; (5) $(L, H, L, H, L, ..., H, L, H, H, ..., H)$.

Behavior that seems "contrarian"—that is, not following the herd—in a model with independent signals, can in fact be rationalized in a model with Markovian signals.

7 Conclusion

Social learning imposes a social cost. Economic agents rely too much on learning from others' observed behavior and under-weigh or even neglect their own information. This has a cascading effect. If my choices only imperfectly (or not at all) reflect my information, then others who observe these choices will not be able to infer much from doing so. This can lead to choices that are too similar relative to a social optimum. This is not only an inefficient use of available information for the agents, but may be sub-optimal for wider society as well. If all banks hold similar portfolios as a result of "herding", then this is not only inefficient for the banks, but the lack of diversity in investments is too risky for society—a shock to one sector can trigger widespread bank failures. Intuition then suggests that the more correlated agents information is, the more correlated their actions will be. In this paper, we have argued that while this not true in general. To do so, we have developed a tractable model of social learning that allows for the possibility that agents' information is correlated—perhaps because of common sources. Such a model can be used to examine economically relevant questions that could not be examined using existing models.

A Appendix: Reduction Property

Lemma A.1 (Reduction) *Consider a signal history $\boldsymbol{s}^t$ that begins with $s_1 = \ell$. Then $\lambda(\boldsymbol{s}^t)$ must equal one of: (i) $\lambda(\ell^{d+1}h)$; (ii) $\lambda(\ell h^{d+1})$; (iii) $\lambda(\ell^{d+1})$; or (iv) $\lambda(\ell h^{d+1}\ell)$ where, in each case, d is some non-negative integer.*

Proof. First, consider a signal history $\boldsymbol{s}^t$ such that $s_1 = \ell$ and $s_t = h$. Any such history can be written as

$$\boldsymbol{s}^t = \ell^{k_1}h^{j_1}\dots\ell^{k_p}h^{j_p}$$

where each $k_i \geq 1$ and $j_i \geq 1$ and $\sum_{i=1}^p k_i + \sum_{i=1}^p j_i = t$. We then have,

$$\begin{aligned}
&\lambda\left(\ell^{k_1}h^{j_1}\dots\ell^{k_p}h^{j_p}\right)\\
=\;& \frac{\rho}{1-\rho}\underbrace{\frac{1-q}{q}}_{s_1=\ell}\underbrace{\left(\frac{\beta}{\alpha}\right)^{k_1-1}}_{\ell\to\ell}\underbrace{\frac{1-\beta}{1-\alpha}}_{\ell\to h}\underbrace{\left(\frac{\alpha}{\beta}\right)^{j_1-1}}_{h\to h}\times\underbrace{\frac{1-\alpha}{1-\beta}}_{h\to\ell}\underbrace{\left(\frac{\beta}{\alpha}\right)^{k_2-1}}_{\ell\to\ell}\underbrace{\frac{1-\beta}{1-\alpha}}_{\ell\to h}\underbrace{\left(\frac{\alpha}{\beta}\right)^{j_2-1}}_{h\to h}\times\dots\\
&\dots\times\underbrace{\frac{1-\alpha}{1-\beta}}_{h\to\ell}\underbrace{\left(\frac{\beta}{\alpha}\right)^{k_p-1}}_{\ell\to\ell}\underbrace{\frac{1-\beta}{1-\alpha}}_{\ell\to h}\underbrace{\left(\frac{\alpha}{\beta}\right)^{j_p-1}}_{h\to h}
\end{aligned}$$

In the expression above, notice that the probability ratio of each $\ell \to h$ transition is just the inverse of a $h \to \ell$ transition. Similarly, the probability ratio of each $\ell \to \ell$ transition is just the inverse of a $h \to h$ transition. This implies that the formula above reduces to

$$\lambda\left(\ell^{k_1}h^{j_1}\dots\ell^{k_p}h^{j_p}\right) = \frac{\rho}{1-\rho}\frac{1-q}{q}\left(\frac{\beta}{\alpha}\right)^d\frac{1-\beta}{1-\alpha}$$

where $d \equiv \sum_{i=1}^p k_i - \sum_{i=1}^p j_i$. Note that when $d \geq 0$,

$$\lambda\left(\ell^{k_1}h^{j_1}...\ell^{k_p}h^{j_p}\right) = \lambda\left(\ell^{d+1}h\right)$$

and when $d < 0$,

$$\lambda\left(\ell^{k_1} h^{j_1} ... \ell^{k_p} h^{j_p}\right) = \lambda\left(\ell h^{|d|+1}\right)$$

Similarly, consider a signal history $\boldsymbol{s}^t$ such that $s_1 = \ell$ and $s_t = \ell$. Any such history can be rewritten as

$$\boldsymbol{s}^t = \ell^{k_1} h^{j_1} ... \ell^{k_p}$$

and in a manner similar to that above, we obtain

$$\lambda\left(\ell^{k_1} h^{j_1} ... \ell^{k_p}\right) = \frac{\rho}{1-\rho} \frac{1-q}{q} \left(\frac{\beta}{\alpha}\right)^d$$

where $d \equiv \sum_{i=1}^{p} k_i - \sum_{i=1}^{p-1} j_i - 1$. As above, note that when $d \geq 0$,

$$\lambda\left(\ell^{k_1} h^{j_1} ... \ell^{k_p}\right) = \lambda\left(\ell^{d+1}\right)$$

and if $d < 0$,

$$\lambda\left(\ell^{k_1} h^{j_1} ... \ell^{k_p}\right) = \lambda\left(\ell h^{|d|+1} \ell\right)$$

■

The reduction lemma has an important consequence.

Lemma A.2 *Consider two signal histories* $\boldsymbol{s}^t$ *and* $\overline{\boldsymbol{s}}^\tau$ *(of possibly different lengths) which begin with* ℓ. *If* $\lambda(\boldsymbol{s}^t) = \lambda(\overline{\boldsymbol{s}}^\tau)$, *then* $s_t = \overline{s}_\tau$, *that is, they must have the same last signal.*

Proof. The proof of Lemma A.1 demonstrates that if the last signal $s_t = h$, then $\lambda(\boldsymbol{s}^t)$ is either of the form $\lambda\left(\ell^{d+1} h\right)$ or $\lambda\left(\ell h^{d+1}\right)$. Similarly, if the last signal $s_t = \ell$, then $\lambda(\boldsymbol{s}^t)$ is either of the form $\lambda\left(\ell^{d+1}\right)$ or $\lambda\left(\ell h^{d+1} \ell\right)$. Thus, $s_t = h$ if and only if the last signal in its canonical form is also h. The same is true of $\overline{\boldsymbol{s}}^\tau$. ■

B Appendix: Equilibrium

Here we first prove Proposition 3.1 that shows that the public beliefs resulting from the strategy $\boldsymbol{\sigma}^*$ take on a finite set of values, that is, those in $\Lambda(m,n)$.

Proof of Proposition 3.1. The proof is by induction on t. First, $B_2^* \in \Lambda(m,n)$ since it is either $\lambda(h) \in \mathcal{A}$ or $\lambda(\ell) \in \mathcal{T}$.

Suppose, as the induction hypothesis, that in all periods $\tau \leq t$, $B^*_\tau \in \Lambda(m,n) = \mathcal{A} \cup \mathcal{T}$. We will argue that $B^*_{t+1} \in \Lambda(m,n)$ as well.

Case 1. In some period $\tau \leq t$, $B^*_\tau \in \mathcal{A}$.

Now the strategy $\boldsymbol{\sigma}^*$ specifies that in period τ the agent ignore his own signal s_τ. Thus the action a_τ carries no information about s_τ and so the new public belief $B^*_{\tau+1}$ is the same as B^*_τ. Thus, if B^*_τ is ever in an absorbing state it remains so thereafter and hence $B^*_{t+1} \in \mathcal{A}$.

Case 2. In all periods $\tau \leq t$, $B^*_\tau \in \mathcal{T}$.

We first claim that if in any period $\tau \leq t$, $B^*_\tau \in \mathcal{T}$, then $B^*_{\tau+1} = \lambda(\boldsymbol{s}^\tau)$. This is because if $B^*_\tau \in \mathcal{T}$ is transient it must have been transient in all preceding periods and so the actions in the preceding periods must have been revealed. In particular, $B^*_{t+1} = \lambda(\boldsymbol{s}^t)$.

Now, as shown in Appendix A, $\lambda(\boldsymbol{s}^{t-1})$ can be reduced to one of the canonical forms: (i) $\lambda\left(\ell^{d+1}h\right)$; (ii) $\lambda\left(\ell h^{d+1}\right)$; (iii) $\lambda\left(\ell^{d+1}\right)$; and (iv) $\lambda\left(\ell h^{d+1}\ell\right)$ where, in each case, d is some positive integer. Moreover, since $B^*_t = \lambda(\boldsymbol{s}^{t-1}) \in \mathcal{T}$, in cases (i) and (iii), $d < n$ and cases (ii) and (iv), $d < m$.

Notice that by definition

$$\lambda\left(\boldsymbol{s}^t\right) = \lambda\left(\boldsymbol{s}^{t-1}\right) \times \frac{\mathbb{P}^H\left[s_t \mid s_{t-1}\right]}{\mathbb{P}^L\left[s_t \mid s_{t-1}\right]}$$

where $\lambda(\boldsymbol{s}^{t-1})$ is equal to one of the canonical forms (Lemma A.1) with $d < m$ or $d < n$ (because by the induction hypothesis $\lambda(\boldsymbol{s}^{t-1})$ is transient).

Suppose for concreteness the canonical form is $\lambda\left(\ell^{d+1}h\right)$, that is $\lambda(\boldsymbol{s}^{t-1}) = \lambda\left(\ell^{d+1}h\right)$ for some $d < n$. Then

$$\begin{aligned} B^*_{t+1} &= \lambda\left(\boldsymbol{s}^t\right) \\ &= \lambda\left(\boldsymbol{s}^{t-1}\right) \times \frac{\mathbb{P}^H\left[s_t \mid s_{t-1}\right]}{\mathbb{P}^L\left[s_t \mid s_{t-1}\right]} \\ &= \lambda\left(\ell^{d+1}h\right) \times \frac{\mathbb{P}^H\left[s_t \mid s_{t-1}\right]}{\mathbb{P}^L\left[s_t \mid s_{t-1}\right]} \end{aligned}$$

Now Lemma A.2 says that if $\lambda(\boldsymbol{s}^{t-1}) = \lambda\left(\ell^{d+1}h\right)$, then $s_{t-1} = h$. Thus,

$$B^*_{t+1} = \lambda\left(\ell^{d+1}h\right) \times \frac{\mathbb{P}^H\left[s_t \mid s_{t-1} = h\right]}{\mathbb{P}^L\left[s_t \mid s_{t-1} = h\right]}$$

There are two cases to consider. If $s_t = h$, then the public belief $B^*_{t+1} = \lambda\left(\ell^{d+1}h^2\right)$, which reduces to $\lambda\left(\ell^d h\right) \in \Lambda(m,n)$. On the other hand, if $s_t = \ell$, then $B^*_{t+1} = \lambda\left(\ell^{d+1}h\ell\right)$ which reduces to $\lambda\left(\ell^{d+1}\right) \in \Lambda(m,n)$. Thus in both cases, $B^*_{t+1} \in \Lambda(m,n)$.

A similar argument applies to the other three canonical forms. ■

C Appendix: Total Variation of Ψ

The *total variation* of a function $f : [a,b] \to \mathbb{R}$ is

$$TV(f) = \sup_{\mathcal{P}} \sum_{i=0}^{n_P - 1} \left| f(x_{i+1}) - f(x_i) \right|$$

where the supremum runs over the set of all partitions

$$\mathcal{P} = \left\{ \{x_0, x_1, ..., x_{n_P}\} : a = x_0 < x_1 < ... < x_{n_P - 1} < x_{n_P} = b \right\}$$

of $[a,b]$ into sub-intervals.

In this appendix, we show that $TV(\Psi) = \infty$.

We will construct a particular partition $P_N = \{x_0, x_1, ..., x_N\}$ of $\left[\frac{1}{\chi}, 1\right]$ as follows. Let $x_0 = \frac{1}{\chi}$, $x_1 = r$, $x_2 = r^{\frac{1}{2}}, ..., x_n = r^{\frac{1}{n}}, ..., x_N = r^{\frac{1}{N}}$, $x_{N+1} = 1$.

We first show

Lemma C.1 *If $\gamma = r^{\frac{1}{n}}$, then the m corresponding to γ satisfies*

$$n\delta \leq m < n\delta + 1$$

where $\delta = \ln\left(\frac{1}{r\chi}\right) / \ln r > 0$.

Proof. From the definition of m,

$$\left(\frac{1}{r\chi}\right)^{\frac{1}{m-1}} < \gamma \leq \left(\frac{1}{r\chi}\right)^{\frac{1}{m}}$$

The second inequality is equivalent to

$$\begin{aligned} m &\geq \frac{\ln\left(\frac{1}{r\chi}\right)}{\ln\gamma} \\ &= n\frac{\ln\left(\frac{1}{r\chi}\right)}{\ln r} \\ &= n\delta \end{aligned}$$

since, by definition, $\ln\gamma = \frac{1}{n}\ln r$.

Similarly, the first inequality is equivalent to

$$\begin{aligned} m-1 &< n\frac{\ln\left(\frac{1}{r\chi}\right)}{\ln r} \\ &= n\delta \end{aligned}$$

■

By definition, for all γ such that $r^{\frac{1}{n-1}} \leq \gamma < r^{\frac{1}{n}}$

$$\Psi(\gamma) = \rho(1-q)\left(r + \frac{r\chi\gamma^m(1-\gamma^n) + \gamma^n(\chi - \gamma^m)}{\chi - \gamma^{m+n}}\right)$$

where we know from Lemma C.1 that $m < n\delta + 1$. Now the fact that Ψ is decreasing in m implies that

$$\Psi(\gamma) \geq \rho(1-q)\left(r + \frac{r\chi\gamma^{n\delta+1}(1-\gamma^n) + \gamma^n\left(\chi - \gamma^{n\delta+1}\right)}{\chi - \gamma^{n\delta+1+n}}\right)$$

and so,

$$\lim_{\gamma \to r^{\frac{1}{n}}} \Psi(\gamma) \geq \rho(1-q)\, r\left(1 + \frac{\chi r^{\delta+\frac{1}{n}}(1-r) + \left(\chi - r^{\delta+\frac{1}{n}}\right)}{\chi - r r^{\delta+\frac{1}{n}}}\right)$$

For $\gamma = r^{\frac{1}{n}}$,

$$\Psi\left(r^{\frac{1}{n}}\right) = \rho(1-q)\left(r + \frac{r\chi\gamma^m\left(1 - r r^{\frac{1}{n}}\right) + r r^{\frac{1}{n}}(\chi - \gamma^m)}{\chi - \gamma^m r r^{\frac{1}{n}}}\right)$$

where we know from Lemma C.1 that $n\delta \leq m$. Now the fact that Ψ is decreasing in

m implies that

$$\Psi\left(r^{\frac{1}{n}}\right) \leq \rho\left(1-q\right) r\left(1+\frac{\chi r^{\delta}\left(1-r r^{\frac{1}{n}}\right)+r^{\frac{1}{n}}\left(\chi-r^{\delta}\right)}{\chi-r^{\delta} r r^{\frac{1}{n}}}\right)$$

Now define

$$\begin{aligned} D_n &= \lim_{\gamma \uparrow r^{\frac{1}{n}}} \Psi\left(\gamma\right)-\Psi\left(r^{\frac{1}{n}}\right) \\ &\geq \rho\left(1-q\right) r \chi\left(1-r^{\delta}\right)\left(\frac{1-r^{\frac{1}{n}}}{\chi-r^{\delta+\frac{1}{n}+1}}\right) \end{aligned}$$

So

$$\lim_{n \to \infty} n D_n \geq \rho\left(1-q\right) r \chi\left(1-r^{\delta}\right) \lim_{n \to \infty} n\left(\frac{1-r^{\frac{1}{n}}}{\chi-r^{\delta+\frac{1}{n}+1}}\right) > 0$$

By the limit comparison test, since $\sum \frac{1}{n}$ diverges and $\lim_{n\to\infty} nD_n > 0$, $\sum D_n$ also diverges.

References

[1] Banerjee, A. V. (1992). "A Simple Model of Herd Behavior." *Quarterly Journal of Economics,* 107 (3), 797–817.

[2] Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy*, 100 (5), 992–1026.

[3] Bikhchandani, S., Hirshleifer, D., O. Tamuz and Welch, I. (2024). "Information Cascades and Social Learning." *Journal of Economic Literature*, 62 (3), 1040–1093.

[4] Chamley, C. P. (2004). *Rational Herds.* Cambridge University Press.

[5] Feller, W. (1966). *An Introduction to Probability Theory and Its Applications*, Volume II, Wiley.

[6] Graham, J. R. (1999). "Herding Among Investment Newsletters: Theory and Evidence." *Journal of Finance*, 54 (1), 237–268.

[7] Grinstead, C. M. and L. J. Snell (1997). *Introduction to Probability*, American Mathematical Society.

[8] Jácome, L. I., and E. W. Nier (2012). "Macroprudential policy: Protecting the Whole." IMF Finance and Development, 2012.

[9] Kartik, N., S. Lee, T. Liu and D. Rappoport (2024). "Beyond Unbounded Beliefs: How Preferences and Information Interplay in Social Learning." *Econometrica*, 92 (4), 1033–1062.

[10] Moscarini, G., M. Ottaviani, and L. Smith (1998). "Social Learning in a Changing World." *Economic Theory*, 11 (3), 657–665.

[11] Rosenberg, D. and N. Vielle (2022). "On the Efficiency of Social Learning." *Econometrica*, 87(6), pp. 2141–2168.

[12] Schaal, E. and M. Taschereau-Dumouchel (2023). "Herding through Booms and Busts." *Journal of Economic Theory*, 210, p. 105669.

[13] Scharfstein, D. and J. Stein (1990). "Herd Behavior and Investment." *American Economic Review*, 80 (3), 465–479.

[14] Sato, H. and K. Shimizu (2025). "Value of Information in Social Learning." arXiv 2503.05015.

[15] Smith, L. and P. Sørensen (2000). "Pathological Outcomes of Observational Learning." *Econometrica*, 68(2), pp. 371–398.

[16] Smith, L. and P. Sorensen (2011). "Observational Learning." *The New Palgrave Dictionary of Economics Online Edition*, 29–52.

[17] Wang, Wanying (2024). "The Emergence of Fads in a Changing World," Working Paper, Monash University.

[18] Wu, J. (2021). "Correlated Information Cascades." Available at SSRN: https://ssrn.com/abstract=3951879.