Learning by Slow Unlearning: An Unlearning Algorithm for Social Learning

Jaden Yang Chen* Yuyi Li†

December 21, 2025

Abstract

This paper studies a sequential social learning problem in which society follows a power-Bayesian updating protocol that governs how public information is weighted over time. We show that incorrect information cascades can be avoided by gradually "unlearning" public information. Under mild regularity conditions, a correct information cascade emerges almost surely in the limit if and only if the weight on public information decays more slowly than $1/i$, where i denotes the decision order. If public information is discounted too quickly, no information aggregation occurs in the limit and individuals asymptotically rely only on private signals; if it is never fully discarded, both correct and incorrect cascades persist. From an algorithmic perspective, the results suggest that recommendation rules with time-decaying weights can mitigate persistent herding while preserving long-run information aggregation.

JEL Classification: D81, D83, C72

Keywords: social learning; information cascades; unlearning algorithm; belief dynamics; stochastic approximation

*University of North Carolina at Chapel Hill, E-mail: yangch@unc.edu

†University of North Carolina at Chapel Hill, E-mail: yuyil123@ad.unc.edu

1 Introduction

Individuals often learn from observing others. In many economic settings, ranging from financial markets to technology adoption, such learning unfolds sequentially: individuals act in turn, each observing predecessors' actions but not their private signals. In the classical Bayesian models, an *information cascade* occurs almost surely: eventually, agents disregard their own signals and imitate others, even when those actions are suboptimal (Banerjee, 1992; Bikhchandani et al., 1992).[1] This learning failure arises because early actions exert a disproportionate influence on later decisions, crowding out subsequent private information.

This paper considers a setting in which individuals collectively follow a non-Bayesian updating protocol and shows that it can successfully avoid incorrect information cascades. In particular, we introduce an *unlearning algorithm*, in which individual i raises the likelihood of the observed public history to a power $\alpha_i \in [0, +\infty)$.[2] When $\alpha_i = 1$, the algorithm coincides with standard Bayesian learning; when $\alpha_i < 1$, individuals underweight public information. A decaying α_i enables society to gradually "unlearn" public information over time. We show that such unlearning mitigates the excessive influence of early actions in the standard cascade model. When α_i decays at a rate slower than $1/i$, the following results hold:

1. Correct limit cascade: A *correct* information cascade occurs almost surely in the limit. This stands in contrast to the classical literature, in which both correct and incorrect cascades arise with positive probability.

2. Belief dynamics and convergence rate: The public belief enters and exits cascade regions infinitely often. It oscillates around the boundary of the correct-cascade region with shrinking amplitude and converges to that boundary (Figure 3a). We further show that the (local) convergence rate in L^2 is approximately $O(\alpha_i^{1/2})$.

3. Asymptotic actions: The expected fraction of correct actions exceeds $1 - \delta$ in the limit, where δ is the probability of the strongest signal favoring the incorrect state. With continuous signals ($\delta = 0$), our unlearning algorithm (with slight modification) drives asymptotically almost *all* individuals to choose the correct action.[3]

The condition for correct limit cascades is tight. Under the power-rate assumption for α_i, a correct limit cascade emerges almost surely *if and only if* $\alpha_i \to 0$ more slowly than $1/i$. If α_i decays faster than $1/i$, public information is discarded too quickly, and individuals follow only their private signals in the limit. Conversely, if α_i does not converge to 0 (e.g., converges

[1] See also Chamley (2004) and Easley and Kleinberg (2010) for expositions.

[2] We use the term *algorithm* because beliefs are updated cooperatively at the societal level; it is therefore more natural to view agents as following a common protocol rather than applying isolated individual rules.

[3] More precisely, we show that the fraction of correct actions converges to 1 in probability.

to a positive constant or diverges), a minimum weight on public information persists and both correct and incorrect cascades occur in the limit, as in the standard model.

This paper makes several contributions to the social learning literature. (I) It proposes a *simple algorithm* that overturns the classical cascade result: by allowing society to unlearn public information at a controlled speed, the analysis shows that incorrect cascades can be avoided without requiring unbounded signals or designated sacrificial agents. (II) The analysis uncovers a *new learning dynamic*. Unlike traditional models—which generate either path-dependent herding or complete learning—the unlearning algorithm yields an almost-sure correct cascade in which beliefs converge to an interior limit (i.e., the boundary of the correct-cascade region). Society therefore converges to the correct action, yet learning remains incomplete because sufficiently strong private signals can still induce deviations. (III) The paper develops a *methodological approach* for studying learning dynamics when the martingale property fails. By employing tools from stochastic approximation, the analysis is able to characterize the non-martingale belief dynamics generated by the non-Bayesian updating process, particularly for the continuous-signal case. (IV) Finally, from a practical perspective, the results also speak to *algorithm design*: they suggest that a time-decaying recommendation algorithm can systematically improve long-run information aggregation, a point developed further in Section 7.

Related Literature.

A natural method to mitigate incorrect cascades is the *sacrificial-lambs algorithm*: infinitely many agents are designated "lambs" who ignore social information and act solely on their private signals, thereby injecting infinitely many independent observations and delivering complete learning for the non-lambs (e.g., Peres et al., 2020; Cheng et al., 2018).[4] Our approach differs along several dimensions. First, in terms of mechanism, we do not designate special cohorts or exclude anyone from social information; instead, every individual follows the same unlearning update that gradually down-weights public history via a position-dependent weight α_i.[5] Second, in terms of learning outcomes, the two approaches differ both belief-wise and action-wise.

- *Belief-wise*: Sacrificial-lamb schemes achieve complete learning for non-lambs—their beliefs converge to a point mass on the true state in the limit—while lambs themselves

[4]The sacrificial-lamb idea appears already in classical social learning papers such as Banerjee (1992) and Bikhchandani et al. (1992). The recent works cited above show how it can be used to obtain asymptotic learning and also study convergence rates.

[5]Notably, in many settings, instructing or compelling some agents to ignore observable information may be impractical due to legal or ethical concerns. Our approach thus offers a softer, more defensible intervention for algorithm design in social learning.

do not learn. In contrast, under our algorithm, society achieves a correct limit cascade for all individuals; learning remains incomplete, as beliefs converge to an interior limit.

- *Action-wise*: Both approaches can deliver asymptotically correct actions for almost all agents, but under different requirements. In sacrificial-lamb schemes, the density of lambs must vanish over time, whereas under our algorithm, correctness requires a thin-tailed (or continuous) signal structure.

Finally, previous papers focus on binary signals, whereas this paper considers general signal structures—covering both discrete and continuous signals—and provides a comprehensive characterization of the corresponding asymptotic dynamics and convergence rates.

The paper's approach of underweighting public information is also related to the literature showing that overconfidence can enhance social learning. Bernardo and Welch (2001) analyze a sequential social learning model in which a proportion of agents are overconfident and underweight public information; they show that moderate overconfidence improves social learning outcomes. Arieli et al. (2025) study a model with unbounded signals where individuals misinterpret their predecessors' data-generating processes; they demonstrate that mild condescension can lead to efficient learning, a result stronger than complete learning in the standard unbounded-signal model.[6] These papers assume a *homogeneous* confidence level. In contrast, our model allows the confidence parameter α_i to vary across individuals, and the analysis shows that not only the *level* of confidence but also the *rate* at which confidence changes plays a crucial role in social learning.[7] Our assumption of heterogeneous weights on public information is also related to Goeree et al. (2006), who show that complete learning can occur even with bounded signals when individuals have idiosyncratic tastes for actions; preference diversity ensures that all actions occur with positive probability in every state, continuously revealing information.

This paper also contributes to the broader literature on non-Bayesian social learning. A prominent strand studies naive or heuristic updating rules, such as DeGroot learning and related models (DeGroot, 1974; DeMarzo et al., 2003; Golub and Jackson, 2010).[8] In this

[6] Efficient learning means that the expected number of incorrect actions is finite, as defined in Rosenberg and Vieille (2019). While Smith and Sørensen (2000) show that complete learning occurs with unbounded signals, it does not necessarily ensure efficiency.

[7] As noted in a recent survey by Bikhchandani et al. (2024): "A possible direction for future research is understanding conditions under which overconfidence harms social learning instead of helping. For example, if overconfidence were growing rapidly with later agents, all agents would act based only on their private signals and there would be no information aggregation." Proposition 7 directly echoes their conjecture in the quote.

[8] See also Ellison and Fudenberg (1993, 1995), Bala and Goyal (1998), Epstein et al. (2010), Jadbabaie et al. (2012), Guarino and Jehiel (2013), Arieli et al. (2021), Molavi et al. (2018), Eyster and Rabin (2010, 2014), Dasaratha and He (2020), Chen (2019), Frick et al. (2024).

paper, individuals can also be interpreted as using a heuristic rule, assigning varying weights to public information based on their position in the sequence. A relevant paper is Molavi et al. (2018), who axiomatize the log-linear rule and show that if individuals' weights on neighbors decay more slowly than $1/t$—the same threshold identified in this paper—then complete learning occurs. Their environment, however, differs in key respects. Individuals repeatedly exchange beliefs, and information cascades do not arise. Moreover, in their setting complete learning still obtains when weights are constant, whereas in our model a correct limit cascade *requires* that the weight on public information decay to zero.

2 Model

Preliminaries. The setup follows the traditional sequential social learning framework. The true state of the world is $\theta \in \{1, 0\}$. Without loss of generality, assume that the true state is 0. An infinite sequence of individuals $I = \{1, 2, 3, \dots\}$ act in an exogenously given order, each choosing an action from $A = \{1, 0\}$ in an attempt to match the state. Individual i receives payoff 1 if her action a_i matches the true state and 0 otherwise.

Information Structure. Individuals do not know the true state and begin with a flat prior π_0. Before acting, individual i privately receives a signal s_i from a signal space S. Signals are i.i.d. conditional on θ and have conditional distribution G^θ. Following the convention, we work with the *normalized signal* $\lambda(s) = \frac{dG^1(s)}{dG^0(s)}$, which is the likelihood ratio induced by signal s. Let F^θ denote the conditional c.d.f. of the normalized signal. We assume $co\left(\text{supp}\left(F^\theta\right)\right) = [1/\gamma, \gamma]$, where $\gamma \in (1, +\infty)$, so signals are informative ($\gamma > 1$) and bounded ($\gamma < +\infty$). [9] This paper allows both discrete and continuous signal structures: in the discrete case, $\text{supp}\left(F^\theta\right)$ contains finitely many points; in the continuous case, we further assume that F^θ is continuously differentiable.

Learning Algorithm. All individuals follow some non-Bayesian algorithm to update beliefs. Individual i's posterior is

$$\pi_i\left(\theta|\lambda_i, h_i\right) = \frac{f^\theta\left(\lambda_i\right) \times p_i\left(\theta|h_i\right)}{\sum_{\theta' \in \Theta} f^{\theta'}\left(\lambda_i\right) \times p_i\left(\theta'|h_i\right)},$$

where: (i) $f^\theta\left(\lambda_i\right)$ is the p.d.f. or p.m.f. of the normalized signal λ_i under state θ; and (ii) $p_i\left(\theta|h_i\right)$ is the *public belief*, i.e., the belief based solely on public history, h_i. Thus, π_i is the standard Bayesian update of the public belief based on private information.

[9] The assumption of signals being informative is to avoid triviality. We also focus on bounded signals because cascades do not arise with unbounded signals (Smith and Sørensen, 2000).

The non-standard feature of the model lies in how public beliefs are formed. Individuals follow a *power-Bayesian algorithm*, defined by

$$p_i(\theta|h_i) = \frac{p_F^{\alpha_i}(\theta|h_i)}{\sum_{\theta' \in \Theta} p_F^{\alpha_i}(\theta'|h_i)}, \tag{1}$$

where: (i) $p_F(h_i|\theta)$ is the *Bayesian posterior* of θ based on h_i; and (ii) $\alpha_i \geq 0$ is the *Bayesian weight* used by individual i. Intuitively, the power-Bayesian rule transforms the Bayesian public belief by raising it to a position-dependent exponent α_i. When $\alpha_i = 1$, it corresponds to Bayesian learning. When $\alpha_i < 1$, individuals flatten extreme likelihood ratios and thereby discount the accumulated public information; when $\alpha_i > 1$, they amplify it. The sequence (α_i) therefore governs the rate at which society "unlearns" or "reinforces" the information encoded in past actions. This exponent-based transformation can be viewed as a simple algorithmic adjustment of the society's confidence in public history; Section 7 gives a concrete connection to recommendation systems.

Recursive Form. The update in (1) can be expressed in the recursive form

$$z_{i+1} = \beta_{i+1} \times z_i + \alpha_{i+1} \times \log\left(\frac{p_F(a_i|h_i, \theta = 1)}{p_F(a_i|h_i, \theta = 0)}\right), \tag{2}$$

where: (i) $z_i = \log\left(\frac{p_i(\theta=1|h_i)}{p_i(\theta=0|h_i)}\right)$ denotes the public *log likelihood ratio* (LLR), (ii) $\beta_{i+1} = \alpha_{i+1}/\alpha_i$ can be interpreted as the discount factor that individual $i+1$ applies to the accumulated public information; and (iii) the final term is the LLR contribution from observing action a_i, i.e., the new public information. In words, next period's public LLR z_{i+1} equals a discounted version of the current LLR z_i, plus the new public information generated by action a_i, weighted by the Bayesian weight α_{i+1}.

Remark 1 (Power-Bayesian Algorithm). The power-Bayesian algorithm in (1) is a special case of *generalized Bayesian updating* (Bissiri et al., 2016), in which the likelihood contribution of historical data is raised to an exponent. It is closely related to *power priors* in statistics, which discount past information by applying a tuning parameter to the likelihood (Ibrahim and Chen, 2000). In economics, the formulation is connected to *Grether's models* (Grether, 1980, 1992), where individuals differentially weight base rates and sample evidence, generating systematic departures from Bayesian updating.[10]

[10] A large experimental literature documents similar deviations from Bayesian weighting when combining private and public information: see, for example, Anderson and Holt (1997), Weizsäcker (2010), Duffy et al. (2021), and the recent findings in Conlon et al. (2022). Surveys include Morin et al. (2021) and Benjamin (2019).

2.1 Assumptions

The goal of the paper is to understand how asymptotic learning outcomes depend on the *weighting scheme* $\boldsymbol{\alpha} = \{\alpha_i\}_{i=1}^{\infty}$. To simplify the discussion, we impose the following regularity conditions.

Assumption 1 (Monotonicity)**.** *α_i is monotonic in i when i is sufficiently large.*

Assumption 2 (Power Rate)**.** *As $i \to +\infty$,*

$$\alpha_i \sim c \times i^{-\rho},$$

for some constant $c > 0$ and $\rho \in \mathbb{R}$. When $\rho = 0$, we additionally assume that $|\alpha_i - c_0| \sim c \times i^{-q}$ for some $c_0, q > 0$.

Assumption 1 rules out sequences in which α_i oscillates infinitely often. Assumption 2 is the central assumption: it requires that the Bayesian weights eventually follow a power-law rate of decay (or growth). Together, these assumptions allow us to reduce the characterization of asymptotic learning to a single parameter ρ. Three cases are particularly important: (i) when $\rho > 0$, $\alpha_i \to 0$, so tail individuals eventually ignore public information; (ii) when $\rho < 0$, $\alpha_i \to +\infty$, so tail individuals place increasingly large weight on public information; and (iii) when $\rho = 0$, α_i converges to a positive constant.

3 Decision Criterion and Learning Concepts

This section characterizes individuals' decision rule in the sequential social learning environment and introduces the notions of information cascades and limit cascades.

3.1 Decision Rule

We begin by describing each individual's optimal action using the public likelihood ratio.

Proposition 1 (Characterization of Decision Rule)**.** *For all $i \in I$ and all feasible h_i,*

$$a_i = \begin{cases} 1, & \lambda_i > 1/x_i \\ 0, & \lambda_i < 1/x_i \end{cases}, \tag{3}$$

where $x_i \equiv \exp(z_i)$ is individual i's public likelihood ratio based on history h_i.

Proof. Under the power-Bayesian update,

$$\frac{\pi_i\left(1|\lambda_i, h_i\right)}{\pi_i\left(0|\lambda_i, h_i\right)} = \frac{f^1\left(\lambda_i\right)}{f^0\left(\lambda_i\right)} \times \left[\frac{p_F\left(h_i|\theta=1\right)}{p_F\left(h_i|\theta=0\right)}\right]^{\alpha_i} = \lambda_i \times x_i.$$

Thus, individual i chooses action 1 (or 0) if the ratio is larger (or smaller) than 1. □

When an indifference arises, we adopt a tie-breaking rule under which individuals follow the public likelihood ratio: they choose action 1 whenever $x_i \geq 1$.[11] From Proposition 1, the public likelihood ratio x_i serves as a sufficient statistic for the learning problem. Therefore, despite the complexity brought by the Bayesian weight, the learning problem is still tractable: it suffices to keep track of x_i to understand the learning dynamics.

3.2 Information Cascades and Limit Cascades

When the public likelihood ratio x_i exceeds the highest possible signal likelihood ratio γ or falls below the smallest one $1/\gamma$, individual i's action no longer depends on her private signal. This motivates the following definitions.

Definition 1. We say that: (i) individual i is in an *information cascade* if $x_i \in C \equiv [0, 1/\gamma] \cup [\gamma, +\infty]$, and (ii) a *limit cascade* occurs if $x_\infty \in C$, where $x_\infty = \lim_{i\to+\infty} x_i$, assuming that a limit exists.

An information cascade occurs when the public likelihood ratio is so extreme—either very large or very small—that no private signal can overturn the public belief. We refer to the set C as the *cascade set*, and its complement set $D \equiv (1/\gamma, \gamma)$ as the *noncascade set*. We also define $C_1 = [\gamma, +\infty]$ and $C_0 = [0, 1/\gamma]$, where C_a is the cascade set of action a. If $x_i \in C_a$, then individual i chooses action a regardless of her private signal.

Unlike the classical Bayesian model, under the power-Bayesian algorithm, the occurrence of a cascade at time i does not necessarily imply that all subsequent individuals also cascade. Hence the more relevant concept is a *limit cascade*, in which public information eventually dominates private signals. We say that a *correct limit cascade* occurs if $x_\infty \in C_0$, meaning that the public likelihood ratio converges to the cascade set of the correct action.

4 Discussion of Learning Dynamics

In this section, we discuss the learning dynamics under the power-Bayesian algorithm and explain how they differ from those in classical Bayesian models. By Proposition 1, the public

[11] That is, individuals choose the action favored by public belief in the knife-edge case. This assumption simplifies the exposition but does not affect any of the asymptotic results.

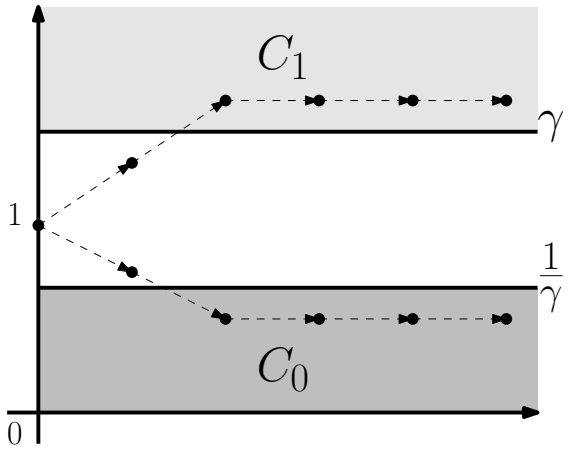

Figure 1: Bayesian Case: Path-Dependent Cascades

likelihood ratios $\{x_i\}$ satisfy the recursive expression:

$$x_{i+1} = x_i^{\frac{\alpha_{i+1}}{\alpha_i}} \times \left[\frac{\phi^1(a_i, x_i)}{\phi^0(a_i, x_i)}\right]^{\alpha_{i+1}}, \tag{4}$$

where

$$\phi^\theta(a, x) = p_F(a|h, \theta) = \begin{cases} 1 - F^\theta\left(\frac{1}{x}\right), & \text{if } a = 1, \\ F^\theta\left(\frac{1}{x}\right), & \text{if } a = 0, \end{cases}$$

represents the probability of taking action a in state θ given the public likelihood ratio x.

4.1 Bayesian Benchmark

When $\alpha_i = 1$ for all i, the power-Bayesian model reduces to the standard Bayesian model. In this case, equation (4) simplifies to

$$x_{i+1} = x_i \times \frac{\phi^1(a_i, x_i)}{\phi^0(a_i, x_i)}. \tag{5}$$

A key property of x_i in this setting is that it is a *martingale*.[12] By the martingale convergence theorem, x_i almost surely converges to a limit x_∞.

Lemma 1 (Bayesian Benchmark)**.** *If $\alpha_i = 1$ for all i, a limit cascade occurs almost surely, and both correct and incorrect cascades occur with positive probability.*

Proof. Note x_∞ must be a stationary point of (5), so it must satisfy $\frac{\phi^1(a_\infty, x_\infty)}{\phi^0(a_\infty, x_\infty)} = 1$, where a_∞ is the limit action. The only possibility is that $x_\infty \in C$. [13] Moreover, the martingale

[12] To be more precise, it is a martingale under the probability measure conditional on the true state.

[13] This is because $F^0(x) > F^1(x)$ for $x \in (1/\gamma, \gamma)$, so any point in the noncascade set is not stationary.

property implies $\mathbb{E}(x_\infty) = x_0 \in (1/\gamma, \gamma)$, so x_∞ cannot always exceed γ or remain below $1/\gamma$. Thus both correct and incorrect cascades occur with positive probability. □

Furthermore, when signals are discrete (or continuous with sufficiently heavy tails), an information cascade occurs in finite time. Once $x_i \in C$, all subsequent individuals ignore their private signals, and society becomes trapped in a permanent cascade (Figure 1). Such cascades are path-dependent: depending on early signals, society may enter either a correct or an incorrect cascade, each with positive probability.

For some continuous signal distributions, an information cascade may never occur in finite time, yet herding still arises almost surely, and the eventual herd behavior is also path-dependent; see Herrera and Hörner (2012) and Smith et al. (2021).

4.2 Power-Bayesian Case

Under the power-Bayesian algorithm, the learning dynamics governed by equation (4) differ fundamentally from those in the Bayesian benchmark.

First, x_i is generally *not* a martingale under the power-Bayesian rule. The martingale convergence theorem therefore does not apply, and neither the existence nor the characterization of a limit is immediate. Even if x_i converges, its limit cannot be identified by the usual stationary-point argument. For example, if $\alpha_i = 1/i$ and $x_i \to x_\infty$, then

$$\left[\frac{\phi^1(a_\infty, x_\infty)}{\phi^0(a_\infty, x_\infty)}\right]^{\alpha_\infty} = 1, \tag{6}$$

where $\alpha_\infty = 0$. Since any positive number raised to the power 0 equals 1, *every* x_∞ satisfies this condition. Thus, unlike the Bayesian case, stationary points provide no information about the limit.

Second, information cascades no longer imply the cessation of information aggregation. Suppose individual i is in a cascade, so $x_i \in C$. Then, by equation (4),

$$x_{i+1} = x_i^{\alpha_{i+1}/\alpha_i}. \tag{7}$$

In the Bayesian model, where $\alpha_i \equiv 1$, this yields $x_{i+1} = x_i$ and no further information aggregation occurs. Under the power-Bayesian algorithm, however, the relative change in the Bayesian weight affects the next likelihood ratio (Figure 2):

- If $\alpha_{i+1} > \alpha_i$, x_{i+1} moves deeper into the cascade set;
- if $\alpha_{i+1} < \alpha_i$, x_{i+1} moves closer to the boundary of the cascade set or exits it entirely.

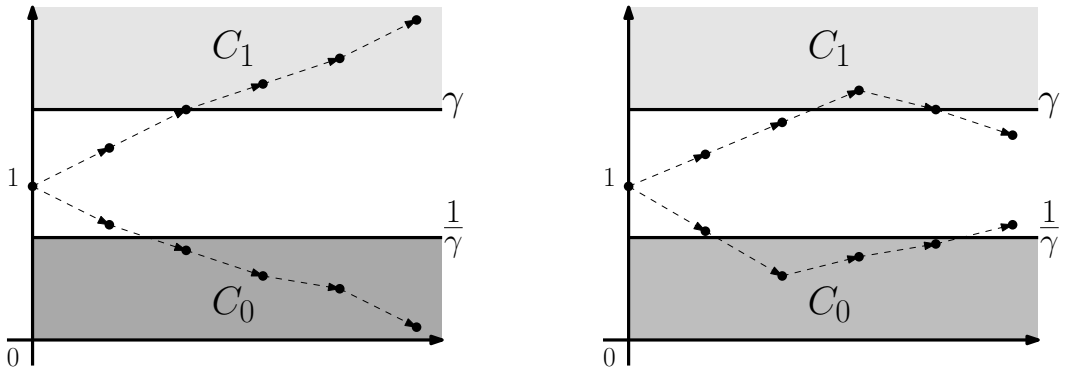

(a) Increasing α_i (b) Decreasing α_i

Figure 2: Dynamics with Different α_i

In the latter case, once x_i re-enters the noncascade region, society resumes learning from private information. Thus, society may experience multiple rounds of halting and restarting information aggregation. This dynamic makes the analysis of the power-Bayesian algorithm considerably more complex compared to the Bayesian case.

5 Main Results: Slow Unlearning and Correct Cascades

This section characterizes learning dynamics under the power-Bayesian algorithm. Throughout, we evaluate events under the objective probability measure $\mathbb{P}$ (i.e., conditional on the true state being 0). In the classical sequential learning model with bounded signals, both correct and incorrect cascades arise with positive probability. The results below show that these conclusions change fundamentally once individuals follow the unlearning algorithm.

5.1 Discrete Signals

We begin with the case of discrete signals.

Theorem 1. *Suppose signals are discrete. A correct limit cascade occurs $\mathbb{P}$-almost surely if $\rho \in (0,1)$ and only if $\rho \in (0,1]$. Moreover, during a correct cascade, $x_i \to 1/\gamma$ $\mathbb{P}$-almost surely.*

The theorem provides an almost necessary and sufficient condition for the emergence of a correct limit cascade.[14] Two insights are central. First, the requirement $\rho > 0$ means that $\alpha_i \to 0$: the Bayesian weights must eventually vanish. A correct limit cascade therefore

[14]When $\rho = 1$, we have $\alpha_i \sim c/i$ for some $c > 0$. In this boundary case, whether a correct limit cascade arises depends on the value of c.

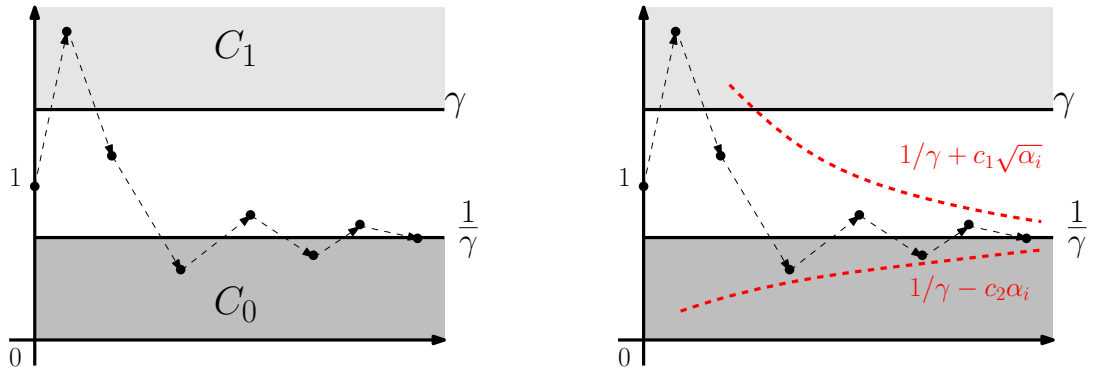

(a) Correct Limit Cascades (b) Convergence Rate

Figure 3: Emergence of Correct Limit Cascades

requires society to discard *all* public information in the limit. Second, the condition $\rho \leq 1$ means that this decline must be sufficiently slow; in particular, α_i cannot fall faster than $1/i$. Together, these conditions imply that correct cascades arise only when society gradually—but not too quickly—unlearns past social information.

The theorem further implies that the public likelihood ratio converges to the boundary of the correct-cascade region, $1/\gamma$. Figure 3a illustrates the learning trajectory: the sequence x_i repeatedly enters and exits the cascade region before eventually stabilizing near $1/\gamma$. This behavior is fundamentally different from the Bayesian benchmark (Figure 1), where the asymptotic learning outcome is path-dependent.

The next proposition characterizes the rate at which x_i converges to $1/\gamma$.

Proposition 2 (Convergence Rate)**.** *Suppose $\rho \in (0,1)$, and signals are discrete. Then, for every $\varepsilon > 0$,*

$$\mathbb{E}\left[(x_i - 1/\gamma)^2 \, 1_{\{x_i \notin C_1\}}\right] = O\left(\alpha_i^{1-\varepsilon}\right). \tag{8}$$

Moreover, whenever $x_i \in C_0$, we have $1/\gamma - x_i = O(\alpha_i)$.

Proposition 2 says that conditional on that x_i stays outside of the incorrect cascade set—which happens with probability 1 in the limit—the expected squared distance of x_i from the correct boundary $1/\gamma$ decays at a rate faster than $\alpha_i^{1-\varepsilon}$ for any $\varepsilon > 0$.[15] Roughly speaking, for large i, the trajectory x_i lies between $1/\gamma \pm c_1\sqrt{\alpha_i}$ with probability close to 1, where $c_1 > 0$ is a constant. The second part of the proposition sharpens this picture: if x_i enters the correct cascade set, then $x_i \geq 1/\gamma - c_2\alpha_i$ for some $c_2 > 0$. Taken together, these bounds imply that the sequence x_i eventually fluctuates within a narrowing band between $1/\gamma - c_2\alpha_i$ and $1/\gamma + c_1\sqrt{\alpha_i}$, as illustrated in Figure 3b.

[15] When $\rho < 1/2$, one may take $\varepsilon = 0$.

The next proposition analyzes long-run behavior of actions.

Proposition 3 (Asymptotic Actions)**.** *Suppose $\rho \in (0,1)$ and signals are discrete. Let $\vartheta_i \equiv \frac{1}{i}\sum_{j=1} 1_{\{a_j=0\}}$ denote the fraction of correct actions. Then,*

$$\liminf_{i\to\infty} \mathbb{E}[\vartheta_i] \geq \liminf_{i\to\infty} \mathbb{P}\left(a_i = 0\right) \geq F^0\left(\gamma^-\right),$$

where $F^0\left(\gamma^-\right) \equiv \lim_{x\uparrow\gamma} F^0\left(x\right)$.

Proposition 3 provides a lower bound for the asymptotic fraction and probability of correct actions. It follows almost directly from the dynamics of x_i—that is, it oscillates around $1/\gamma$ and eventually converges to it in the limit (Figure 3). Therefore, when i is large, there are two possibilities. First, $x_i \leq 1/\gamma$. In this case, an information cascade occurs, and individual i chooses action 0 with probability 1. Second, $x_i > 1/\gamma$ and very close to $1/\gamma$. In this case, individual i is in the noncascade set and chooses action 1 if and only if she receives the strongest signal, γ. Thus, the lower bound for the asymptotic probability of taking correct actions is $F^0\left(\gamma^-\right)$, the probability of *not* receiving γ. This lower bound also applies to the expected fraction of correct actions.

The proposition further shows that the performance of the unlearning algorithm depends on the tail behavior of F^0. If the strongest signal γ occurs with very small probability (i.e., F^0 has a thin right tail), then $\mathbb{E}[\vartheta_i]$ is close to 1. However, when signals are discrete, this fraction is always strictly below 1. This naturally suggests that with continuous signals, one might obtain $\mathbb{E}[\vartheta_i] \to 1$; this case is analyzed in the next subsection.

5.2 Continuous Signals

We now analyze learning dynamics when signals are continuous. We first establish conditional convergence results parallel to those in Section 5.1. We then study unconditional convergence and show that, with continuous signals, society can achieve $\mathbb{E}[\vartheta_i] \to 1$.

5.2.1 Conditional Convergence

The following theorem establishes conditional convergence to the correct cascade set:

Theorem 2. *Suppose signals are continuous. Let Q be any compact subset of $C_1^c = [0,\gamma)$. Conditional on the event that x_i visits Q infinitely often, a correct limit cascade occurs $\mathbb{P}$-almost surely if $\rho \in (0,1)$ and only if $\rho \in (0,1]$. Moreover, during a correct cascade, $x_i \to 1/\gamma$ almost surely.*

Theorem 2 parallels Theorem 1 in showing that a correct limit cascade occurs almost surely if $\rho \in (0,1)$ and only if $\rho \in (0,1]$. However, it additionally requires that x_i visits a compact set Q outside the incorrect cascade set C_1 infinitely often. This condition arises from the Kushner-Clark lemma in stochastic approximation (Kushner and Clark, 1978). Infinite returns to a compact subset of $[0,\gamma)$ ensure that the process repeatedly enters regions where the drift is well behaved (i.e., bounded away from zero or already within the correct cascade set). See Section 5.3 for further discussion of the method.

Remark 2 (Discrete vs. Continuous Signals). Because α_i is decreasing, we already know that $x_i < \gamma$ infinitely often. In the discrete-signal case, each action conveys a minimum amount of information whenever a cascade is absent, which ensures that x_i enters Q infinitely often almost surely. With continuous signals, by contrast, x_i may drift toward the boundary γ, where actions carry vanishing informational content, making it more delicate to show that the process returns to any compact subset of $[0,\gamma)$.

The following proposition characterizes the rate of conditional convergence.

Proposition 4 (Convergence Rate). *Suppose $\rho \in (0,1)$, and signals are continuous with $f^\theta(\gamma) > 0$. Then, for every $\varepsilon > 0$,*

$$\mathbb{E}\left[(x_i - 1/\gamma)^2 1_{\{x_i \in Q\}}\right] = O\left(\alpha_i^{1-\varepsilon}\right)$$

for any compact set $Q \subset C_1^c$. Moreover, when $x_i \in C_0$, we have $1/\gamma - x_i = O(\alpha_i)$.

This result parallels the discrete-signal case but differs in two respects. First, convergence is established only on a compact set $Q \subset C_1^c$, which reflects the same technical requirement in Theorem 2. Second, we require $f^\theta(\gamma) > 0$. This assumption guarantees that the dynamics of x_i near $1/\gamma$ admit a first-order linear approximation. If instead $f^\theta(\gamma) = 0$, the increments of x_i near $1/\gamma$ depend on higher-order terms, and the convergence rate can differ substantially.

We also obtain a parallel statement for asymptotic actions.

Corollary 1 (Asymptotic Actions). *Suppose signals are continuous. Let Q be any compact subset of C_1^c. Conditional on the event that x_i visits Q infinitely often, both $\mathbb{E}[\vartheta_i] \to 1$ and $\mathbb{P}(a_i = 0) \to 1$ as $i \to +\infty$.*

This corollary follows from the same reasoning as in the discrete case: conditional on infinitely many returns to a compact subset of $[0,\gamma)$, the public likelihood ratio converges to $1/\gamma$ almost surely. As a result, the expected fraction of correct actions, as well as the probability of taking the correct action, converges to 1. The proof is analogous to Proposition 3 and is therefore omitted.

5.2.2 Global Convergence Under ε-Truncated Algorithms

A central technical challenge is verifying that the process $\{x_i\}$ returns to a compact subset of $[0, \gamma)$ infinitely often. As noted earlier, establishing it directly is difficult: as x_i approaches the boundary point $1/\gamma$, its increment becomes arbitrarily small, so many classical tools for proving infinite returns fail.[16]

To resolve this difficulty, we introduce an *ε-truncated algorithm*. This procedure forces the belief process to return to an interior compact region infinitely often, but only rarely; the frequency of truncations vanishes over time. This allows us to derive unconditional convergence results, both for beliefs and for actions.

Definition 2 (ε-Truncated Algorithm)**.** Define the truncation operator

$$T_\varepsilon(x) = \begin{cases} \gamma - \varepsilon & x \in (\gamma - \varepsilon, \gamma) \\ 1/\gamma + \varepsilon & x \in (1/\gamma, 1/\gamma + \varepsilon) \\ x & \text{otherwise} \end{cases}.$$

Let $I^* = \{i_1, i_2, ...\}$ be an infinite subset of individuals. For each $i_k \in I^*$, the update rule is:

$$x_{i_k+1} = \begin{cases} \mathcal{A}(x_{i_k}, a_{i_k}) & \text{with prob. } 1 - p_k \\ T_\varepsilon(\mathcal{A}(x_{i_k}, a_{i_k})) & \text{with prob. } p_k \end{cases},$$

where $\mathcal{A}(x, a)$ denotes the power-Bayesian algorithm and $p_k \to 0$ while $\sum_{k=1}^{\infty} p_k = +\infty$. We refer to this as the *ε-truncated power-Bayesian algorithm*.

Figure 4 illustrates the perturbed algorithm. The operator $T_\varepsilon(x)$ pushes beliefs away from the boundary layers $(\gamma - \varepsilon, \gamma)$ and $(1/\gamma, 1/\gamma + \varepsilon)$ toward the interior points $\gamma - \varepsilon$ and $1/\gamma + \varepsilon$, respectively (Figure 4a). Note that only individuals in $i \in I^*$ are eligible for truncation, and even there truncation occurs with probability p_k. When truncated, the next period belief x_{i_k+1} is a truncated version of the power-Bayesian update $\mathcal{A}(x_{i_k}, a_{i_k})$; see Figure 4b.

The condition $\sum_{k=1}^{\infty} p_k = +\infty$ guarantees that truncation occurs infinitely often, ensuring that x_i returns to the compact interval $[1/\gamma + \varepsilon, \gamma - \varepsilon]$ infinitely many times. The condition $p_k \to 0$ ensures that truncations become rare, so their long-run effect is negligible.

We now state the main result for continuous signals.

[16] To ensure infinite visits, it is often required that the perturbation around the stationary point ($1/\gamma$ in our case) is bounded away from zero with strictly positive probability, e.g., Pemantle (1990).

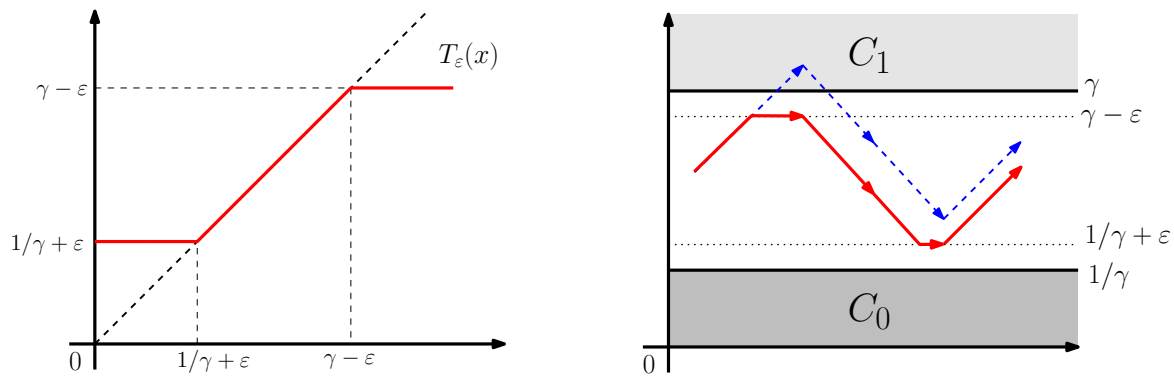

(a) ε-Truncation (b) Truncated Dynamics

Figure 4: ε-Truncated Algorithms

Theorem 3 (Limit Cascades and Actions)**.** *Suppose signals are continuous and $\rho \in (0,1)$. For any $\varepsilon > 0$, there exists an ε-truncated power-Bayesian algorithm such that: (i) for every $\delta > 0$,*

$$\lim_{n\to+\infty} \frac{1}{n}\sum_{i=1}^{n} \mathbb{P}\left(|x_i - 1/\gamma| > \delta\right) = 0,$$

and, (ii) for asymptotic actions,

$$\lim_{i\to+\infty} \mathbb{E}[\vartheta_i] = \lim_{n\to+\infty} \frac{1}{n}\sum_{i=1}^{n} \mathbb{P}\left(a_i = 0\right) = 1,$$

which further implies that $\vartheta_i \to 1$ in probability.

Theorem 3(i) parallels Theorem 1 by establishing that a correct limit cascade emerges: asymptotically, almost all beliefs lie arbitrarily close to $1/\gamma$. Part (ii) then shows that the unlearning algorithm achieves asymptotically correct behavior: the fraction of correct actions converges to one in probability.

5.3 Methodology and Intuition

This subsection outlines the methodology and provides intuition behind our main results. For clarity, we focus first on the case in which $\rho \in (1/2, 1)$ and the signal distribution F^θ is continuously differentiable; other cases are discussed at the end of the subsection.

Rewrite the Dynamics. From (2) and (4), the public LLR z_i satisfies

$$z_{i+1} = z_i + \alpha_{i+1}\left[h\left(z_i\right) + M_{i+1} + \varepsilon_i\right], \tag{9}$$

where: (i) $h(z_i) = \mathbb{E}[\ln(\phi^1/\phi^0)|\mathcal{F}_i]$ is the drift, i.e., the expected change in the public LLR, (ii) M_{i+1} is a martingale difference, and (iii) the approximation error $\varepsilon_i \to 0$. Thus the learning process is a stochastic approximation (SA) algorithm with state-dependent drift.

Stochastic Approximation. Ignoring the vanishing term $\varepsilon_i = 0$, recursion (9) becomes a classical Robbins-Monro algorithm (Robbins and Monro, 1951). When $\rho \in (1/2, 1)$, the step size α_i satisfies the standard SA conditions:[17]

$$\sum_i \alpha_i = \infty \text{ and } \sum_i \alpha_i^2 < +\infty.$$

Under these conditions, the recursion is well-approximated by the ODE

$$\dot{z}(t) = h(z(t)).$$

Because F^θ is continuously differentiable, the drift h is Lipschitz continuous, and the ODE admits a unique solution.

Analysis of the Deterministic Dynamics $z(t)$. Let $D^* \equiv (\ln(\gamma), \ln(1/\gamma))$, the log-transformed noncascade region. The drift satisfies

$$h(z) \begin{cases} < 0 & \text{when } z \in D^* \\ = 0 & \text{when } z \notin D^* \end{cases} \tag{10}$$

because actions are informative when $z_i \in D^*$ (so Jensen's inequality implies $h(z_i) < 0$), and uninformative during cascades, when $z_i \notin D^*$.

Lemma 2 (Attraction Area). *If $z_0 \in D^* \cup \{z^*\}$, then $z(t) \to z^* \equiv \ln(1/\gamma)$ as $t \to +\infty$.*

Thus z^* is the global attractor of the ODE on the noncascade region. This provides heuristic intuition for Theorem 1: if the stochastic recursion tracks this ODE, then $x_i = e^{z_i}$ should converge to $1/\gamma$, meaning that the process converges to the correct limit cascade.

From $z(t)$ to Stochastic Dynamics z_i. The connection between the ODE convergence and the convergence of z_i follows from Kushner and Clark (1978):

[17] The condition $\sum_i \alpha_i = \infty$ ensures that the "algorithmic time" $t_i = \sum_{j \le i} \alpha_i \to +\infty$ as $i \to +\infty$, allowing the process to follow long-run dynamics. The condition $\sum_i \alpha_i^2 < \infty$ implies $\alpha_i \to 0$ and ensures that the cumulative noise $\sum \alpha_i M_i$ remains finite.

Lemma 3 (Kushner-Clark)**.** *Conditional on that z_i visits a compact set $Q \subset D^* \cup \{z^*\}$ infinitely often, $z_i \to z^*$ almost surely.*

Theorem 2 is a direct implication of this lemma. However, as noted earlier, in the continuous-signal case the drift $h(z) \to 0$ as $z \to \ln \gamma$, so the process may drift toward the incorrect boundary, and showing infinite returns to a compact $Q \subset D^*$ becomes delicate.

In contrast, when signals are discrete, the drift h on D^* is bounded away from zero. Actions therefore generate minimum-information shocks that continually push the process back into the attraction area. This yields infinite returns to any compact $Q \subset D^*$, and hence unconditional convergence $z_i \to z^*$. This provides a more refined intuition for Theorem 1.

Other Cases. The argument above assumes $\rho \in (1/2, 1)$ and continuous signals, but the approach extends more generally. First, for general ρ, we use SA results from Ljung (1977) and Benveniste et al. (1990), which require only

$$\sum_i \alpha_i = \infty \text{ and } \sum_i \alpha_i^p < +\infty \text{ for some } p > 1,$$

conditions satisfied by all power-decay schemes with $\rho \in (0, 1)$. Second, for discrete signals, differentiability fails, and standard SA tools do not apply directly. Instead, we use a different argument to show that $1/\gamma$ is globally attracting. The key step, mentioned above, is to prove that x_i returns infinitely often to a compact subset of the attraction region of $1/\gamma$.[18]

6 Learning Dynamics When $\rho \leq 0$ or $\rho > 1$

To complete the characterization of learning dynamics, this section studies the remaining cases in which $\rho \leq 0$ or $\rho > 1$. Unless explicitly stated otherwise, the results in this section apply to both discrete and continuous signal structures.

6.1 Path-Dependent Cascades

We begin with the case $\rho \leq 0$.

Proposition 5. *If $\rho \leq 0$, a limit cascade occurs $\mathbb{P}$-almost surely, and both correct and incorrect limit cascades occur with strictly positive $\mathbb{P}$-probability.*

The condition $\rho \leq 0$ encompasses two cases: (i) $\rho < 0$, in which the Bayesian weight diverges to infinity; (ii) $\rho = 0$, in which the Bayesian weight converges to a non-zero constant,

[18] In particular, the proof shows that x_i visits the window $[1/\gamma, 1/\gamma + \varepsilon]$ infinitely often for any $\varepsilon > 0$.

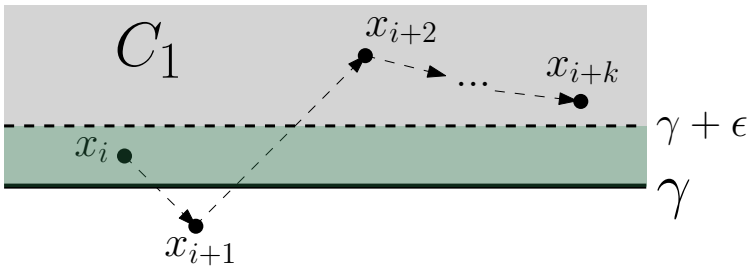

Figure 5: Power-Bayesian With $\rho = 0$ and Decreasing α_i

but it may increase or decrease with i. Proposition 5 establishes that if society does not eventually discard all public history, a limit cascade must occur, and the resulting cascade may be either correct or incorrect.

When signals are discrete, a stronger result holds: an information cascade occurs in finite time and persists forever. That is, there exists a finite index I such that $x_i \in C$ for all $i \geq I$, mirroring the classic cascade results of Bikhchandani et al. (1992).

Corollary 2. *If signals are discrete and $\rho \leq 0$, then an information cascade occurs and persists forever $\mathbb{P}$-almost surely. Both correct and incorrect cascades occur with strictly positive $\mathbb{P}$-probability.*

This result follows from the fact that when α_i does not converge to zero, there exists a uniform positive lower bound on the weight placed on public information. As a result, finitely many identical actions suffice to trigger a cascade, as in the classical model. To show that such a cascade persists forever, consider two cases.

First, if α_i is weakly increasing, then once a cascade forms, all subsequent individuals place at least as much weight on public information as the agent who initiated the cascade, ensuring that the cascade persists.

Second, if α_i is decreasing, the public belief may temporarily exit the cascade region. However, when i is large, any such exit occurs arbitrarily close to the cascade boundary, since α_i has nearly stabilized. Because the Bayesian weight remains bounded away from zero, a single cascade action by the next agent is sufficient to push the public belief back into the interior of the cascade region, after which it becomes trapped (see Figure 5). Based on this argument, we can further show that the public belief almost surely becomes absorbed into a (possibly incorrect) cascade in the long run.

The Case of Constant α_i

A special case of Proposition 5 is when α_i is constant.

Proposition 6. *Suppose $\alpha_i = \alpha > 0$ for all i. Then: (i) A limit cascade occurs $\mathbb{P}$-almost surely, and both correct and incorrect cascades occur with strictly positive $\mathbb{P}$-probability; (ii) letting $p(\alpha)$ denote the probability of a correct cascade, we have $p(\alpha) \to 1$ as $\alpha \to 0$.*

When α_i is constant, the public likelihood ratio x_i is a martingale, and the learning dynamics coincide with those of the standard Bayesian model illustrated in Figure 1. Part (i) therefore follows directly from Proposition 5. Part (ii) provides a simple comparative static. As the common weight on public information decreases, the probability of a correct limit cascade approaches one. Intuitively, a smaller α increases the number of identical actions required to trigger a cascade, allowing society to aggregate more information before herding begins. As α becomes arbitrarily small, learning remains informative for a long time, making incorrect cascades increasingly unlikely.

6.2 No Information Aggregation

We now turn to the case in which the Bayesian weights decay too rapidly.

Definition 3. We say that *no information aggregation* occurs in the limit if $x_i \to 1$.

By Proposition 1, convergence of x_i to 1 implies that, asymptotically, individuals rely solely on their private signals. In this case, public information no longer affects the decisions of tail individuals, and social learning fails to aggregate information across agents. We have the following result.

Proposition 7. *If $\rho > 1$, no information aggregation occurs in the limit $\mathbb{P}$-almost surely.*

Proposition 7 states that if the society eventually discards public information too aggressively, social learning breaks down asymptotically. Intuitively, if the weight placed on public information decays too quickly, newly generated public information cannot accumulate fast enough to influence future beliefs. As a result, individuals increasingly ignore social signals and act almost entirely on their private information. Figure 6 illustrates this behavior. The trajectory of x_i is trapped between two deterministic bounds that both converge to 1.[19] As a result, we have $x_i \to 1$, and no information aggregation is achieved.

6.3 Discussion

Controlled Speed of Unlearning. The previous discussion highlights the delicate balance required for successful information aggregation. Excessive reliance on public information prevents individuals from effectively conveying their private information, potentially

[19] As shown in the appendix, the upper curve is $\gamma^{i\alpha_{i+1}}$ and the lower curve is $1/\gamma^{i\alpha_{i+1}}$. When $\rho > 1$, we have $i\alpha_{i+1} \to 0$ as $i \to +\infty$, so both curves converge to 1 in the limit.

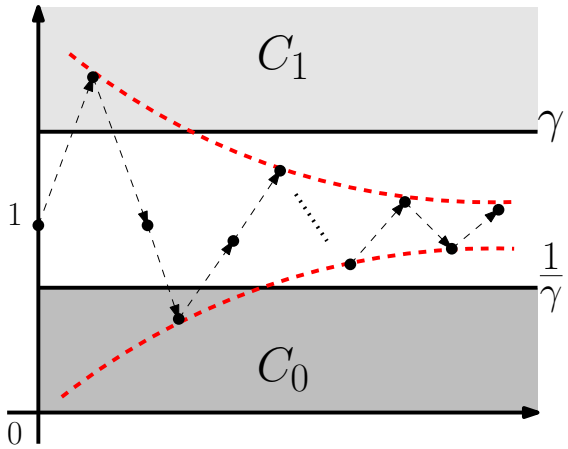

Figure 6: Power-Bayesian With $\rho > 1$

trapping society in an incorrect informational cascade (Proposition 5). Conversely, insufficient reliance on public information can also harm social learning, as the speed of discarding public information outpaces the speed of information accumulation, leading to no information aggregation in the limit (Proposition 7). Thus, successful information aggregation can only be achieved when society carefully regulates the rate at which it discards public information.

Learning Incompleteness. The result of almost-surely correct cascades is different from the path-dependent patterns in standard models. That being said, it still maintains several features of cascade results—notably, learning is *incomplete* in the sense that individuals' posteriors converge to a non-extreme belief in the limit. Recall that $x_i \to 1/\gamma$, so individuals' posteriors satisfy:

$$\pi_i\left(0|\lambda_i, h_i\right) \to \pi_\infty\left(0|\lambda_i\right) = \frac{\gamma}{\gamma + \lambda_i},$$

so individuals never achieve full confidence in the true state: their subjective beliefs remain strictly less than 1. Notably, the limit belief $\pi_\infty\left(0|\lambda_i\right) \geq 1/2$ for all λ_i between $1/\gamma$ and γ. This indicates that, in the limit, individuals believe the true state is the most likely state for almost all signals. Therefore, although learning is incomplete, asymptotic actions are correct under almost all private signals.

Information Aggregation. Although learning is incomplete, information accumulation in social learning continues indefinitely. Due to the infinite cycles of exiting and re-entering the cascade set, society reveals an unbounded amount of information over time. More precisely, consider a *Bayesian outside observer*. For such an observer, the belief about the true state

satisfies:

$$\text{As } i \to +\infty : \quad \pi^B\left(0|h_i\right) \to 1 \quad \mathbb{P}-\text{a.s.}$$

In other words, a Bayesian outside learner can completely learn the true state based on the public information. The co-existence of incomplete learning and continuous information aggregation is another interesting feature of our model. This feature is also present in sequential social learning models with continuous signals (Smith and Sørensen, 2000; Herrera and Hörner, 2012). In those models, learning remains incomplete because the speed of information aggregation decreases over time. As a result, even though information continues to accumulate, the informativeness of society becomes bounded in the limit. In contrast, in our model, this phenomenon arises because the weight on public information decays to zero, yet the speed of information accumulation remains steady, allowing society to achieve unbounded informativeness in the Bayesian sense asymptotically.

7 Implications for Algorithm Design

This paper has implications for the design of recommendation systems and online decision-support algorithms that aggregate past users' behavior. To illustrate the connection, consider a sequential environment with binary states and actions, $\Theta = A = \{0, 1\}$, in which individuals arrive one by one and aim to match their action to the true state. Each individual i observes two pieces of information:

1. a private signal λ_i as before, and

2. a recommendation produced by an algorithm that summarizes past actions.

Importantly, individuals need not observe the entire action history. Instead, they rely on the algorithm's output, which aggregates previous users' choices into a simple posterior-like score. This informational structure mirrors many online platforms, such as product ratings on Amazon or Yelp, where users do not perfectly see all past decisions, but instead condition their behavior on a summary statistic generated by the platform.

7.1 The Algorithm and Implications

Recommendation Algorithm. Formally, a *recommendation algorithm* is a mapping

$$\mathcal{A} : \cup_{i=1}^{\infty} \{0, 1\}^i \to \Delta\left(\Theta\right),$$

which assigns each feasible action history h_i to a belief over states. For each individual i, the recommendation $\mathcal{A}(h_i)(\theta)$ can be interpreted as the platform's summary assessment of the likelihood of state θ based on past actions. The benchmark model in this paper corresponds to the following recommendation rule:

$$\mathcal{A}(h_i)(\theta) \propto p_F^{\alpha_i}(\theta|h_i), \tag{11}$$

where $p_F(\theta|h_i)$ is the Bayesian posterior based on the public history, and $\alpha_i > 0$ determines how heavily the algorithm relies on accumulated public information. When $\alpha_i = 1$, the algorithm reports the Bayesian posterior; when $\alpha_i < 1$, it downweights public information; and when α_i decreases over time, it implements the unlearning algorithm analyzed in this paper.

Individual Updating. Suppose individuals are *obedient* and take the recommendation at face value. They incorporate their private signals using Bayes' rule, so their posterior beliefs satisfy

$$\pi_i(\theta|h_i, \lambda_i) \propto f^\theta(\lambda_i) \times \mathcal{A}(h_i)(\theta).$$

Under this interpretation, the induced belief dynamics coincide exactly with those analyzed in the main text. Consequently, all of our results apply directly to this algorithmic environment. It is worth noting that full obedience is a strong benchmark assumption; we discuss this assumption and possible relaxations in the next subsection.

Implications. A central implication of the analysis is that recommendation algorithms can benefit from placing a *time-decaying weight* on aggregated public information. When the weight on past actions declines gradually—specifically, at a sub-reciprocal rate—public information is slowly unlearned. This prevents permanent herding while still allowing information aggregation to continue over time, yielding correct limit cascade in the model.

From a design perspective, such an algorithm balances two forces: it tempers the influence of early, potentially misleading actions, while preserving enough influence of past behavior to aggregate dispersed private information. In practice, this logic aligns with systems that discount older content and emphasize more recent signals, thereby mitigating incorrect herding without eliminating social learning altogether.[20]

[20]For example, content-ranking and recommendation systems often prioritize more recent actions or observations and gradually downweight older ones (e.g., recency-based ranking in search engines, news feeds, or online content platforms).

7.2 Discussion

The previous discussion relies on some assumptions, which we discuss below.

Full Obedience. We assume that individuals fully trust and follow the recommendation algorithm. This assumption captures environments in which users face limited time, information, or cognitive resources, and therefore rely on a platform's recommendation as a summary statistic rather than reconstructing the entire inference problem themselves.[21] Empirically, many studies document that people often defer to algorithmic advice—even when the algorithm is imperfect—due to automation bias and perceived algorithmic authority (Suresh et al., 2020; Krügel et al., 2022). The full-obedience case thus provides a natural benchmark for isolating the informational role of the algorithm. Once this benchmark is understood, the analysis can be extended in several directions: individuals may partially trust the algorithm,[22] strategically interpret recommendations, or obedience could be supported through appropriate incentives or platform design. We leave these extensions to future work.

No Observation of Private Signals. The algorithm observes only past actions but not private signals. This distinction between observable behavior and latent private information is standard in recommender-systems design (Hu et al., 2008). It also reflects many online environments—such as product purchases, ratings, likes, or flagging decisions—where platforms observe what users do but not their underlying private assessments. One may wonder whether outcomes could be improved if users directly reported their private information to the algorithm. While this is true in principle, in practice privacy concerns often prevent platforms from accessing fine-grained user data even when technically feasible (Goldfarb and Tucker, 2012). As a result, algorithms often infer information from observed choices alone, mirroring the informational structure in sequential social learning.

Specific Algorithmic Form. The algorithm takes the power-Bayesian form with decaying α_i. A natural extension is to consider a more general framework in which individuals assign weights to public information not only based on their position in the decision sequence, but also on the positions of their predecessors. Such an approach would capture settings

[21] More broadly, this can be interpreted as an extreme case of limited attention or bounded rationality, in which users economize on cognitive effort by relying on a platform-provided summary rather than reconstructing the inference problem from scratch.

[22] For example, one can consider a more general formulation in which

$$\pi_i\left(\theta|h_i,\lambda_i\right) \propto f^{\theta}\left(\lambda_i\right) \times \psi\left[\mathcal{A}\left(h_i\right)\right](\theta),$$

where $\psi : \Delta(\Theta) \to \Delta(\Theta)$ represents how individuals process or distort the algorithm's output. A simple specification is $\psi(p) \propto p^{\beta}$,where $\beta > 1$ captures over-trust and $\beta < 1$ captures skepticism.

where individuals place greater weight on the actions of predecessors who are temporally or socially closer to them. Additional extensions could involve stochastic Bayesian weights, where weights vary randomly, or endogenous Bayesian weights, where individuals strategically decide how much weight to assign to each predecessor. While these richer formulations offer greater flexibility, we view the simple, time-decaying algorithm studied in this paper as the essential first step in understanding the effect of unlearning algorithms in social learning.

Uniform Arrival Rate. For analytical tractability, the model assumes a uniform arrival rate, with one agent arriving in each period. In practice, user participation on platforms is often uneven or clustered. When arrivals are non-uniform, the effective speed at which public information accumulates changes, which in turn may call for adjusting the decay rate of the unlearning algorithm. We believe, however, that the core trade-off identified in this paper continues to govern learning outcomes. Extending the analysis to endogenous or irregular arrival processes is a natural direction for future work.

8 Conclusion

This paper introduces a power-Bayesian learning framework in which the weight placed on public information varies systematically with an agent's position in the sequence. We show that, even with bounded private signals, correct information cascades can arise almost surely when public information is discarded gradually at a controlled rate. By contrast, discarding public history too quickly, or failing to sufficiently discount it in the limit, undermines information aggregation. These results also carry implications for the design of recommendation algorithms: applying a slowly decaying weight to past actions can mitigate incorrect herding while preserving long-run learning. The analysis focuses on a stylized sequential environment with common weighting schemes and a fixed state. Extending the framework to richer informational structures, endogenous weighting, or dynamic environments is a promising direction for future research.

A Stochastic approximation theory

We first introduce some useful results in stochastic approximation theory tailored to the context of this paper. Most of the following results can be found in Borkar (2023) and Benveniste et al. (1990).

A.1 Robbins–Monro (RM) Algorithm

Let $\{\lambda_n\}$ be a sequence of bounded i.i.d. random variables, e.g., the signals in our context. Suppose that $\{z_n\}$ is generated by the following adaptive algorithm

$$z_{n+1} = z_n + \alpha_{n+1} \left[H\left(z_n, \lambda_n\right) + \varepsilon_n \right],$$

where $z_n \in \mathbb{R}$ with z_1 fixed, and $\varepsilon_n = \varepsilon_n\left(z_n, \lambda_n\right)$ captures the (possibly biased) noise. Another way to express the dynamics is

$$z_{n+1} = z_n + \alpha_{n+1} \left[h\left(z_n\right) + M_{n+1} + \varepsilon_n \right], \tag{12}$$

where: (i) $h\left(z_n\right) = \mathbb{E}\left(H\left(z_n, \lambda_n\right) | \mathcal{F}_n\right)$ is the conditional expectation based on $\mathcal{F}_n = \sigma\left(\lambda_1, ..., \lambda_{n-1}\right)$, and (ii) $M_{n+1} = H\left(z_n, \lambda_n\right) - h\left(z_n\right)$ is a martingale difference sequence.

Assumption 3. *$h : \mathbb{R} \to \mathbb{R}$ is Lipschitz continuous.*

Assumption 4. *$\{M_n\}$ satisfies $\mathbb{E}\left(M_{n+1}^2 | \mathcal{F}_n\right) \leq K_0\left(1 + z_n^2\right)$ for some $K_0 < \infty$.*

Assumption 5. *$\{\alpha_n\}$ satisfies $\sum \alpha_n = +\infty$ and $\sum \alpha_n^2 < \infty$.*

Assumption 6. *$\sup_n |\varepsilon_n| < K < +\infty$ almost surely and $\varepsilon_n \to 0$ almost surely.*

This is the Robbins–Monro algorithm (Robbins and Monro, 1951) with the slight twist being the noise ε_n term, which could be biased. Assumptions 3 to 5 are standard assumptions in stochastic approximation. Assumption 6 ensures that the noise term ε_n vanishes in the limit. We have the following result.

Theorem 4. *Under Assumptions 3 to 6, suppose the ODE*

$$\dot{z}\left(t\right) = h\left(z\left(t\right)\right)$$

has an attractor z^* whose domain of attraction is D^*.[23] *On the event that the trajectory $\{z_n\}$ visits a compact set $Q \subset D^*$ infinitely often, $z_n \to z^*$ $\mathbb{P}$-almost surely.*

[23] The domain of attraction is the set of initial conditions z such that $z\left(t\right) \to z^*$.

Proof. See Corollary 2.3 in Borkar (2023) for results without the noise. The presence of a vanishing noise term does not affect the approximation (see Section 2.2 in Borkar (2023)). □

A.2 RM Algorithm With General Stepsize

Notice that if $\alpha_i \sim c \times i^{-\rho}$ as $i \to +\infty$, Assumption 5 is only applicable for $1/2 < \rho \leq 1$. Below are two alternative assumptions:

Assumption 7. $\{\alpha_n\}$ *satisfies* $\sum \alpha_n = +\infty$ *and* $\sum \alpha_n^p < \infty$ *for some* $p > 1$.

Assumption 8. $\sup_n |\varepsilon_n/\alpha_n| \leq K < \infty$ *almost surely.*

Assumption 7 is more general than Assumption 5, in particular it accommodates cases where $\rho \leq 1/2$. Assumption 8 requires ε_n to decay to zero at least as fast as α_n.

Theorem 5 (General Steps)**.** *Suppose Assumptions 3, 4, 7 and 8 hold. The result in Theorem 4 still holds.*

Proof. See Corollary 6 in Chapter 2 of Benveniste et al. (1990). □

B Learning Dynamics in the SA Form

Now, we express this paper's learning dynamics in the form of (12).

Lemma 4 (Stochastic Approximation Form)**.** *Let* $z_i = \ln(x_i)$*. Then, we have*

$$\forall i: \quad z_{i+1} = z_i + \alpha_{i+1}\left[h(z_i) + M_{i+1} + \varepsilon_i\right],$$

where (i) $h(z_i) = \mathbb{E}\left(\ln\left(\frac{\phi^1(a_i,x_i)}{\phi^0(a_i,x_i)}\right)|\mathcal{F}_i\right)$*, (ii)* $M_{i+1} = \ln\left(\frac{\phi^1(a_i,x_i)}{\phi^0(a_i,x_i)}\right) - h(z_i)$*, and (iii)* $\varepsilon_i = \left(\frac{1}{\alpha_i} - \frac{1}{\alpha_{i+1}}\right) z_i$*.*

Proof. It comes directly from the dynamics of $\{x_i\}$. □

Lemma 5 (Lipschitz Continuity)**.** *Suppose* F^θ *is continuously differentiable, then* $h: \mathbb{R} \to \mathbb{R}$ *is Lipschitz continuous.*

Proof. Define a function $g: (1/\gamma, \gamma) \to (1/\gamma, \gamma)$, where

$$g(x) = \left(1 - F^0(x)\right) \ln\left(\frac{1 - F^1(x)}{1 - F^0(x)}\right) + F^0(x) \ln\left(\frac{F^1(x)}{F^0(x)}\right).$$

It can be verified that

$$g'(x) = f^0(x) \ln\left(\frac{F^1(x)(1-F^0(x))}{F^0(x)(1-F^1(x))}\right) + f^1(x)\left[\frac{F^0(x)}{F^1(x)} - \frac{1-F^0(x)}{1-F^1(x)}\right].$$

Notice that $\frac{F^1(x)}{F^0(x)}, \frac{1-F^1(x)}{1-F^0(x)} \in [1/\gamma, \gamma]$, $F^1(x) \in [0,1]$. Since F^0 is continuously differentiable, f^0 is also bounded on $[1/\gamma, \gamma]$. This implies that $|g'(x)| \leq M$ for some finite M. Notice that

$$h(z) = \begin{cases} g(e^{-z}) & z \in D^* \\ 0 & z \notin D^* \end{cases},$$

where $D^* = (-\ln(\gamma), \ln(\gamma))$. Since $h'(z) = -g'(e^{-z})e^{-z}$, so h also has bounded derivative on D^*. Thus, h is also Lipschitz continuous on D^*, and let L_0 be the Lipschitz constant. We can show that L_0 is also the Lipschitz constant of h over the entire $\mathbb{R}$. Suppose $z_1 \in D^*$ and $z_2 \notin D^*$, then there exists a boundary point $d \in \{-\ln\gamma, \ln\gamma\}$ that is between z_1 and z_2. Therefore,

$$\frac{|h(z_1) - h(z_2)|}{|z_1 - z_2|} = \frac{|h(z_1)|}{|z_1 - z_2|} \leq \frac{|h(z_1)|}{|z_1 - d|} = \frac{|h(z_1) - h(d)|}{|z_1 - d|} \leq L_0.$$

The cases when $z_1, z_2 \in D^*$ and $z_1, z_2 \notin D^*$ trivially satisfy the previous inequality. Hence, h is globally Lipschitz continuous. □

Lemma 6 (Boundedness)**.** *If $\alpha_i \to 0$ and $\frac{1}{\alpha_{i+1}} - \frac{1}{\alpha_i} \to 0$, then z_i, M_i and ε_i are all bounded.*

Proof. (i) Since signals are bounded, $\frac{\phi^1(a_i,x_i)}{\phi^0(a_i,x_i)} \in [1/\gamma, \gamma]$, so both h and M_i are bounded; (ii) Now we show $\{z_i\}$ is bounded. If $\{z_i\}$ remains in the noncascade set, then it is obviously bounded. If $|z_i| \geq \ln(\gamma)$, then $|z_{i+1}| = \frac{\alpha_{i+1}}{\alpha_i}|z_i| \leq |z_i|$, so we only need to show that the first-entry into the cascade set is bounded. Suppose i is such that $|z_{i-1}| < \ln(\gamma)$ and $|z_i| \geq \ln(\gamma)$, then

$$|z_i| \leq |z_{i-1}| + \alpha_i \times \left(\overline{h} + \overline{M}\right) + \left|\left(\frac{1}{\alpha_{i-1}} - \frac{1}{\alpha_i}\right) z_{i-1}\right|,$$

where $\overline{h}$ and $\overline{M}$ are upper bounds of h and M_i. Note that the right-hand side is bounded since $|z_{i-1}| \leq \ln(\gamma)$ and both α_i and $\frac{1}{\alpha_{i-1}} - \frac{1}{\alpha_i}$ are bounded. □

Remark 3. Suppose $\alpha_i \sim c \times i^{-\rho}$ with $\rho \in (0,1)$, then: (i) $\alpha_i \to 0$ holds, and (ii) $\frac{1}{\alpha_i} - \frac{1}{\alpha_{i+1}} \sim -\frac{\rho}{c} i^{\rho-1}$, so $\frac{1}{\alpha_{i-1}} - \frac{1}{\alpha_i} \to 0$ also holds.

C Proofs

The proofs are organized in an order different from the main text. We first prove the results for continuous signals (Theorem 2 and Theorem 3), since they align more closely with standard form of stochastic approximation. We then prove all results for discrete signals.

C.1 Proof of Theorem 2

Proof. We first check all conditions of stochastic approximation hold. Since F^θ is continuously differentiable on $[1/\gamma, \gamma]$, h is Lipschitz continuous by Lemma 5, so Assumption 3 holds. When $\rho \in (0,1)$, both z and signals are bounded, M is also bounded by Lemma 6 and remark 3, so Assumption 4 holds as well.

To check the conditions for noise ε_i. We consider the following two cases:

- Suppose $\rho \in (1/2, 1)$, then Assumption 5 holds. $\varepsilon_i = \left(\frac{1}{\alpha_i} - \frac{1}{\alpha_{i+1}}\right) z_i \approx C\left(i^\rho - (i+1)^\rho\right) z_i$. Since z_i is bounded, $i^\rho - (i+1)^\rho \sim -\rho i^{\rho-1}$, we have $\varepsilon_i \approx -C\rho i^{\rho-1} z_i \to 0$, so Assumption 6 holds, and hence the conditions of Theorem 4 hold.

- Suppose $\rho \in (0, 1/2]$, then Assumption 7 holds. Moreover, $\varepsilon_i/\alpha_i \approx -\frac{C\rho}{c} i^{2\rho-1} z_i = O(1)$ when $\rho \leq 1/2$. Assumption 8 also holds, thus the conditions of Theorem 5 hold.

Now, we examine the properties of the following ODE

$$\dot{z}(t) = h(z(t)) \begin{cases} < 0 & \text{when } z \in (-\ln(\gamma), \ln(\gamma)) \\ = 0 & \text{when } z \notin (-\ln(\gamma), \ln(\gamma)) \end{cases}.$$

Let $D^* = (-\ln(\gamma), \ln(\gamma))$ and $z^* = -\ln(\gamma)$. Notice that for any $z(0) \in D^*$, $z(t) \to z^*$, so D^* is the domain of attraction. Theorem 2 then comes directly from Theorems 4 and 5. □

C.2 Proof of Theorem 3

The proof proceeds in the following three steps:

Step 1: Approximate the Untruncated Dynamics Using ODE

Definitions. Let $t_i = \sum_{j=1}^{i} \alpha_j$, which is referred to as the *actual time* of the algorithm for individual i. Let $i[t]$ denote the minimum i such that $t_i \geq t$, called the *t-th threshold individual*. We let

$$I^* = \left\{ i\left[2^{k-1}\right] : k = 1, 2, 3, ... \right\}.$$

For individual $i_k \equiv i\left[2^{k-1}\right]$, if $x_{i_k} \notin [1/\gamma+\varepsilon, \gamma-\varepsilon]$, truncation happens with probability p_k. If truncation happens, the next period public belief $x_{i_k+1} \in \{1/\gamma+\varepsilon, \gamma-\varepsilon\}$, so it escapes the δ-neighborhood of $1/\gamma$ for small δ. Since $\sum p_k = \infty$, such escaping happens almost surely—therefore, it is impossible that x_i converges almost surely to $1/\gamma$ as in the discrete case. The goal of the proof is not to establish almost sure convergence, but to show that the fraction of x_i outside of $B_\delta(1/\gamma)$ is zero in the limit.

Stochastic Approximation. Suppose $x_{i_k+1} \in \{1/\gamma+\varepsilon, \gamma-\varepsilon\}$. Consider the following ODE

$$\dot{z}_k(t) = h(z_k(t)), \text{where } z_k(0) = \ln(x_{i_k+1}).$$

Since $z_k(0)$ is in the attraction area of $\ln(1/\gamma)$, we have $z_k(t) \to \ln(1/\gamma)$ as $t \to +\infty$. Thus, for any $\delta > 0$, there exists $T < \infty$ such that for all k, we have:

$$|z_k(t) - \ln(1/\gamma)| < \delta/2 \quad \forall t \geq T. \tag{13}$$

Notice that when $i \in [i_k+1, i_{k+1}]$, the dynamics of z_i are untruncated as in the standard case. The theory of stochastic approximation implies that

$$\mathbb{P}\left(\sup_{i\in[i_k+1,i_{k+1}]} \left|z_i - z_k\left(t_i - 2^{k-1}\right)\right| > \delta/2\right) \leq Q_k,$$

for some sequence $\{Q_k\}$ where $Q_k \to 0$ as $k \to +\infty$.[24]

Step 2: Prove That $\frac{1}{n}\sum_{i=1}^{n} \mathbb{P}(|z_i - \ln(1/\gamma)| > \delta) \to 0$.

Main Idea. We first explain the main idea of the proof. For each k, we define

$$I_k \equiv \{i_k+1, ..., i_{k+1}\} = \left\{i\left[2^{k-1}\right]+1, ..., i\left[2^k\right]\right\}.$$

We divide I_k into two parts

$$I_k^0 \equiv \left\{i\left[2^{k-1}\right]+1, ..., i\left[2^{k-1}+T\right]\right\}, \text{ and } I_k^1 \equiv \left\{i\left[2^{k-1}+T\right]+1, ..., i\left[2^k\right]\right\}.$$

[24] In fact, $Q_k = C\sum_{i\geq i_k} \alpha_i^p$, where $C > 0$ is a constant and p is such that $\sum_i \alpha_i^p < \infty$ (see Corollary 5 in Chapter 2 of Benveniste et al. (1990)).

When $i \in I_k^1$, we have $t_i - 2^{k-1} > T$, so we have $\left|z_k\left(t_i - 2^{k-1}\right) - \ln(1/\gamma)\right| < \delta/2$ by (13), which further implies

$$\sup_{i \in I_k^1} \mathbb{P}\left(|z_i - \ln(1/\gamma)| > \delta\right) \leq \mathbb{P}\left(\sup_{i \in I_k^1} |z_i - \ln(1/\gamma)| > \delta\right) \leq Q_k.$$

Note that as $k \to +\infty$, the probability of $|z_i - \ln(1/\gamma)| > \delta$ vanishes for all $i \in I_k^1$. The only case where $|z_i - \ln(1/\gamma)| > \delta$ can happen with non-vanishing probability is when $i \in I_k^0$. The idea of the whole proof relies on that when k is large, the size of I_k^0 is negligible relative to that of I_k^1, so the pattern in I_k^1 dominates.

Formal Argument. We now provide formal arguments. For any n, suppose $t_n \in T_K = (2^{K-1}, 2^K]$. Then, it can be written as $n = i\left[2^{K-1}\right] + m$ where $0 < m < i\left[2^K\right] - i\left[2^{K-1}\right]$. We can divide the first n individuals into the following groups:

$$\{1, 2, ..., n\} = I_0 \cup I_1 \cup ... \cup I_{K-1} \cup \left\{i\left[2^{K-1}\right] + 1, ..., n\right\},$$

where $I_0 = \{1, ..., i[1]\}$. We further define $I^0 \equiv \cup_{k=1}^K I_k^0$ and $I^1 \equiv \cup_{k=1}^K I_k^1$. Let $P_i \equiv \mathbb{P}\left(|z_i - \ln(1/\gamma)| > \delta\right)$, then

$$\frac{1}{n}\sum_{i=1}^n P_i \leq \frac{1}{n}\sum_{i \in I_0} P_i + \frac{1}{n}\sum_{i \in I^0} P_i + \frac{1}{n}\sum_{i \in I^1} P_i$$

The first term $\frac{1}{n}\sum_{i \in I_0} P_i \to 0$ since $\sum_{i \in I_0} P_i$ is fixed. We now consider the next two terms:

$$\frac{1}{n}\sum_{i \in I^0} P_i \leq \frac{1}{n} \times |I^0| = \frac{\sum_{k=1}^K i\left[2^{k-1} + T\right] - i\left[2^{k-1}\right]}{i\left[2^{K-1}\right] + m} \leq \frac{\sum_{k=1}^K i\left[2^{k-1} + T\right] - i\left[2^{k-1}\right]}{i\left[2^{K-1}\right]}.$$

Recall that $i\left[2^{k-1}\right] = \min\left\{i : \sum_{j=1}^i \alpha_j \geq 2^{k-1}\right\}$ and $\alpha_i \sim c \times i^{-\rho}$ with $\rho \in (0, 1)$. Then, $i\left[2^{k-1}\right] \approx C_1 \cdot 2^{\frac{k-1}{1-\rho}}$ for some constant C_1 when k is large. Therefore,

$$\frac{1}{n}\sum_{i \in I^0} P_i \leq \frac{\sum_{k=1}^K i\left[2^{k-1} + T\right] - i\left[2^{k-1}\right]}{i\left[2^{K-1}\right]} \approx C_2 \cdot 2^{-K} \to 0 \quad \text{as } K \to +\infty,$$

where $C_2 > 0$ is a constant. For the third term:

$$\frac{1}{n}\sum_{i \in I^1} P_i \leq \frac{1}{|I^1|}\sum_{i \in I^1} P_i.$$

Suppose $i \in I_k^1$, then $P_i \leq Q_k$, so

$$\frac{1}{|I^1|}\sum_{i\in I^1} P_i \leq \frac{1}{|I^1|}\left(\sum_{i\in I_1^1} Q_1 + ... \sum_{i\in I_K^1} Q_K\right) \to 0 \quad \text{as } K \to +\infty,$$

where the limit comes from Cauchy's limit theorem and the fact that $Q_k \to 0$ as $k \to +\infty$. As a consequence, $\frac{1}{n}\sum_{i=1}^n P_i \to 0$.

Step 2: Prove That the Expected Fraction $\mathbb{E}[\vartheta_n] \to 1$.

Next, we estimate the expected fraction of correct actions. Note that for all $\delta > 0$,

$$\mathbb{P}(a_i = 0) \;\geq\; \mathbb{P}\big(\{|z_i - \ln(1/\gamma)| \leq \delta\} \cap \{\lambda_i < \gamma e^{-\delta}\}\big).$$

Taking the time average, we obtain

$$\mathbb{E}[\vartheta_n] = \frac{1}{n}\sum_{i=1}^n \mathbb{P}\left(a_i = 0\right) \geq F^0\left(\gamma e^{-\delta}\right) \times \frac{1}{n}\sum_{i=1}^n \mathbb{P}\left(|z_i - \ln\left(1/\gamma\right)| \leq \delta\right).$$

Since $\frac{1}{n}\sum_{i=1}^n \mathbb{P}\left(|z_i - \ln\left(1/\gamma\right)| \leq \delta\right) \to 1$, we have

$$\liminf \mathbb{E}[\vartheta_n] = \liminf \frac{1}{n}\sum_{i=1}^n \mathbb{P}\left(a_i = 0\right) \geq F^0\left(\gamma e^{-\delta}\right)\left(1 + o\left(1\right)\right).$$

The argument holds for all $\delta > 0$, so

$$\lim_{n\to+\infty} \mathbb{E}[\vartheta_n] = \lim_{n\to+\infty} \frac{1}{n}\sum_{i=1}^n \mathbb{P}\left(a_i = 0\right) = 1.$$

By Markov's inequality, $\mathbb{P}\left(|\vartheta_n - 1| > \epsilon\right) \leq \frac{1-\mathbb{E}[\vartheta_n]}{\epsilon} \to 0$ for any $\epsilon > 0$, so $\vartheta_n \to 1$ in probability.

C.3 Proof of Theorem 1

When signals are discrete, standard tools of stochastic approximation are not longer available, so we use different techniques to prove the case where signals are discrete. Notice that when signals are discrete, tie cases are not probability-zero event. Throughout the discussion, we impose a tie-breaking rule that individual chooses action 0 when indifferent, but this choice

is nonessential.[25]

Preliminaries

Suppose that the support $\Lambda = \{\lambda^1, \lambda^2, \ldots, \lambda^n\}$ with $\gamma = \lambda^1 > \lambda^2 > \cdots > \lambda^n = 1/\gamma$. We assume $1 \in \Lambda$ for convenience.[26] We divide the interval $D = (1/\gamma, \gamma)$ into $n-1$ disjoint zones:

$$Z_1 = (1/\gamma, 1/\lambda^2], Z_2 = (1/\lambda^2, 1/\lambda^3], ..., Z_{n-2} = (1/\lambda^{n-2}, 1/\lambda^{n-1}], Z_{n-1} = (1/\lambda^{n-1}, \gamma),$$

and let $Z_0 = [0, 1/\gamma]$ and $Z_n = [\gamma, +\infty]$. Let $\mathcal{T} = \{1, 2, ..., n-1\}$, and let $\mathcal{T}^- = \{\tau \in \mathcal{T} : Z_\tau \subset (0,1)\}$ and $\mathcal{T}^+ = \{\tau \in \mathcal{T} : Z_\tau \subset (1, \infty)\}$, which are indexes of zones below and above 1, respectively. We have the following facts.

Fact 1 (Equilibrium Strategies)**.** *Suppose $x_i \in Z_\tau$ where $\tau \in \mathcal{T}$, then*

$$a_i = \begin{cases} 1 & \text{if } \lambda_i \in S_1^\tau \equiv \{\lambda^\tau, ..., \gamma\} \\ 0 & \text{if } \lambda_i \in S_0^\tau \equiv \{1/\gamma, \ldots, \lambda^{\tau+1}\} \end{cases},$$

where S_a^τ represents the set of signals such that individuals will choose action a when $x_i \in Z_\tau$.

Fact 2 (Dynamics of x_i)**.** *Suppose $x_i \in Z_\tau$ where $\tau \in \mathcal{T}$, then*

$$\ln(x_{i+1}) = \frac{\alpha_{i+1}}{\alpha_i} \ln(x_i) + \alpha_{i+1} \times X_i^\tau,$$

where

$$X_i^\tau \equiv \sum_{a \in A} \ln\left[\frac{\mathbb{P}^1(S_a^\tau)}{\mathbb{P}^0(S_a^\tau)}\right] \times 1_{\{\lambda_i \in S_a^\tau\}},$$

where $\mathbb{P}^\theta(S_a^\tau)$ denotes the probability that $\lambda_i \in S_a^\tau$ under F^θ. Here, X_i^τ represents the increment of log public likelihood ratio when $x_i \in Z_\tau$.

The formal proof proceeds in the following three steps.

[25] Different tie-breaking rules only affect our definitions of Z_k below. Under the assumption that individuals choose action 0 in tie cases, Z_k's in the noncascade region take the form of $(a, b]$. In other cases, if we impose a tie-breaking rule that individuals choose action 1 in tie cases, then it takes the form of $[a, b)$. If we impose a tie-breaking rule that individuals follow the public belief (as in the main text), then Z_k takes the form of $[a, b]$ below 1 and (a, b) above 1 (also we need to impose a tie-breaking rule right at 1), which is a bit tedious in notation. Notice that although the tie-breaking rule we adopt in the appendix is different from that in the main text, the asymptotic results are not affected.

[26] This is also without loss of generality. Suppose $1 \notin \Lambda$, we simply let the probability on 1 be zero. The sole purpose is that for any zone Z_i, it is either above 1 or below 1.

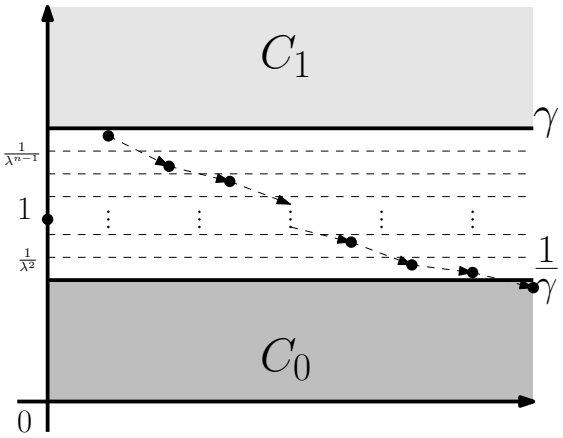

Figure 7: General Signal Space

Step 1: $\mathbb{P}\left(x_i \leq \frac{1}{\gamma} \textbf{ happens infinitely often}\right) = 1$**.**

The proof idea is to show that x_i can transit in the following trajectory:

$$Z_n \to Z_{n-1} \to Z_{n-2} \to ... \to Z_1 \to Z_0,$$

with every step happening with some probability bounded away from 0, as illustrated in Figure 7. In other words, x_i can transit from C_1 to C_0 with strictly positive probability.

For each $i \in I$, and $\tau \in \mathcal{T}$, we define an event

$$\mathcal{E}_i^{\delta,\tau} \equiv \left\{\mathcal{M}_i^k\left(X^\tau\right) < -\delta : \forall k \geq i\right\}, \text{ where } \delta > 0,$$

where $\mathcal{M}_i^k\left(X^\tau\right) \equiv \frac{1}{k}\sum_{j=i}^{i+k-1} X_j^\tau$ represents the k-moving average starting from individual i. The event means that the moving average starting from i is always less than $-\delta$. Since X_i^τ's are i.i.d., applications of Hoeffding's inequality yield that there exists $\delta, \zeta > 0$ such that $\mathbb{P}\left(\mathcal{E}_i^{\delta,\tau}\right) > \zeta$ for all i and $\tau \in \mathcal{T}$ (See Lemma 16 in Appendix D).

Lemma 7 (Bound for Transition Probability)**.** *There exists $I > 0$ such that for all $i \geq I$ and all $\tau \in \mathcal{T}$, we have*

$$\mathbb{P}\left(x_{i+k} \in \cup_{\tau'<\tau} Z_{\tau'} \text{ for some } k \geq 0 | x_i \in Z_\tau\right) \geq \mathbb{P}\left(\mathcal{E}_i^{\delta,\tau}\right) > \zeta > 0.$$

Besides, conditional on that $x_i \in Z_\tau$ and on event $\mathcal{E}_i^{\delta,\tau}$, there almost surely exists some $k \geq 0$ such that $x_{i+k} \in \cup_{\tau'<\tau} Z_{\tau'}$.

Proof. Suppose that $x_i, x_{i+1}, ..., x_{i+k-1} \in Z_\tau$, then we have:

$$\ln(x_{i+k}) = \frac{\alpha_{i+k}}{\alpha_i} \ln(x_i) + \alpha_{i+k} k \times \mathcal{M}_i^k(X^\tau).$$

Consider the following two cases:

Case 1: $\tau \in \mathcal{T}^+$**.** In this case, $x_i > 1$, so $\ln(x_i) > 0$. Let I_1 be the cutoff such that α_i is monotonic after $i \geq I_1$. When $i \geq I_1$, $\frac{\alpha_{i+k}}{\alpha_i} \leq 1$, so

$$\begin{aligned}\ln(x_{i+k}) &\leq \ln(x_i) + \alpha_{i+k} k \times \mathcal{M}_i^k(X^\tau) \\ &\leq \ln(x_i) + \alpha_{i+k} k \times (-\delta) \quad \text{on event } \mathcal{E}_i^{\delta,\tau}.\end{aligned}$$

Therefore, we have $\ln(x_{i+k}) < \ln(x_i)$ on event $\mathcal{E}_i^{\delta,\tau}$. In other words, x_{i+k} is always smaller than x_i on event $\mathcal{E}_i^{\delta,\tau}$, so it cannot increase above the upper boundary of Z_τ. On the other hand, x_{i+k} cannot stay in Z_τ forever; otherwise, we have $\alpha_{i+k} k \times (-\delta) \to -\infty$ as $k \to +\infty$, which implies $x_{i+k} \to 0$, leading to a contradiction. As a consequence, there must be some $K < \infty$ such that $x_{i+K} < \lambda^\tau$, i.e., x_{i+k} must enter some $Z_{\tau'}$ with $\tau' < \tau$.

Case 2: $\tau \in \mathcal{T}^-$. In this case, $\ln(x_i) < 0$, implying that $\frac{\alpha_{i+k}}{\alpha_i} \ln(x_i) \geq \ln(x_i)$, so we cannot use the same approach as in Case 1. We notice that

$$\begin{aligned}\ln(x_{i+k}) &= \frac{\alpha_{i+k}}{\alpha_i} \ln(x_i) + \alpha_{i+k} k \times \mathcal{M}_i^k(X^\tau) \\ &= \ln(x_i) + \alpha_{i+k} k \left[\mathcal{M}_i^k(X^\tau) - \left(\frac{1}{k\alpha_{i+k}} - \frac{1}{k\alpha_i} \right) \ln(x_i) \right].\end{aligned}$$

As $i \to +\infty$, we have

$$\mathcal{M}_i^k(X^\tau) - \left(\frac{1}{k\alpha_{i+k}} - \frac{1}{k\alpha_i} \right) \ln(x_i) \approx \mathcal{M}_i^k(X^\tau) - C \left(\frac{(i+k)^\rho - i^\rho}{k} \right) \ln(x_i).$$

Notice that $g(i,k) = \frac{(i+k)^\rho - i^\rho}{k}$ is a decreasing function in k, so $g(i,k) \leq \lim_{k\to 0} g(i,k) = \rho i^{\rho-1}$. Note that $\mathcal{M}_i^k(X^\tau) < -\delta$, and $\rho i^{\rho-1}$ can be made arbitrarily small when i is large. Besides, $|\ln(x_i)| < L_\tau$. Therefore, we can find I_2 such that $\mathcal{M}_i^k(X^\tau) - \left(\frac{1}{k\alpha_{i+k}} - \frac{1}{k\alpha_i} \right) \ln(x_i) < 0$ for all $i \geq I_2$ and $k \geq 0$. Therefore, we have $\ln(x_{i+k}) < \ln(x_i)$ on event $\mathcal{E}_i^{\delta,\tau}$ when $i \geq I_2$. Following the argument in Case 1, we can show that x_{i+k} must enter some $Z_{\tau'}$ with $\tau' < \tau$ as well.

Combining these two cases, we obtain Lemma 7. □

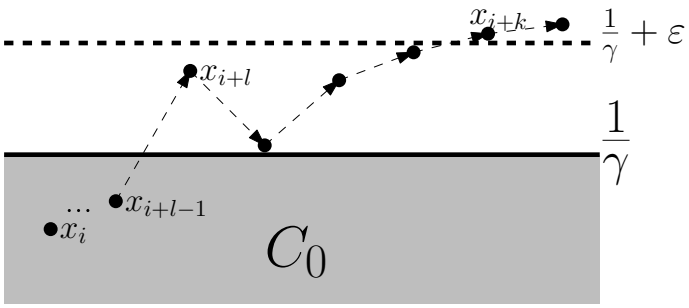

Figure 8: Illustration of an Upcrossing

Since α_i is decreasing, $x_i \in (1/\gamma, \gamma)$ infinitely often. Lemma 7 implies that

$$\forall i: \quad \mathbb{P}\left(x_{i+k} \leq 1/\gamma \text{ for some } k \geq 0 | \mathcal{F}_i\right) \geq \zeta^n > 0.$$

By Levy's extension of Borel-Cantelli Lemma, we have $x_i \leq \frac{1}{\gamma}$ happens infinitely often with probability 1.

Step 2: For large i, $\mathbb{P}\left(x_{i+k} < \frac{1}{\gamma}+\varepsilon \text{ for all } k \geq 0 | x_i \leq \frac{1}{\gamma}\right) > 0$ for all small $\varepsilon > 0$.

Next, we show that x_i will remain in the ε-neighborhood of C_0 with positive probability. We first introduce the concept of the upcrossing:

Definition 4. A chain $x_{j:j+k} \equiv \{x_j, x_{j+1}, ..., x_{j+k}\}$ consititutes an *upcrossing* from a to b if $x_{j-1} \leq a$, $x_j, ..., x_{j+k-1} \in (a, b)$ and $x_{j+k} \geq b$.

Notice that for x_i to escape from the ε-neighborhood of C_0, it must experience an upcrossing from $\frac{1}{\gamma}$ to $\frac{1}{\gamma}+\varepsilon$ during which it never falls back into C_0 (illustrated in Figure 8). We are interested in small ε, so we assume $1/\gamma + \varepsilon < 1/\lambda^2$. Next, we show that such upcrossing cannot occur infinitely often.

For any $s, L \geq 0$, we refer to the upcrossing $x_{s:s+L} = \{x_s, ..., x_{s+L}\}$ from $1/\gamma$ to $1/\gamma+\varepsilon$ as (s, L)-upcrossing, where s denotes the *starting point*, and L denotes the *length* of this upcrossing. We define

$$\Omega(s, L) = \sum_{j=s}^{s+L-1} X_j^1, \text{ where } X_j^1 = \sum_{a \in A} \ln\left[\frac{\mathbb{P}^1(S_a^1)}{\mathbb{P}^0(S_a^1)}\right] \times 1_{\{\lambda_j \in S_a^1\}}$$

where $\Omega(s, L)$ denotes the *increment of the public log-likelihood ratio* during the (s, L)-upcrossing.

Lemma 8 (Increment of Public LLR)**.** *There exists $A, B > 0$ such that $\Omega(s, L) > A/\alpha_s - B/\alpha_{s+L}$ when s is sufficiently large.*

Proof. We first note that when an (s, L)-upcrossing occurs,

$$\ln(x_{s+L}) = \frac{\alpha_{s+L}}{\alpha_s}\ln(x_s) + \alpha_{s+L} \times \Omega(s, L), \quad \text{where } \Omega(s, L) = \sum_{j=s}^{s+L-1} X_j^1.$$

Rearranging terms, we have

$$\begin{aligned}\Omega(s, L) &= \frac{1}{\alpha_{s+L}}\ln(x_{s+L}) - \frac{1}{\alpha_s}\ln(x_s) = \frac{1}{\alpha_s}\ln\left(\frac{1}{x_s}\right) - \frac{1}{\alpha_{s+L}}\ln\left(\frac{1}{x_{s+L}}\right) \\ &\geq \frac{1}{\alpha_s}\ln\left(\frac{1}{1/\gamma + \varepsilon/2}\right) - \frac{1}{\alpha_{s+L}}\ln\left(\frac{1}{1/\gamma+\varepsilon}\right) \equiv \frac{1}{\alpha_s} \times A - \frac{1}{\alpha_{s+L}} \times B,\end{aligned}$$

where $A = \ln\left(\frac{1}{1/\gamma+\varepsilon/2}\right) > B = \ln\left(\frac{1}{1/\gamma+\varepsilon}\right) > 0$. The inequality employs the fact that $1 > x_{s+L} \geq 1/\gamma + \varepsilon$ and $x_s < 1/\gamma + \varepsilon/2 < 1$ when s is large.[27] □

Lemma 9 (Probability of (s, L)-Upcrossing)**.** *Let $P(s, L)$ denote the probability of an (s, L)-upcrossing. When s is sufficiently large, we have:*

$$P(s, L) \leq \exp\left(-\bar{C}s^\rho\right) \quad \textit{when } L \leq L(s),$$

and

$$P(s, L) \leq \exp\left(-\tilde{C}L\right) \quad \textit{when } L > L(s),$$

where $\bar{C}, \tilde{C} > 0$ and $L(s)$ is some threshold L depending on s.

Proof. When s is sufficiently large, Lemma 8 implies that

$$P(s, L) \leq \mathbb{P}\left(\Omega(s, L) \geq \frac{1}{\alpha_s} \times A - \frac{1}{\alpha_{s+L}} \times B\right) = \mathbb{P}\left(\sum_{j=1}^{L} X_j^1 \geq C_1 s^\rho - C_2(s+L)^\rho\right),$$

where $C_1 = A \times c > C_2 = B \times c > 0$ are constants for some $c > 0$.
Denote

$$R(s, L) \equiv C_1 s^\rho - C_2(s+L)^\rho = s^\rho\left[C_1 - C_2\left(1 + \frac{L}{s}\right)^\rho\right],$$

then

$$R(s, L) \geq 0 \quad \text{if and only if} \quad C_1 - C_2\left(1 + \frac{L}{s}\right)^\rho \geq 0.$$

That is,

$$\frac{L}{s} \leq \left(\frac{C_1}{C_2}\right)^{\frac{1}{\rho}} - 1.$$

[27] Recall that s is the first individual during an upcrossing. The stepsize goes to zero as $i \to +\infty$, so x_s is very close to $1/\gamma$ when i is large.

Fix $\kappa \in \left(0, \left(\frac{C_1}{C_2}\right)^{\frac{1}{\rho}} - 1\right)$. Whenever $\frac{L}{s} \leq \kappa$, we have $R(s,L) \geq 0$. Define $L(s) \equiv \lfloor \kappa s \rfloor$. Next, we consider the bounds on $\mathbb{P}\left(\sum_{j=1}^{L} X_j^1 \geq R(s,L)\right)$ for cases where $1 \leq L \leq L(s)$ and $L > L(s)$ separately.

Case 1: $1 \leq L \leq L(s)$**.** In this case, $R(s,L) \geq 0$. First, X^1 is bounded and can be included in a set $[m, M]$. Second, $R(s,L) - L\mathbb{E}\left(X_j^1\right) > 0$ because $\mathbb{E}\left(X_j^1\right) < 0$. Then, we can apply Hoeffding's inequality and obtain:

$$\begin{aligned}
\mathbb{P}\left(\sum_{j=1}^{L} X_j^1 \geq R(s,L)\right) &= \mathbb{P}\left(\sum_{j=1}^{L} X_j^1 - L\mathbb{E}\left(X_j^1\right) \geq R(s,L) - L\mathbb{E}\left(X_j^1\right)\right) \\
&\leq \exp\left(-\frac{2\left(R(s,L) - L\mathbb{E}\left(X_j^1\right)\right)^2}{L(M-m)^2}\right) \\
&= \exp\left(-\frac{2\left(R(s,L) + L\mu\right)^2}{L(M-m)^2}\right),
\end{aligned}$$

where $\mu = |\mathbb{E}\left(X_j^1\right)| = -\mathbb{E}\left(X_j^1\right) > 0$. Since $1 \leq L \leq L(s) = \lfloor \kappa s \rfloor$, then

$$\frac{L}{s} \leq \kappa \Rightarrow 1 + \frac{L}{s} \leq 1 + \kappa \Rightarrow \left(1 + \frac{L}{s}\right)^{\rho} \leq (1+\kappa)^{\rho} \Rightarrow C_1 - C_2\left(1 + \frac{L}{s}\right)^{\rho} \geq C_1 - C_2(1+\kappa)^{\rho}.$$

Therefore, for all $1 \leq L \leq L(s)$,

$$R(s,L) = s^{\rho}\left[C_1 - C_2\left(1 + \frac{L}{s}\right)^{\rho}\right] \geq Cs^{\rho},$$

where $C = C_1 - C_2(1+\kappa)^{\rho} > 0$. Hence,

$$R(s,L) + L\mu \geq Cs^{\rho} + L\mu > 0,$$

implying

$$\frac{\left(R(s,L) + L\mu\right)^2}{L} \geq \frac{\left(Cs^{\rho} + L\mu\right)^2}{L} = \frac{C^2 s^{2\rho}}{L} + \mu^2 L + 2\mu C s^{\rho}.$$

Consider this as a function of continuous $L > 0$. Then, it is minimized when $L = \frac{Cs^{\rho}}{\mu}$. This yields:

$$\min_{1 \leq L \leq L(s)} \frac{\left(R(s,L) + L\mu\right)^2}{L} \geq 4\mu C s^{\rho}.$$

Thus,

$$P(s,L) \leq \mathbb{P}\left(\sum_{j=1}^{L} X_j^1 \geq R(s,L)\right) \leq \exp\left(-\frac{8\mu C s^\rho}{(M-m)^2}\right) = \exp\left(-\bar{C}s^\rho\right),$$

where $\bar{C} = \frac{8\mu C}{(M-m)^2} > 0$.

Case 2: $L > L(s)$**.** That is, $L > \lfloor \kappa s \rfloor$, so $s < \frac{L}{\kappa}$. We then have

$$C_2(s+L)^\rho < C_2\left(\frac{1}{\kappa}+1\right)^\rho L^\rho.$$

Thus,

$$\begin{aligned} R(s,L) + L\mu &= C_1 s^\rho - C_2(s+L)^\rho + L\mu \\ &> C_1 s^\rho - C_2\left(\frac{1}{\kappa}+1\right)^\rho L^\rho + L\mu \\ &= C_1 s^\rho + L\left(\mu - C_2\left(\frac{1}{\kappa}+1\right)^\rho L^{\rho-1}\right). \end{aligned}$$

Since $L^{\rho-1} \to 0$ as $L \to \infty$ for all $0 < \rho < 1$, then there exists some threshold L_0 such that for all $L \geq L_0$, $C_2\left(\frac{1}{\kappa}+1\right)^\rho L^{\rho-1} \leq \frac{\mu}{2}$. Therefore, for all $L \geq \max\{L_0, L(s)\}$,

$$R(s,L) + L\mu > C_1 s^\rho + \frac{\mu}{2}L > \frac{\mu}{2}L.$$

When s is sufficiently large, $L(s) = \lfloor \kappa s \rfloor \geq L_0$, and therefore,

$$R(s,L) + L\mu > \frac{\mu}{2}L \quad \text{for all } L > L(s).$$

From Hoeffding's inequality, we get that when s is sufficiently large,

$$\begin{aligned} P(s,L) &\leq \mathbb{P}\left(\sum_{j=1}^{L} X_j^1 + L\mu \geq R(s,L) + L\mu\right) \\ &\leq \mathbb{P}\left(\sum_{j=1}^{L} X_j^1 + L\mu \geq \frac{\mu}{2}L\right) \\ &\leq \exp\left(-\frac{2\left(\frac{\mu}{2}L\right)^2}{L(M-m)^2}\right) \\ &= \exp\left(-\tilde{C}L\right), \end{aligned}$$

where $\tilde{C} = \frac{\mu^2}{2(M-m)^2} > 0$. □

Let $P(s)$ denote the probability of the event that an upcrossing starts at s. Lemma 9 implies that when s is large,

$$\begin{aligned} P(s) &= \mathbb{P}\left(\cup_{L=1}^{\infty}\left\{\sum_{j=1}^{L} X_j^1 \geq C_1 s^\rho - C_2 (s+L)^\rho\right\}\right) \\ &\leq \sum_{L=1}^{L(s)} \mathbb{P}\left(\sum_{j=1}^{L} X_j^1 \geq R(s,L)\right) + \sum_{L>L(s)} \mathbb{P}\left(\sum_{j=1}^{L} X_j^1 \geq R(s,L)\right). \end{aligned}$$

For the first term:

$$\sum_{L=1}^{L(s)} \mathbb{P}\left(\sum_{j=1}^{L} X_j^1 \geq R(s,L)\right) \leq \sum_{L=1}^{L(s)} \exp\left(-\bar{C} s^\rho\right) \leq \kappa s \exp\left(-\bar{C} s^\rho\right).$$

For the second term:

$$\begin{aligned} \sum_{L>L(s)} \mathbb{P}\left(\sum_{j=1}^{L} X_j^1 \geq R(s,L)\right) &\leq \sum_{L>L(s)} \exp\left(-\tilde{C} L\right) \\ &= \exp\left(-\tilde{C}\left(L(s)+1\right)\right) \sum_{k=0}^{\infty} \exp\left(-\tilde{C} k\right) \\ &= \frac{\exp\left(-\tilde{C}\left(L(s)+1\right)\right)}{1-\exp\left(-\tilde{C}\right)}. \end{aligned}$$

Since $L(s)+1 > \kappa s$, then

$$\sum_{L>L(s)} \mathbb{P}\left(\sum_{j=1}^{L} X_j^1 \geq R(s,L)\right) \leq \frac{1}{1-e^{-\tilde{C}}} e^{-\tilde{C}\kappa s}.$$

Therefore,

$$\sum_{s=1}^{\infty} P(s) \leq \sum_{s=1}^{s^*} P(s) + \sum_{s=s^*+1}^{\infty} \kappa s \exp\left(-\bar{C} s^\rho\right) + \sum_{s=s^*+1}^{\infty} \frac{1}{1-e^{-\tilde{C}}} e^{-\tilde{C}\kappa s} < \infty,$$

where s^* is finite but large such that $L(s) \geq L_0$ for all $s > s^*$. So, the probability that an upcrossing from $1/\gamma$ to $1/\gamma + \varepsilon$ occurs infinitely often is zero by Borel-Cantelli lemma. In

other words,

$$\mathbb{P}\left(x_{i+k} < \frac{1}{\gamma} + \varepsilon \text{ for all } k \geq 0 | x_i \leq \frac{1}{\gamma}\right) \geq 1 - \sum_{s=i}^{\infty} P(s) \to 1 \text{ as } i \to +\infty,$$

so x_i remains below $1/\gamma + \varepsilon$ with a strictly positive probability when i is large.

Step 3: x_i converges to $\frac{1}{\gamma}$ almost surely.

Let $G_\varepsilon \equiv \left\{x_i < \frac{1}{\gamma} + \varepsilon \text{ for all } i \text{ sufficiently large}\right\}$. From the last step, we know that there exists $\sigma > 0$ such that $\mathbb{P}(G_\varepsilon | \mathcal{F}_i) > \sigma$. By Levy's 0-1 law, $\mathbb{P}(G_\varepsilon) = 1$. This implies that $\limsup_{i \to +\infty} x_i < \frac{1}{\gamma} + \varepsilon$ almost surely. On the other hand, we also have $\liminf x_i = 1/\gamma$, since: (i) the stepsize α_i decreases to zero, and (ii) x_i keeps increasing once it falls below $1/\gamma$, so the smallest x_i—which is obtained at the first entry of C_0—becomes arbitrarily close to $1/\gamma$ in the limit. In summary,

$$\frac{1}{\gamma} = \liminf_{i \to +\infty} x_i \leq \limsup_{i \to +\infty} x_i < \frac{1}{\gamma} + \varepsilon.$$

Since $\varepsilon > 0$ can be arbitrarily small, we have $x_i \to \frac{1}{\gamma}$ almost surely.

C.4 Proof of Propositions 2 and 4

Recall that

$$z_{i+1} = z_i + \alpha_{i+1} \left[h(z_i) + M_{i+1} + \varepsilon_i\right].$$

Since $\varepsilon_i = \left(\frac{1}{\alpha_i} - \frac{1}{\alpha_{i+1}}\right) z_i \approx -C\rho i^{\rho-1} z_i$, we can rewrite the process as:

$$z_{i+1} = z_i + \alpha_{i+1} \left[h(z_i) + M_{i+1}\right] + \alpha_{i+1}^{1+\eta} \xi_i, \tag{14}$$

where $\eta \in (0, 1/\rho - 1]$ and ξ_i is bounded. Define $e_i = z_i - \ln(1/\gamma)$. We have the following lemmas:

Lemma 10. *$|e_i| \times 1_{\{z_i \leq \ln(1/\gamma)\}} = O(\alpha_i)$ when i is large.*

Proof. Suppose $z_{i-1} > \ln(1/\gamma)$ but $z_i \leq \ln(1/\gamma)$. Then, $H(z_i, \lambda_i) = h(z_i) + M_{i+1} = 0$, and

$$z_{i+1} - z_i \sim -\rho/i.$$

Let's then consider the minimum value of z_i:

$$\begin{aligned} z_i &= \frac{\alpha_i}{\alpha_{i-1}} z_{i-1} + \alpha_i H\left(z_{i-1}, \lambda_{i-1}\right) \\ &\geq \left(1 - C_2 \frac{1}{i-1}\right) z_{i-1} - C_1 \alpha_i, \text{ where } C_{1,} C_2 > 0 \\ &\geq \left(1 - C_2 \frac{1}{i-1}\right) \ln\left(1/\gamma\right) - C_1 \alpha_i \\ &= \ln\left(1/\gamma\right) - \left(C_1 \alpha_i - C_2 \frac{1}{i-1} \ln\left(\gamma\right)\right) > \ln\left(1/\gamma\right) - C_1 \alpha_i, \end{aligned}$$

Suppose $z_{i+1}, ..., z_{i+k} \in C_0$, then

$$z_{i+k} = \frac{\alpha_{i+k}}{\alpha_i} z_i \geq \left(\frac{i}{i+k}\right)^{\rho} \left[\ln\left(1/\gamma\right) - C_1 \alpha_i\right].$$

Let $K\left(i\right)$ be the maximum steps for z_i to stay in C_0. It must satisfy

$$\left(\frac{i}{i+K\left(i\right)}\right)^{\rho} \left[\ln\left(1/\gamma\right) - C_1 \alpha_i\right] \leq \ln\left(1/\gamma\right).$$

Therefore

$$\left(\frac{i}{i+K\left(i\right)}\right)^{\rho} \geq \frac{\ln\left(\gamma\right)}{\ln\left(\gamma\right) + C_1 \alpha_i}.$$

This implies

$$\left(1 + \frac{K\left(i\right)}{i}\right)^{\rho} \leq C' \alpha_i + 1 \Rightarrow K\left(i\right) \leq \left(\left[C' \alpha_i + 1\right]^{1/\rho} - 1\right) i \Rightarrow K\left(i\right) = O\left(i^{1-\rho}\right),$$

where C' is a constant. For any $k \in \{0, 1, ..., K\left(i\right)\}$, we have $k/i = O\left(i^{-\rho}\right) = O\left(\alpha_i\right)$. Also note that for all $z_{i+k} \in C_0$,

$$e_{i+k} = z_{i+k} - \ln\left(1/\gamma\right) \geq \left(\frac{i}{i+k}\right)^{\rho} \left[\ln\left(1/\gamma\right) - C_1 \alpha_i\right] - \ln\left(1/\gamma\right),$$

which implies that

$$\begin{aligned} \left|e_{i+k}\right| &\leq \left|\left(1 - \rho \frac{k}{i} + o\left(\frac{k}{i}\right)\right) \left[\ln\left(1/\gamma\right) - C_1 \alpha_i\right] - \ln\left(1/\gamma\right)\right| \\ &\leq \left|\left[-\rho \frac{k}{i} + o\left(\frac{k}{i}\right)\right] \ln\left(1/\gamma\right)\right| + C_1 \alpha_i \\ &= C'' \alpha_i, \end{aligned}$$

where the last equality comes from $\frac{k}{i} = O\left(i^{-\rho}\right)$. So, we have $|e_{i+k}| = O\left(\alpha_i\right)$ for any k. That is, $|e_i| = O\left(\alpha_i\right)$ for all i such that $z_i \in C_0$. □

Lemma 11. *When signals are discrete, there exists $\delta > 0$ such that*

$$h\left(z\right) \leq -\delta\left(z - \ln\left(1/\gamma\right)\right) \tag{15}$$

for all $z < \ln\left(\gamma\right)$. When signals are continuous with $f^0\left(\gamma\right) > 0$, for all $\varepsilon > 0$, there exists $\delta > 0$ such that (15) holds for $z \leq \ln\left(\gamma\right) - \varepsilon$.

Proof. When signals are discrete, there exists some constant $c > 0$ such that $h\left(z\right) \leq -c$ for all $z \in \left(\ln\left(1/\gamma\right), \ln\left(\gamma\right)\right)$ and $h = 0$ if z is less than $\ln\left(1/\gamma\right)$. The lemma holds if we choose δ to be sufficiently small.

When signals are continuous, we have $h\left(z\right) < 0$ for all $z \in \left(\ln\left(1/\gamma\right), \ln\left(\gamma\right)\right)$, and $h\left(z\right)$ diminishes to zero as $z \to \ln\left(1/\gamma\right)$. Notice that

$$h'\left(\ln\left(1/\gamma\right)\right) = -\gamma f^0\left(\gamma\right)\left(\gamma - 1 - \ln\left(\gamma\right)\right) < 0,$$

because $f^1\left(\gamma\right)/f^0\left(\gamma\right) \equiv \gamma$. So, when $z \in \left(\ln\left(1/\gamma\right), \ln\left(1/\gamma\right) + \varepsilon_1\right)$ where ε_1 is small, we have:

$$h\left(z\right) = h'\left(\ln\left(1/\gamma\right)\right)\left(z - \ln\left(1/\gamma\right)\right) + o\left(z - \ln\left(1/\gamma\right)\right),$$

which satisfies (15) for some small $\delta > 0$. Since h is continuous, so the minimum of $\frac{h(z)}{z - \ln(1/\gamma)}$ is also strictly negative on $[\ln\left(1/\gamma\right) + \varepsilon_1, \ln\left(\gamma\right) - \varepsilon]$, so (15) holds for all $z \leq \ln\left(\gamma\right) - \varepsilon$. □

Lemma 12. *For any compact set $Q \subset \left[\ln\left(1/\gamma\right), \ln\left(\gamma\right)\right)$ and any $\epsilon > 0$, we have:*

$$\mathbb{E}\left(e_i^2 1_{\{z_i \in Q\}}\right) = O\left(\alpha_i^{1-\epsilon}\right)$$

when i is sufficiently large. When signals are discrete, we can let $Q = \left[\ln\left(1/\gamma\right), \ln\left(\gamma\right)\right)$.

Proof. From (14), we have:

$$\begin{aligned} e_{i+1}^2 = {} & e_i^2 + 2\alpha_{i+1} e_i h\left(z_i\right) + 2\alpha_{i+1} e_i M_{i+1} \\ & + 2e_i \alpha_{i+1}^{1+\eta} \xi_i + \alpha_{i+1}^2 \left[h\left(z_i\right) + M_{i+1} + \alpha_{i+1}^{\eta} \xi_i\right]^2. \end{aligned}$$

If $z_i \in Q$, then by Lemma 11,

$$e_{i+1}^2 \leq \left(1 - 2\delta\alpha_{i+1} + D_1 \alpha_{i+1}^2\right) e_i^2 + \left(2\alpha_{i+1} M_{i+1} + 2\alpha_{i+1}^{1+\eta} \xi_i\right) e_i + H\left(z_i, \lambda_i\right)^2 \alpha_{i+1}^2 + \alpha_{i+1}^{2+2\eta} \xi_i^2$$

where $\delta, D_1 > 0$. So,

$$\mathbb{E}\left[e_{i+1}^2 1_{\{z_{i+1}\in Q\}}\right] \leq \left(1 - 2\delta\alpha_{i+1} + D_1\alpha_{i+1}^2\right)\mathbb{E}\left[e_i^2 1_{\{z_i\in Q\}}\right] + D_2 \min\left\{\alpha_{i+1}^{1+\eta}, \alpha_{i+1}^2\right\},$$

where $D_2 > 0$.

To estimate the convergence rate, we follow the approach of Benveniste et al. (1990).[28] Suppose that if we can construct sequence $\{u_n\}$ with $u_n = \lambda\alpha_n^\beta$ that satisfies

$$u_{n+1} \geq \left(1 - 2\delta\alpha_{i+1} + D_1\alpha_{i+1}^2\right)u_n + D_2 \min\left\{\alpha_{i+1}^{1+\eta}, \alpha_{i+1}^2\right\} \tag{16}$$

for all large n. Then we can fix a large N such that $\mathbb{E}\left[e_N^2 1_{\{z_N\in Q\}}\right] \leq \lambda\alpha_N^\beta$. Then

$$\begin{aligned}\mathbb{E}\left[e_{N+1}^2 1_{\{z_{N+1}\in Q\}}\right] &\leq \left(1 - 2\delta\alpha_{N+1} + D_1\alpha_{N+1}^2\right)\mathbb{E}\left[e_N^2 1_{\{z_N\in Q\}}\right] + D_2 \min\left\{\alpha_{N+1}^{1+\eta}, \alpha_{N+1}^2\right\} \\ &\leq \lambda\alpha_{N+1}^\beta,\end{aligned}$$

and by induction, we have $\mathbb{E}\left[e_n^2 1_{\{z_n\in Q\}}\right] \leq \lambda\alpha_n^\beta$ for all $n \geq N$.

Now we show that such sequence $\{u_n\}$ exists. The sequence must satisfy the following condition:

$$\lambda\alpha_{n+1}^\beta \geq \left(1 - 2\delta\alpha_{n+1} + D_1\alpha_{n+1}^2\right)\lambda\alpha_n^\beta + D_2 \min\left\{\alpha_{n+1}^{1+\eta}, \alpha_{n+1}^2\right\},$$

which can be rewritten as

$$\lambda\left(\alpha_{n+1}^\beta - \alpha_n^\beta + \alpha_{n+1}\alpha_n^\beta\left(2\delta - D_1\alpha_{n+1}\right)\right) \geq D_2 \min\left\{\alpha_{n+1}^{1+\eta}, \alpha_{n+1}^2\right\}.$$

That is, we want to see if there exists β such that

$$\frac{\lambda}{D_2} \times \frac{\alpha_{n+1}^\beta - \alpha_n^\beta + \alpha_{n+1}\alpha_n^\beta\left(2\delta - D_1\alpha_{n+1}\right)}{\min\left\{\alpha_{n+1}^{1+\eta}, \alpha_{n+1}^2\right\}} \geq 1.$$

Suppose $\eta > 1$, then the left-hand side becomes

$$\frac{\lambda}{D_2} \times \frac{\alpha_{n+1}^\beta - \alpha_n^\beta + \alpha_{n+1}\alpha_n^\beta\left(2\delta - D_1\alpha_{n+1}\right)}{\alpha_{n+1}^{1+\eta}.} = \frac{\lambda}{D_2} \times \left[\frac{\alpha_{n+1}^\beta - \alpha_n^\beta}{\alpha_{n+1}^{1+\eta}} + \frac{\alpha_n^\beta\left(2\delta - D_1\alpha_{n+1}\right)}{\alpha_{n+1}^\eta}\right] \to +\infty$$

[28] In Part II, Section 1.10, the authors employ a similar construction of a sequence u_n in the proofs of Theorems 22 and 24.

for any $\beta \leq 1$, so the sequence exists. Similarly, suppose $\eta \leq 1$, then $\min\left\{\alpha_{n+1}^{1+\eta}, \alpha_{n+1}^{2}\right\} = \alpha_{n+1}^{2}$, and

$$\frac{\lambda}{D_2} \times \frac{\alpha_{n+1}^{\beta} - \alpha_n^{\beta} + \alpha_{n+1}\alpha_n^{\beta}\left(2\delta - D_1\alpha_{n+1}\right)}{\alpha_{n+1}^{2}} \to +\infty$$

for any $\beta < 1$. In summary, when n is large, we have:

$$\mathbb{E}\left[e_n^2 1_{\{z_n \in Q\}}\right] = O\left(\alpha_n^{\beta}\right)$$

for any $\beta < 1$. □

Propositions 2 and 4 then follow directly from Lemmas 10 and 12.

C.5 Proof of Proposition 3

Proof. From Theorem 1, $x_i \to 1/\gamma$ almost surely. When x_i stays below $1/\gamma$, $\mathbb{P}\left(a_i = 0\right) = 1$. When x_i stays above $1/\gamma$ and sufficiently close to it, $\mathbb{P}\left(a_i = 0\right) = F^0\left(\gamma^-\right)$ since only the individual who receives the strongest opposite signal would choose action 1. Therefore, $\liminf_i \mathbb{P}\left(a_i = 0\right) \geq F^0\left(\gamma^-\right)$, thus

$$\liminf_i \mathbb{E}\vartheta_i = \liminf_i \frac{1}{i}\sum_{j=1} \mathbb{P}\left(a_i = 0\right) \geq \liminf_i \mathbb{P}\left(a_i = 0\right) \geq F^0\left(\gamma^-\right),$$

which establishes the proposition. □

C.6 Proof of Proposition 5

Case 1: $\rho < 0$**.**

Since $\alpha_i \to +\infty$, if x_∞ is a limit point of $\{x_i\}$, we must have $\frac{\phi^1(a_\infty, x_\infty)}{\phi^0(a_\infty, x_\infty)} = 1$, i.e., $x_\infty \in C$.[29] Besides, we can show that both $x_\infty \in C_0$ and $x_\infty \in C_1$ happen with positive probability. To see that, consider the case where $a_1 = a_2 = ... = 1$. Then

$$\ln\left(x_i\right) = \alpha_i \times \sum_{j=1}^{i-1} \ln\left[\frac{\phi^1\left(a_j, x_j\right)}{\phi^0\left(a_j, x_j\right)}\right] \geq \alpha_i \times \ln\left[\frac{\phi^1\left(a_1, x_1\right)}{\phi^0\left(a_1, x_1\right)}\right] = \alpha_i \times \ln\left[\frac{1 - F^1\left(1\right)}{1 - F^0\left(1\right)}\right],$$

where the inequality comes from that $\frac{\phi^1(a_1,x_1)}{\phi^0(a_1,x_1)} \geq 1$ when $a_j = 1$. Since $\alpha_i \to +\infty$ and $\frac{1-F^1(1)}{1-F^0(1)} > 1$, there exists some $I < +\infty$ such that $x_I \geq \gamma$, i.e., an action-1 cascade occurs

[29] Suppose otherwise, then there exists $\varepsilon > 0$ such that $x_i \in (1/\gamma + \varepsilon, \gamma - \varepsilon)$ infinitely often. Note that the absolute value of $X_i = \ln\left[\frac{\phi^1(a_i,x_i)}{\phi^0(a_i,x_i)}\right]$ is bounded away from zero. When α_{i+1} is very large, x_{i+1} will enter the interior of the cascade set and never go out, contradicting $x_\infty \notin C$.

almost surely. Here, we can choose I adequately large such that α_i is increasing after I, so x_i is also increasing, and hence it stays in C^1 forever. The probability of I individuals taking action 1 is positive, meaning that $x_\infty \in C^1$ happens with positive probability. The argument for $x_\infty \in C^0$ is symmetric.

Case 2: $\rho = 0$**.**

We rewrite the dynamics of x_i as follows:

$$\ln(x_{i+1}) = \ln(x_i) + \alpha_{i+1} \ln\left[\frac{\phi^1(a_i, x_i)}{\phi^0(a_i, x_i)}\right] + \xi_{i+1}, \quad \text{where } \xi_{i+1} = \frac{\alpha_{i+1} - \alpha_i}{\alpha_i} \ln(x_i).$$

Define $\sigma_i \equiv \sum_{j=2}^{i} \xi_j$ and $w_i \equiv \ln(x_i) - \sigma_i$, then $\{w_i\}$ satisfies the following:

$$w_{i+1} = w_i + \alpha_{i+1} \times \ln\left[\frac{\phi^1(a_i, x_i)}{\phi^0(a_i, x_i)}\right],$$

since $\mathbb{E}\left[\ln\left[\frac{\phi^1(a_i,x_i)}{\phi^0(a_i,x_i)}\right] | \mathcal{F}_i\right] \leq 0$, $\{w_i\}$ is a supermartingale. When $\rho = 0$, $\alpha_i \to c \in (0, +\infty)$, so x_i is bounded by some constant $M_0 < +\infty$. Therefore,

$$|\sigma_i| = \left|\sum_{j=2}^{i} \frac{\alpha_j - \alpha_{j-1}}{\alpha_{j-1}} \ln(x_{j-1})\right| \leq M_0 \times \sum_{i=2}^{\infty} \left|\frac{\alpha_i - \alpha_{i-1}}{\alpha_{i-1}}\right|.$$

Since $|\alpha_i - c|$ has the same order as i^{-q}, we have $\left|\frac{\alpha_i - \alpha_{i-1}}{\alpha_{i-1}}\right|$ has the same order as i^{-q-1} when i is large. Therefore, $\sum_{i=2}^{\infty} \left|\frac{\alpha_i - \alpha_{i-1}}{\alpha_{i-1}}\right|$ is finite, so σ_i is also bounded. Since $\{w_i\}$ is a bounded supermartingale, it converges almost surely. Note that $\{\sigma_i\}$ is absolute convergent (hence convergent), so x_i also converges almost surely. Since $\alpha_i \to c > 0$, the limit point x_∞ can only be in the cascade set, so a limit cascade occurs almost surely.

The supermartingale property implies that $\mathbb{E}\ln(x_\infty) \leq \ln(x_1) = 0$, so $x_\infty \in C_0$ with positive probability. Next we show that $x_\infty \in C_1$ also with positive probability. Choose a constant ψ such that $\psi\alpha_i > 1$ for all i. Notice that $\{x_i\}$ satisfies:

$$x_{i+1}^{\psi} = x_i^{\psi} \times \left[\frac{\phi^1(a_i, x_i)}{\phi^0(a_i, x_i)}\right]^{\psi\alpha_{i+1}} \times \exp(\psi\xi_{i+1}).$$

Define $y_i = x_i^{\psi} / \exp(\psi\sigma_i)$, so $\{y_i\}$ satisfies

$$y_{i+1} = y_i \times \left[\frac{\phi^1(a_i, x_i)}{\phi^0(a_i, x_i)}\right]^{\psi\alpha_{i+1}}.$$

Since $\psi\alpha_{i+1} > 1$, Jensen's inequality implies that

$$\mathbb{E}\left(\left[\frac{\phi^1(a_i,x_i)}{\phi^0(a_i,x_i)}\right]^{\psi\alpha_{i+1}}|\mathcal{F}_i\right) \geq \left(\mathbb{E}\left(\frac{\phi^1(a_i,x_i)}{\phi^0(a_i,x_i)}|\mathcal{F}_i\right)\right)^{\psi\alpha_{i+1}} = 1,$$

so $\{y_i\}$ is a submartingale. Since $\exp(\psi\sigma_i)$ is bounded, y_i also converges almost surely, and we must have $\mathbb{E}y_\infty \geq y_1 = 1$. This implies that $x_\infty \in C_1$ with positive probability.

C.7 Proof of Corollary 2

Lemma 13. *There exists $K < \infty$ and $\delta > 0$ such that $\mathbb{P}(x_{i+K} \in C | x_i \notin C) \geq \delta$ for all i.*

Proof. Suppose $x_i \notin C$, then the strongest signal γ (or $1/\gamma$) will induce individual i to choose action 1 (or 0). Given $x_i \geq 1$ and that $a_i = ... = a_{i+k-1} = 1$, we have

$$\begin{aligned} x_{i+k} &= x_i^{\frac{\alpha_{i+k}}{\alpha_i}} \times \left[\prod_{j=i}^{i+k-1}\frac{\phi^1(1,x_j)}{\phi^0(1,x_j)}\right]^{\alpha_{i+k}} \geq 1 \times \left[\prod_{j=i}^{i+k-1}\frac{\phi^1(1,x_j)}{\phi^0(1,x_j)}\right]^{\underline{\alpha}} \\ &\geq \left(\frac{1-F^1(1/\gamma)}{1-F^0(1/\gamma)}\right)^{k\underline{\alpha}} \text{ if } x_i, ..., x_{i+k-1} \notin C, \end{aligned}$$

where $\underline{\alpha} = \inf_i \alpha_i > 0$, and the first inequality comes from $x_i \geq 1$, and the second inequality comes from

$$\frac{\phi^1(1,x_j)}{\phi^0(1,x_j)} = \frac{1-F^1(1/x_j)}{1-F^0(1/x_j)}$$

is decreasing in x_j (see Lemma A.1 in Smith and Sørensen (2000)), so we have

$$\frac{\phi^1(1,x_j)}{\phi^0(1,x_j)} \geq \frac{1-F^1(1/\gamma)}{1-F^0(1/\gamma)} \equiv \Omega_1 > 1.$$

Therefore, after at most $K_1 = \lceil \frac{1}{\underline{\alpha}} \log_{\Omega_1}^{\gamma} \rceil$ steps, we have $x_{i+K_1} \geq \gamma$, so a cascade of action 1 is triggered. Symmetrically, given $x_i \leq 1$ and that $a_i = ... = a_{i+k-1} = 0$, we have

$$\begin{aligned} x_{i+k} &= x_i^{\frac{\alpha_{i+k}}{\alpha_i}} \times \left[\prod_{j=i}^{i+k-1}\frac{\phi^1(0,x_j)}{\phi^0(0,x_j)}\right]^{\alpha_{i+k}} \\ &\leq 1 \times \left[\prod_{j=i}^{i+k-1}\frac{\phi^1(0,x_j)}{\phi^0(0,x_j)}\right]^{\underline{\alpha}} \leq \Omega_0^{k\underline{\alpha}} \text{ if } x_i, ..., x_{i+k-1} \notin C, \end{aligned}$$

where $\Omega_0 = \frac{F^1(1)}{F^0(1)} < 1$. Then, after at most $K_0 = \lceil \frac{1}{\underline{\alpha}} \log_{\Omega_0}^{1/\gamma} \rceil$ steps, we have $x_{i+K_0} \leq 1/\gamma$, so a cascade of action 0 is triggered. Note than when $x_i \notin C$, a sequence of consecutive

$K = \max\{K_1, K_2\}$ signals γ (or $1/\gamma$) will trigger a sequence of action 1 (or 0). Let the probability of realizing K signals being δ, then we prove the lemma. □

Lemma 14. *If α_i is increasing, an information cascade occurs and persists forever $\mathbb{P}$-almost surely, and both correct and incorrect cascades occur with strictly positive probability.*

Proof. Suppose α_i is increasing. When $x_i \in C$, then

$$x_{i+1} = x_i^{\frac{\alpha_{i+1}}{\alpha_i}} \times \left[\frac{\phi^1(a_i, x_i)}{\phi^0(a_i, x_i)}\right]^{\alpha_{i+1}} = x_i^{\frac{\alpha_{i+1}}{\alpha_i}} \begin{cases} \geq x_i & x_i \in C_1 \\ \leq x_i & x_i \in C_0 \end{cases},$$

so $x_{i+1} \in C$ as well, which implies that the public LLR will remain in the cascade set. Lemma 13 implies that the probability of entering a cascade set is uniformly bounded away from 0. Standard Borel-Cantelli arguments show that an information cascade occurs almost surely. Finitely many signals can trigger a cascade, so both cascades will emerge with positive probability. □

Lemma 15. *If $\rho = 0$ and α_i is decreasing, the result in Lemma 14 still holds.*

Proof. The nontrivial part is that with decreasing α_i, we can no longer guarantee that $x_i \in C$ implies $x_{i+1} \in C$. However, we notice that

$$x_{i+1} - x_i = x_i^{\frac{\alpha_{i+1}}{\alpha_i}} - x_i \to 0 \text{ as } i \to +\infty.$$

Therefore, when i is adequately large, even if x_{i+1} falls outside of the cascade set, it is very close to the boundary, γ or $1/\gamma$. Without loss, suppose $x_{i+1} > \gamma - \varepsilon$, where $\varepsilon > 0$ can be arbitrarily small. Suppose that $\lambda_{i+1} = \gamma$. When i is sufficiently large, we have:

$$x_{i+2} = x_{i+1}^{\frac{\alpha_{i+2}}{\alpha_{i+1}}} \times \left[\frac{1 - F^1(1/\gamma)}{1 - F^0(1/\gamma)}\right]^{\alpha_{i+2}} \approx x_{i+1} \times \left[\frac{1 - F^1(1/\gamma)}{1 - F^0(1/\gamma)}\right]^{\underline{\alpha}} > (\gamma - \varepsilon) \times \Omega_1^{\underline{\alpha}}.$$

Therefore, we have:

$$x_{i+2} - \gamma > (\Omega_1^{\underline{\alpha}} - 1)\gamma - \varepsilon\Omega_1^{\underline{\alpha}} \equiv \Delta.$$

Note that $\Omega_1^{\underline{\alpha}} > 1$, so we can choose $\varepsilon > 0$ such that $\Delta > 0$, then $x_{i+2} > \gamma + \Delta$, and

$$x_{i+3} = x_{i+2}^{\frac{\alpha_{i+3}}{\alpha_{i+2}}} > (\gamma + \Delta)^{\frac{\alpha_{i+3}}{\alpha_{i+2}}}.$$

Since $\alpha_i \to \underline{\alpha} > 0$, for all $\xi > 0$, there exists $I < \infty$ such that $|\alpha_m/\alpha_n - 1| < \xi$ whenever $m, n > I$. We choose $\xi = 1 - \log_{\gamma+\Delta}^{\gamma}$, then when i is sufficiently large,

$$x_{i+3} > (\gamma + \Delta)^{\frac{\alpha_{i+3}}{\alpha_{i+2}}} > (\gamma + \Delta)^{1-\xi} = \gamma,$$

so $x_{i+3} \in C_1$. Suppose $x_{i+3}, ..., x_{i+k-1} \in C_1$, then

$$x_{i+k} = x_{i+2}^{\frac{\alpha_{i+k}}{\alpha_{i+2}}} > (\gamma + \Delta)^{\frac{\alpha_{i+k}}{\alpha_{i+2}}} > (\gamma + \Delta)^{1-\xi} = \gamma,$$

so x_i stays inside the cascade set forever. In summary, when i is sufficiently large, even if x_{i+1} falls outside of C, another signal can bring it back to the cascade set and make it stay there forever. Formally, we define ϵ-*cascade sets* as follows

$$C_0^\epsilon \equiv [0, 1/\gamma + \epsilon], \quad C_1^\epsilon \equiv [\gamma + \epsilon, +\infty],$$

and $C^\epsilon \equiv C_0^\epsilon \cup C_1^\epsilon$. From the previous discussion, there exists some $\epsilon > 0$ such that when i is sufficiently large, one of the following three cases holds:

1. If $x_i \in C^\epsilon$, then x_{i+k} remains in the cascade set C for all $k \geq 0$;
2. If $x_i \in C$, but $x_{i+1} \notin C$, then a single signal will bring x_{i+2} to C^ϵ;
3. If $x_i \notin C$, finitely many signals will bring x_i to C, and we have either case 1 or 2.

In summary, there exists $K < \infty$ such that when i is sufficiently large, the probability of x_{i+K} entering C and staying forever is bounded away from 0. Standard 0-1 arguments yield the result in Lemma 14. □

C.8 Proof of Proposition 6

Proof. Suppose $\alpha_i = \alpha$ for all i. Then, the public likelihood ratio $l_i = x_i^{1/\alpha}$ is a martingale, so it almost surely converges to a finite random variable l_∞. Therefore, x_i also converges to some limit x_∞ almost surely, and we have

$$\mathbb{E}(x_\infty^{1/\alpha}) = \mathbb{E}(l_\infty) = \lim_{i \to +\infty} \mathbb{E}(l_i) = 1,$$

where the interchange of limit and expectation comes from the boundedness of l_i. Note that we must have $x_\infty \in C_0$ or $x_\infty \in C_1$ in the limit. Therefore,

$$1 = \int_{\{x_\infty \in C_0\}} x_\infty^{1/\alpha} d\mathbb{P} + \int_{\{x_\infty \in C_1\}} x_\infty^{1/\alpha} d\mathbb{P} \geq 0 + \mathbb{P}(x_\infty \in C_1) \times \gamma^{1/\alpha},$$

so $\mathbb{P}(x_\infty \in C_1) \leq \gamma^{-1/\alpha}$. From Proposition 5, an information cascade almost surely occurs when $\rho = 0$, so

$$p(\alpha) = \mathbb{P}(x_\infty \in C_0) \geq 1 - \gamma^{-1/\alpha}.$$

That is, the probability of a correct cascade is greater than $1 - \gamma^{-1/\alpha}$. As $\alpha \to 0$, we have $1 - \gamma^{-1/\alpha} \to 1$, thus $p(\alpha) \to 1$. □

C.9 Proof of Proposition 7

Proof. For any i, we have

$$x_{i+1} = x_1^{\frac{\alpha_{i+1}}{\alpha_1}} \left[\prod_{k=1}^{i} \frac{\phi^1(a_k, x_k)}{\phi^0(a_k, x_k)}\right]^{\alpha_{i+1}} = \left[\prod_{k=1}^{i} \frac{\phi^1(a_k, x_k)}{\phi^0(a_k, x_k)}\right]^{\alpha_{i+1}},$$

where the equality comes from the flat prior assumption, so $x_1 = 1$.[30] Since $\lambda_i \in [1/\gamma, \gamma]$, and

$$\phi^\theta(a, x) = \begin{cases} 1 - F^\theta\left(\frac{1}{x}\right), & a = 1 \\ F^\theta\left(\frac{1}{x}\right), & a = 0 \end{cases},$$

so we have $\frac{\phi^1(a_i, x_i)}{\phi^0(a_i, x_i)} \in \left[\frac{1}{\gamma}, \gamma\right]$ for all i, and hence

$$x_{i+1} = \left[\prod_{k=1}^{i} \frac{\phi^1(a_k, x_k)}{\phi^0(a_k, x_k)}\right]^{\alpha_{i+1}} \in \left[\frac{1}{\gamma^{i\alpha_{i+1}}}, \gamma^{i\alpha_{i+1}}\right].$$

Since $\rho > 1$, we have $\lim_{i \to +\infty} i\alpha_{i+1} = 0$, which implies that $x_i \to 1$ as $i \to +\infty$. Thus, in the limit, all individuals will only act according to their own signal, i.e., they choose action 1 whenever $\lambda_i > 1$ and action 0 otherwise, so there is no information aggregation. □

D A Maximum Inequality

In this section, we present a maximum inequality that we used in the proof of Theorem 1:

Lemma 16. *Suppose that $X_1, X_2, \ldots$ are bounded i.i.d. random variables. Then,*

$$\forall \varepsilon > 0: \quad \mathbb{P}\left(\overline{X}_n - \mathbb{E}(X) \leq \varepsilon \text{ for all } n \geq 1\right) > 0,$$

[30] This is non-essential. For general x_1, we have $x_1^{\frac{\alpha_{i+1}}{\alpha_1}} \approx 1$ when i is large enough, and the same argument applies.

where $\overline{X}_n$ denotes the sample mean of random variables $X_1, ..., X_n$.

Proof. Fix any $\varepsilon > 0$. From Hoeffding's inequality, we have:

$$\mathbb{P}\left(\overline{X}_n - \mathbb{E}(X) > \varepsilon\right) \leq \exp(-Cn) \text{ for some } C > 0.$$

Therefore,

$$\sum_n \mathbb{P}\left(\overline{X}_n - \mathbb{E}(X) > \varepsilon\right) \leq \sum_n \exp(-Cn) < \infty,$$

which implies $\mathbb{P}\left(\overline{X}_n - \mathbb{E}(X) > \varepsilon \text{ i.o.}\right) = 0$ from the Borel–Cantelli lemma. Therefore,

$$\mathbb{P}\left(\cup_{k=1}^{\infty}\left\{\overline{X}_n - \mathbb{E}(X) > \varepsilon \text{ for all } n \geq k\right\}\right) = 1.$$

The continuity from below of the probability measure implies that there exists some k such that

$$\mathbb{P}\left(\overline{X}_n - \mathbb{E}(X) \leq \varepsilon \text{ for all } n \geq k\right) > 0, \tag{17}$$

and we choose k to be the smallest number that satisfies the above inequality.

Suppose $k = 1$, Lemma 16 holds automatically. Next we want to show that k cannot be greater than 1. We proceed by contradiction. Suppose that $k > 1$, and we define:

$$E = \{\boldsymbol{x} \in [a,b]^{\infty} : \overline{x}_n - \mathbb{E}(X) \leq \varepsilon \text{ for all } n \geq k\} \subset \mathbb{R}^{\infty},$$

where $\overline{x}_n = \frac{1}{n}\sum_{i=1}^{n} x_i$ and $[a,b]$ is the range of X. For any sequence $\boldsymbol{x} = (x_1, x_2, ...) \in \mathbb{R}^{\infty}$, we define $\text{Proj}_{k+}(x) \equiv (x_k, x_{k+1}, ...)$, which is the subsequence of x starting from the k-th term. We further define:

$$H = \left\{\text{Proj}_{k+}(\boldsymbol{x}) : \boldsymbol{x} \in E\right\}$$

and

$$G = \left\{(x_1, .., x_{k-1}) \in [a,b]^{k-1} : \overline{x}_n - \mathbb{E}(X) \leq \varepsilon \text{ for all } n = 1, 2, ..., k-1\right\}.$$

For any $g = (g_1, ..., g_{k-1}) \in G$ and $h = (h_1, h_2, ...) \in H$, we construct a new sequence:

$$\boldsymbol{x} = (g_1, g_2, ..., g_{k-1}, h_1, h_2, ...) \in G \times H.$$

By construction, whenever $n \leq k-1$, we have:

$$\overline{x}_n - \mathbb{E}(X) = \overline{g}_n - \mathbb{E}(X) \leq \varepsilon.$$

Suppose that $n \geq k$. Then, by the definition of H, h is associated with some $(\hat{x}_1, ..\hat{x}_{k-1}) \in$

$[a,b]^{k-1}$ such that the combined sequence $\hat{\boldsymbol{x}} = (\hat{x}_1, ..\hat{x}_{k-1}, h)$ satisfies: (i) $\overline{\hat{x}_n} - \mathbb{E}(X) \leq \varepsilon$ for all $n \geq k$, and (ii) $\overline{\hat{x}}_{k-1} - \mathbb{E}(X) > \varepsilon$. Therefore, we have:

$$\begin{aligned}\overline{x}_n - \mathbb{E}(X) &= \frac{k-1}{n}\overline{x}_{k-1} + \frac{1}{n}\sum_{i=k}^{n} x_i \\ &< \frac{k-1}{n}\overline{\hat{x}}_{k-1} + \frac{1}{n}\sum_{i=k}^{n} x_i = \overline{\hat{x}_n} - \mathbb{E}(X) \leq \varepsilon,\end{aligned}$$

where the first inequality comes from the fact that $\overline{x}_{k-1} - \mathbb{E}(X) \leq \varepsilon$, and hence $\overline{x}_{k-1} < \overline{\hat{x}}_{k-1}$. To sum up, whenever $\boldsymbol{x} \in G \times H$, we must have $\overline{x}_n - \mathbb{E}(X) \leq \varepsilon$ for all $n \geq 1$. From the independence assumption,

$$\begin{aligned}\mathbb{P}(X_1, X_2, ... \in G \times H) &= \mathbb{P}(X_1, X_2, ...X_{k-1} \in G) \times \mathbb{P}(X_k, X_{k+1}, ... \in H) \\ &\geq \mathbb{P}(X_1, X_2, ...X_{k-1} \in G) \times \mathbb{P}(X_1, X_2, ... \in E).\end{aligned}$$

Note that G is a non-empty set, and it is finitely dimensional, so $\mathbb{P}(X_1, X_2, ...X_{k-1} \in G) > 0$. From (17), we know that $\mathbb{P}(X_1, X_2, ... \in E) > 0$. Therefore, $\mathbb{P}(X_1, X_2, ... \in G \times H) > 0$, implying that $\overline{x}_n - \mathbb{E}(X) \leq \varepsilon$ for all $n \geq 1$ is a positive probability event, contradicting the presumption that $k > 1$. So, the lemma is proved.[31] □

[31] Intuitively, since later sample averages depend on early observations only through their sum, lowering early realizations can only decrease all subsequent averages, which allows a tail property (i.e., the sample average remaining below a given threshold) to be extended to the whole sequence.

Bibliography

Anderson, Lisa R and Charles A Holt (1997) "Information cascades in the laboratory," *American Economic Review*, 847–862. 6

Arieli, Itai, Yakov Babichenko, and Manuel Mueller-Frank (2021) "Sequential naive learning," *arXiv preprint arXiv:2101.02897*. 4

Arieli, Itai, Yakov Babichenko, Stephan Müller, Farzad Pourbabaee, and Omer Tamuz (2025) "The Hazards and Benefits of Condescension in Social Learning," *Theoretical Economics*, 20 (1). 4

Bala, Venkatesh and Sanjeev Goyal (1998) "Learning from neighbours," *The review of economic studies*, 65 (3), 595–621. 4

Banerjee, Abhijit V (1992) "A simple model of herd behavior," *Quarterly Journal of Economics*, 107 (3), 797–817. 2, 3

Benjamin, Daniel J (2019) "Errors in probabilistic reasoning and judgment biases," *Handbook of Behavioral Economics: Applications and Foundations 1*, 2, 69–186. 6

Benveniste, Albert, Michel Métivier, and Pierre Priouret (1990) *Adaptive Algorithms and Stochastic Approximations*: Springer-Verlag. 18, 26, 27, 30, 44

Bernardo, Antonio E and Ivo Welch (2001) "On the evolution of overconfidence and entrepreneurs," *Journal of Economics & Management Strategy*, 10 (3), 301–330. 4

Bikhchandani, Sushil, David Hirshleifer, Omer Tamuz, and Ivo Welch (2024) "Information cascades and social learning," *Journal of Economic Literature*, 62 (3), 1040–1093. 4

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992) "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, 100 (5), 992–1026. 2, 3, 19

Bissiri, Pier Giovanni, Chris C Holmes, and Stephen G Walker (2016) "A general framework for updating belief distributions," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 78 (5), 1103–1130. 6

Borkar, Vivek S (2023) *Stochastic approximation: a dynamical systems viewpoint*: Springer. 26, 27

Chamley, Christophe (2004) *Rational herds: Economic models of social learning*: Cambridge University Press. 2

Chen, Jaden Yang (2019) "Sequential learning under informational ambiguity," *Available at SSRN 3480231*. 4

Cheng, Yu, Wade Hann-Caruthers, and Omer Tamuz (2018) "A deterministic protocol for sequential asymptotic learning," in *2018 IEEE International Symposium on Information Theory (ISIT)*, 1735–1738, IEEE. 3

Conlon, John J, Malavika Mani, Gautam Rao, Matthew W Ridley, and Frank Schilbach (2022) "Not learning from others,"Technical report, National Bureau of Economic Research. 6

Dasaratha, Krishna and Kevin He (2020) "Network structure and naive sequential learning," *Theoretical Economics*, 15 (2), 415–444. 4

DeGroot, Morris H (1974) "Reaching a consensus," *Journal of the American Statistical association*, 69 (345), 118–121. 4

DeMarzo, Peter M, Dimitri Vayanos, and Jeffrey Zwiebel (2003) "Persuasion bias, social influence, and unidimensional opinions," *Quarterly Journal of Economics*, 118 (3), 909–968. 4

Duffy, John, Ed Hopkins, and Tatiana Kornienko (2021) "Lone wolf or herd animal? Information choice and learning from others," *European Economic Review*, 134, 103690. 6

Easley, David and Jon Kleinberg (2010) *Networks, crowds, and markets: Reasoning about a highly connected world*: Cambridge university press. 2

Ellison, Glenn and Drew Fudenberg (1993) "Rules of thumb for social learning," *Journal of political Economy*, 101 (4), 612–643. 4

——— (1995) "Word-of-mouth communication and social learning," *The Quarterly Journal of Economics*, 110 (1), 93–125. 4

Epstein, Larry G, Jawwad Noor, and Alvaro Sandroni (2010) "Non-bayesian learning," *The BE Journal of Theoretical Economics*, 10 (1), 0000102202193517041623. 4

Eyster, Erik and Matthew Rabin (2010) "Naive herding in rich-information settings," *American Economic Journal: Microeconomics*, 2 (4), 221–43. 4

——— (2014) "Extensive imitation is irrational and harmful," *The Quarterly Journal of Economics*, 129 (4), 1861–1898. 4

Frick, Mira, Ryota Iijima, and Yuhta Ishii (2024) "Welfare comparisons for biased learning," *American Economic Review*, 114 (6), 1612–1649. 4

Goeree, Jacob K, Thomas R Palfrey, and Brian W Rogers (2006) "Social learning with private and common values," *Economic theory*, 28, 245–264. 4

Goldfarb, Avi and Catherine Tucker (2012) "Privacy and innovation," *Innovation policy and the economy*, 12 (1), 65–90. 24

Golub, Benjamin and Matthew O Jackson (2010) "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, 2 (1), 112–49. 4

Grether, David M (1980) "Bayes rule as a descriptive model: The representativeness heuristic," *The Quaterly Journal of Economics*, 95 (3), 537–557. 6

——— (1992) "Testing Bayes rule and the representativeness heuristic: Some experimental evidence," *Journal of Economic Behavior & Organization*, 17 (1), 31–57. 6

Guarino, Antonio and Philippe Jehiel (2013) "Social learning with coarse inference," *American Economic Journal: Microeconomics*, 5 (1), 147–74. 4

Herrera, Helios and Johannes Hörner (2012) "A necessary and sufficient condition for information cascades." 10, 22

Hu, Yifan, Yehuda Koren, and Chris Volinsky (2008) "Collaborative filtering for implicit feedback datasets," in *2008 Eighth IEEE international conference on data mining*, 263–272, Ieee. 24

Ibrahim, Joseph G and Ming-Hui Chen (2000) "Power prior distributions for regression models," *Statistical Science*, 46–60. 6

Jadbabaie, Ali, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi (2012) "Non-Bayesian social learning," *Games and Economic Behavior*, 76 (1), 210–225. 4

Krügel, Sebastian, Andreas Ostermaier, and Matthias Uhl (2022) "Zombies in the loop? Humans trust untrustworthy AI-advisors for ethical decisions," *Philosophy & Technology*, 35 (1), 17. 24

Kushner, Harold Joseph and Dean S Clark (1978) *Stochastic approximation methods for constrained and unconstrained systems*: New York: Springer. 14, 17

Ljung, Lennart (1977) "Analysis of recursive stochastic algorithms," *IEEE transactions on automatic control*, AC-22 (4), 233–257. 18

Molavi, Pooya, Alireza Tahbaz-Salehi, and Ali Jadbabaie (2018) "A theory of non-Bayesian social learning," *Econometrica*, 86 (2), 445–490. 4, 5

Morin, Olivier, Pierre Olivier Jacquet, Krist Vaesen, and Alberto Acerbi (2021) "Social information use and social information waste," *Philosophical Transactions of the Royal Society B*, 376 (1828), 20200052. 6

Pemantle, Robin (1990) "Nonconvergence to unstable points in urn models and stochastic approximations," *The Annals of Probability*, 698–712. 15

Peres, Yuval, Miklós Z Rácz, Allan Sly, and Izabella Stuhl (2020) "How fragile are information cascades?" *The Annals of Applied Probability*, 30 (6), 2796–2814. 3

Robbins, Herbert and Sutton Monro (1951) "A stochastic approximation method," *The annals of mathematical statistics*, 400–407. 17, 26

Rosenberg, Dinah and Nicolas Vieille (2019) "On the efficiency of social learning," *Econometrica*, 87 (6), 2141–2168. 4

Smith, Lones and Peter Sørensen (2000) "Pathological outcomes of observational learning," *Econometrica*, 68 (2), 371–398. 4, 5, 22, 47

Smith, Lones, Peter Norman Sørensen, and Jianrong Tian (2021) "Informational herding, optimal experimentation, and contrarianism," *Review of Economic Studies*, 88 (5), 2527–2554. 10

Suresh, Harini, Natalie Lao, and Ilaria Liccardi (2020) "Misplaced trust: Measuring the interference of machine learning in human decision-making," in *Proceedings of the 12th ACM conference on web science*, 315–324. 24

Weizsäcker, Georg (2010) "Do we follow others when we should? A simple test of rational expectations," *American Economic Review*, 100 (5), 2340–2360. 6