

THE EMERGENCE OF FADS IN A CHANGING WORLD

WANYING HUANG[†]

ABSTRACT. We study how fads emerge under social learning in a changing environment. We consider a simple sequential social learning model where rational agents arrive in order, each acting only once, and the underlying unknown state constantly evolves. Each agent receives a private signal, observes all past actions of others, and chooses an action to match the current state. Because the state changes over time, cascades cannot last forever, and actions also fluctuate. We show that despite the rise of temporary information cascades, in the long run, actions change more often than the state. This result provides a theoretical foundation for faddish behavior in which people often change their actions more frequently than necessary.

1. INTRODUCTION

The term “fad” refers to transient behavior that rapidly rises and fades in popularity. Specifically, these rapid behavioral shifts cannot be fully attributed to changes in the underlying fundamentals. For example, in macroeconomics, there are boom-and-bust business cycles that cannot be explained by changes in the underlying economy.¹ Similarly, in finance, it has long been documented that price deviations from an asset’s intrinsic value can stem from speculative bubbles and fads (Camerer, 1989; Aggarwal and Rivoli, 1990). Although the phenomenon of fads is widely observed in many economic activities, the question of how and why they emerge remains unresolved. In this paper, we model fads as excessive behavioral fluctuations relative to changes in the underlying environment and demonstrate how they can emerge under social learning in a dynamic setting.

The pioneering work in the social learning literature (Banerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992, hereafter referred to as BHW) shows that, under appropriate conditions, an *information cascade* will always occur. This is the event in which social information overwhelms agents’ private information, leading them to follow the actions of their predecessors even if their own private information suggests otherwise. At the same time, because these cascades are typically formed based on limited information,

Date: November 19, 2024.

[†]Monash University. Email: kate.huang@monash.edu.

I am indebted to Omer Tamuz for his continued support and encouragement. I am grateful (in alphabetical order) to Krishna Dasaratha, Sumit Goel, Wade Hann-Caruthers, Pawel Janas, Aldo Lucia, Thomas Palfrey, Luciano Pomatto and Jean-Laurent Rosenthal for their helpful comments. All errors are my own.

¹See a recent study by Schaal and Taschereau-Dumouchel (2021) modeling business cycles through the lens of social learning.

they are fragile to small shocks over time. This fragility of cascades is briefly discussed in BHW, where it is shown that the possibility of a one-time change in the underlying environment can lead to “*seemingly whimsical swings in mass behavior without obvious external stimulus*”—a phenomenon they refer to as fads. Inspired by this idea, we formally define fads and study their long-term behavior.

While BHW present an early idea of a fad, they mainly focus on learning in a fixed environment where fads cannot recur indefinitely. In contrast, the recurrence of fads is possible in a changing environment, a setting that has recently attracted some attention (see, e.g., [Dasaratha, Golub, and Hak, 2020](#); [Lévy, Pęski, and Vieille, 2022](#)). Indeed, this is an important setting to study, as many applications of social learning—such as investment, employment, cultural norms, and technological advancement—often operate in a dynamic environment.

In this paper, we study a canonical model of social learning (i.e. a binary state with symmetric and informative binary signals) with a slight twist: in each period, the underlying state switches with a small (and symmetric) probability.² We focus on the long-term behavior of agents, who arrive sequentially and learn from observing the past actions of others as well as their private signals. Each agent acts once and obtains a positive payoff if their action matches the current state. As the underlying state evolves, the optimal action fluctuates as well. The questions we aim to address are: How frequently do actions change? And more specifically, do they change more or less often compared to the underlying state?

Note that in this dynamic environment, the occurrence of information cascades is no longer guaranteed. Intuitively, if the underlying state changes too frequently, the social information extracted from past actions contains little information about the current state, as it primarily reflects the previous state, which is likely to be different from the current one. Indeed, as shown in [Moscarini, Ottaviani, and Smith \(1998\)](#), information cascades occur only if the state changes relatively slowly, and even then, they can only last temporarily. This is because, once a cascade forms, there is no new inflow of information. Consequently, the social information on which the cascade was built depreciates over time and it becomes less relevant to the current agent. Eventually, this social information becomes weak enough so that agents will revert to using their private signals and adjust their actions accordingly. Thus, unlike in a static environment where cascades last indefinitely, when the state evolves slowly, cascades can still arise but they cannot persist indefinitely.

We henceforth focus on the setting in which the state evolves slowly so that temporary cascades can arise; otherwise, all action fluctuations are driven solely by fluctuations in

²This is also known as a simple two-state Markov process. See other studies of social learning with a Markovian state, e.g., [Moscarini, Ottaviani, and Smith \(1998\)](#); [Hirshleifer and Welch \(2002\)](#) and [Lévy, Pęski, and Vieille \(2022\)](#). Our model is mostly similar to that in [Moscarini, Ottaviani, and Smith \(1998\)](#) except for the tie-breaking rule.

private signals, which are clearly more volatile than the state. In this setting, the question of whether actions change more or less frequently than the state is no longer clear. On the one hand, agents sometimes ignore their private signals due to cascades and therefore do not change their actions even when the state changes.³ On the other hand, because the signals are not perfectly informative about the state, agents sometimes change actions unnecessarily. Building on the idea of excessive behavioral changes, we say that *fads emerge* if there are more action changes than state changes. Our main result (Theorem 1) shows that even with the occurrence of temporary cascades, fads emerge almost surely in the long run—that is, actions change excessively relative to the state over time.

For example, consider a private signal that matches the current state 80 percent of the time. When the state changes once every 100 periods on average, we demonstrate that it takes fewer than sixty-one periods for agents to change their actions (see Proposition 1). As a result, the long-term frequency of action changes must exceed that of state changes, leading to the emergence of fads in the long run. It is worth emphasizing that in this model, agents are fully rational, and fads arise from their desire to match the ever-changing state based on the information they have, rather than from any payoff externalities between agents or heuristic-driven behavior.

The proof strategy behind the long-term emergence of fads is as follows. First, for any fixed signal precision and probability of state change, there exists a maximum length for each temporary information cascade. Consequently, although the presence of temporary cascades prolongs action inertia, this effect is limited by their bounded length. Furthermore, once agents exit a cascade, they only need one opposing signal to change their actions. Together, these observations allow us to establish a lower bound on the probability of an action change, which in turn provides an upper bound for the expected time between these changes. We then show that this upper bound is less than the expected time between state changes, implying that action changes occur more frequently than state changes on average. Finally, by translating the expected time between changes for both the state and the action into their long-term frequency of changes, we conclude that fads emerge in the long run.

1.1. Related Literature. This paper is closely related to a small stream of studies on social learning in a changing state. In BHW, they briefly discuss the case where a one-time shock to the state could break the cascade, even though that shock may never be realized. They provide a numerical example in which the probability of an action change is at least 87% higher than the probability of a state change (see their Result 4), which aligns with our main result. Later, Moscarini, Ottaviani, and Smith (1998) show that if the underlying state evolves in every period and is sufficiently persistent, an information cascade must arise, but it only lasts temporarily, i.e., it must end in finite time. Our

³The symmetry of binary states further amplifies this effect: consider a scenario in which the state has changed an even number of times, say twice. However, an agent in a cascade would mistakenly perceive the state as unchanged and therefore have no reason to change their action.

work builds on their model but with a different focus. Instead of analyzing the short-term patterns of information cascades, such as the conditions under which they arise or end, we ask: in the long run, should one expect more volatility in actions or in the state?

In a setting where a single agent repeatedly receives private signals, [Hirshleifer and Welch \(2002\)](#) examine the effect of *memory loss*—a situation in which the agent only recalls past actions but not past signals—on the continuity of the agent’s behavior. They analyze the equilibrium behavior of a five-period stylized model and show that in a relatively stable environment, memory loss induces excessive action inertia compared to a full-recall regime. In contrast, in a more volatile environment, memory loss results in excessive action impulsiveness.⁴ Different from their work—which examines how an agent’s ability to recall past actions and signals affects action fluctuations—our study investigates the long-term fluctuations in the actions of agents who observe past actions in a changing environment and compares these fluctuations to those of the underlying state.

Among a few more recent studies that consider a dynamic state, the efficiency of learning has been a primary focus of study. For example, [Frongillo, Schoenebeck, and Tamuz \(2011\)](#) consider a specific environment in which the underlying state follows a random walk with non-Bayesian agents who use different linear rules when updating. Their main result is that the equilibrium updating weights may be Pareto suboptimal, causing inefficiency in learning.⁵ In a similar but more general environment, [Dasaratha, Golub, and Hak \(2020\)](#) show that having sufficiently diverse network neighbors with different signal distributions improves learning. Their idea is that having diverse signals enables agents to extract the most relevant information from the old and confounded data, thereby achieving higher efficiency in information aggregation.

In a setup similar to ours, a recent study by [Lévy, Pęski, and Vieille \(2022\)](#) considers the welfare implication of a dynamic state. In their model, agents observe a random subsample drawn from all past actions and then decide whether to acquire private signals that are potentially costly. These model generalizations allow them to highlight the trade-off between learning efficiency and responsiveness to environmental changes in maximizing equilibrium welfare. In contrast, we assume that agents observe the full history of past actions and there is no cost associated with obtaining their private signals. We consider this canonical sequential learning model without further complications as our focus is on comparing the long-term relative frequency of action and state changes—a question that turns out to be nontrivial even in this simple setup.

⁴Intuitively, as volatility of the environment increases, past actions become less relevant to the current state. At some point, this information weakens enough so that the amnesiac agent would always follow her latest signal, but the full-recall agent may not do so at this point. Hence, there is an increase in the probability of an action change due to amnesia.

⁵See more studies in the computer science literature, e.g., [Acemoglu, Nedic, and Ozdaglar \(2008\)](#); [Shahrampour, Rakhlin, and Jadbabaie \(2013\)](#) that consider a dynamic environment with non-Bayesian agents.




FIGURE 1. An illustration of processes (θ_t) and (a_t) .

2. MODEL

We follow the setup from [Moscarini, Ottaviani, and Smith \(1998\)](#) closely. Time is discrete, and the horizon is infinite, i.e., $t \in \mathbb{N}_+ = \{1, 2, \dots\}$. There is a binary state $\theta_t \in \{+1, -1\}$ that evolves over time. Specifically, the state evolves according to a Markov chain with a symmetric transition probability $\varepsilon \in (0, 1)$:

$$\mathbb{P}[\theta_{t+1} \neq i | \theta_t = i] = \varepsilon, \text{ for } i \in \{+1, -1\}.$$

For simplicity, we assume that both states are equally likely at the beginning of time. Note that this uniform distribution is also the stationary distribution of θ_t .

A sequence of short-lived agents indexed by time t arrive in order, each acting once by choosing an action $a_t \in \{+1, -1\}$. For each agent t , she obtains a payoff of one if her action matches the current state, i.e., $a_t = \theta_t$ and zero otherwise. Before choosing an action, she receives a binary private signal $s_t \in \{+1, -1\}$ and observes the history of all past actions made by her predecessors, $h_{t-1} = (a_1, \dots, a_{t-1})$. Conditional on the entire sequence of states, the private signals (s_t) are independent, and each s_t follows a Bernoulli distribution $B_{\theta_t}(\alpha)$ where $\alpha \in (1/2, 1)$ is the probability of matching the current state:

$$\mathbb{P}[s_t = i | \theta_t = i] = \alpha, \text{ for } i \in \{+1, -1\}.$$

We assume throughout that the state is *sufficiently persistent*: for any signal precision $\alpha \in (1/2, 1)$, the probability of a state change $\varepsilon \in (0, \alpha(1 - \alpha))$. Under this assumption, [Moscarini, Ottaviani, and Smith \(1998\)](#) show that information cascades can occur, but they only last temporarily. Equivalently, one can think of this assumption as follows: in every period, with probability $2\varepsilon \in (0, 2\alpha(1 - \alpha))$ the state will be redrawn from the set $\{+1, -1\}$ with equal probability. Thus, the probability of a state change is equal to $\varepsilon \in (0, \alpha(1 - \alpha))$.

At any time t , the sequence of events is as follows. First, the agent arrives and observes the history of all past actions, h_{t-1} . Second, the state θ_{t-1} transitions to θ_t with a probability ε of switching. After the state transitions, agent t receives a private signal s_t that matches the current state θ_t with probability α . Finally, she chooses an action a_t that maximizes the probability of matching θ_t , conditional on (h_{t-1}, s_t) the information available to her. A graphical illustration of the sequence of events is shown in [Figure 1](#).

2.1. **Fads.** Given that each agent aims to match the current state, as the state evolves, the best action to take also fluctuates. BHW informally discuss the idea of faddish behavior as a situation where action changes occur more frequently than state changes. In other words, fads represent scenarios where there are *excessive* action changes relative to the state. To formalize this idea, we denote the fraction of time periods $t \leq n$ for which $a_t \neq a_{t+1}$ and $\theta_t \neq \theta_{t+1}$ by

$$\mathcal{Q}_a(n) := \frac{1}{n} \sum_{t=1}^n \mathbb{1}(a_t \neq a_{t+1}) \quad \text{and} \quad \mathcal{Q}_\theta(n) := \frac{1}{n} \sum_{t=1}^n \mathbb{1}(\theta_t \neq \theta_{t+1}), \text{ respectively.}$$

Formally, we say that *fads emerge at time $n + 1$* if

$$\mathcal{Q}_a(n) > \mathcal{Q}_\theta(n). \tag{2.1}$$

Multiplying both sides of (2.1) by n , the emergence of fads at time $n + 1$ means that actions have changed more frequently than the state by time $n + 1$.

2.2. **Agents' Beliefs.** Let $p_t := \mathbb{P}[\theta_t = +1 | h_{t-1}, s_t]$ denote the *posterior belief* of agent t that the state is positive after observing (h_{t-1}, s_t) the pair of action history and private signal. The log-likelihood ratio (LLR) of agent t 's posterior beliefs of the state being $+1$ over the state being -1 is

$$L_t = \log \frac{p_t}{1 - p_t} = \log \frac{\mathbb{P}[\theta_t = +1 | h_{t-1}, s_t]}{\mathbb{P}[\theta_t = -1 | h_{t-1}, s_t]}.$$

We call L_t the *posterior LLR* at time t . By Bayes' rule, the posterior LLR at time t is equal to

$$L_t = \log \frac{\mathbb{P}[\theta_t = +1 | h_{t-1}]}{\mathbb{P}[\theta_t = -1 | h_{t-1}]} + \log \frac{\mathbb{P}[s_t | \theta_t = +1, h_{t-1}]}{\mathbb{P}[s_t | \theta_t = -1, h_{t-1}]}.\tag{2.2}$$

We refer to the first term in (2.2) as the *public LLR* at time t and denote it by ℓ_t . This is the log-likelihood ratio of the *public belief* of agent t after observing only the history of actions h_{t-1} , which we denote by $q_t := \mathbb{P}[\theta_t = +1 | h_{t-1}]$. Since the private signal is independent of the history of actions conditional on the current state, the second term in (2.2) reduces to the LLR induced by the signal itself, which is equal to $c_\alpha := \log \frac{\alpha}{1-\alpha}$ if $s_t = +1$ and $-c_\alpha$ if $s_t = -1$. Therefore, depending on the realization of the private signal, the posterior LLR at time t evolves as follows:

$$L_t = \begin{cases} \ell_t - c_\alpha & \text{if } s_t = -1, \\ \ell_t + c_\alpha & \text{if } s_t = +1. \end{cases} \tag{2.3}$$

2.3. **Agents' Behavior.** The optimal action for agent t is the action that maximizes her expected payoff conditional on the information available to her:

$$a_t \in \arg \max_{a \in \{-1, +1\}} \mathbb{P}[\theta_t = a | h_{t-1}, s_t].$$

Thus $a_t = +1$ if $L_t > 0$ and $a_t = -1$ if $L_t < 0$. When $L_t = 0$, agent t is indifferent between both actions. We assume that in this case, she would follow what her immediate predecessor did in the previous period, i.e., $a_t = a_{t-1}$.⁶ This tie-breaking rule differs from the one used in Moscarini, Ottaviani, and Smith (1998), where indifferent agents are assumed to follow their own private signals. We make this assumption so that any action changes are driven by agents' strict preference for one action over another, rather than by the specific choice of the tie-breaking rule.

2.4. Cascade and Learning Regions. As mentioned before, an *information cascade* occurs when the social information inferred from others' past actions outweighs an agent's private signal, leading the agent to disregard their private information. From (2.3), we see that when $|\ell_t| > c_\alpha$, the sign of L_t —and therefore the optimal action for agent t —is determined solely by the sign of ℓ_t , regardless of s_t . Thus, in this case, $a_t = +1$ if $\ell_t > c_\alpha$ and $a_t = -1$ if $\ell_t < -c_\alpha$.

Similarly, when $|\ell_t| = c_\alpha$, the tie breaking rule at indifference implies that agent t will choose the same action as agent $t - 1$, regardless of the realization of s_t . Furthermore, note that a_{t-1} is equal to the sign of ℓ_t .⁷ Thus, we refer to the region of the public LLR where $|\ell_t| \geq c_\alpha$ as the *cascade region*. Conversely, when $|\ell_t| < c_\alpha$, agent t chooses the action according to her private signal (i.e., $a_t = s_t$), and we refer to this region as the *learning region*.

3. RESULTS

3.1. A Benchmark. As a benchmark, we briefly discuss the case in which each short-lived agent only observes her own private signal but not the actions of her predecessors. In this scenario, agent t simply follows her private signal s_t , as it is her only source of information about the state.⁸ Consequently, in the long run, agents' actions fluctuate as frequently as their private signals. By the strong law of large numbers,

$$\lim_{n \rightarrow \infty} \mathcal{Q}_a(n) = \mathbb{P}[s_t \neq s_{t+1}] \quad \text{almost surely.}$$

A straightforward calculation then shows that

$$\mathbb{P}[s_t \neq s_{t+1}] = (1 - \alpha^2)(1 - \varepsilon) + \alpha^2\varepsilon.$$

First, observe that the above probability is strictly greater than ε , which is the long-term frequency of state changes (see details in Section 4.2). This is intuitive since in

⁶Our results do not depend on this assumption and are robust to any tie-breaking rule that is common knowledge.

⁷To see this, suppose without loss of generality that $\ell_t = c_\alpha$, so $\text{sign}(\ell_t) = +1$. If $s_t = +1$, then $L_t = \ell_t + c_\alpha > c_\alpha$, which leads to $a_t = +1$. If $s_t = -1$, then $L_t = \ell_t - c_\alpha = 0$ and by the tie-breaking rule, $a_t = a_{t-1}$. In either case, $a_t = a_{t-1}$, which is equal to $\text{sign}(\ell_t) = +1$. This is because if $a_{t-1} = -1$, it would follow from $\ell_t = c_\alpha$ that $\ell_{t-1} > c_\alpha$, which implies $a_{t-1} = +1$. A contradiction.

⁸This case is behaviorally equivalent to a scenario in which the state lacks sufficient persistence, thereby preventing temporary information cascades and causing agents to rely solely on their private signals.

this benchmark, agents always follow their signals, so actions exhibit the same level of volatility as the signals, which are more volatile than the state itself. Additionally, as the private signal becomes less noisy (i.e., as $\alpha \rightarrow 1$), the long-run frequency of action changes converges to that of state changes. This is because, as signals become increasingly precise, unnecessary changes in actions decrease. In the limit, when signals are perfectly informative, actions change as often as the state in the long run.

3.2. Public Actions. Now, we turn to our main setting, where each short-lived agent observes both her private signal and the actions of her predecessors. Since the state is sufficiently persistent, information cascades arise temporarily, during which agents mimic their predecessors rather than respond to their private signals. As a result, compared to the benchmark, action changes are less frequent. Our main result shows that, despite these periods of cascade-driven inaction, fads still almost surely emerge in the long run. Recall that in (2.1) we defined the emergence of fads at time $n + 1$ as a higher relative frequency of action changes compared to state changes.

Theorem 1. *For any signal precision $\alpha \in (1/2, 1)$ and probability of state change $\varepsilon \in (0, \alpha(1 - \alpha))$, fads emerge in the long run almost surely:*

$$\lim_{n \rightarrow \infty} Q_a(n) > \lim_{n \rightarrow \infty} Q_\theta(n) \quad \text{almost surely.}$$

Thus Theorem 1 shows that in the long run, fads are guaranteed to emerge under social learning, even when the underlying state evolves slowly and there are temporary information cascades during which agents do not change their actions. As an example, consider a private signal that matches the current state 80 percent of the time. When the probability of a state change is 1 percent, the state changes on average once every 100 periods. Meanwhile, the average time between action changes is strictly less than 61 periods.⁹ This implies that, over time, action changes will occur more frequently than state changes, leading to faddish behavior.

The idea behind the proof of Theorem 1 is as follows. Intuitively, as the state evolves, the older social information on which a cascade is built becomes less relevant to the current agent. Consequently, agents periodically stop mimicking past actions and begin responding to their private information. This periodic responsiveness to private signals drives fluctuations in actions since agents at these times are susceptible to opposing news. Formally, we show that once the agent's public belief exits the cascade region, their action either changes or the public belief re-enters the same cascade region. We upper bound the probability of the latter event, thus providing an upper bound for the expected time between action switches (Proposition 1). We then compare this upper bound to the expected time between state changes and show that the former is strictly less than the latter. Finally, our main result (Theorem 1) translates the expected times between action

⁹This follows from Proposition 1 in Section 4.2 by substituting $\alpha = 0.8$ and $\varepsilon = 0.01$ into $M(\alpha, \varepsilon)$, yielding $M(0.8, 0.01) \approx 60.7$.

$\alpha \setminus \varepsilon$	0.05	0.1	0.2
0.51	16,766 (5,081)	28,564 (10,055)	42,128 (20,024)
0.75	15,240 (5,100)	26,149 (10,034)	–
0.9	14,252 (5,096)	–	–

TABLE 1. The numerical simulations show the number of action and state changes (values in parentheses) across different values of α and $\varepsilon \in (0, \alpha(1 - \alpha))$ over 100,000 periods.

and state changes into their long-run frequencies and concludes that actions change more frequently than the state.

Note that action changes are not independent events, and consequently, we cannot directly establish the connection between the expected time between changes and its long-run frequency using the standard law of large numbers. To address this, we study the process of the (random) time interval between action changes, which as shown in Lemma 2 in the appendix, has well-defined moments. In addition, since the state changes over time, the agent’s public belief also ceases to be a martingale—an important tool in analyzing the long-term outcome of learning in a fixed state model.¹⁰ Nevertheless, it is still a Markov process. Given the Markov property, we analyze the transitional patterns of the public belief across different regions (see Lemma 1) and use them to study the expected time between action switches.

We assume a sufficiently persistent state for two reasons. First, as discussed before, when the state is not sufficiently persistent—so that temporary information cascades never arise—agents would always follow their signals and change their actions accordingly. Hence, as in our benchmark, action clearly changes more frequently than the state. Second, even with a sufficiently persistent state, it is a priori unclear whether the state or the action would change more often. As the likelihood of state changes decreases, action changes also slow down. This is because when state changes are less frequent, past actions become more informative about the current state; consequently, cascades tend to last longer, leading to extended periods of action inertia. Our main result suggests that, over time, this prolonged action inertia is eventually overtaken by action impulsiveness, resulting in excessive changes in actions relative to the state.

3.3. Numerical Simulations. To illustrate our main result, we simulate the frequencies of action and state changes under different values of signal precision (α) and state volatility (ε) over 100,000 periods. Table 1 displays the results of these simulations.

These simulations confirm our main result: for all pairs of parameter values considered, we see that action changes are more frequent than state changes. Next, we explore how the frequencies of action and state changes vary across different parameter values. In the first column of Table 1, when the probability of state change is fixed at 0.05, the

¹⁰For example, it is essential in proving asymptotic learning (Smith and Sørensen, 2000) for unbounded signals.

frequency of action changes decreases as signal precision increases, aligning with the intuition that more precise signals help reduce unnecessary action changes. However, even with high signal precision ($\alpha = 0.9$), actions still change at least twice as often as the state, underscoring the persistence of faddish behavior despite increasingly informative signals. Conversely, as state volatility increases, both action and state changes become more frequent, as shown in the first row of Table 1. Yet, the ratio of action to state changes declines from approximately 3.3 to 2.1, suggesting that state volatility directly impacts the frequency of state changes more than that of action changes.

4. ANALYSIS

In this section, we analyze how the agent’s public belief evolves in both the learning and cascade regions. These dynamics allow us to establish an upper bound on the expected time between action changes. We then compare this upper bound to the expected time between state changes and show that the former is strictly less than the latter.

4.1. The Belief Dynamics.

Cascade Region. It is well-known that in this model, if the state is fixed ($\varepsilon = 0$), an information cascade will be triggered and, once triggered, will last indefinitely. This is because once the agent’s public belief enters the cascade region, it remains there, as all subsequent agents face the same problem as the initial agent who started the cascade (Banerjee, 1992; Bikhchandani et al., 1992). Since signals are binary and imperfectly informative, the resulting cascade is formed based on limited information and thus can be incorrect with positive probability.¹¹

However, if the state is changing ($\varepsilon > 0$), the behavior of the agent’s public belief becomes more complex. To see this, consider the case where the public LLR at time t satisfies $\ell_t \geq |c_\alpha|$, and suppose t is the time at which the public LLR first enters the cascade region from the learning region. In this case, agent t follows the action of her immediate predecessor, so a_t contains no additional information about θ_t beyond what a_{t-1} provides. Meanwhile, between time t and $t+1$, the state may change with probability ε . Since θ_t follows a Markov chain, conditional on θ_t , the history h_t provides no further information about θ_{t+1} . Thus, while ℓ_t remains in the cascade region, the corresponding public belief updates deterministically as follows:

$$\begin{aligned} q_{t+1} &= \mathbb{P}[\theta_{t+1} = +1|h_t] = (1 - \varepsilon)q_t + \varepsilon(1 - q_t) \\ &= (1 - 2\varepsilon)q_t + (2\varepsilon)\frac{1}{2}. \end{aligned} \tag{4.1}$$

¹¹More generally, in a fixed state model with non-binary signals, whether agents eventually all choose the correct action depends on whether the private signals are unbounded or bounded. Our case with binary signals and $\alpha \in (1/2, 1)$ is a special case of bounded signals.

Equivalently, we can write it in terms of the public LLR:

$$\ell_{t+1} = \log \frac{q_{t+1}}{1 - q_{t+1}} = \log \frac{(1 - \varepsilon)e^{\ell_t} + \varepsilon}{1 - \varepsilon + \varepsilon e^{\ell_t}}. \quad (4.2)$$

From (4.1), we observe that q_{t+1} tends toward $1/2$, so eventually, ℓ_{t+1} exits the cascade region. Intuitively, having a changing state depreciates the value of older social information, as actions observed in earlier periods become less relevant to the current agent. Consequently, after a finite number of periods, the agent's public belief gradually converges to $1/2$, so that a cascade supported by this belief eventually ceases. This is the main insight from Moscarini, Ottaviani, and Smith (1998), where they show that information cascades, if they arise, can only be temporary under a changing state.

Learning Region. Next, we consider the learning region where $|\ell| < c_\alpha$. If the state is fixed ($\varepsilon = 0$), then the agent at time t simply follows their private signal: $a_t = s_t$. As a result, at time $t + 1$, the agent's public belief coincides with their posterior belief:

$$q_{t+1} = \mathbb{P}[\theta = +1 | h_{t-1}, a_t] = \mathbb{P}[\theta = +1 | h_{t-1}, s_t] = p_t.$$

Hence, the corresponding log-likelihood ratios also coincide, i.e., $\ell_t = L_t$, and so ℓ_t also evolves according to (2.3).

In contrast, if the state changes with probability $\varepsilon > 0$ in every period, upon observing the latest history, each agent needs to consider the possibility that the state may have changed after the most recent action was taken. However, neither the learning nor the cascade region is affected by a changing state as the state only transitions after the history of past actions is observed. It follows from Bayes' rule that

$$\begin{aligned} \ell_{t+1} &= \log \frac{q_{t+1}}{1 - q_{t+1}} = \log \frac{\mathbb{P}[\theta_{t+1} = +1 | h_{t-1}, a_t]}{\mathbb{P}[\theta_{t+1} = -1 | h_{t-1}, a_t]} \\ &= \log \frac{\sum_{i \in \{-1, +1\}} \mathbb{P}[\theta_{t+1} = +1, a_t | h_{t-1}, \theta_t = i] \cdot \mathbb{P}[\theta_t = i | h_{t-1}]}{\sum_{i \in \{-1, +1\}} \mathbb{P}[\theta_{t+1} = -1, a_t | h_{t-1}, \theta_t = i] \cdot \mathbb{P}[\theta_t = i | h_{t-1}]}. \end{aligned} \quad (4.3)$$

Since θ_t follows a Markov chain and $a_t = s_t$ in the learning region, conditional on θ_t , both θ_{t+1} and a_t are independent of the history h_{t-1} and of each other. Therefore, the public LLR evolves as follows:

$$\ell_{t+1} = \begin{cases} \log \frac{(1-\varepsilon)\alpha e^{\ell_t} + \varepsilon(1-\alpha)}{\varepsilon\alpha e^{\ell_t} + (1-\varepsilon)(1-\alpha)} := f_+(\ell_t) & \text{if } s_t = +1, \\ \log \frac{(1-\varepsilon)(1-\alpha) e^{\ell_t} + \varepsilon\alpha}{\varepsilon(1-\alpha) e^{\ell_t} + (1-\varepsilon)\alpha} := f_-(\ell_t) & \text{if } s_t = -1. \end{cases} \quad (4.4)$$

Note that both f_+ and f_- are strictly increasing. This means that conditional on receiving any signal, the agent with a higher public belief has a higher posterior belief. In addition, it is easy to check that $f_+(\ell) > \ell$ and $f_-(\ell) < \ell$. This indicates that, relative to her prior, an agent's posterior belief increases when she receives a positive signal and decreases when she receives a negative signal.

From (4.4), we see that the magnitude difference between ℓ_t and ℓ_{t+1} depends on both the realization of the private signal s_t and the current value of ℓ_t . The following lemma summarizes the transitional patterns of the public LLR when it is in the learning region. At any time t , we say that an action is *opposing* to the current public belief if $a_t \neq \text{sign}(\ell_t)$ and *supporting* otherwise. The following lemma is in spirit close to the *overturning principle* (Smith and Sørensen, 2000), but it applies to a changing state.

Lemma 1. *For any ℓ_t such that $|\ell_t| < c_\alpha$, the following two conditions hold.*

- (i) $a_t \neq \text{sign}(\ell_t)$ implies that $\text{sign}(\ell_{t+1}) = -\text{sign}(\ell_t)$.
- (ii) $a_t = a_{t+1} = \text{sign}(\ell_t)$ implies that $|\ell_{t+2}| \geq c_\alpha$.

The first part of this lemma states that a single opposing action is sufficient to overturn the sign of the public LLR. The second part indicates that initiating a cascade requires at most two supporting actions. Intuitively, because the public LLR in the learning region tends to be moderate, it is sensitive to opposing evidence; at the same time, although the public belief adjusts more conservatively due to the possibility of a changing state, consecutive observations of supporting evidence are still sufficient to trigger a cascade.

Another important observation is that, regardless of whether the state is fixed or changing, the process (ℓ_t) forms a Markov chain.¹² In the case of a fixed state, the state space of this Markov chain is finite since the magnitude difference between ℓ_t and ℓ_{t+1} is a constant for any given signal precision. However, in the case of a changing state, the state space becomes infinite, as these magnitude differences also depend on the current value of ℓ_t .¹³ This poses a significant challenge in finding its stationary distribution, which is required to calculate the exact expected time between sign switches. We circumvent this problem by providing an upper bound to this expected time instead.

4.2. Expected Time Between Switches. We first calculate the expected time between state changes. Since θ_t follows a simple two-state Markov chain with a symmetric transition probability ε , the expected time between state changes is inversely proportional to the likelihood of a state change. To illustrate this, let x represent the expected time between state changes. Then x satisfies the following equation:

$$x = \varepsilon + (1 - \varepsilon)(1 + x),$$

which implies that $x = 1/\varepsilon$. That is, a higher likelihood of state changes corresponds to a shorter average time between changes.

To calculate the expected time between action changes, note that while a_t is not a Markov chain, it is a function of a Markov chain; more specifically, we have $a_t = \text{sign}(\ell_{t+1})$. However, as discussed earlier, this Markov chain is complicated—with infinitely many possible values and different transition probabilities—making it difficult to directly analyze

¹²This is because conditional on the state θ_t , the private signal s_t is independent of ℓ_τ , for any $\tau < t$.

¹³In fact, in almost all cases, two consecutive opposing signals do not exactly offset each other, i.e., $f_+(f_-(\ell)) \neq \ell$ and vice versa.

the expected time between its sign switches. Therefore, we establish an upper bound, which also bounds the expected time between action changes. To do so, consider the maximum length of any cascade. Recall that such a maximum exists because the public belief in the cascade region slowly converges toward uniformity. Meanwhile, for any given signal precision and probability of a state change, no cascade can last longer than the one starting at $f_+(c_\alpha)$, which is the supremum of the public LLR. Hence, from (4.2), we can calculate a tight upper bound on the length of any cascade.

Following Moscarini, Ottaviani, and Smith (1998), we denote this upper bound by $K(\alpha, \varepsilon)$ where¹⁴

$$K(\alpha, \varepsilon) = \frac{\log(1 - 2\alpha(1 - \alpha))}{\log(1 - 2\varepsilon)}.$$

It is straightforward to see that the upper bound $K(\alpha, \varepsilon)$ decreases in both α and ε . As private signals become less precise, cascades contain more information relative to private signals, potentially extending the duration of a cascade. Meanwhile, as the state becomes less volatile, temporary cascades last longer because social information depreciates at a lower rate. Together, this suggests that prolonged cascades may result from either less precise private signals or a more stable environment.

Now, we are ready to provide an upper bound for the expected time between sign switches of the public LLR. For any positive integer $i = 1, 2, \dots$, let \mathcal{T}_i denote the random time at which the public LLR switches its sign for the i -th time, with $\mathcal{T}_0 = 0$. Let $\mathcal{D}_i = \mathcal{T}_i - \mathcal{T}_{i-1}$ denote the (random) time elapsed between the $(i - 1)$ -th and i -th sign switch.

Proposition 1. *For any positive integers $i \geq 2$, conditional on the public LLR switching its sign for the $(i - 1)$ -th time, the expected time to the next sign switch satisfies*

$$\mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] < 1 + \frac{K(\alpha, \varepsilon)}{2\alpha(1 - \alpha)} =: M(\alpha, \varepsilon).$$

This proposition states that, on average, the public LLR is expected to change its sign at least once every $M(\alpha, \varepsilon)$ periods. For example, when $\alpha = 0.8$ and $\varepsilon = 0.01$, $M(0.8, 0.01)$ is approximately 61, indicating that the public LLR experiences at least one sign switch every 61 periods. Furthermore, notice that $M(\alpha, \varepsilon)$ decreases in α , so $M(1/2, \varepsilon)$ is the maximum upper bound for any $\varepsilon \in (0, \alpha(1 - \alpha))$. Intuitively, when

¹⁴For completeness, we provide a similar calculation of $K(\alpha, \varepsilon)$ as in Section 3.B of Moscarini, Ottaviani, and Smith (1998). Fix any arbitrary $\alpha \in (1/2, 1)$ and $\varepsilon \in (0, \alpha(1 - \alpha))$. Denote m as the supremum of public belief, where $m = \frac{(1-\varepsilon)\alpha^2 + \varepsilon(1-\alpha)^2}{\alpha^2 + (1-\alpha)^2}$. Since the public belief in a cascade evolves deterministically according to (4.1), after h periods, the public belief starting at m equals

$$g(h) := \varepsilon \sum_{i=1}^{h-1} (1 - 2\varepsilon)^i + (1 - 2\varepsilon)^h m.$$

This implies that after spending h periods in the cascade region, any public belief would have a value strictly lower than $g(h)$. Thus, whenever $g(h) \leq \alpha$, or equivalently $(1 - 2\varepsilon)^{h+1} \leq 1 - 2\alpha(1 - \alpha)$, the public LLR with value $g(h)$ must have exited the cascade region. Hence, the maximum number of periods that the public LLR can stay in the cascade region is $\frac{\log(1 - 2\alpha(1 - \alpha))}{\log(1 - 2\varepsilon)}$.

private signals are only weakly informative, agents rely more on social information. As a result, information cascades are more likely to arise, and so is action inertia.

We thus illustrate the proof idea of Proposition 1 using a weakly informative signal. Suppose that $\alpha = 1/2 + \delta$ for a small $\delta > 0$. Since $K(\alpha, \varepsilon)$ decreases in α , $K(\frac{1}{2}, \varepsilon)$ is the greatest upper bound for the length of any cascade as δ approaches zero. Upon exiting a cascade, the probability that the public LLR switches its sign is about $1/2$, since the agent, who follows her private signal, receives either a positive or negative signal with nearly equal probability. Therefore, the expected time between sign switches is bounded from above by a geometric distribution:

$$1 + \sum_{i=1}^{\infty} \frac{i}{2^i} K(\frac{1}{2}, \varepsilon) = 1 + \frac{2 \log 2}{-\log(1 - 2\varepsilon)} = M(1/2, \varepsilon),$$

and $M(1/2, \varepsilon)$ decreases in ε . Thus, for sufficiently small ε , $M(1/2, \varepsilon) \approx (\log 2)/\varepsilon$, which is strictly less than $1/\varepsilon$, the expected time between state changes.¹⁵ Given the one-to-one mapping between the agent's action and their public LLR, Proposition 1 immediately implies the following result.

Corollary 2. *The expected time between action changes is strictly less than the expected time between state changes.*

That is, on average, actions change more quickly than the state. For instance, when the probability of a state change is equal to 0.05, the state changes every twenty periods on average. In comparison, the maximum time for an action change is less than fourteen periods. Building on this, Theorem 1 formally shows the connection between the expected time between action changes and their long-run frequency, demonstrating that a shorter expected time implies a higher frequency.

5. CONCLUSION

We study the long-term behavior of agents who receive a private signal and observe the past actions of their predecessors in a changing environment. As the state evolves, the best action to take also fluctuates. We show that, in the long run, the frequency of action changes exceeds that of state changes, indicating that fads emerge under social learning in a dynamic environment. This result holds even with the occurrence of temporary information cascades in which agents simply mimic the action of their predecessors.

One may wonder if the main result is driven by the high frequency of action changes when the posterior belief is around $1/2$. Accordingly, we can further restrict the definition of fads to exclude consecutive action changes, i.e., cases where $a_t \neq a_{t-1}$ and $a_{t-1} \neq a_{t-2}$. Simulation results show that even under this more restricted definition of fads, actions still change more frequently than the state. For example, with $\alpha = 0.75$, $\varepsilon = 0.05$,

¹⁵See Claim 3 in the appendix.

and a total of 100,000 periods, action changes approximately 8,150 times, which is more frequent than the number of state changes, which is about 5,100 times.

An interesting extension would be to study the frequency of action changes for a single long-lived agent who repeatedly receives private signals about a changing state. We conjecture that, in this case, action changes would be less frequent than in our model, where only past actions are observable, but still more frequent than state changes. Intuitively, removing noisy observations of others' private signals (i.e., their actions) could reduce unnecessary action changes. If so, this would highlight the role of observational learning in accelerating action fluctuations, particularly in a slowly evolving environment. We leave this for future research, as it would require a different proof technique from the approach used here.

There are several possible avenues for future research. Recall that Proposition 1 implies that $M(\alpha, \varepsilon)$ is an upper bound to the expected time between action changes. One could ask whether this upper bound $M(\alpha, \varepsilon)$ is tight, and if so, for any finite time N , whether the number of action changes would be close to $N/M(\alpha, \varepsilon)$. Based on the simulation results, we conjecture that it is not a tight bound. E.g., we let $\alpha = 0.9$ and $\varepsilon = 0.05$, and $N = 100,000$. Since $M(0.9, 0.05) \approx 11.5$, it implies that within these hundred thousand periods, the action should change at least about 8700 times. However, our numerical simulation shows that the action changes about 14,200 times, almost double the number suggested by $M(0.9, 0.05)$. Furthermore, our simulations suggest that as the private signal becomes less informative and the state changes more slowly, i.e., when α approaches $1/2$ and ε approaches 0 at the same rate, the ratio between the frequency of action changes and state changes approaches a constant that is close to 4. This suggests that achieving a very accurate understanding of fads in this regime might be possible.

REFERENCES

- ACEMOGLU, D., A. NEDIC, AND A. OZDAGLAR (2008): "Convergence of rule-of-thumb learning rules in social networks," in *2008 47th IEEE Conference on Decision and Control*, IEEE, 1714–1720.
- AGGARWAL, R. AND P. RIVOLI (1990): "Fads in the initial public offering market?" *Financial management*, 45–57.
- 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," *Journal of Political Economy*, 992–1026.
- CAMERER, C. (1989): "Bubbles and fads in asset prices," *Journal of Economic Surveys*, 3, 3–41.

- DASARATHA, K., B. GOLUB, AND N. HAK (2020): “Learning from neighbors about a changing state,” *Available at SSRN 3097505*.
- FELLER, W. (1966): *An introduction to probability theory and its applications, vol 2*, John Wiley & Sons.
- FRONGILLO, R. M., G. SCHOENEBECK, AND O. TAMUZ (2011): “Social learning in a changing world,” in *International Workshop on Internet and Network Economics*, Springer, 146–157.
- HIRSHLEIFER, D. AND I. WELCH (2002): “An economic approach to the psychology of change: Amnesia, inertia, and impulsiveness,” *Journal of Economics & Management Strategy*, 11, 379–421.
- LÉVY, R., M. PEŠKI, AND N. VIEILLE (2022): “Stationary social learning in a changing environment,” *arXiv preprint arXiv:2201.02122*.
- MOSCARINI, G., M. OTTAVIANI, AND L. SMITH (1998): “Social learning in a changing world,” *Economic Theory*, 11, 657–665.
- SCHAAL, E. AND M. TASCHEREAU-DUMOUCHEL (2021): “Herding Through Booms and Busts,” *CEPR Discussion Paper No. DP16368*.
- SHAHRAMPOUR, S., S. RAKHLIN, AND A. JADBABAIE (2013): “Online learning of dynamic parameters in social networks,” *Advances in Neural Information Processing Systems*, 26.
- SMITH, L. AND P. SØRENSEN (2000): “Pathological outcomes of observational learning,” *Econometrica*, 68, 371–398.

APPENDIX A. PROOFS

Proof of Lemma 1. Fix any arbitrary $\alpha \in (1/2, 1)$ and $\varepsilon \in (0, \alpha(1 - \alpha))$. Consider first the case where $\ell_t \in (0, c_\alpha)$. For part (i), suppose $a_t = -1$. Given that $\ell_t \in (0, c_\alpha)$ is in the learning region, we have $s_t = a_t = -1$. Since $f_-(\cdot)$ in (4.4) is strictly increasing, and $f_-(c_\alpha) = 0$, we have $f_-(\ell_t) < 0$ for all $\ell_t \in (0, c_\alpha)$. Thus, it follows from (4.4) that $\ell_{t+1} = f_-(\ell_t)$ and $\text{sign}(f_-(\ell_t)) = -1 = -\text{sign}(\ell_t)$. For part (ii), it suffices to show that $f_+(f_+(0)) \geq c_\alpha$ as $f_+(\cdot)$ is strictly increasing. Let $c_u = f_+^{-1}(c_\alpha)$ denote the threshold at which exactly one positive signal is required for the public LLR to enter the cascade region on the positive action. This threshold is given by

$$c_u = \log \frac{(1 - \alpha)(\alpha - \varepsilon)}{\alpha(1 - \alpha - \varepsilon)} \in (0, c_\alpha).$$

Note that $f_+(0) > c_u$ for all $\varepsilon \in (0, \alpha(1 - \alpha))$, and thus $f_+(f_+(0)) > f_+(c_u)$. By definition, $f_+(f_+(0)) > f_+(c_u) = f_+(f_+^{-1}(c_\alpha)) = c_\alpha$, as required. The case where $\ell_t \in (-c_\alpha, 0)$ follows an analogous argument. \square

Proof of Proposition 1. Fix any arbitrary $\alpha \in (1/2, 1)$, $\varepsilon \in (0, \alpha(1 - \alpha))$ and some positive integer $i \geq 2$. Recall that \mathcal{T}_{i-1} is the time at which the public LLR changes its sign for the $i - 1$ -th time and $\mathcal{D}_i = \mathcal{T}_i - \mathcal{T}_{i-1}$. Consider $\ell_{\mathcal{T}_{i-1}} > 0$ and there are three disjoint intervals

for the value of $\ell_{\mathcal{T}_{i-1}}$: (i) $[c_u, c_\alpha]$ where $c_u = f_+^{-1}(c_\alpha)$; (ii) $(0, c_u)$ and (iii) $[c_\alpha, f_+(c_\alpha))$. We will show that for all cases, $\mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] < 1 + \frac{K(\alpha, \varepsilon)}{2\alpha(1-\alpha)}$, where $K(\alpha, \varepsilon)$ is the upper bound on the length of any cascade.

Denote by $\pi(\ell)$ the probability of receiving a positive signal conditional on the public LLR being ℓ .¹⁶ By the law of total probability,

$$\pi(\ell) = \alpha \cdot \frac{e^\ell}{1 + e^\ell} + (1 - \alpha) \cdot \frac{1}{1 + e^\ell} = \frac{1 + \alpha(e^\ell - 1)}{1 + e^\ell}, \quad (\text{A.1})$$

and it is strictly increasing in ℓ . Thus,

$$\bar{\pi} := \sup_{\ell \in (0, c_\alpha)} \pi(\ell) = 1 - 2\alpha(1 - \alpha).$$

Let $\kappa(\ell)$ denote the length of a positive cascade triggered by receiving a positive signal conditional on the public LLR being ℓ . Let $\mathcal{L}(\ell)$ represent the value of the public LLR after exiting the cascade region for the first time. We use $\lfloor K(\alpha, \varepsilon) \rfloor \geq 1$ to denote the greatest integer less than or equal to $K(\alpha, \varepsilon)$.

Case (i). Suppose $\ell_{\mathcal{T}_{i-1}} \in [c_u, c_\alpha]$. By part (i) of Lemma 1, since $\ell_{\mathcal{T}_{i-1}}$ is in the learning region, one opposing signal is sufficient to change the sign of $\ell_{\mathcal{T}_{i-1}}$. Thus, the expected time to the next sign switch is

$$\mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] = 1 - \pi(\ell_{\mathcal{T}_{i-1}}) + \pi(\ell_{\mathcal{T}_{i-1}}) \left(\kappa(\ell_{\mathcal{T}_{i-1}}) + \mathbb{E}[\mathcal{D}_i | \mathcal{L}(\ell_{\mathcal{T}_{i-1}})] \right),$$

and by the definition of $\bar{\pi}$, we have

$$\mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] < 1 - \bar{\pi} + \bar{\pi} \left(\lfloor K(\alpha, \varepsilon) \rfloor + \mathbb{E}[\mathcal{D}_i | \mathcal{L}(\ell_{\mathcal{T}_{i-1}})] \right). \quad (\text{A.2})$$

Note that there are two possible cases for the value of $\mathcal{L}(\ell_{\mathcal{T}_{i-1}})$: either $\mathcal{L}(\ell_{\mathcal{T}_{i-1}}) \in [c_u, c_\alpha]$ or $\mathcal{L}(\ell_{\mathcal{T}_{i-1}}) \in (0, c_u)$. If it is the former case, then taking the supremum on both sides of (A.2) and rearranging,

$$\sup_{c_u \leq \ell_{\mathcal{T}_{i-1}} < c_\alpha} \mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] \leq 1 + \frac{\bar{\pi}}{1 - \bar{\pi}} \lfloor K(\alpha, \varepsilon) \rfloor. \quad (\text{A.3})$$

If it is the latter case, by the definition of $\bar{\pi}$,

$$\mathbb{E}[\mathcal{D}_i | \mathcal{L}(\ell_{\mathcal{T}_{i-1}})] < 1 - \bar{\pi} + \bar{\pi} \left(1 + \mathbb{E}[\mathcal{D}_i | f_+(\mathcal{L}(\ell_{\mathcal{T}_{i-1}}))] \right).$$

Substituting the above inequality into (A.2),

$$\mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] < 1 - \bar{\pi} + \bar{\pi} \left(\lfloor K(\alpha, \varepsilon) \rfloor + 1 - \bar{\pi} + \bar{\pi} \left(1 + \mathbb{E}[\mathcal{D}_i | f_+(\mathcal{L}(\ell_{\mathcal{T}_{i-1}}))] \right) \right).$$

Since $f_+(\cdot)$ is strictly increasing, by part (ii) of Lemma 1, $f_+(\mathcal{L}(\ell_{\mathcal{T}_{i-1}})) \in [c_u, c_\alpha]$. Thus, taking the supremum on both sides and rearranging,

$$\sup_{c_u \leq \ell_{\mathcal{T}_{i-1}} < c_\alpha} \mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] \leq \frac{1 - \bar{\pi} + (\lfloor K(\alpha, \varepsilon) \rfloor + 1)\bar{\pi}}{1 - \bar{\pi}^2} \leq 1 + \frac{\bar{\pi}}{1 - \bar{\pi}} \lfloor K(\alpha, \varepsilon) \rfloor,$$

¹⁶For ease of notation, we suppress its dependence on α .

where the second inequality holds since $\lfloor K(\alpha, \varepsilon) \rfloor \geq 1$.

Case (ii). Suppose $\ell_{\mathcal{T}_{i-1}} \in (0, c_u)$. By part (i) of Lemma 1 and the definition of $\bar{\pi}$, the expected time to the next sign switch is bounded above:

$$\mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] < (1 - \bar{\pi}) + \bar{\pi}(1 + \mathbb{E}[\mathcal{D}_i | f_+(\ell_{\mathcal{T}_{i-1}})]).$$

Since $f_+(\cdot)$ is strictly increasing, part (ii) of Lemma 1 implies that $f_+(\ell_{\mathcal{T}_{i-1}}) \in [c_u, c_\alpha)$. It then follows from (A.3) that

$$\mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] < \frac{\bar{\pi}^2(\lfloor K(\alpha, \varepsilon) \rfloor - 1) + 1}{1 - \bar{\pi}}. \quad (\text{A.4})$$

Case (iii). Suppose $\ell_{\mathcal{T}_{i-1}} \in [c_\alpha, f_+(c_\alpha))$. In this case, after at most $\lfloor K(\alpha, \varepsilon) \rfloor$ periods, the public LLR initiated at $\ell_{\mathcal{T}_{i-1}}$ would have exited the cascade region. Hence, the expected time to the next sign switch is bounded above:

$$\mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] \leq \lfloor K(\alpha, \varepsilon) \rfloor + \mathbb{E}[\mathcal{D}_i | \mathcal{L}(\ell_{\mathcal{T}_{i-1}})].$$

Again, there are two possible cases for $\mathcal{L}(\ell_{\mathcal{T}_{i-1}})$: either $\mathcal{L}(\ell_{\mathcal{T}_{i-1}}) \in [c_u, c_\alpha)$ or $\mathcal{L}(\ell_{\mathcal{T}_{i-1}}) \in (0, c_u)$. If it is the former case, it follows from (A.3) that

$$\mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] < 1 + \frac{1}{1 - \bar{\pi}} \lfloor K(\alpha, \varepsilon) \rfloor. \quad (\text{A.5})$$

If it is the latter case, then it follows from (A.4) that

$$\begin{aligned} \mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] &< \lfloor K(\alpha, \varepsilon) \rfloor + \frac{\bar{\pi}^2(\lfloor K(\alpha, \varepsilon) \rfloor - 1) + 1}{1 - \bar{\pi}} \\ &= \lfloor K(\alpha, \varepsilon) \rfloor + 1 + \bar{\pi} + \frac{\bar{\pi}^2}{1 - \bar{\pi}} \lfloor K(\alpha, \varepsilon) \rfloor \leq 1 + \frac{1}{1 - \bar{\pi}} \lfloor K(\alpha, \varepsilon) \rfloor. \end{aligned}$$

Now, note that the maximum of these three upper bounds given in (A.3) to (A.5) is $1 + \frac{1}{1 - \bar{\pi}} \lfloor K(\alpha, \varepsilon) \rfloor$. Furthermore, by definition, $\lfloor K(\alpha, \varepsilon) \rfloor \leq K(\alpha, \varepsilon)$. Therefore, we conclude that

$$\mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}} > 0] < 1 + \frac{K(\alpha, \varepsilon)}{2\alpha(1 - \alpha)}.$$

The case where $\ell_{\mathcal{T}_{i-1}} < 0$ follows from an analogous argument. \square

Claim 3. $M(1/2, \varepsilon) < 1/\varepsilon$ for all $\varepsilon \in (0, 1/4)$.

Proof. Recall that

$$M(1/2, \varepsilon) = \sup_{\alpha \in (1/2, 1)} M(\alpha, \varepsilon) = 1 + \frac{2 \log 2}{-\log(1 - 2\varepsilon)},$$

and rearranging, we have $M(1/2, \varepsilon) < 1/\varepsilon$ if

$$2 \log 2 < -\left(\frac{1}{\varepsilon} - 1\right) \log(1 - 2\varepsilon)$$

since $\varepsilon \in (0, 1/4)$. By the L'Hôpital's rule,

$$\lim_{\varepsilon \rightarrow 0} -\left(\frac{1}{\varepsilon} - 1\right) \log(1 - 2\varepsilon) = \lim_{\varepsilon \rightarrow 0} 2 \frac{(1 - \varepsilon)^2}{1 - 2\varepsilon} = 2 > 2 \log 2.$$

Since $-(\frac{1}{\varepsilon} - 1) \log(1 - 2\varepsilon)$ is strictly increasing in ε , we conclude that $M(1/2, \varepsilon) < 1/\varepsilon$ for all $\varepsilon \in (0, 1/4)$. \square

Since $M(\alpha, \varepsilon)$ decreases in α , it follows from Claim 3 and Proposition 1 that for all $\alpha \in (1/2, 1)$ and $\varepsilon \in (0, \alpha(1 - \alpha))$,

$$M(\alpha, \varepsilon) < 1/\varepsilon. \quad (\text{A.6})$$

Before proceeding to the proof of Theorem 1, we provide the following lemma, which will be useful. This lemma shows that the process (\mathcal{D}_i) has well-defined moments. In particular, it implies that there is a finite uniform upper bound to its second moment $\mathbb{E}[\mathcal{D}_i^2]$, which is required to apply the standard martingale convergence theorem. Intuitively, since any cascade must end after $K(\alpha, \varepsilon)$ periods, the probability that \mathcal{D}_i is larger than some finite periods decreases exponentially fast, and so \mathcal{D}_i must have finite moments.

Lemma 2. *For every $r \in \{1, 2, \dots\}$ there is a constant c_r that depends on α and ε such that for all i , $\mathbb{E}[|\mathcal{D}_i|^r] < c_r$. I.e., each moment of \mathcal{D}_i is uniformly bounded, independently of i .*

Proof. Fix any arbitrary $\alpha \in (1/2, 1)$, $\varepsilon \in (0, \alpha(1 - \alpha))$ and some positive integer $i \geq 2$. Suppose that $\ell_{\mathcal{T}_{i-1}} > 0$ and so $\mathcal{D}_i = \mathcal{T}_i - \mathcal{T}_{i-1}$ is the time elapsed from a positive public LLR to a negative one. For any $n \geq 2$, we denote the minimum number (which may not be an integer) of temporary cascades required for $\mathcal{D}_i > n$ by

$$k(n) := \max \left\{ \frac{n-1}{\lfloor K(\alpha, \varepsilon) \rfloor}, 1 \right\}.$$

Recall that $\bar{\pi}$ is the supremum of the probability of receiving a positive signal conditional on the public LLR being ℓ for all $\ell \in (0, c_\alpha)$. By part (ii) of Lemma 1, for any $n \geq 2$, the probability of the event $\{\mathcal{D}_i > n\}$ is bounded above:

$$\mathbb{P}[\mathcal{D}_i > n] < \bar{\pi}^{2+(\lfloor k(n) \rfloor - 1)}.$$

Since \mathcal{D}_i is a positive random variable, it follows that for any $p > 0$,

$$\begin{aligned} \lim_{n \rightarrow \infty} n^p \mathbb{P}[\mathcal{D}_i > n] &= \lim_{n \rightarrow \infty} \frac{n^p}{1/\mathbb{P}[\mathcal{D}_i > n]} \\ &< \lim_{n \rightarrow \infty} \frac{n^p}{(1/\bar{\pi})^{1+\lfloor k(n) \rfloor}} = 0. \end{aligned} \quad (\text{A.7})$$

For any $r \geq 1$, the r -th moment of $|\mathcal{D}_i|$ satisfies

$$\begin{aligned}\mathbb{E}[|\mathcal{D}_i|^r] &= \int_0^\infty \mathbb{P}[|\mathcal{D}_i|^r > t] dt \\ &< 1 + \int_1^\infty \mathbb{P}[\mathcal{D}_i > y] r y^{r-1} dy \\ &= 1 + \sum_{n=1}^\infty \int_n^{n+1} \mathbb{P}[\mathcal{D}_i > y] r y^{r-1} dy \\ &< 1 + \sum_{n=1}^\infty \mathbb{P}[\mathcal{D}_i > n] r (n+1)^{r-1},\end{aligned}$$

where the second inequality follows from a change of variable $y = t^{1/r}$. Since (A.7) implies that $\mathbb{P}[\mathcal{D}_i > n] < Cn^{-p}$ for some nonnegative constant C , it follows that for any $p > r$,

$$\begin{aligned}\mathbb{E}[|\mathcal{D}_i|^r] &< 1 + rC \sum_{n=1}^\infty \frac{(n+1)^{r-1}}{n^p} \\ &< 1 + r2^{r-1}C \sum_{n=1}^\infty \frac{1}{n^{p-r+1}} < \infty,\end{aligned}$$

which holds for all i . Hence, for every $r \in \{1, 2, \dots\}$, there exists a constant $c_r = 1 + r2^{r-1}C \sum_{n=1}^\infty \frac{1}{n^{p-r+1}}$, independently of i , that uniformly bounds $\mathbb{E}[|\mathcal{D}_i|^r]$. \square

Proof of Theorem 1. Fix any arbitrary $\alpha \in (1/2, 1)$ and $\varepsilon \in (0, \alpha(1-\alpha))$. Since θ_t follows a two-state Markov chain with a symmetric transition probability ε , $(\mathbb{1}(\theta_1 \neq \theta_2), \mathbb{1}(\theta_2 \neq \theta_3), \dots)$ is a sequence of i.i.d. random variables. By the strong law of large numbers,

$$\lim_{n \rightarrow \infty} \mathcal{Q}_\theta(n) := \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n \mathbb{1}(\theta_t \neq \theta_{t+1}) = \mathbb{P}[\theta_t \neq \theta_{t+1}] = \varepsilon \quad \text{almost surely.}$$

Let $\Phi = (\mathcal{F}_1, \mathcal{F}_2, \dots)$ be the filtration where each $\mathcal{F}_i = \sigma(\mathcal{D}_1, \dots, \mathcal{D}_i)$ and thus $\mathcal{F}_j \subseteq \mathcal{F}_i$ for any $j \leq i$. So the process $(\mathcal{D}_1, \mathcal{D}_2, \dots)$ is adapted to Φ since each \mathcal{D}_i is \mathcal{F}_i -measurable. By Proposition 1 and (A.6), there exists $\delta = 1/\varepsilon - M(\alpha, \varepsilon) > 0$ such that for all $i \geq 2$,

$$\mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}] < 1/\varepsilon - \delta.$$

By the law of iterated expectation and the Markov property of the public LLR,

$$\mathbb{E}[\mathcal{D}_i | \mathcal{F}_{i-1}] = \mathbb{E}[\mathbb{E}[\mathcal{D}_i | \ell_{\mathcal{T}_{i-1}}, \mathcal{F}_{i-1}] | \mathcal{F}_{i-1}] < 1/\varepsilon - \delta. \quad (\text{A.8})$$

Let $X_i = \mathcal{D}_i - \mathbb{E}[\mathcal{D}_i | \mathcal{F}_{i-1}]$ for all $i \geq 2$ and since $\mathcal{F}_{i-1} \subseteq \mathcal{F}_i$, each X_i is \mathcal{F}_i -measurable. Denote a partial sum of the process $(X_i)_{i \geq 2}$ by

$$Y_n = X_2 + \frac{1}{2}X_3 + \dots + \frac{1}{n-1}X_n.$$

By definition, $\mathbb{E}[X_i|\mathcal{F}_{i-1}] = 0$ for all $i \geq 2$. Since each Y_{n-1} is \mathcal{F}_{n-1} -measurable,

$$\mathbb{E}[Y_n|\mathcal{F}_{n-1}] = \mathbb{E}\left[\sum_{i=2}^n \frac{1}{i-1} X_i|\mathcal{F}_{n-1}\right] = Y_{n-1} + \frac{1}{n-1} \mathbb{E}[X_n|\mathcal{F}_{n-1}] = Y_{n-1},$$

and so the process $(Y_n)_n$ forms a martingale.

By Lemma 2 and (A.8), both $\mathbb{E}[\mathcal{D}_i^2]$ and $\mathbb{E}[\mathcal{D}_i|\mathcal{F}_{i-1}]$ are uniformly bounded. Therefore, $\mathbb{E}[X_i^2]$ is also uniformly bounded for all $i \geq 2$. Furthermore, since $\mathbb{E}[X_i X_j] = 0$ for any $i \neq j$, it then follows that for all $n \geq 2$,

$$\mathbb{E}[Y_n^2] = \sum_{i=2}^n \frac{1}{(i-1)^2} \mathbb{E}[X_i^2] < \infty.$$

By the martingale convergence theorem, Y_n converges almost surely. It then follows from Kronecker's lemma that¹⁷

$$\lim_{n \rightarrow \infty} \frac{1}{n-1} (X_2 + \cdots + X_n) = 0 \quad \text{almost surely.}$$

Thus, by the definition of X_i , we can write

$$\lim_{n \rightarrow \infty} \frac{1}{n-1} \sum_{i=2}^n \mathcal{D}_i = \lim_{n \rightarrow \infty} \frac{1}{n-1} \sum_{i=2}^n \mathbb{E}[\mathcal{D}_i|\mathcal{F}_{i-1}] \quad \text{almost surely.}$$

It follows from (A.8) that

$$\lim_{n \rightarrow \infty} \frac{1}{n-1} \sum_{i=2}^n \mathcal{D}_i \leq 1/\varepsilon - \delta < 1/\varepsilon \quad \text{almost surely.} \quad (\text{A.9})$$

Since $a_t = \text{sign}(\ell_{t+1})$ for all $t \geq 2$,

$$\lim_{n \rightarrow \infty} \mathcal{Q}_a(n) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n \mathbb{1}(a_t \neq a_{t+1}) = \lim_{n \rightarrow \infty} \frac{1}{n-1} \sum_{t=2}^n \mathbb{1}(\text{sign}(\ell_{t+1}) \neq \text{sign}(\ell_{t+2})).$$

By the definitions of \mathcal{T}_i and \mathcal{D}_i ,

$$\frac{1}{n-1} \sum_{t=2}^n \mathbb{1}(\text{sign}(\ell_{t+1}) \neq \text{sign}(\ell_{t+2})) = \frac{n-1}{\mathcal{T}_n - \mathcal{T}_1} = \frac{n-1}{\sum_{i=2}^n \mathcal{D}_i}.$$

Hence, we conclude from (A.9) that

$$\lim_{n \rightarrow \infty} \mathcal{Q}_a(n) = \lim_{n \rightarrow \infty} \frac{n-1}{\sum_{i=2}^n \mathcal{D}_i} > \varepsilon = \lim_{n \rightarrow \infty} \mathcal{Q}_\theta(n) \quad \text{almost surely.}$$

□

¹⁷This result is also known as the strong law for martingales (See p.238, Feller (1966, Theorem 2)).