ACTION VOLATILITY IN A CHANGING WORLD

WANYING HUANG

ABSTRACT. We study how volatile agents' behavior is in a social learning environment with a changing state. 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 actions change less often than private signals but, despite the presence of temporary information cascades, still more often than the underlying state in the long run.

1. INTRODUCTION

In many economic environments, individuals rely on the observed decisions of others to guide their own choices. At the same time, the conditions they are learning about—such as market demand, the value of new technologies, or organizational constraints—often change over time. In these situations, learning is not simply a matter of identifying a fixed state of the world, but of adjusting actions as new information arrives and circumstances evolve. This raises a natural question: does behavior adjust more or less frequently than the underlying conditions themselves? Understanding this relative frequency of adjustment is important in many domains in which behavioral volatility, or the lack thereof, can generate persistent structural mismatches and decouple market outcomes from underlying fundamentals.

As an illustrative example, consider a university student choosing a field to major in. In addition to her own private research, she can infer future labor-market prospects, such as wages and working conditions, by observing the choices of recent graduates. However, this social information is only partially relevant, since the relative attractiveness of different fields evolves over time. As a result, the popularity of certain majors may fluctuate, and these fluctuations may not perfectly reflect underlying changes in market demand, potentially creating a wedge between labor supply and demand. Similarly, in financial markets, when professional traders decide whether to buy or sell a stock, they often rely on past trading volumes to infer an asset's intrinsic value, which itself changes over time.

Date: February 15, 2026.
Department of Economics, Monash University. Email: kate.huang@monash.edu
Acknowledgements. This work is based on the third chapter of my PhD dissertation at Caltech and was previously circulated with the title "The emergence of fads in a changing world". I am indebted to Omer Tamuz for his continued support and encouragement. I would like to thank Arthur Campbell, Krishna Dasaratha, Keaton Ellis, 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.

Consequently, trading behavior that adjusts too slowly or too quickly to fundamentals can generate price deviations from intrinsic value, giving rise to speculative bubbles and fads (Camerer, 1989; Aggarwal and Rivoli, 1990).

To understand the extent to which actions fluctuate in these settings, we consider a canonical sequential social learning model in which agents aim to match a binary state. Unlike the standard fixed-state framework (Banerjee, 1992; Bikhchandani et al., 1992), the state evolves over time: in each period, it switches with a small, symmetric probability.[1] Agents arrive sequentially, each acting once and receiving a positive payoff if her action matches the current state. Before making a decision, each agent receives a private binary signal that is informative about the current state and observes the past actions of her predecessors. As the underlying state evolves, the optimal action fluctuates as well. The question we ask is: how frequently do actions change, and more specifically, do they change more or less often than the underlying state and the private signals that agents receive?

We focus on a setting in which the underlying state is relatively stable, so that the information inferred from past actions remains relevant and accumulates over time.[2] This allows information cascades to arise, in which social information overwhelms agents' private signals and leads them to mimic their predecessors even when their own signals suggest otherwise. However, as shown in Moscarini et al. (1998), these cascades can only be temporary: as the state evolves, the social information sustaining a cascade depreciates over time and becomes progressively less relevant to the current agent. Eventually, once this social information becomes sufficiently weak, the cascade breaks and agents resume responding to their private signals. This is in contrast to a static environment, where information cascades, once formed, can persist indefinitely.

Given the presence of information cascades, one might expect actions to fluctuate less frequently than the underlying state, since agents ignore their private signals during cascades and may therefore fail to adjust their actions when the state changes.[3] However, because cascades are temporary and signals are not perfectly informative about the state, agents may also switch actions unnecessarily. Therefore, the question of whether actions change more or less often than the state is *a priori* unclear.

[1]This is also known as a simple two-state Markov process. There are only a handful of papers studying social learning with a changing state; see, for example, Moscarini, Ottaviani, and Smith (1998); Hirshleifer and Welch (2002) and Lévy, Pęski, and Vieille (2024). Our model is most similar to that in Moscarini, Ottaviani, and Smith (1998), except for the tie-breaking rule. See a more detailed discussion in the literature review section.

[2]If, instead, the state were to evolve rapidly, past actions would quickly become outdated and convey little information. In that case, agents would rely on the most recent information, namely, their private signal. Action fluctuations would then be driven by signal noise, which is more volatile than the state itself.

[3]The symmetry of the binary state further amplifies this effect. For example, suppose the state changes an even number of times, say twice. An agent in a cascade may then incorrectly infer that the state is unchanged and thus have no reason to change her action.

Our main result (Theorem 1) answers this question by establishing a strict ordering between the volatility of actions, the underlying state, and private signals. We show that, despite the presence of temporary information cascades, actions fluctuate strictly more frequently than the underlying state but less frequently than private signals in the long run. In other words, while information cascades reduce excessive responsiveness to noisy signals, they do not eliminate excess action fluctuation relative to the state. It is worth emphasizing that, in this model, agents are fully rational, and these excessive action switches arise from their incentives to match the evolving state given their information, rather than from payoff externalities or heuristic behavior.

For example, consider a regime in which the state changes once every 100 periods on average and each private signal matches the current state with probability 0.8. In this case, signals change approximately once every three periods. We show that the expected time between action changes is strictly greater than three, but it is strictly less than sixty-one periods (see the upper bound in Proposition 2), which is much shorter than the average duration between state changes. As a result, the long-term frequency of action changes is lower than that of signal changes but still exceeds that of state changes, leading to faddish behavior in the long run.

The idea behind the proof of Theorem 1 is as follows. Intuitively, when agents are in a cascade, they simply mimic the actions of their predecessors, which dampens action volatility relative to that of the signal. However, this effect is limited by the bounded duration of cascades. Once a cascade ends, agents again respond to their private information and are therefore susceptible to noisy or conflicting signals. These periodic episodes of renewed responsiveness then generate excess action fluctuation relative to the state. Formally, to compare action volatility to that of the state and signals, we establish upper bounds on the expected times between cascade entries and action switches. The bound on cascade entries implies that cascades occur with positive frequency in the long run, thereby reducing action volatility relative to signals. Meanwhile, the bound on action switches is strictly lower than the expected time between state changes, implying that action switches occur more frequently than state changes on average. Taken together, these results imply that, in the long run, action volatility is lower than that of signals but higher than that of the state.

Related Literature. Despite an extensive literature on social learning,[4] only a small strand of papers considers environments with a changing state. Nevertheless, in the pioneering work of Bikhchandani et al. (1992), the idea of a changing state is briefly mentioned in their discussion of the fragility of information cascades. Through a numerical example, they show how a one-time shock to the state could break a cascade, even if such a shock is never realized, and that the probability of an action change is at least

[4]See Bikhchandani et al. (2024) for a comprehensive survey of recent developments in the social learning literature.

87% higher than the probability of a state change (see their Result 4), which aligns with our main result. They describe this phenomenon as "*seemingly whimsical swings in mass behavior without obvious external stimulus*" and refer to it as fads.

While Bikhchandani et al. (1992) presents an early notion 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. Indeed, as shown later in Moscarini, Ottaviani, and Smith (1998), if the underlying state evolves in every period but remains relatively stable, information cascades must arise, though they can only be temporary, i.e., it must end in finite time. Our work builds on their model and asks the following natural question: how does the long-run volatility of actions compare to that of signals and the underlying 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.[5] Different from their work—which examines how an agent's ability to recall past actions and signals affects action fluctuations in the short-term—our study investigates the long-term fluctuations in the actions of agents who observe past actions 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, 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.[6] In a similar but more general environment, Dasaratha, Golub, and Hak (2023) show that having sufficiently diverse network neighbors with different signal distributions improves learning efficiency, as such diversity allows agents to extract the most relevant information from old and confounded data.

In a setup similar to ours, a recent study by Lévy, Pęski, and Vieille (2024) 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

[5] Intuitively, 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.

[6] See 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.

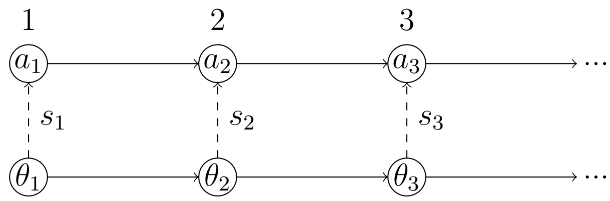

FIGURE 1. An illustration of processes (θ_t) and (a_t).

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 focus on this canonical sequential learning model without further complications as our objective is to compare the long-term relative frequency of action, signal, and state changes—a question that turns out to be nontrivial even in this simple setting.

2. Model

We follow the setup from Moscarini, Ottaviani, and Smith (1998) closely. Time is discrete, and the horizon is infinite, i.e., $t \in \mathbb{N}_+ = \{1, 2, \ldots\}$. 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 $\varepsilon \in (0, 1)$:

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

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, $h_{t-1} = (a_1, \ldots, a_{t-1})$. 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:

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

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 then 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 the information available to her, (h_{t-1}, s_t). A graphical illustration of the sequence of events is shown in Figure 1.

Throughout, we assume that the state is *relatively stable*: for any signal precision $\alpha \in (1/2, 1)$, the probability of a state change is $\varepsilon \in (0, \alpha(1-\alpha))$. Equivalently, one can think of this assumption as follows: in every period, with probability $2\varepsilon \in (0, 2\alpha(1-\alpha))$, the state is redrawn from the set $\{+1, -1\}$ with equal probability. Thus, the probability of a state change is equal to $\varepsilon \in (0, \alpha(1-\alpha))$. As shown in Moscarini et al. (1998), this assumption ensures information cascades can arise but only temporarily.

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

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

We call L_t the *posterior LLR* at time t. By Bayes' rule, it can be written as the sum of two terms:

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

We refer to the first term as the *public LLR* at time t and denote it by ℓ_t. This is the log-likelihood ratio of agent t's *public belief*, which is derived from the history of past actions h_{t-1} and is given by $q_t := \mathbb{P}[\theta_t = +1 \mid h_{t-1}]$. Since the private signal is conditionally independent of past actions given the current state, the second term in the above sum simplifies to the LLR induced by the signal alone. We define this value as $c_\alpha := \log \frac{\alpha}{1-\alpha}$. Thus, the second term equals c_α if $s_t = +1$ and $-c_\alpha$ if $s_t = -1$. Consequently, depending on the realization of the private signal, the posterior LLR evolves as follows:

$$L_t = \ell_t + c_\alpha \cdot s_t \tag{2.1}$$

The optimal action for agent t is the action that maximizes her expected payoff conditional on the information available to her:

$$a_t \in \operatorname*{arg\,max}_{a \in \{-1,+1\}} \mathbb{P}[\theta_t = a \mid h_{t-1}, s_t].$$

It thus follows from (2.1) that $a_t = +1$ 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 follows the action taken by her immediate predecessor, i.e., $a_t = a_{t-1}$.[7] This tie-breaking rule differs from that 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.

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

Cascade and Learning Regions. As mentioned before, an *information cascade* occurs when the social information inferred from past actions outweighs an agent's private signal, causing the agent to disregard her private information. From (2.1), one can see that this occurs precisely when $|\ell_t| > c_\alpha$. In this case, the sign of L_t—and therefore the optimal action for agent t—is determined solely by the sign of ℓ_t, regardless of the realization of s_t. Thus, when $\ell_t > c_\alpha$, agent t always chooses $a_t = +1$, and when $\ell_t < -c_\alpha$, she always chooses $a_t = -1$.

Similarly, when $|\ell_t| = c_\alpha$, the tie-breaking rule at indifference implies that agent t chooses the same action as agent $t-1$, regardless of the realization of s_t. Moreover, one can verify that in this case a_{t-1} equals the sign of ℓ_t.[8] We therefore refer to the region where $|\ell_t| \geq c_\alpha$ as the *cascade region*, in which actions are fully determined by social information. Conversely, when $|\ell_t| < c_\alpha$, the private signal can overturn the public belief, and thus agent t follows her private signal, i.e., $a_t = s_t$. We refer to this region as the *learning region*.

3. State and Signal Volatility

To study the long-run volatility of actions, it is useful to first examine the volatility of the underlying state and signals. 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 a state change corresponds to a shorter average time between changes. We denote the fraction of periods $t \leq n$ in which the state changes by

$$\mathcal{Q}_\theta(n) = \frac{1}{n}\sum_{t=1}^{n} \mathbb{1}(\theta_t \neq \theta_{t+1}).$$

By the strong law of large numbers,

$$\lim_{n\to\infty} \mathcal{Q}_\theta(n) = \varepsilon \quad \text{almost surely.}$$

Equivalently, an average *state duration*—defined as the number of consecutive periods during which the state remains unchanged—of $1/\varepsilon$ implies that the state changes in a fraction ε of periods in the long run. Similarly, we denote the fraction of periods $t \leq n$

[8] To see this, suppose without loss of generality that $\ell_t = c_\alpha$, so $\mathrm{sign}(\ell_t) = +1$. If $s_t = +1$, then $L_t = \ell_t + c_\alpha > c_\alpha$, which leads to $a_t = +1$. If instead $s_t = -1$, then $L_t = \ell_t - c_\alpha = 0$ and by the tie-breaking rule, $a_t = a_{t-1}$. In either case, $a_{t-1} = +1$ because $\ell_t = c_\alpha$ implies that either $\ell_{t-1} > c_\alpha$, in which case agent $t-1$ is in the cascade region and chooses $+1$, or $\ell_{t-1} < c_\alpha$ but a_{t-1} must be $+1$ since $\ell_t = c_\alpha$.

in which the signal changes by

$$\mathcal{Q}_s(n) = \frac{1}{n}\sum_{t=1}^{n} \mathbb{1}(s_t \neq s_{t+1}).$$

Since each signal matches the current state with probability α, it also follows from the strong law of large numbers that

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

and a direct calculation yields

$$\mathbb{P}[s_t \neq s_{t+1}] = \varepsilon + 2(1-2\varepsilon)\alpha(1-\alpha). \tag{3.1}$$

As is evident from (3.1), this probability is strictly greater than ε, and so signals switch more often than the state in the long run. This is simply because signals are not perfectly informative and thus exhibit additional volatility beyond that of the underlying state. As signals become more precise ($\alpha \to 1$), this excess volatility vanishes; in the limit of perfectly informative signals, signal changes occur as often as state changes.

4. Results

We now turn to the volatility of agents' actions. Recall that each agent receives a private signal about the current state and observes the history of past actions. Analogous to the volatility measures defined in Section 3, we denote the fraction of periods $t \leq n$ in which the action changes by

$$\mathcal{Q}_a(n) = \frac{1}{n}\sum_{t=1}^{n} \mathbb{1}(a_t \neq a_{t+1}).$$

Since the state is relatively stable, agents periodically enter temporary information cascades, during which they stop responding to their private signals and instead mimic the actions of their predecessors. As a result, one would expect actions to be less volatile than signals. A natural question, then, is whether this cascade-induced inertia is strong enough to align the volatility of actions with that of the underlying state. Our main result answers this question by establishing a strict ordering among the long-run frequencies of signal, action, and state switches.

Theorem 1. *In the long run, actions are less volatile than signals but more volatile than the state:*

$$\lim_{n\to\infty} \mathcal{Q}_s(n) > \lim_{n\to\infty} \mathcal{Q}_a(n) > \lim_{n\to\infty} \mathcal{Q}_\theta(n) \quad \textit{almost surely.}$$

Thus, Theorem 1 shows that while information cascades reduce action volatility relative to that of signals, agents nevertheless switch actions more frequently than the underlying state. As an example, consider a regime in which the state changes with probability 1% per period and private signals match the current state with probability 80%. In this case, the state changes once every 100 periods on average, whereas signals change roughly once every three periods.[9] Meanwhile, the long-run average duration between action changes lies strictly between these two cases: it is strictly greater than three periods and bounded above by 61 periods.[10] Taken together, this implies that, in the long run, action switches occur less frequently than signal switches but more frequently than state switches.

The next result, which can be obtained as a corollary of Theorem 1, further quantifies the extent to which actions are more volatile than the underlying state by providing a lower bound on their relative frequencies.

Corollary 1. *The long-run frequency of action switches is at least* $1/(\log 2 + \varepsilon)$ *times that of state switches:*

$$\lim_{n\to\infty} \mathcal{Q}_a(n) \geq \frac{1}{\log 2 + \varepsilon} \lim_{n\to\infty} \mathcal{Q}_\theta(n) \quad \textit{almost surely.}$$

In particular, as the state becomes more stable, this corollary implies that the long-run frequency of action switches is at least about 44 percent higher than the frequency of state changes.

The idea behind the proof of Theorem 1 is as follows. Intuitively, when agents are in a cascade, they simply mimic the actions of their predecessors, which dampens action volatility relative to that of signals. As the state evolves, however, the social information sustaining a cascade becomes less relevant to the current agent. Consequently, cascades are guaranteed to break after some finite period, and agents resume responding to their private signals. These episodes of renewed responsiveness then generate excessive action fluctuations, as agents at such times are particularly susceptible to opposing information.

Before formalizing this idea, note that because action changes are not independent events, we cannot directly relate the expected time between changes to the corresponding long-run frequency using the standard law of large numbers. Nevertheless, an agent's action is a function of her public belief, which itself is a Markov process.[11] We will utilize this Markov property to study how the public belief transitions across different regions and use these transitions to derive bounds on the expected time between action switches.

Formally, the proof of Theorem 1 proceeds in three steps. First, we show that the time between consecutive entries of the public belief process into the cascade region is uniformly bounded (Proposition 1). This implies that cascade entries—during which signals may change while actions do not—occur with positive frequency in the long run, thereby reducing action switches relative to signal switches. Second, we show that once the agent's public belief exits the cascade region, her action either changes or the public belief re-enters the same cascade region. By upper bounding the probability of the latter event,

[9]To see this, substituting $\alpha = 0.8$ and $\varepsilon = 0.01$ into (3.1) yields a probability of approximately 0.324.

[10]The upper bound follows from Proposition 2 by substituting $\alpha = 0.8$ and $\varepsilon = 0.01$ into $M(\alpha, \varepsilon)$, yielding $M(0.8, 0.01) \approx 61$.

[11]Since the state changes over time, the agent's public belief ceases to be a martingale, which is an important tool in analyzing long-term learning outcomes in fixed-state models (Smith and Sørensen, 2000).

we obtain an upper bound on the expected time between action switches (Proposition 2). Third, we compare this bound to the expected time between state changes and show that the former is strictly less than the latter. Theorem 1 then combines these bounds to compare the long-run frequencies of action, signal, and state switches, concluding that actions are less volatile than signals but more volatile than the state.

5. 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 bounds on the expected time between action changes, which will be useful in proving our main result. We provide a proof sketch of Theorem 1 at the end of the section.

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. Since signals are binary and imperfectly informative, the resulting cascade is formed based on limited information and thus can be incorrect with positive probability.[12]

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:

$$q_{t+1} = \mathbb{P}[\theta_{t+1} = +1 \mid h_t] = (1-\varepsilon)q_t + \varepsilon(1-q_t) = (1-2\varepsilon)q_t + (2\varepsilon)\frac{1}{2}. \tag{5.1}$$

From (5.1), we observe that q_{t+1} tends toward $1/2$, and so the public belief eventually exits the cascade region. As mentioned before, 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.

[12] More 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.

Learning Region. Next, we consider the learning region where $|\ell_t| < c_\alpha$. 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:

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

Hence, the corresponding log-likelihood ratios also coincide, i.e., $\ell_t = L_t$, and so ℓ_t also evolves according to (2.1).

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. Note that 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. By Bayes' rule, the public LLR at time $t+1$ satisfies

$$\begin{aligned} \ell_{t+1} &= \log \frac{\mathbb{P}[\theta_{t+1} = +1 \mid h_{t-1}, a_t]}{\mathbb{P}[\theta_{t+1} = -1 \mid h_{t-1}, a_t]} \\ &= \log \frac{\sum_{i \in \{-1,+1\}} \mathbb{P}[\theta_{t+1} = +1, a_t \mid h_{t-1}, \theta_t = i] \cdot \mathbb{P}[\theta_t = i \mid h_{t-1}]}{\sum_{i \in \{-1,+1\}} \mathbb{P}[\theta_{t+1} = -1, a_t \mid h_{t-1}, \theta_t = i] \cdot \mathbb{P}[\theta_t = i \mid h_{t-1}]}. \end{aligned} \tag{5.2}$$

Since θ_t follows a Markov chain and, in the learning region, the agent's action coincides with her signal ($a_t = s_t$), it follows that conditional on θ_t, both θ_{t+1} and a_t are independent of the history h_{t-1} and of each other. Therefore, if $s_t = +1$, the public LLR evolves according to

$$\ell_{t+1} = \log \frac{(1-\varepsilon)\alpha \mathrm{e}^{\ell_t} + \varepsilon(1-\alpha)}{\varepsilon \alpha \mathrm{e}^{\ell_t} + (1-\varepsilon)(1-\alpha)}.$$

If $s_t = -1$, it evolves as

$$\ell_{t+1} = \log \frac{(1-\varepsilon)(1-\alpha)\mathrm{e}^{\ell_t} + \varepsilon\alpha}{\varepsilon(1-\alpha)\mathrm{e}^{\ell_t} + (1-\varepsilon)\alpha}.$$

In sum, when ℓ_t is in the learning region, its evolution in (5.2) can be written as

$$\ell_{t+1} = G_\varepsilon(\ell_t + c_\alpha \cdot s_t),$$

where

$$G_\varepsilon(x) := \log \frac{(1-\varepsilon)\mathrm{e}^x + \varepsilon}{1 - \varepsilon + \varepsilon \mathrm{e}^x}. \tag{5.3}$$

When ℓ_t is in the cascade region, the update rule in (5.1) similarly reduces, in terms of the log-likelihood ratio, to

$$\ell_{t+1} = G_\varepsilon(\ell_t).$$

Note that $G_\varepsilon(\cdot)$ is strictly increasing. Thus, in the learning region, belief updating is monotone in both the public belief and the realized signal: $G_\varepsilon(\ell + c_\alpha) > \ell > G_\varepsilon(\ell - c_\alpha)$. Moreover, as seen in (5.1), because the state is evolving, belief updating is attenuated toward the uniform belief while preserving its direction: for all $x > 0$, $0 < G_\varepsilon(x) < x$ and for all $x < 0$, $0 > G_\varepsilon(x) > x$.

From above, 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.[13] 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.

Lemma 1. *For any ℓ_t such that $|\ell_t| < c_\alpha$, the following two conditions hold.*
(i) $a_t \neq sign(\ell_t)$ implies that $sign(\ell_{t+1}) = -sign(\ell_t)$.
(ii) $a_t = a_{t+1} = sign(\ell_t)$ implies that $|\ell_{t+2}| > c_\alpha$.

The first part of this lemma states that when the public belief is in the learning region, 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, observing consecutive supporting evidence is 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.[14] 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. 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 bounds to this expected time instead.

Bounds on Expected Durations. To bound the expected time between the sign switches of the public LLR, consider first the maximum length of any cascade. Such a maximum exists because the public belief in the cascade region slowly converges toward uniformity. Moreover, for any given signal precision and probability of a state change, no cascade can last longer than the one starting at $G_\varepsilon(2c_\alpha)$, as this is the supremum of the public LLR. Using the evolution of the public LLR in (5.3), one can therefore calculate a tight upper bound on the length of any cascade, which we denote by

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

For completeness, we provide in the appendix (see Lemma 2) a derivation of $K(\alpha, \varepsilon)$ analogous to the calculation in Section 3.B of Moscarini, Ottaviani, and Smith (1998). Notice that $K(\alpha, \varepsilon)$ decreases in both α and ε. As private signals become less precise,

[13]This lemma is in spirit close to the *overturning principle* (Smith and Sørensen, 2000), but it applies to a changing state.

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

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. Taken together, this suggests that prolonged cascades may result from either less precise private signals or a more stable environment.

Building on the cascade-length bound, we now establish two critical upper bounds. The first concerns the expected time between consecutive entries into the cascade region. Because actions remain constant within the cascade region while private signals continue to fluctuate, bounding this inter-entry time is essential for comparing the long-run frequency of action switches with that of signal switches. The second bound concerns the expected time between sign switches of the public LLR. Since there is a one-to-one mapping between agents' actions and the sign of the public LLR—more specifically, $a_t = \text{sign}(\ell_{t+1})$—this bound is a key ingredient in comparing the frequency of action switches with that of state switches.

Formally, for each integer $i = 1, 2, \ldots$, let $\mathcal{T}_i$ and $\mathcal{C}_i$ denote the (random) times at which the public LLR switches sign and enters the cascade region from the learning region, respectively, for the i-th time. We adopt the conventions $\mathcal{T}_0 = 0$ and $\mathcal{C}_0 = 0$. Define $\mathcal{D}_i := \mathcal{T}_i - \mathcal{T}_{i-1}$ as the duration between the $(i-1)$-th and i-th sign switches. Similarly, define the inter-entry times by $\mathcal{J}_i = \mathcal{C}_i - \mathcal{C}_{i-1}$.

The next proposition shows that the expected time between successive entries of the public LLR process into the cascade region is uniformly bounded.

Proposition 1. *For any integer $i \geq 2$, conditional on the public LLR entering the cascade region from the learning region for the $(i-1)$-th time, the expected time until the next entry into the cascade region satisfies*

$$\mathbb{E}[\mathcal{J}_i \mid \ell_{\mathcal{C}_{i-1}}] \leq \lfloor K(\alpha, \varepsilon) \rfloor + 8.$$

The proof of Proposition 1 relies on the observation that the public belief process alternates between the cascade and learning regions. Accordingly, the bound consists of two components: the maximum duration spent inside a cascade and the expected time spent in the learning region. Since two supporting signals suffice to trigger a cascade (Lemma 1), the latter is uniformly bounded above.

Next, we derive an upper bound on the expected time between sign switches of the public belief, providing an upper bound on the expected time between action switches.

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

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

This result implies 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. Moreover, notice that $M(\alpha, \varepsilon)$ is strictly decreasing in α (see Claim 1 in the appendix). As a result, $M(1/2, \varepsilon)$ provides the maximal upper bound for any $\varepsilon \in (0, \alpha(1-\alpha))$.

The proof idea behind Proposition 2 is as follows. Since $K(\alpha, \varepsilon)$ decreases in α, the greatest upper bound on the length of any cascade is obtained as $\alpha \to 1/2$, in which case $K(1/2, \varepsilon) = -\log 2 / \log(1 - 2\varepsilon)$. For such weakly informative signals, once a cascade ends and the process returns to the learning region, agents follow their private signals, which can be either positive or negative with nearly equal probability. As a result, upon exiting a cascade, the probability that the public LLR switches its sign also approaches 1/2. Since each cascade entry can last at most $K(1/2, \varepsilon)$ periods, the time between sign switches of the public LLR is stochastically dominated by the product of $K(1/2, \varepsilon)$ and the number of cascade entries before a sign switch occurs. The latter is geometrically distributed with success probability approximately 1/2. Therefore, the expected time until the next sign switch of the public LLR is bounded above by

$$1 + \sum_{i=1}^{\infty} \frac{i}{2^i} K(1/2, \varepsilon) = 1 + \frac{2 \log 2}{-\log(1 - 2\varepsilon)} = M(1/2, \varepsilon). \tag{5.4}$$

Proof Sketch of Theorem 1. We end this section by providing a proof sketch of Theorem 1. First, to see that action volatility is strictly lower than signal volatility, observe that upon each entry into a cascade, there is a uniformly positive probability that signals switch while actions remain constant. Since the expected time between successive entries is uniformly bounded above (Proposition 1), it follows that signals fluctuate more frequently than actions in the long run. Second, to see why action volatility is strictly higher than state volatility, consider a weakly informative signal where α is close to 1/2. By Proposition 2, $M(1/2, \varepsilon)$ is the largest upper bound on the expected time between sign switches of the public LLR. Using a standard approximation for (5.4), for ε small enough, $M(1/2, \varepsilon)$ can be approximated by $(\log 2)/\varepsilon$, which is strictly less than $1/\varepsilon$, the expected time between state changes.[15] Hence, on average, it takes less time for the action to switch than for the state, implying that in the long run, actions switch more often than the state.

6. Numerical Simulations

While Theorem 1 establishes a strict ordering of long-run action volatility relative to signals and the state, it leaves open the question of magnitude: how large are the gaps between these frequencies of change? Moreover, to what extent is the excess volatility of actions driven by the informational coarseness of observing past actions rather than

[15] In fact, in Claim 2 in the appendix, we show that this strict inequality $M(1/2, \varepsilon) < 1/\varepsilon$ holds for all $\varepsilon \in (0, 1/4)$. Since $M(\alpha, \varepsilon)$ is strictly decreasing in α, it follows that $M(\alpha, \varepsilon) < 1/\varepsilon$ for all $\varepsilon \in (0, \alpha(1-\alpha))$.

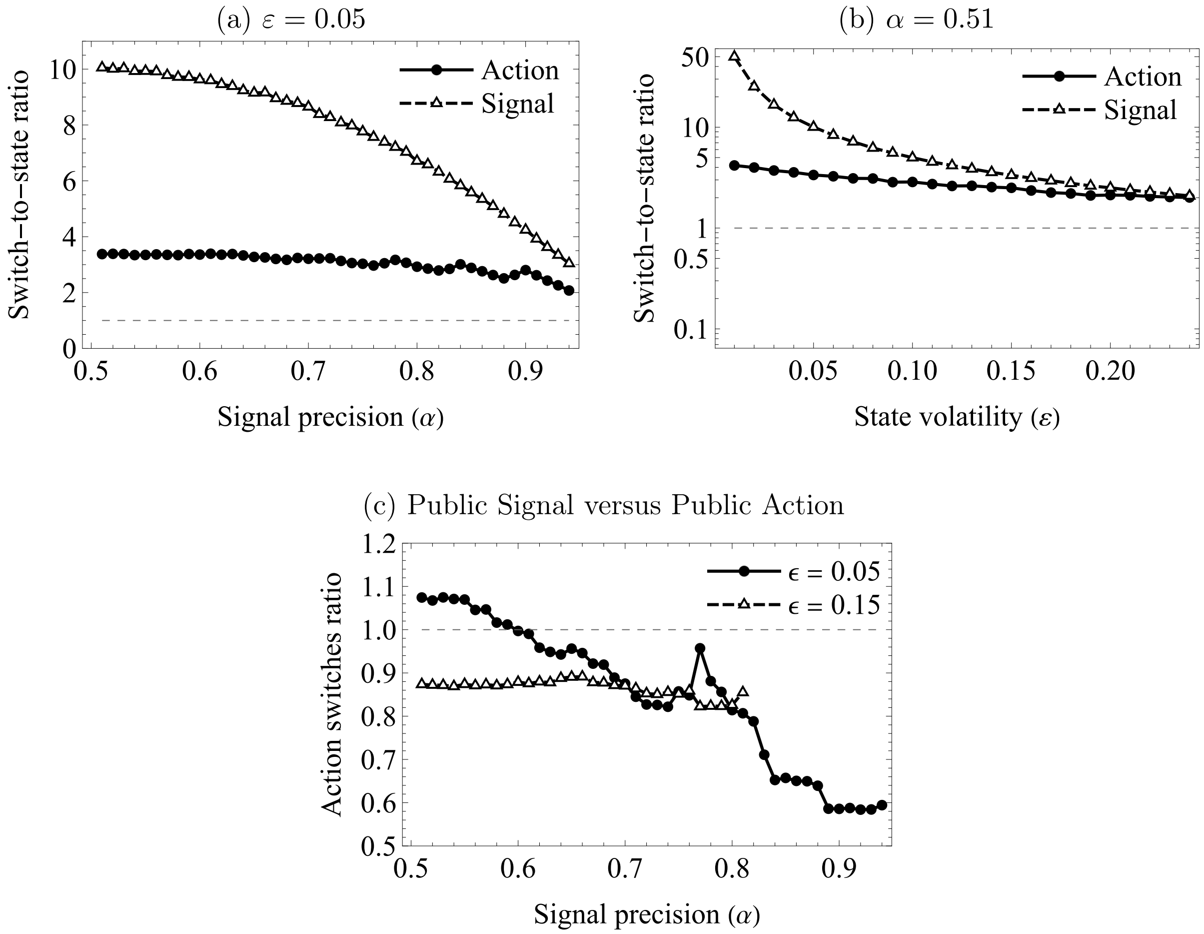

FIGURE 2. On the top left: the ratios between action (signal) switches and state switches as functions of signal precision, with state volatility fixed at 5%. The top right figure plots the same ratios as functions of state volatility, with signal precision fixed at 0.51. On the bottom: the ratio of action switch frequencies when past signals are public to those when only past actions are public as a function of signal precision α for $\varepsilon \in \{0.05, 0.15\}$. All dashed horizontal lines correspond to a ratio of one.

past signals? To shed light on these questions, we use numerical simulations to quantify action volatility under different parameter values and information regimes.

We simulate the frequencies of signal, action, and state changes over one million periods for different values of signal precision (α) and state volatility (ε) under the assumption that the state is relatively stable. In addition to our main setting, in which agents observe past actions, we also consider an alternative regime in which all past signals are publicly observable. We refer to the former as the *public-action* setting and the latter as the *public-signal* setting. We then compute (i) the ratio of action switches to state changes, (ii) the ratio of signal switches to state changes, and (iii) the ratio of action-switch frequencies across the two regimes.

The results are depicted in Figure 2. On the top left, we fix state volatility at 0.05 and plot the ratios of action switches (when past actions are observable) to state switches and of signal switches to state switches while varying signal precision. As signal precision increases, both action and signal switches become less frequent relative to the state. However, the reduction is much more pronounced for signal switches than for action switches: as α increases, the two ratios converge; nevertheless, even at high signal precision ($\alpha = 0.9$), actions still switch at about twice the rate of state changes, exceeding the lower bound suggested in Corollary 1. On the top right, we fix signal precision at 0.51 and vary state volatility. As state volatility increases, both signals and actions switch more frequently. However, as shown in the figure, the action-to-state and signal-to-state ratios decline, indicating that state changes accelerate faster than either action or signal changes.

On the bottom of Figure 2, we plot the ratio of action-switch frequencies in the public-signal setting to those in the public-action setting, as a function of signal precision for different levels of state volatility. Thus, a ratio above one indicates that action volatility is higher in the public-signal setting; a ratio below one indicates the opposite. A priori, it is unclear which regime induces greater action volatility. On the one hand, because past actions are noisy, binary summaries of agents' private signals—which are themselves noisy observations of the state—replacing actions with signals could reduce unnecessary action switches. On the other hand, when all past signals are publicly observable, agents never become stuck in cascades, so their beliefs remain responsive to each new signal, which may in turn lead to more frequent action switching.

Indeed, as shown in the figure, replacing actions with signals sometimes reduces long-run action volatility and sometimes exacerbates it. For example, when the state is very stable ($\varepsilon = 0.05$) and signals are very imprecise ($\alpha = 0.51$), action volatility is higher in the public-signal setting than in the public-action setting. Intuitively, when the state is sufficiently stable, past information depreciates slowly, so cascades last longer, which dampens action changes in the public-action setting. Meanwhile, in this stable environment, if past signals were publicly observable, then agents would never stop incorporating new signals, and such persistent responsiveness would generate more action switches under very noisy signals. In contrast, we see that when the state is moderately volatile ($\varepsilon = 0.15$), the informational advantage of observing precise past signals outweighs the stabilizing effect of temporary cascades. In this case, action volatility is lower in the public-signal setting than in the public-action setting across all levels of signal precision.[16] These simulations suggest that the region of the parameter space in which action volatility is higher in the public-signal setting is relatively small.

[16]The assumption of a relatively stable state, $\varepsilon \in (0, \alpha(1-\alpha))$, implies that for any fixed ε, signal precision is bounded above by $\bar{\alpha}(\varepsilon) = (1 + \sqrt{1-4\varepsilon})/2$. Hence, if $\varepsilon = 0.15$, the maximum signal precision is approximately 0.82.

7. 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, actions switch less frequently than signals but more frequently than the state itself. This result holds despite the presence of temporary information cascades in which agents simply mimic their predecessors' actions.

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 measurement of action volatility 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 restrictive counting, actions still change more frequently than the state. For example, with $\alpha = 0.75$, $\varepsilon = 0.05$, and a total of $100,000$ periods, action changes occur approximately 8,150 times, compared to the number of state changes, which is about 5,100 times.

There are several possible avenues for future research. As seen in Section 6, action volatility need not be uniformly higher or lower in the public-signal setting than in the public-action setting; rather, the ranking depends on the underlying parameters. It would therefore be interesting to characterize the regions in which each information structure generates higher action volatility. It is also natural to study the corresponding long-run mistake probability under these information regimes. We leave these questions for future work, as they currently seem to be beyond what is technically tractable.

Another open question concerns the tightness of the bound in Proposition 2. In particular, one could ask whether this upper bound $M(\alpha, \varepsilon)$ is tight, and if so, whether for any finite time N, the number of action changes is well approximated by $N/M(\alpha, \varepsilon)$. Our simulation results suggest that this is not the case. For example, when $\alpha = 0.9$ and $\varepsilon = 0.05$, and $N = 100,000$, we have $M(0.9, 0.05) \approx 11.5$, which would predict at least about 8,700 action changes. In contrast, our simulations exhibit approximately 14,200 action changes—almost twice as many as implied by the bound. Furthermore, the simulations suggest that as the private signal becomes less informative and the state evolves 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 action volatility 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, O. Tamuz, and I. Welch (2024): “Information cascades and social learning,” *Journal of Economic Literature*, 62, 1040–1093.

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 (2023): “Learning from neighbours about a changing state,” *Review of Economic Studies*, 90, 2326–2369.

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. Pęski, and N. Vieille (2024): “Stationary social learning in a changing environment,” *Econometrica*, 92, 1939–1966.

Moscarini, G., M. Ottaviani, and L. Smith (1998): “Social learning in a changing world,” *Economic Theory*, 11, 657–665.

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. Omitted Proofs from the main text

Proof of Lemma 1. Consider the case where $\ell_t \in (0, c_\alpha)$. For part (i), if $a_t = -1 \neq \text{sign}(\ell_t)$, then given that ℓ_t is in the learning region, we have $s_t = a_t = -1$. Since G_ε in (5.3) is strictly increasing, and $G_\varepsilon(0) = 0$, we have that $\ell_{t+1} = G_\varepsilon(\ell_t - c_\alpha) < 0$ for all $\ell_t \in (0, c_\alpha)$. Thus, $\text{sign}(\ell_{t+1}) = -1 = -\text{sign}(\ell_t)$.

For part (ii), it suffices to show that $G_\varepsilon(G_\varepsilon(0 + c_\alpha) + c_\alpha) \geq c_\alpha$ since G_ε is strictly increasing. Let c_u be the unique solution to $G_\varepsilon(c_u + c_\alpha) = c_\alpha$, and so it is the threshold at which exactly one positive signal is required for the public LLR to enter the cascade region on the positive action. A direct calculation yields that

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

Note that $G_\varepsilon(c_\alpha) > c_u$ for all $\varepsilon \in (0, \alpha(1-\alpha))$, and since G_ε is strictly increasing, $G_\varepsilon(G_\varepsilon(c_\alpha) + c_\alpha) > G_\varepsilon(c_u + c_\alpha) = c_\alpha$. Hence, we conclude that

$$G_\varepsilon(G_\varepsilon(c_\alpha) + c_\alpha) > c_\alpha,$$

as required. The case where $\ell_t \in (-c_\alpha, 0)$ follows an analogous argument. $\square$

Lemma 2. *The maximum length of any temporary information cascade is given by*

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

Proof. Since $G_\varepsilon(2c_\alpha)$ is the supremum of the public LLR, the corresponding supremum of the public belief, m, is given by

$$m = \frac{(1-\varepsilon)\alpha^2 + \varepsilon(1-\alpha)^2}{\alpha^2 + (1-\alpha)^2}.$$

Within the cascade region, the public belief evolves deterministically according to (5.1). Starting from an initial value x, iterating this update rule h times yields

$$g(h, x) := \varepsilon \sum_{i=0}^{h-1} (1-2\varepsilon)^i + (1-2\varepsilon)^h x.$$

Consequently, after spending h periods in the cascade region, any public belief must be strictly less than $g(h, m)$, where summing the geometric series yields

$$\begin{aligned} g(h, m) &= \varepsilon \left[\frac{1-(1-2\varepsilon)^h}{2\varepsilon} \right] + (1-2\varepsilon)^h m \\ &= \frac{1}{2} + (1-2\varepsilon)^h \left(m - \frac{1}{2} \right). \end{aligned}$$

The cascade terminates when the belief drops to α. Thus, we examine the condition $g(h, m) \le \alpha$, which is equivalent to

$$(1-2\varepsilon)^h \left(m - \frac{1}{2} \right) \le \alpha - \frac{1}{2}. \tag{A.1}$$

Substituting the expression for m, we obtain

$$m - \frac{1}{2} = \frac{(1-\varepsilon)\alpha^2 + \varepsilon(1-\alpha)^2}{\alpha^2 + (1-\alpha)^2} - \frac{1}{2} = \frac{(1-2\varepsilon)(2\alpha-1)}{2(\alpha^2 + (1-\alpha)^2)}.$$

Plugging this back into (A.1) yields

$$(1-2\varepsilon)^h \cdot \frac{(1-2\varepsilon)(2\alpha-1)}{2(\alpha^2 + (1-\alpha)^2)} \le \frac{2\alpha-1}{2}.$$

Since $\alpha > 1/2$, we can divide both sides by $\frac{2\alpha-1}{2}$. Noting that the denominator $\alpha^2 + (1-\alpha)^2 = 1 - 2\alpha(1-\alpha)$, the inequality simplifies to

$$(1-2\varepsilon)^{h+1} \leq 1 - 2\alpha(1-\alpha).$$

Solving for h yields the maximum duration of a temporary information cascade:

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

□

Claim 1. *The upper bound $M(\alpha,\varepsilon) = 1 + \frac{K(\alpha,\varepsilon)}{2\alpha(1-\alpha)}$ is strictly decreasing in α.*

Proof. Let $x(\alpha) = 2\alpha(1-\alpha)$. For $\alpha \in (1/2, 1)$, $x(\alpha)$ takes values in $(0, 1/2)$ and is strictly decreasing. Using this change of variable, we write

$$M(x,\varepsilon) = 1 + \frac{1}{\log(1-2\varepsilon)} \cdot f(x),$$

where $f(x) = \frac{\log(1-x)}{x}$. Note that $\frac{1}{\log(1-2\varepsilon)} < 0$ for $\varepsilon \in (0, 1/2)$. Differentiating $f(x)$ with respect to x yields

$$f'(x) = \frac{1}{x^2}\left[-\frac{x}{1-x} - \log(1-x)\right].$$

Let $g(x) = -\frac{x}{1-x} - \log(1-x)$. Since $g(0) = 0$ and $g'(x) = \frac{-x}{(1-x)^2} < 0$ for $x > 0$, it follows that $g(x) < 0$ for all $x \in (0, 1/2)$. Consequently, $f'(x) < 0$, meaning that $f(x)$ is strictly decreasing in x. Finally, applying the chain rule, we have

$$\frac{dM}{d\alpha} = \frac{1}{\log(1-2\varepsilon)} \cdot f'(x) \cdot x'(\alpha).$$

Since $\frac{1}{\log(1-2\varepsilon)} < 0$, $f'(x) < 0$, and $x'(\alpha) < 0$, the product is negative. Thus, $M(\alpha, \varepsilon)$ is strictly decreasing in α. □

We introduce the following notation, which will be useful in the proofs of Proposition 1 and 2. Let $\pi(\ell)$ denote the probability of receiving a positive signal conditional on the public LLR being ℓ.[17] By the law of total probability,

$$\pi(\ell) = \alpha \cdot \frac{\mathrm{e}^{\ell}}{1+\mathrm{e}^{\ell}} + (1-\alpha)\cdot\frac{1}{1+\mathrm{e}^{\ell}} = \frac{1+\alpha(\mathrm{e}^{\ell}-1)}{1+\mathrm{e}^{\ell}},$$

which is strictly increasing in ℓ. Therefore,

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

Note that $\pi(0) = 1/2$. Consequently, for any $\ell \neq 0$, the probability of receiving a supporting signal (i.e., $\pi(\ell)$ if $\ell > 0$ and $1 - \pi(\ell)$ if $\ell < 0$) is strictly greater than 1/2.

[17] For ease of notation, we suppress its dependence on α.

Proof of Proposition 1. We decompose $\mathcal{J}_i$ into the time spent in the cascade region and the subsequent time spent in the learning region. By Lemma 2, the former is deterministically bounded by $\lfloor K(\alpha, \varepsilon) \rfloor$.

For the latter, part (ii) of Lemma 1 implies that two consecutive supporting actions suffice to re-enter the cascade region. Moreover, in the learning region $a_t = s_t$, and the probability of a supporting signal is at least $1/2$. Therefore, conditional on any history upon being in the learning region, the probability of getting two consecutive supporting actions over the next two periods is at least $(1/2)^2 = 1/4$. Hence the waiting time from entering the learning region to the next cascade entry is stochastically dominated by $2G$ where G is a geometric random variable with success probability $1/4$. Since $\mathbb{E}[2G] = 8$, we obtain the result. □

Proof of Proposition 2. Fix a positive integer $i \geq 2$. Recall that $\mathcal{T}_{i-1}$ denotes the time at which the public LLR changes sign for the $(i-1)$-th time, and we defined the duration of the i-th run by $\mathcal{D}_i = \mathcal{T}_i - \mathcal{T}_{i-1}$. We will consider the case $\ell_{\mathcal{T}_{i-1}} > 0$ as the case where $\ell_{\mathcal{T}_{i-1}} < 0$ follows from an analogous argument. Let c_u, as in the proof of Lemma 1, be the unique solution to $G_\varepsilon(c_u + c_\alpha) = c_\alpha$.

There are three disjoint intervals for the value of $\ell_{\mathcal{T}_{i-1}}$: (i) $[c_u, c_\alpha)$; (ii) $(0, c_u)$, or (iii) $[c_\alpha, G_\varepsilon(2c_\alpha))$. We will show that in each case

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

where $K(\alpha, \varepsilon)$ is the upper bound on the length of any cascade.

To this end, let $\kappa(\ell)$ denote the length of a positive cascade triggered by receiving a positive signal conditional on the public LLR being ℓ. Let $\mathcal{L}(\ell)$ denote the value of the public LLR upon first exiting the cascade region. Finally, let $K := \lfloor K(\alpha, \varepsilon) \rfloor \geq 1$ 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 can be written as

$$\mathbb{E}[\mathcal{D}_i \mid \ell_{\mathcal{T}_{i-1}}] = 1 - \pi(\ell_{\mathcal{T}_{i-1}}) + \pi(\ell_{\mathcal{T}_{i-1}}) \Big(\kappa(\ell_{\mathcal{T}_{i-1}}) + \mathbb{E}[\mathcal{D}_i \mid \mathcal{L}(\ell_{\mathcal{T}_{i-1}})] \Big).$$

Using $\pi(\ell) \leq \bar{\pi}$ for $\ell \in (0, c_\alpha)$ from (A.2) and $\kappa(\ell) < K$, we obtain the bound

$$\mathbb{E}[\mathcal{D}_i \mid \ell_{\mathcal{T}_{i-1}}] < 1 - \bar{\pi} + \bar{\pi} \Big(K + \mathbb{E}[\mathcal{D}_i \mid \mathcal{L}(\ell_{\mathcal{T}_{i-1}})] \Big). \tag{A.3}$$

Now, there are two possibilities for $\mathcal{L}(\ell_{\mathcal{T}_{i-1}})$. If $\mathcal{L}(\ell_{\mathcal{T}_{i-1}}) \in [c_u, c_\alpha)$, then taking the supremum on both sides and rearranging yields

$$\sup_{c_u \leq \ell_{\mathcal{T}_{i-1}} < c_\alpha} \mathbb{E}[\mathcal{D}_i \mid \ell_{\mathcal{T}_{i-1}}] \leq 1 + \frac{\bar{\pi} K}{1 - \bar{\pi}}. \tag{A.4}$$

If $\mathcal{L}(\ell_{\mathcal{T}_{i-1}}) \in (0, c_u)$, then one more positive signal pushes the LLR into $[c_u, c_\alpha)$. Again, using $\pi(\ell) \leq \bar{\pi}$, we get

$$\mathbb{E}[\mathcal{D}_i \mid \mathcal{L}(\ell_{\mathcal{T}_{i-1}})] < 1 - \bar{\pi} + \bar{\pi}\Big(1 + \mathbb{E}[\mathcal{D}_i \mid G_\varepsilon(\mathcal{L}(\ell_{\mathcal{T}_{i-1}}) + c_\alpha)]\Big).$$

Substituting the above inequality into (A.3) yields

$$\mathbb{E}[\mathcal{D}_i \mid \ell_{\mathcal{T}_{i-1}}] < 1 - \bar{\pi} + \bar{\pi}\Big(K + 1 - \bar{\pi} + \bar{\pi}\big(1 + \mathbb{E}[\mathcal{D}_i \mid G_\varepsilon(\mathcal{L}(\ell_{\mathcal{T}_{i-1}}) + c_\alpha)]\big)\Big)$$

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

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

which is less than the upper bound in (A.4) since $K \geq 1$.

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:

$$\begin{aligned} \mathbb{E}[\mathcal{D}_i \mid \ell_{\mathcal{T}_{i-1}}] &< (1 - \bar{\pi}) + \bar{\pi}(1 + \mathbb{E}[\mathcal{D}_i \mid G_\varepsilon(\ell_{\mathcal{T}_{i-1}} + c_\alpha)]) \\ &< (1 - \bar{\pi}) + \bar{\pi}(1 + 1 + \frac{\bar{\pi}}{1 - \bar{\pi}}K) \\ &= \frac{\bar{\pi}^2(K-1) + 1}{1 - \bar{\pi}} \end{aligned} \tag{A.5}$$

where the second inequality follows from the fact that $G_\varepsilon(\ell_{\mathcal{T}_{i-1}} + c_\alpha) \in [c_u, c_\alpha)$ and (A.4).

Case (iii). Suppose $\ell_{\mathcal{T}_{i-1}} \in [c_\alpha, G_\varepsilon(2c_\alpha))$. In this case, after at most K 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:

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

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.4) that

$$\mathbb{E}[\mathcal{D}_i \mid \ell_{\mathcal{T}_{i-1}}] < 1 + \frac{1}{1 - \bar{\pi}}K. \tag{A.6}$$

If it is the latter case, then it follows from (A.5) that

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

Now, note that the maximum of these three upper bounds given in (A.4) to (A.6) is $1 + \frac{1}{1-\bar{\pi}}K$. Furthermore, by definition, $K \leq K(\alpha, \varepsilon)$. Therefore, we conclude that

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

$\square$

Claim 2. *For all $\alpha \in (1/2, 1)$ and $\varepsilon \in (0, \alpha(1-\alpha))$, we have*

$$M(\alpha, \varepsilon) < 1/\varepsilon.$$

Proof. Since $M(\alpha, \varepsilon)$ is strictly decreasing in α (Claim 1), for any $\alpha > 1/2$, it is bounded above by its limit as $\alpha \to 1/2$:

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

The condition $M(1/2, \varepsilon) < 1/\varepsilon$ is equivalent to

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

Since $\varepsilon < \alpha(1-\alpha)$ and α can be arbitrarily close to $1/2$, we restrict attention to $\varepsilon \in (0, 1/4)$. By the L'Hôpital's rule,

$$\lim_{\varepsilon \to 0} -(\frac{1}{\varepsilon} - 1) \log(1 - 2\varepsilon) = \lim_{\varepsilon \to 0} 2 \frac{(1-\varepsilon)^2}{1 - 2\varepsilon} = 2.$$

Since $2 > 2 \log 2$ and $-(\frac{1}{\varepsilon} - 1) \log(1 - 2\varepsilon)$ is strictly increasing in ε, the above inequality holds for all $\varepsilon \in (0, 1/4)$. Thus we conclude that for all $\alpha \in (1/2, 1)$ and $\varepsilon \in (0, \alpha(1-\alpha))$,

$$M(\alpha, \varepsilon) < M(1/2, \varepsilon) < 1/\varepsilon.$$

$\square$

The following lemma will be useful in the proof of Theorem 1. It establishes 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 3. *For every $r \in \{1, 2, \ldots\}$ 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 $p > 0$,

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

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 $r \in \{1, 2, \ldots\}$, 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 $\mathbb{E}[|\mathcal{D}_i|^r]$. □

The following lemma will also be useful in the proof of Theorem 1. Define the number *wasted signal switch* up to time n by

$$W_n = \sum_{t=1}^n \mathbb{1}(s_t \neq s_{t+1}) \cdot \mathbb{1}(a_t = a_{t+1} = s_{t+1}).$$

That is, W_n counts times at which the signal switches but the action does not, and after the signal switch the action already matches the new signal. Recall that $n \cdot \mathcal{Q}_a(n) = \sum_{t=1}^n \mathbb{1}(a_t \neq a_{t+1})$ is the total number of action switches by time n. Likewise, $n \cdot \mathcal{Q}_s(n)$ is the total number of signal switches by time n.

Lemma 4. *For every $n \geq 1$,*

$$n \cdot \mathcal{Q}_a(n) \leq n \cdot \mathcal{Q}_s(n) - W_n + 1. \tag{A.8}$$

Proof. Partition $\{1, \ldots, n+1\}$ into S_n+1 maximal signal runs, denoted by $\mathcal{R}_k = [\tau_k, \tau_{k+1}-1]$, where $\tau_1 = 1$ and $\{\tau_k\}_{k=2}^{S_n+1}$ are the times of signal switches. Within any run $\mathcal{R}_k$, the signal s_t is constant, say $s_t = s \in \{-1, +1\}$.

We first show that on any run $\mathcal{R}_k$, the action a_t switches at most once. Suppose $a_u = s$ for some $u \in \mathcal{R}_k$. If the next period is in the learning region, $a_{u+1} = s_{u+1} = s$. If it is in the cascade region, the public LLR evolves deterministically according to (5.3), i.e., $\ell_{u+2} = G_\varepsilon(\ell_{u+1})$; since G_ε preserves the sign, and $\operatorname{sign}(\ell_{u+1}) = a_u = s$, it follows that $a_{u+1} = \operatorname{sign}(\ell_{u+2}) = s$. By induction, once a_t aligns with the signal s, it remains constant for the remainder of the run. Thus, the only possible action switch within $\mathcal{R}_k$ is a single transition from $-s$ to s.

Next, consider the signal switch at $t = \tau_k - 1$ $(k > 1)$ which initiates the run $\mathcal{R}_k$. Let $s_{t+1} = s$. There are three cases: (i) if $a_t \neq a_{t+1}$, this signal switch contributes exactly one action switch at time t; (ii) if $a_t = a_{t+1} = s$, then by the argument above, the action remains constant throughout $\mathcal{R}_k$, and thus it is counted in W_n; (iii) if $a_t = a_{t+1} = -s$, the action may switch at most once within $\mathcal{R}_k$. Hence, among the S_n signal switches, each can generate at most one action switch, except those counted by W_n, which generate none. Finally, the initial signal run (before the first signal switch) can contribute at most one action switch, so we obtain $n \cdot \mathcal{Q}_a(n) \leq n \cdot \mathcal{Q}_s(n) - W_n + 1$. $\square$

Now, we are ready to prove the main theorem.

Proof of Theorem 1

We divide the proof of Theorem 1 into two parts. The first part proves the lower bound. The second part proves the upper bound.

Part (i). We first prove that $\lim_{n\to\infty} \mathcal{Q}_a(n) > \lim_{n\to\infty} \mathcal{Q}_\theta(n)$ almost surely. Recall that $\mathcal{D}_i$ is the (random) duration of the i-th run of ℓ, where a *run* is defined as the number of consecutive periods during which the public LLR has the same sign. Let $\Phi = (\mathcal{F}_1, \mathcal{F}_2, \ldots)$ be the filtration where each $\mathcal{F}_i = \sigma(\mathcal{D}_1, \ldots, \mathcal{D}_i)$ and thus $\mathcal{F}_j \subseteq \mathcal{F}_i$ for any $j \leq i$. The process $(\mathcal{D}_1, \mathcal{D}_2, \ldots)$ is adapted to Φ since each $\mathcal{D}_i$ is $\mathcal{F}_i$-measurable. By Proposition 2 and Claim 2, there exists $\delta = 1/\varepsilon - M(\alpha, \varepsilon) > 0$ such that for all $i \geq 2$,

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

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

$$\mathbb{E}[\mathcal{D}_i \mid \mathcal{F}_{i-1}] = \mathbb{E}[\mathbb{E}[\mathcal{D}_i \mid \ell_{\mathcal{T}_{i-1}}, \mathcal{F}_{i-1}] \mid \mathcal{F}_{i-1}] < 1/\varepsilon - \delta. \tag{A.9}$$

Let $X_i = \mathcal{D}_i - \mathbb{E}[\mathcal{D}_i \mid \mathcal{F}_{i-1}]$ for all $i \geq 2$ and since $\mathcal{F}_{i-1} \subseteq \mathcal{F}_i$, each X_i is $\mathcal{F}_i$-measurable. Define the partial sum process as

$$Y_K = X_2 + \frac{1}{2}X_3 + \cdots + \frac{1}{K-1}X_K.$$

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

$$\mathbb{E}[Y_K \mid \mathcal{F}_{K-1}] = \mathbb{E}[\sum_{i=2}^{K} \frac{1}{i-1} X_i \mid \mathcal{F}_{K-1}] = Y_{K-1} + \frac{1}{K-1}\mathbb{E}[X_K \mid \mathcal{F}_{K-1}] = Y_{K-1},$$

and so the process $(Y_K)_{K\geq 2}$ forms a martingale.

By Lemma 3 and (A.9), both $\mathbb{E}[\mathcal{D}_i^2]$ and $\mathbb{E}[\mathcal{D}_i \mid \mathcal{F}_{i-1}]$ 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 $K \geq 2$,

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

By the martingale convergence theorem, Y_K converges almost surely as $K \to \infty$. It then follows from Kronecker's lemma that[18]

$$\lim_{K\to\infty} \frac{1}{K-1}(X_2 + \cdots X_K) = 0 \quad \text{almost surely.}$$

Substituting the definition of X_i, we have

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

Using the bound from (A.9), this implies

$$\lim_{K\to\infty} \frac{1}{K-1}\sum_{i=2}^{K} \mathcal{D}_i \leq 1/\varepsilon - \delta < 1/\varepsilon \quad \text{almost surely.} \tag{A.10}$$

To connect this to the action volatility, recall that $a_t = \operatorname{sign}(\ell_{t+1})$ for all $t \geq 2$, and $\mathcal{T}_1 < \mathcal{T}_2 < \cdots$ denote the successive times at which the public LLR switches its sign. Let $K_n = \sum_{t=1}^{n} \mathbb{1}(\operatorname{sign}(\ell_t) \neq \operatorname{sign}(\ell_{t+1}))$ denote the number of sign switches occurring up to time n. Then

$$\mathcal{T}_{K_n} - \mathcal{T}_1 = \sum_{i=2}^{K_n} \mathcal{D}_i,$$

and since $\mathcal{T}_{K_n} \leq n < \mathcal{T}_{K_n+1}$,

$$\lim_{n\to\infty} \mathcal{Q}_a(n) = \lim_{n\to\infty} \frac{K_n}{n} = \lim_{K\to\infty} \frac{K}{\mathcal{T}_K} = \lim_{K\to\infty} \frac{K}{\mathcal{T}_1 + \sum_{i=2}^{K} \mathcal{D}_i}.$$

Dividing the numerator and denominator by $K-1$, using $\mathcal{T}_1/(K-1) \to 0$ as $K \to \infty$, and applying the bound from (A.10), we conclude that

$$\lim_{n\to\infty} \mathcal{Q}_a(n) \geq \frac{1}{1/\varepsilon - \delta} > \varepsilon = \lim_{n\to\infty} \mathcal{Q}_\theta(n) \quad \text{almost surely.}$$

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

Part (ii). We now prove that $\lim_{n\to\infty} \mathcal{Q}_s(n) > \lim_{n\to\infty} \mathcal{Q}_a(n)$ almost surely. Dividing (A.8) in Lemma 4 by n yields

$$\mathcal{Q}_a(n) \le \mathcal{Q}_s(n) - \frac{W_n}{n} + \frac{1}{n},$$

so it suffices to show that

$$\lim_{n\to\infty} \frac{W_n}{n} > 0 \quad \text{almost surely.} \tag{A.11}$$

Recall that $\mathcal{C}_1 < \mathcal{C}_2 < \cdots$ are the successive times at which the public LLR enters the cascade region, and $\mathcal{J}_i = \mathcal{C}_i - \mathcal{C}_{i-1}$ (with $\mathcal{C}_0 := 0$) are the inter-entry times. Let $\mathcal{H}_i = \sigma(\mathcal{J}_1, \ldots, \mathcal{J}_i)$ and define the martingale differences $X_i = \mathcal{J}_i - \mathbb{E}[\mathcal{J}_i \mid \mathcal{H}_{i-1}]$. As shown in the proof of Proposition 1, $\mathcal{J}_i$ is stochastically dominated by a constant plus a geometric random variable. Consequently, $\mathcal{J}_i$ has uniformly bounded second moments. Therefore, using the bound in Proposition 1 and following a similar argument that led to (A.10), we have

$$\lim_{M\to\infty} \frac{1}{M} \sum_{i=1}^{M} \mathcal{J}_i \le \lfloor K(\alpha, \varepsilon) \rfloor + 8 \quad \text{almost surely.}$$

Since $\sum_{i=1}^{M} \mathcal{J}_i = \mathcal{C}_M$, it follows that

$$\lim_{M\to\infty} \frac{M}{\mathcal{C}_M} = \lim_{n\to\infty} \frac{N_{\mathcal{C}}(n)}{n} \ge \frac{1}{\lfloor K(\alpha, \varepsilon) \rfloor + 8} \quad \text{almost surely,} \tag{A.12}$$

where $N_{\mathcal{C}}(n) = \max\{i : \ \mathcal{C}_i \le n\}$ denotes the number of cascade entries by time n.

Next, we show that each cascade entry produces a wasted signal switch with uniformly positive probability. For each $i \ge 1$, define the event

$$F_i = \{s_{\mathcal{C}_i} = -a_{\mathcal{C}_i}, s_{\mathcal{C}_i+1} = a_{\mathcal{C}_i}\}.$$

On this event, clearly, $s_{\mathcal{C}_i} \neq s_{\mathcal{C}_i+1}$. By definition, at time $\mathcal{C}_i$, the public LLR is in the cascade region, and so $a_{\mathcal{C}_i}$ depends only on public action history but not $s_{\mathcal{C}_i}$. The condition $s_{\mathcal{C}_i+1} = a_{\mathcal{C}_i}$ then ensures that $a_{\mathcal{C}_i+1} = a_{\mathcal{C}_i}$ regardless of the region at time $\mathcal{C}_i + 1$. Thus, every occurrence of F_i contributes to the wasted switch count W_n, i.e,

$$W_n \ge \sum_{i=1}^{N_{\mathcal{C}}(n)} \mathbb{1}(F_i). \tag{A.13}$$

We now lower bound the probability of F_i. Let $\mathcal{G}_i = \sigma(s_1, \ldots, s_{\mathcal{C}_i-1})$ and note that $a_{\mathcal{C}_i}$ is $\mathcal{G}_i$-measurable. Since the state is Markov, regardless of the history, $\mathbb{P}[\theta_{\mathcal{C}_i+1} = \theta_{\mathcal{C}_i} \mid \mathcal{G}_i] = 1 - \varepsilon$. Conditional on this event, the signals $s_{\mathcal{C}_i}, s_{\mathcal{C}_i+1}$ are independent draws with accuracy α. Hence, conditional on $\mathcal{G}_i$, we have

$$\mathbb{P}[F_i \mid \theta_{\mathcal{C}_i+1} = \theta_{\mathcal{C}_i}, \mathcal{G}_i] = \alpha(1-\alpha),$$

and hence by the law of total probability,

$$\mathbb{P}[F_i \mid \mathcal{G}_i] \ge (1-\varepsilon)\alpha(1-\alpha).$$

Applying the strong law of large numbers for bounded martingale differences yields

$$\lim_{M\to\infty} \frac{1}{M}\sum_{i=1}^{M} \mathbb{1}(F_i) \geq (1-\varepsilon)\alpha(1-\alpha) \quad \text{almost surely.}$$

Combining this with (A.12) and (A.13), we obtain

$$\begin{aligned}
\lim_{n\to\infty} \frac{W_n}{n} &\geq \lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{N_{\mathcal{C}}(n)} \mathbb{1}(F_i) \\
&= \left(\lim_{n\to\infty} \frac{N_{\mathcal{C}}(n)}{n}\right) \cdot \left(\lim_{M\to\infty} \frac{1}{M}\sum_{i=1}^{M} \mathbb{1}(F_i)\right) \\
&\geq \frac{(1-\varepsilon)\alpha(1-\alpha)}{\lfloor K(\alpha,\varepsilon)\rfloor + 8} > 0 \quad \text{almost surely,}
\end{aligned}$$

as required by (A.11). This concludes the proof of Theorem 1.

Proof of Corollary 1. Following the same martingale argument, the long-run frequency of action switches converges almost surely to the inverse of the limiting expected run duration:

$$\lim_{n\to\infty} \mathcal{Q}_a(n) = \lim_{K\to\infty} \frac{1}{\frac{1}{K-1}\sum_{i=2}^{K} \mathbb{E}[\mathcal{D}_i \mid \mathcal{F}_{i-1}]} \quad \text{almost surely.}$$

We now bound this expectation. By Proposition 2 and the Markov property, $\mathbb{E}[\mathcal{D}_i \mid \mathcal{F}_{i-1}] < M(\alpha,\varepsilon)$. Since $M(\alpha,\varepsilon)$ is strictly decreasing in α (Claim 1), it is uniformly bounded by the worst case $\alpha = 1/2$, and so

$$\lim_{n\to\infty} \mathcal{Q}_a(n) \geq \frac{1}{M(1/2,\varepsilon)} \quad \text{almost surely.}$$

Using the inequality $-\log(1-x) \geq x$ for $x \in (0,1)$, we have $-\log(1-2\varepsilon) \geq 2\varepsilon$. Applying this to the definition of $M(1/2,\varepsilon)$ gives

$$M(1/2,\varepsilon) = 1 + \frac{2\log 2}{-\log(1-2\varepsilon)} \leq 1 + \frac{2\log 2}{2\varepsilon} = \frac{\varepsilon + \log 2}{\varepsilon}.$$

Substituting this into the above inequality, we have

$$\lim_{n\to\infty} \mathcal{Q}_a(n) \geq \frac{\varepsilon}{\varepsilon + \log 2} \quad \text{almost surely.}$$

Since $\lim_{n\to\infty} \mathcal{Q}_\theta(n) = \varepsilon > 0$ almost surely, dividing this limit yields the desired claim:

$$\frac{\lim_{n\to\infty} \mathcal{Q}_a(n)}{\lim_{n\to\infty} \mathcal{Q}_\theta(n)} \geq \frac{1}{\log 2 + \varepsilon} \quad \text{almost surely.}$$

□