

THE EMERGENCE OF FADS IN A CHANGING WORLD

WANYING HUANG[†]

ABSTRACT. We study how fads emerge from social learning in a changing environment. We consider a simple sequential learning model in which rational agents arrive in order, each acting only once, and the underlying unknown state is constantly evolving. Each agent receives a private signal, observes all past actions of others, and chooses an action to match the current state. Since the state changes over time, cascades cannot last forever, and actions fluctuate too. We show that in the long run, actions change more often than the state. This describes many real-life faddish behaviors in which people often change their actions more frequently than what is necessary.

1. INTRODUCTION

The term “fad” describes transient behavior that rises and fades quickly in popularity. In particular, these fast changes in behavior cannot be explained entirely by changes in the fundamentals. For example, in macroeconomics, there are boom-and-bust business cycles that cannot be pinned down by changes in the underlying economy.¹ It has been long documented in finance that price deviation from the asset’s intrinsic values can stem from speculative bubbles and fads (Camerer, 1989; Aggarwal and Rivoli, 1990). While the phenomenon of fads is widely observed in many economic activities, the question of how and why fads emerge has yet to be resolved in the literature. In this paper, we show how fads can arise from social learning in an ever-changing environment.

The pioneering work in the social learning literature (Banerjee, 1992; Bikhchandani et al., 1992) (BHW thereafter) shows that under appropriate conditions, *information cascades* always occur – this is the event in which the social information swamps agents’ private information so that agents would follow others’ action regardless of their private signals. However, as discussed in BHW, this long-run cascading outcome is also fragile to small shocks. For example, the possibility of a one-time change in the underlying state could cause “seemingly whimsical swings in mass behavior without obvious external stimulus” for which they to as fads. Inspired by BHW’s original idea, we introduce a formal definition of fads and study their long-term behavior.

Date: March 13, 2023.

[†]California Institute of Technology. Email: whhuang@caltech.edu.

I am indebted to my advisor Omer Tamuz for his continued support and encouragement. I am grateful (in alphabetical order) to Krishna Dasaratha, 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 interesting study by Schaal and Taschereau-Dumouchel (2021) modeling business cycles through the lens of social learning.

While BHW present an early idea of fads, 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 in the literature (see, e.g., [Dasaratha et al., 2020](#); [Lévy et al., 2022](#)). More importantly, a model with a changing state allows us to understand whether behavior under observational learning is more or less volatile than the underlying state. In other words, does social learning exaggerate the fluctuations in behavior and cause faddish behavior, or does it have the opposite effect?

To answer this question, we study the canonical model of social learning (i.e., a binary state with symmetric binary signals) with a slight twist: with a small (and symmetric) probability, the underlying state switches in every period.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 her action matches the current state. As the underlying state evolves, the best action to take also fluctuates. The question we ask is: compared to state changes, how frequently do actions change? On the one hand, agents sometimes ignore their private signals because of information cascades, and thus they do not change their actions even when the state changes. On the other hand, because signals are noisy, agents sometimes change actions unnecessarily. We say *fads* emerge if there are more action changes than state changes, and our main result shows that fads do emerge in the long run. We stress that in our model, fads arise from rational agents' desire to match the ever-changing state instead of any payoff externalities between agents or heuristics or irrationality of the agents.

More specifically, the slowly evolving state in our model allows cascades to arise, but they can only be temporary so that agents cannot herd on a single action forever ([Moscarini et al., 1998](#)). Intuitively, older social information is depreciated over time as it becomes less relevant to the current agent, and eventually, agents will return to utilize their private signals and change their actions. Nevertheless, the question of how often agents change their actions compared to state changes remains unclear. At first glance, one may expect that actions would change less often than the state as temporary cascades prevent agents from following their private signals and thus reduce the volatility in actions. The symmetry of binary states also amplifies such an effect: imagine the state has changed an even number of times, say twice, while agents in a cascade could mistakenly treat the state as unchanged, so they would have no reason to change their actions.

Perhaps surprisingly, our main result (Theorem 1) shows that even if there are temporary cascades, actions change more often than the state in the long run. In other words,

²This is known as a simple two-state Markov process. See other studies of social learning in a changing environment that also consider a Markovian environment, e.g., [Moscarini et al. \(1998\)](#); [Hirshleifer and Welch \(2002\)](#) and [Lévy et al. \(2022\)](#). Our model is mostly close to [Moscarini et al. \(1998\)](#) except for the tie-breaking rule. We provide a detailed literature review on these studies next.

fads emerge in a changing environment even though it is unlikely that the state has changed at all. For example, consider a private signal that matches the state 80 percent of the time. On average, when the state changes once every hundred periods, we show that agents take less than sixty-one periods to change their actions. Thus, the long-term frequency of action changes must be higher than state changes, resulting in fads. This relatively high frequency of action changes is also in line with the fragility of fads, where small shocks to the system could cause rapid shifts in agents' behavior.³

Our proof strategy behind the emergence of fads is as follows. First, for any fixed signal precision and probability of state change, there exists a maximum length of any cascade. As a result, even though the rise of temporary cascades prolongs action inertia, such an effect on actions is limited by its bounded length. Meanwhile, agents only require one opposing signal to change their actions whenever they have an opportunity, i.e., whenever the public belief exits a cascade. We bound the probability of agents changing actions from below and thus establish an upper bound for the expected time between action changes. We show that this upper bound is less than the expected time of state changes, implying that it takes less time on average for actions to change than the state. Finally, we translate the expected time of changes for both the state and the action into their long-term relative frequency of changes and 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. As mentioned before, BHW show that a one-time shock to the state could break the cascade, even though that shock may never be realized. They provide a numeric example where the probability of an action change is at least 87% higher than the probability of state change (see their Result 4) which is in line with our main result. Later, Moscarini et al. (1998) further explore this idea and show that if the underlying environment is evolving in every period and the state is sufficiently persistent, an information cascade must arise, but it can only be temporary, i.e., it must end in finite time. Our work builds on their model but with a different focus. While they focus on analyzing the short-term behavior of information cascades, e.g., under what conditions do they end or arise, we ask: in the long run, should one expect more action or state changes?

Hirshleifer and Welch (2002) also consider a changing environment in their stylized model (see their §3), but their focus is on examining the effect of *memory loss* on the continuity of behavior of a single agent.⁴ They analyze the equilibrium behavior of a five-period stylized model and show that in a relatively stable environment, memory loss causes the agent to exhibit excess action inertia relative to a full-recall regime, whereas

³As discussed in Bikhchandani et al. (1992, 1998), fads are fragile precisely because they are typically formed based on little information. Thus, different kinds of shocks, such as uncertainty in the underlying state as in our model or the arrival of a better informed agent, etc., would dislodge the previous trend and cause drastic behavioral changes.

⁴The term “memory loss” refers to the case where the agent only recalls past actions but not past signals.

in a volatile environment, memory loss leads to excess action impulsiveness.⁵ However, the authors did not pursue a changing environment when studying the long-term effect of amnesia. We differ from their study by emphasizing long-term behavior in a changing environment.

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 et al. \(2011\)](#) consider a specific dynamic 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 weights may be Pareto suboptimal, causing inefficiency in learning.⁶ In a similar but more general environment, [Dasaratha et al. \(2020\)](#) show that having sufficiently diverse network neighbors with different signal distributions improves learning. Intuitively, having diverse signals allows agents to decipher the most relevant information from the old and confounded information, thus achieving higher efficiency in information aggregation.

A more recent study by [Lévy et al. \(2022\)](#) considers a similar setup to ours, focusing on the implication of a dynamic state on equilibrium welfare. In their model, agents observe a random subsample drawn from all past behaviors 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 the need to be responsive to environmental changes, which reduces equilibrium welfare. In contrast, we assume that all past actions are observable and that the private signals are free of charge. 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.

2. MODEL

We follow the setup from [Moscarini et al. \(1998\)](#) 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 constantly evolves over time. A sequence of agents indexed by time t arrive in order, each acting once by choosing an action $a_t \in \{-1, +1\}$ with a payoff function that depends on the unknown state at time t : $\mathbb{1}(a_t = \theta_t)$, i.e., a positive payoff of one if the action matches the current state and zero otherwise.

Before choosing an action, agent t receives a private signal s_t and observes the history of all past actions made by her predecessors, $h^{t-1} = (a_1, a_2, \dots, a_{t-1})$. Conditional on the entire sequence of states, the private signals s_t are independent, and each s_t has

⁵Intuitively, as the 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.

⁶See more studies in the computer science literature, e.g., [Acemoglu et al. \(2008\)](#); [Shahrampour et al. \(2013\)](#) that consider a dynamic environment with non-Bayesian agents.

a Bernoulli distribution $B_{\theta_t}(\alpha)$ where α is the symmetric probability of matching the current state: $\mathbb{P}[s_t = i | \theta_t = i] = \alpha \in (1/2, 1)$, for $i \in \{-1, +1\}$. We often refer to signal $s = +1$ as an up-signal and to $s = -1$ as a down-signal. Let $\mathcal{I}^t = \{-1, +1\}^{t-1} \times \{-1, +1\}$ be the space of information available to agent t prior to her decision so that $I_t = (h^{t-1}, s_t)$ is an element of \mathcal{I}^t .

For simplicity, we assume that both states are equally likely at the beginning of time, and the state evolves according to a symmetric Markov chain with transition probability ε , i.e., $\mathbb{P}[\theta_{t+1} \neq i | \theta_t = i] = \varepsilon$, for $i \in \{-1, +1\}$. This assumption implies that the stationary distribution of θ_t is uniform. We assume throughout that the state is sufficiently persistent and in particular, $\varepsilon \in (0, \alpha(1 - \alpha))$ so that temporary cascades can arise (Moscarini et al., 1998, Proposition 2). 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 $\{-1, +1\}$ with equal probability. Thus, the probability of a state change is equal to $\varepsilon \in (0, \alpha(1 - \alpha))$.

At any time t , the timing of the events is as follows. First, agent t arrives and observes the history h^{t-1} of all past actions. Second, the state θ_{t-1} transitions to θ_t with probability ε of changing. Then, 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 I_t , the information available to her.

2.1. Fads. Given that each agent aims to match the current state, the best action to take also fluctuates as the state evolves. BHW briefly discuss the idea that agents exhibit faddish behavior if they change their actions more often than the state. Formally, we say that *fads* emerge by time $n + 1$ if the fraction of time periods $t \leq n$ for which $a_t \neq a_{t+1}$ is larger than the fraction of those for which $\theta_t \neq \theta_{t+1}$, i.e.,

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

Multiplying both sides of (1) by n , the emergence of fads by time $n + 1$ implies that actions would have changed more often than the state by time $n + 1$.

2.2. Agents' Beliefs. Let $q_t := \mathbb{P}[\theta_t = +1 | h^{t-1}]$ denote the *public belief* assigned to state $\theta_t = +1$ at time t after observing the history of actions h^{t-1} . Let $p_t := \mathbb{P}[\theta_t = +1 | I_t]$ denote the *posterior belief* assigned to state $\theta_t = +1$ after observing $I_t = (h^{t-1}, s_t)$. Denote the log-likelihood ratio (LLR) of the posterior belief of agent t by

$$L_t := \log \frac{p_t}{1 - p_t} = \log \frac{\mathbb{P}[\theta_t = +1 | I_t]}{\mathbb{P}[\theta_t = -1 | I_t]},$$

and call it the *posterior likelihood* at time t . It follows from Bayes' rule that the posterior likelihood at time t satisfies

$$L_t = \log \frac{\mathbb{P}[s_t | \theta_t = +1, h^{t-1}]}{\mathbb{P}[s_t | \theta_t = -1, h^{t-1}]} + \log \frac{\mathbb{P}[\theta_t = +1 | h^{t-1}]}{\mathbb{P}[\theta_t = -1 | h^{t-1}]} \quad (2)$$

As the private signal is independent of the history conditional on the current state, the first term in (2) reduces to the LLR induced by the signal and it is equal to $c_\alpha := \log \frac{\alpha}{1-\alpha}$ if $s_t = +1$ and $-c_\alpha$ if $s_t = -1$. Denote the second term in (2) by

$$l_t := \log \frac{q_t}{1-q_t} = \log \frac{\mathbb{P}[\theta_t = +1|h^{t-1}]}{\mathbb{P}[\theta_t = -1|h^{t-1}]},$$

the *public likelihood* at time t . Intuitively, anyone who observes all past actions until time $t-1$ can calculate this log-likelihood ratio.

Thus, depending on the realization of the private signals, the posterior likelihood L_t is the sum of the public likelihood l_t and the LLR induced by the private signal at time t :

$$L_t = \begin{cases} l_t - c_\alpha & \text{if } s_t = -1, \\ l_t + c_\alpha & \text{if } s_t = +1. \end{cases} \quad (3)$$

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:

$$a_t \in \arg \max_{a \in \{-1, +1\}} \mathbb{E}[\mathbb{1}(\theta_t = a)|I_t] = \arg \max_{a \in \{-1, +1\}} \mathbb{P}[\theta_t = a|I_t].$$

Thus $a_t = +1$ if $L_t > 0$ and $a_t = -1$ if $L_t < 0$. When agent t is indifferent, i.e., $L_t = 0$, we assume that she would follow what her immediate predecessor did in the previous period, i.e., $a_t = a_{t-1}$.⁷ This assumption ensures that action changes are not due to the specification of the tie-breaking rule but rather due to her strict preference for one action over another.

2.4. Information Cascades and Regions. An *information cascade* is the event in which the past actions of others form an overwhelming influence on agents so that they act independently of the private signals. Specifically, it follows from (3) that the sign of the posterior likelihood L_t is purely determined by the sign of the public likelihood l_t once the absolute value of l_t exceeds c_α . Since the sign of L_t determines the optimal action of agent t , in this case, a_t will also be purely determined by the sign of l_t , independent of the private signal s_t . That is, $a_t = +1$ if $l_t > c_\alpha$ and $a_t = -1$ if $l_t < -c_\alpha$. When $|l_t| < c_\alpha$, agent t chooses the action according to her private signal so that $a_t = s_t$.

When $|l_t| = c_\alpha$, by the tie-breaking rule at indifference, regardless of the private signal that agent t receives, she chooses $a_t = a_{t-1} = \text{sign}(l_t)$.⁸ Thus, we call the region of the public likelihood in which $|l_t| \geq c_\alpha$ the *cascade region* and the region in which $|l_t| < c_\alpha$

⁷This tie-breaking rule differs from the one in Moscarini et al. (1998) where indifferent agents are assumed to follow their own private signals. Our results do not depend on this choice and are robust to any tie-breaking rule that is common knowledge.

⁸Without loss of generality, consider $l_t = c_\alpha$. If $s_t = +1$, $L_t = l_t + c_\alpha > c_\alpha$ so $a_t = +1$. If $s_t = -1$, $L_t = l_t - c_\alpha = 0$ and by the tie-breaking rule, $a_t = a_{t-1}$. To see why $a_{t-1} = \text{sign}(l_t) = +1$, suppose to the contrary that $a_{t-1} = -1$. Given that $l_t = c_\alpha$ and $a_{t-1} = -1$, it must be that $l_{t-1} > c_\alpha$, which implies that $a_{t-1} = +1$. A contradiction.

the *learning region*. We refer to the cascade in which $a = +1$ as an up-cascade and to the cascade in which $a = -1$ as a down-cascade.

3. RESULTS

We now state our main result. Recall that in (1) we define the emergence of fads by some 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, i.e.,*

$$\lim_{n \rightarrow \infty} Q_a(n) > \lim_{n \rightarrow \infty} Q_\theta(n) \quad a.s.$$

Perhaps surprisingly, Theorem 1 shows that even though there are times in which agents stop responding to their private signals, i.e., when cascades arise temporarily, agents who observe their predecessors' past actions still change their actions more often than the state in the long run. In other words, fads can emerge from social learning even though the underlying environment evolves very slowly. For example, consider a private signal with a precision of 0.8. When the probability of state change is equal to 0.01, the state changes once every hundred periods on average. Meanwhile, the average time for the action to change is strictly less than six-one periods.⁹ Thus, in the long run, actions would change strictly more often than the state, resulting in faddish behavior.

The idea behind the proof of Theorem 1 is as follows. Intuitively, as older social information becomes less relevant to the current agent in a changing environment, agents stop cascading on the same action and start periodically aggregating their private information. This periodical information aggregation then drives action fluctuations since agents at these times are easily gullible by opposing news. Formally, we show that once the public belief exits the cascade region, the action either changes or the public belief enters the same cascade region again.¹⁰ We upper bound the probability of the latter event and thus establish an upper bound to the expected time between action switches (Proposition 1). We 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 time between both action and state changes into their long-run relative frequency of changes. Building on this, we conclude that the long-run relative frequency of action changes is higher than that of state changes.

Note that the connection between the expected time of action changes and its long-run relative frequency of changes is not a consequence of the law of large numbers, as these action changes are not independent events. Instead, we circumvent this problem by studying the process of (random) time elapsed between action changes, which has

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

¹⁰It may take more than one up-action for the public belief to enter the same cascade region again.

TABLE 1. The numerical simulations of the number of action and state changes (in parentheses) under different values of α and ε for 100,000 periods.

$\alpha \setminus \varepsilon$	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)	–	–

well-defined moments (see Lemma 2). Also, notice that unlike the fixed-state model, when the state changes over time, the public belief about the current state ceases to be an unconditional martingale.¹¹ Nevertheless, the public belief is still a Markov process with specific transitional patterns (see Lemma 1).

We restrict our model to a persistent state so temporary cascades can arise (Moscarini et al., 1998, Proposition 2). The reason for this restriction is twofold. First, in a world where cascades never arise, agents would follow their signals and change their actions accordingly. Here, action changes are purely driven by the volatility in the state and the noise from the private signals; thus, social learning plays no role in agents' behavior. In contrast, in our model, even with the intervention of social learning, which generates cascades that slow down action changes, we show that in the long run, actions still change more often than the state. This result suggests that the emergence of fads is robust to observational learning. Second, even for a persistent state, it is a priori unclear whether actions or the state changes more often, as the state becomes less likely to change, it also slows down action changes. Intuitively, in this case, past actions become more informative about the current state, and as a result, temporary cascades last longer, prolonging action inertia.

3.1. Numerical Simulations. To complement the asymptotic result of Theorem 1, we simulate the long-run frequencies of action and state changes under different parameter values of α and ε in Table 1. The first observation is that these numerical simulations confirm our main result: for all parameter values considered, the action changes more often than the state. Next, we consider how the magnitude of the action and state changes varies across different parameter values. On the one hand, one can see that as the private signal becomes more precise (α increases), the number of action changes decreases, and so does the ratio between action and state changes. Intuitively, more precise signals reduce unnecessary action changes. On the other hand, as the state becomes more volatile (ε increases), both the action and the state change more frequently. However, the ratio between these changes decreases from around 3.3 to 2.1 (see the first row in Table 1), suggesting that the indirect effect of state volatility on action changes is less significant than its direct effect on state changes.

¹¹The martingale property is an important tool in analyzing the long-run behavior of learning, e.g., it is essential in proving asymptotic learning (Smith and Sørensen, 2000) for unbounded signals.

4. ANALYSIS

In this section, we first analyze how the public likelihood evolves in different regions, i.e., the cascade and learning region. We then compare the expected time between state and action changes and establish an upper bound on the expected time between action changes, which is crucial in proving our main result.

4.1. The Evolution of the Public Likelihood.

4.1.1. Cascade Region. When the state is fixed ($\varepsilon = 0$), it is well-known that the public likelihood stays forever at the value at which it first entered the cascade region, and an incorrect cascade can occur forever with positive probability (Bikhchandani et al., 1992; Banerjee, 1992).

When the state is changing ($\varepsilon > 0$), however, the behavior of the public likelihood changes significantly. To see this, consider first the case where the public likelihood is in the cascade region. Suppose that t is a time at which the public likelihood enters the cascade region from the learning region. Although no agent's actions reveal more information about the state after time t , the state still evolves and changes with probability ε in every period. Since the process $(\theta_t)_t$ is a Markov chain, it follows from the law of total probability that the public belief updates deterministically as follows:

$$\begin{aligned} q_{t+1} &:= \mathbb{P}[\theta_{t+1} = +1 | h^t] = \sum_{\theta_t \in \{-1, +1\}} \mathbb{P}[\theta_{t+1} = +1 | h^t, \theta_t] \mathbb{P}[\theta_t | h^t] \\ &= (1 - \varepsilon)q_t + \varepsilon(1 - q_t), \end{aligned}$$

so that

$$q_{t+1} = (1 - 2\varepsilon)q_t + (2\varepsilon)\frac{1}{2}. \quad (4)$$

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

$$l_{t+1} := \log \frac{q_{t+1}}{1 - q_{t+1}} = \log \frac{(1 - \varepsilon)e^{l_t} + \varepsilon}{1 - \varepsilon + \varepsilon e^{l_t}}. \quad (5)$$

From (4), we see that q_{t+1} tends to $1/2$, and by (5), l_{t+1} moves towards zero over time so that eventually it will exit the cascade region. Intuitively, having a changing state depreciates the value of older information as actions observed in earlier periods become less relevant to the current state. Consequently, after some finite number of periods, the public belief will slowly converge towards uniformity, and thus information cascades built upon this public belief cannot last forever. Indeed, this is the main insight from Moscarini et al. (1998), where they show that information cascades (if they arise) must end in finite time.

4.1.2. Learning Region. When the state is fixed ($\varepsilon = 0$), as the agent's action is informative about her private signal in the learning region, i.e., $a_t = s_t$, the public belief at time

$t + 1$ coincides with the posterior belief at time t :

$$q_{t+1} := \mathbb{P}[\theta_{t+1} = 1 | h^{t-1}, a_t] = \mathbb{P}[\theta_t = 1 | h^{t-1}, s_t] = p_t.$$

Hence, the public likelihood coincides with the posterior likelihood, $l_{t+1} = L_t$ and l_{t+1} evolves according to (3). When the state changes with probability $\varepsilon > 0$ in every period, upon observing the latest history, each agent also needs to consider the possibility that the state may have changed after the latest action was taken. However, neither the learning region nor the cascade region is affected by a changing state as the state only transitions after the history of past actions is observed. By Bayes' rule, the public likelihood at time $t + 1$ in the learning region is

$$\begin{aligned} l_{t+1} &:= \log \frac{q_{t+1}}{1 - q_{t+1}} \\ &= \log \frac{\mathbb{P}[\theta_{t+1} = +1 | h^{t-1}, a_t]}{\mathbb{P}[\theta_{t+1} = -1 | h^{t-1}, a_t]} \\ &= \log \frac{\sum_{\theta_t \in \{-1, +1\}} \mathbb{P}[\theta_{t+1} = +1, a_t | h^{t-1}, \theta_t] \mathbb{P}[\theta_t | h^{t-1}]}{\sum_{\theta_t \in \{-1, +1\}} \mathbb{P}[\theta_{t+1} = -1, a_t | h^{t-1}, \theta_t] \mathbb{P}[\theta_t | h^{t-1}]} \end{aligned} \quad (6)$$

Notice that the process $(\theta_t)_t$ is a Markov chain and $a_t = s_t$ in the learning region, conditioned on θ_t , both θ_{t+1} and a_t are independent of h^{t-1} and independent of each other. Hence, we can write (6) as

$$l_{t+1} = \begin{cases} \log \frac{(1-\varepsilon)\alpha e^{l_t} + \varepsilon(1-\alpha)}{\varepsilon\alpha e^{l_t} + (1-\varepsilon)(1-\alpha)} := f_+(l_t) & \text{if } s_t = +1, \\ \log \frac{(1-\varepsilon)(1-\alpha) e^{l_t} + \varepsilon\alpha}{\varepsilon(1-\alpha) e^{l_t} + (1-\varepsilon)\alpha} := f_-(l_t) & \text{if } s_t = -1. \end{cases} \quad (7)$$

Note that $f_+(l) > l$ and $f_-(l) < l$ and in particular, both $f_+(l)$ and $f_-(l)$ are strictly increasing in l . Intuitively, when an agent starts with a higher prior belief, her posterior belief will be higher upon receiving either an up or down signal. Similarly, an agent's posterior belief will also be higher (lower) than her prior belief upon receiving an up-signal (down-signal).

The following lemma summarizes two distinct patterns of the transition of the public likelihood in the learning region. From (7), one can 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 . At any time t , we say an action is *opposing* to the current public belief if $a_t \neq \text{sign}(l_t)$ and *supporting* otherwise.

Lemma 1. *For any signal precision $\alpha \in (1/2, 1)$ and probability of state change $\varepsilon \in (0, \alpha(1 - \alpha))$, if the public likelihood is in the learning region, then*

- (i) *observing one opposing action is sufficient to overturn its sign;*
- (ii) *at most two supporting actions are required to initiate a cascade.*

This lemma is in spirit close to the *overturning principle* in Sørensen (1996) but adapted to a changing state. Intuitively, the first part of Lemma 1 holds as the public belief in

the learning region is not too extreme and updates monotonically after opposing news. The second part of Lemma 1 holds because although the updating process of the public belief is slowed down by a changing state, observing consecutive good news is sufficient to start a cascade.

Another important observation is that regardless of whether the state is fixed or changing, the process of the public likelihoods $(l_t)_t$ in either case forms a Markov chain.¹² In the case of a fixed state, the state space of $(l_t)_t$ is finite as the magnitude difference between l_t and l_{t+1} is a constant for any fixed signal precision (e.g., c_α in our model). However, in the case of a changing state, the state space of $(l_t)_t$ is infinite as such magnitude differences also depend on the current value of l_t , resulting in many more possible values for l_t .¹³ This poses a significant challenge in finding the stationary distribution for this Markov process. We circumvent this problem by providing an upper bound to the expected time between the sign switches of the public likelihood next.

4.2. Expected Time Between State and Action Changes. We first calculate the expected time between state changes. Since the process $(\theta_t)_t$ is a simple Markov chain, the expected time between state changes is inversely proportional to the probability of state change, which is equal to $1/\varepsilon$. To see this, suppose that the expected time between state changes equals x . As $(\theta_t)_t$ follows a two-state Markov chain with a symmetric transition probability of ε , x satisfies

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

This implies that $x = 1/\varepsilon$. Intuitively, if the state becomes more likely to change in every period, it takes less time to change on average.

In contrast, the question of how long it takes on average for the action to change is more difficult as the process $(a_t)_t$ is not a Markov chain. Nevertheless, a_t is a function of l_{t+1} , namely, the sign of l_{t+1} , which is a Markov chain. However, as discussed before, this Markov chain is complicated as it has infinitely many possible values and different transition probabilities between them. The complexity of $(l_t)_t$ raises difficulty in directly analyzing the expected time between the sign switches of the public likelihood. To overcome this, we provide an upper bound instead and thus obtain an upper bound to the expected time between action changes.

To do so, let us first consider the maximum length of any cascade. Such a maximum exists as the public belief in the cascade region slowly converges towards uniformity. For any signal precision and probability of state change, since no cascade can last longer than the cascade starting at the supremum of the public likelihood $f_+(c_\alpha)$, one can calculate a tight upper bound to the length of any cascade from (5). Following Moscarini et al.

¹²To see that the process $(l_t)_t$ is a Markov chain, notice that conditioned on the state θ_t , the private signal s_t is independent of $l_{t'}$, for any $t' < t$. So its distribution conditioned on l_t is the same as its distribution conditioned on (l_1, \dots, l_{t-1}) which implies that $\mathbb{P}[l_{t+1} = l | l_1, \dots, l_t] = \mathbb{P}[l_{t+1} = l | l_t]$.

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

(1998), we denote this bound by¹⁴

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

Note that this upper bound $K(\alpha, \varepsilon)$ decreases in both α and ε . Intuitively, cascades contain more information relative to private signals as they become less precise. This information will depreciate even slower if the state becomes less volatile, so temporary cascades can last longer.

Given the maximum length of any cascade and the fact that whenever the public likelihood exists a cascade, it would either climb (gradually) back to the cascade region or switch its sign immediately (Lemma 1), one can obtain an upper bound to the expected time between the sign switches, as shown in the next proposition. For all $i = 1, 2, \dots$, denote the random time at which the public likelihood switches its sign for the i -th time by \mathcal{T}_i and let $\mathcal{T}_0 = 0$. Denote the random time elapsed between the $i - 1$ -th and i -th sign switch by $\mathcal{D}_i := \mathcal{T}_i - \mathcal{T}_{i-1}$.

Proposition 1. *For any signal precision $\alpha \in (1/2, 1)$, probability of state change $\varepsilon \in (0, \alpha(1 - \alpha))$ and positive integers $i \geq 2$, conditional on the public likelihood that just switched its sign for the $i - 1$ -th time, the expected time to the next sign switch is strictly bounded above:*

$$\mathbb{E}[\mathcal{D}_i | l_{\mathcal{T}_{i-1}}] < M(\alpha, \varepsilon),$$

where

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

Proposition 1 states that on average, one should expect the public likelihood to change its sign at least once every $M(\alpha, \varepsilon)$ periods. Note that $M(\alpha, \varepsilon)$ also decreases in both α and ε and thus $M(1/2, \varepsilon)$ is the greatest upper bound for any fixed ε . Intuitively, when 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, i.e., $\alpha = 1/2 + \delta$ where δ is strictly positive and close to 0. Denote the maximum length of any cascades as δ approaches zero by $K(\frac{1}{2}, \varepsilon)$. In this case, upon exiting a cascade, the probability of the public likelihood switching its sign is about $1/2$ as it is almost

¹⁴For completeness, we provide a similar calculation of this bound to the one in §3.B of Moscarini et al. (1998). Fix $\alpha \in (1/2, 1)$ and $\varepsilon \in (0, \alpha(1 - \alpha))$. Let m denote the supremum of the public belief and note that $m = \frac{(1-\varepsilon)\alpha^2 + \varepsilon(1-\alpha)^2}{\alpha^2 + (1-\alpha)^2}$. After h periods in a cascade, by (4), the public belief starting at m becomes

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

Since m is the supremum, any public belief after h periods in the cascade region has a value that is strictly less than $g(h)$. Thus, whenever $g(h) \leq \alpha$, or equivalently, $(1 - 2\varepsilon)^{h+1} \leq 1 - 2\alpha(1 - \alpha)$, the public likelihood exits the cascade region. Hence, the public likelihood can never be in the cascade region for more than $\frac{\log(1 - 2\alpha(1 - \alpha))}{\log(1 - 2\varepsilon)}$ periods.

equally likely that the agent, who follows her private signal, will receive an up or down signal. Thus, we can bound the expected time between the sign switches from above by a geometric distribution:

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

which decreases in ε . Since for ε sufficiently small, $M(1/2, \varepsilon) \approx (\log 2)/\varepsilon$ and it is always strictly less than $1/\varepsilon$, which is equal to the expected time between state changes.¹⁵ Since there is a one-to-one mapping between the public likelihood and the agent's action, the following result is an immediate corollary from Proposition 1.

Corollary 2. *For any signal precision $\alpha \in (1/2, 1)$ and probability of state change $\varepsilon \in (0, \alpha(1 - \alpha))$, the expected time between action changes is strictly less than the expected time between state changes.*

That is, on average, actions take less time to change than the state, even for a small probability of state change in which temporary cascades arise. For example, when the probability of state change is equal to 0.05, the state changes every twenty periods on average. In comparison, the maximum average time for the action to change is less than fourteen periods. Our main result (Theorem 1) then builds on this result by connecting the expected time between action changes with its long run 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, agents adjust accordingly so that their actions fluctuate over time. We show that in the long run, the relative frequency of action changes is higher than that of state changes, suggesting fads can emerge from social learning in a changing environment.

One could study the frequency of action changes for a single long-lived agent who repeatedly receives private signals about a changing state. We conjecture that action changes would be less frequent in this case than in our model, where only past actions are observable but still more frequent than the state. Intuitively, the frequency of unnecessary action changes would reduce by shutting down the channel of noisy observations of others' private signals. If this were the case, it would highlight the importance of observational learning in accelerating action fluctuations, especially when the underlying environment is slowly evolving.

One may wonder if the driving force behind our main result is due to the high frequency of action changes when the posterior belief is around $1/2$. Accordingly, we can further restrict the definition of fads to action changes that do not have consecutive changes, i.e., $a_t \neq a_{t-1}$ and $a_{t-1} \neq a_{t-2}$. Simulation results show that actions still change more

¹⁵For any two sequences a_n and b_n , we write $a_n \approx b_n$ if $\frac{a_n}{b_n} \rightarrow 1$ as $n \rightarrow \infty$.

frequently than the state, even under this more restricted definition of fads. For example, for $\alpha = 0.75$, $\varepsilon = 0.05$ and a total of 100,000 periods, the action changes about 8,150 times which is more frequent than the number of state changes, which is about 5,100 times.

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 given finite time N , 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 at least change about 8700 times. However, our numerical simulation shows that the action changes about 14,200 times, which is 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.
- (1998): “Learning from the behavior of others: Conformity, fads, and informational cascades,” *Journal of economic perspectives*, 12, 151–170.
- 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. PEŖSKI, 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*.
- SHAHRAPOUR, 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.
- SØRENSEN, P. N. (1996): “Rational social learning,” Ph.D. thesis, Massachusetts Institute of Technology.

APPENDIX A. PROOFS

Proof of Lemma 1. Fix $\alpha \in (1/2, 1)$, $\varepsilon \in (0, \alpha(1 - \alpha))$. By symmetry, it is without loss of generality to consider the case where $0 < l_t < c_\alpha$.

- (i) First notice that $f_-(c_\alpha) = 0$, and since $f_-(\cdot)$ in (7) is strictly increasing, $f_-(l_t) < 0$, for all $0 < l_t < c_\alpha$. Since l_t is in the learning region, $a_t = s_t$. Hence, $-1 = a_t \neq \text{sign}(l_t) = +1$ implies that $\text{sign}(l_{t+1}) = \text{sign}(f_-(l_t)) = -1$.
- (ii) As $f_+(\cdot)$ is strictly increasing, it suffices to show that $f_+(f_+(0)) \geq c_\alpha$. Note that for all $0 < \varepsilon < \alpha(1 - \alpha)$, $f_+(0) > c_u$ where

$$c_u := f_+^{-1}(c_\alpha) = \log \frac{(1 - \alpha)(\alpha - \varepsilon)}{\alpha(1 - \alpha - \varepsilon)} \in (0, c_\alpha)$$

is the threshold at which exactly one up-signal is required to push the public likelihood into the up-cascade region. Thus, $f_+(f_+(0)) > f_+(c_u)$ and by the definition of c_u , we have $f_+(f_+(0)) > f_+(c_u) = f_+(f_+^{-1}(c_\alpha)) = c_\alpha$, as required. \square

Proof of Proposition 1. Fix $\alpha \in (1/2, 1)$, $\varepsilon \in (0, \alpha(1 - \alpha))$ and some positive integer $i \geq 2$. Denote the greatest integer that is less than or equal to $K(\alpha, \varepsilon)$ by K and note that $K \geq 1$ since $K(\alpha, \varepsilon) \geq 1$. We will provide an upper bound to $\mathbb{E}[\mathcal{D}_i | l_{\mathcal{T}_{i-1}}]$, the expected time between the sign switches of the public likelihood. Consider first the case where $l_{\mathcal{T}_{i-1}} > 0$.

Denote the probability of receiving an up-signal conditioned on the public likelihood being l by $\pi(l)$.¹⁶ By the law of total probability and the fact that signals are independent of the public belief q conditional on the (current) state,

$$\begin{aligned}\pi(l) &= \mathbb{P}[s = +1|q = \frac{e^l}{1+e^l}] \\ &= \sum_{i \in \{+1, -1\}} \mathbb{P}[s = +1|q = \frac{e^l}{1+e^l}, \theta = i] \cdot \mathbb{P}[\theta = i|q = \frac{e^l}{1+e^l}] \\ &= \alpha \cdot \frac{e^l}{1+e^l} + (1-\alpha) \cdot \frac{1}{1+e^l} = \frac{1 + \alpha(e^l - 1)}{1+e^l}.\end{aligned}$$

Since $\partial\pi(l)/\partial l > 0$, the supremum of $\pi(l)$ over all $l \in (0, c_\alpha)$ is equal to $\pi(c_\alpha) = 1 - 2\alpha(1-\alpha)$, and we denote it by $\bar{\pi}$.

Let $\kappa(l)$ denote the length of the cascade initiated by receiving an up-signal conditioned on the public likelihood being l and let $\mathcal{L}(l)$ denote the value of this public likelihood after it first exits the cascade region. We further divide the range of the public likelihood $(0, f_+(c_\alpha))$ into three subregions: (i) $[c_u, c_\alpha)$ — the one up-signal away from the cascade region, where $c_u := f_1^{-1}(c_\alpha)$; (ii) $(0, c_u)$ — the two up-signal away from the cascade region, and (iii) $[c_\alpha, f_+(c_\alpha))$ — the up-cascade region. Next, we obtain an upper bound for each subregion.

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

$$\begin{aligned}\mathbb{E}[\mathcal{D}_i|l_{\mathcal{T}_{i-1}}] &= 1 - \pi(l_{\mathcal{T}_{i-1}}) + \pi(l_{\mathcal{T}_{i-1}}) \left(\kappa(l_{\mathcal{T}_{i-1}}) + \mathbb{E}[\mathcal{D}_i|\mathcal{L}(l_{\mathcal{T}_{i-1}})] \right) \\ &< 1 - \bar{\pi} + \bar{\pi} \left(\kappa(l_{\mathcal{T}_{i-1}}) + \mathbb{E}[\mathcal{D}_i|\mathcal{L}(l_{\mathcal{T}_{i-1}})] \right) \\ &\leq 1 - \bar{\pi} + \bar{\pi} \left(K + \mathbb{E}[\mathcal{D}_i|\mathcal{L}(l_{\mathcal{T}_{i-1}})] \right),\end{aligned}\tag{8}$$

where the second last strict inequality follows from the definition $\bar{\pi}$ and the last inequality follows from the definition of K . Note that there are two possible cases for the value of $\mathcal{L}(l_{\mathcal{T}_{i-1}})$. Case (i): suppose that $\mathcal{L}(l_{\mathcal{T}_{i-1}}) \in [c_u, c_\alpha)$. Then, by taking the supremum on both sides of (8) and rearranging, we obtain

$$\sup_{c_u \leq l_{\mathcal{T}_{i-1}} < c_\alpha} \mathbb{E}[\mathcal{D}_i|l_{\mathcal{T}_{i-1}}] \leq 1 + \frac{K\bar{\pi}}{1-\bar{\pi}}.\tag{9}$$

Case (ii): suppose that $\mathcal{L}(l_{\mathcal{T}_{i-1}}) \in [0, c_u)$. Then, similarly, by the definition of $\bar{\pi}$, we can write

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

¹⁶We suppress the dependence of $\pi(l)$ on α for the ease of notation.

By substituting the above inequality into (8), we have

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

Since $f_+(\mathcal{L}(l_{\mathcal{T}_{i-1}})) \in [c_u, c_\alpha]$ by part (ii) of Lemma 1, we can take the supremum on both sides again and obtain

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

so that (9) holds for all $l_{\mathcal{T}_{i-1}} \in [c_u, c_\alpha]$.

Next, consider $l_{\mathcal{T}_{i-1}} \in (0, c_u)$. By part (i) of Lemma (1) and the definition of $\bar{\pi}$, the expected time to the next sign switch is bounded above by

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

It follows from part (ii) of Lemma 1 that $f_+(l_{\mathcal{T}_{i-1}}) \in [c_u, c_\alpha]$, and thus by (9), we have that for all $l_{\mathcal{T}_{i-1}} \in (0, c_u)$,

$$\begin{aligned} \mathbb{E}[\mathcal{D}_i | l_{\mathcal{T}_{i-1}}] &< (1 - \bar{\pi}) + \bar{\pi}(1 + 1 + \frac{K\bar{\pi}}{1 - \bar{\pi}}) \\ &= \frac{(K - 1)(\bar{\pi})^2 + 1}{1 - \bar{\pi}}. \end{aligned} \tag{10}$$

Finally, consider $l_{\mathcal{T}_{i-1}} \in [c_\alpha, f_+(c_\alpha))$. In this case, after at most K periods, the public likelihood starting at $l_{\mathcal{T}_{i-1}}$ would have exited the cascade region. Hence, the expected time to the next sign switch is bounded above by

$$\mathbb{E}[\mathcal{D}_i | l_{\mathcal{T}_{i-1}}] \leq K + \mathbb{E}[\mathcal{D}_i | \mathcal{L}(l_{\mathcal{T}_{i-1}})].$$

Again, there are two possible cases for the value of $\mathcal{L}(l_{\mathcal{T}_{i-1}})$. Case (i): suppose that $\mathcal{L}(l_{\mathcal{T}_{i-1}}) \in [c_u, c_\alpha]$. It then follows from (9) that

$$\mathbb{E}[\mathcal{D}_i | l_{\mathcal{T}_{i-1}}] < K + 1 + \frac{K\bar{\pi}}{1 - \bar{\pi}} = 1 + \frac{K}{1 - \bar{\pi}}.$$

Case (ii): suppose that $\mathcal{L}(l_{\mathcal{T}_{i-1}}) \in [0, c_u)$. It then follows from (10) that

$$\mathbb{E}[\mathcal{D}_i | l_{\mathcal{T}_{i-1}}] < K + \frac{(K - 1)(\bar{\pi})^2 + 1}{1 - \bar{\pi}} = 1 + K + \bar{\pi} + \frac{K\bar{\pi}^2}{1 - \bar{\pi}} \leq 1 + \frac{K}{1 - \bar{\pi}}.$$

Together, for all $l_{\mathcal{T}_{i-1}} \in [c_\alpha, f_+(c_\alpha))$,

$$\mathbb{E}[\mathcal{D}_i | l_{\mathcal{T}_{i-1}}] < 1 + \frac{K}{1 - \bar{\pi}}. \tag{11}$$

Since the maximum over all three upper bounds from (9) to (11) is $1 + \frac{K}{1 - \bar{\pi}}$, and by the definition of K , $K \leq K(\alpha, \varepsilon)$, this implies that

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

The case where $l_{\tau_i} < 0$ follows from a symmetric argument. Thus, by the law of total probability, the above inequalities also hold unconditionally. \square

A.1. Proof of Theorem 1. Before proving our main theorem, we introduce the following lemma, which provides well-defined moments when working with the process $(\mathcal{D}_i)_i$. Recall that \mathcal{D}_i is the random time elapsed between the $i-1$ -th and i -th sign switch of the public likelihood.

Lemma 2. *Fix any signal precision $\alpha \in (1/2, 1)$ and probability of state change $\varepsilon \in (0, \alpha(1 - \alpha))$. Then, 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 $\alpha \in (1/2, 1)$, $\varepsilon \in (0, \alpha(1 - \alpha))$ and some positive integer $i \geq 2$. Consider without loss of generality the case where \mathcal{D}_i is the (random) time elapsed from a positive public likelihood to the next negative one. Fix any $n \geq 2$, denote the minimum number (may not be an integer) of cascades required for the event $\{\mathcal{D}_i > n\}$ by

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

Let $\bar{\pi} = 1 - 2\alpha(1 - \alpha) \in (1/2, 1)$ denote the supremum of the probability of receiving an up-signal over all $l \in (0, c_\alpha)$ (as seen in the proof of Proposition 1). It follows from Lemma 1 that for any $n \geq 2$, we can bound the probability of the event $\{\mathcal{D}_i > n\}$ by

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

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

$$\begin{aligned} \lim_{n \rightarrow \infty} n^p \mathbb{P}[|\mathcal{D}_i| > n] &= \lim_{n \rightarrow \infty} \frac{n^p}{1/\mathbb{P}[\mathcal{D}_i > n]} \\ &< \lim_{n \rightarrow \infty} \frac{n^p}{(1/\bar{\pi})^{1+\lfloor k(n) \rfloor}} = 0. \end{aligned} \tag{12}$$

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 (12) 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

In particular, this lemma implies that there is a finite uniform upper bound to $\mathbb{E}[\mathcal{D}_i^2]$, the second moment of \mathcal{D}_i , which is required for applying the standard martingale convergence theorem. The intuition behind this lemma is simple: since any cascade must end after $K(\alpha, \varepsilon)$ periods, the probability that \mathcal{D}_i is larger than some finite periods declines exponentially fast, implying that \mathcal{D}_i has finite moments.

Given Lemma 2 and Proposition 1, we are ready to prove our main theorem.

Proof of Theorem 1. Fix $\alpha \in (1/2, 1)$ and $\varepsilon \in (0, \alpha(1-\alpha))$. Since the process $(\theta_t)_t$ 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. It follows from the strong law of large numbers that

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

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 $j \leq i$. Hence, the sequence of random variables $(\mathcal{D}_1, \mathcal{D}_2, \dots)$ is adapted to Φ so that each \mathcal{D}_i is \mathcal{F}_i -measurable. By Proposition 1 and Corollary 2, there exists $\delta = 1/\varepsilon - M(\alpha, \varepsilon) > 0$ such that for all $i \geq 2$,

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

and thus by the law of total expectation and the Markov property of the public likelihood,

$$\mathbb{E}[\mathcal{D}_i | \mathcal{F}_{i-1}] = \mathbb{E}[\mathbb{E}[\mathcal{D}_i | \mathcal{F}_{i-1}, \mathcal{F}_{i-1}] | \mathcal{F}_{i-1}] < 1/\varepsilon - \delta. \quad (13)$$

Let $X_i = \mathcal{D}_i - \mathbb{E}[\mathcal{D}_i | \mathcal{F}_{i-1}]$ for all $i \geq 2$ and denote a partial sum of the process $(X_i)_i$ by

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

By the definition of X_i , $\mathbb{E}[X_i | \mathcal{F}_{i-1}] = 0$ for all $i \geq 2$. Since each Y_{n-1} is \mathcal{F}_{n-1} -measurable, it follows that the process $(Y_n)_n$ forms a martingale:

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

Note that $\mathbb{E}[X_i^2]$ is uniformly bounded for all $i \geq 2$ since both $\mathbb{E}[\mathcal{D}_i^2]$ and $\mathbb{E}[\mathcal{D}_i|\mathcal{F}_{i-1}]$ are uniformly bounded by Lemma 2 and (13). Hence, $\mathbb{E}[Y_n^2] = \sum_{i=2}^n \frac{1}{(i-1)^2} \mathbb{E}[X_i^2] < \infty$ for all n . By the martingale convergence theorem, Y_n converges almost surely, and it then follows from Kronecker's lemma that¹⁷

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

By the definition of X_i , we can write

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

It follows from (13) that

$$\lim_{n \rightarrow \infty} \frac{1}{n-1} \sum_{i=2}^n \mathcal{D}_i \leq 1/\varepsilon - \delta < 1/\varepsilon \quad \text{a.s.} \quad (14)$$

Since there is a one-to-one mapping between the action and the public likelihood, namely, $a_t = \text{sign}(l_{t+1})$ for all $t \geq 2$, and by the definitions of \mathcal{T}_i and \mathcal{D}_i ,

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathcal{Q}_a(n) &= \lim_{n \rightarrow \infty} \frac{1}{n-1} \sum_{t=2}^n \mathbb{1}(\text{sign}(l_{t+1}) \neq \text{sign}(l_{t+2})) \\ &= \lim_{n \rightarrow \infty} \frac{n-1}{\mathcal{T}_n - \mathcal{T}_1} = \lim_{n \rightarrow \infty} \frac{n-1}{\sum_{i=2}^n \mathcal{D}_i}. \end{aligned}$$

Finally, it follows from (14) that

$$\lim_{n \rightarrow \infty} \mathcal{Q}_a(n) > \varepsilon = \lim_{n \rightarrow \infty} \mathcal{Q}_\theta(n) \quad \text{a.s.}$$

□

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