ESSAYS ON LEARNING UNDER MODEL UNCERTAINTY

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Yang Chen

May 2022

ESSAYS ON LEARNING UNDER MODEL UNCERTAINTY

Yang Chen, Ph.D.

Cornell University 2022

The thesis consists of three essays that focus on learning under model uncertainty or relevant topics. They jointly investigate the problem of how individuals learn and make decisions when they do not perfectly understand how to interpret the information they have access to.

The first essay, "Sequential Learning under Informational Ambiguity", introduces model uncertainty into the classic sequential social learning problem. One important phenomenon in the sequential social learning is information cascades. Past research has shown that the occurrence of a cascade depends on the details of people's data-generating processes, so it leaves an open question of whether cascades are prevalent in the social learning. In the essay, I re-examine the problem under the assumption that individuals are ambiguous about others' data-generating process and make decisions according to the max-min criterion. The main result of this research is that, under sufficient ambiguity, an information cascade occurs almost surely for *all* possible data-generating processes. More surprisingly, in many interesting situations, an arbitrarily small amount of ambiguity suffices to generate the results. This suggests that, relative to the presence of ambiguity, the standard literature has focused on a knife-edge case. The key contribution of this paper is to provide an alternative foundation for information cascades by interpreting them as a result of model uncertainty instead of the details of information structures.

The second essay, "Biased Learning under Ambiguous Information", pro-

poses and characterizes a novel updating rule under model uncertainty. In this essay, an agent receives a sequence of signals, but he is ambiguous about the signal-generating process and perceives a set of feasible models for it. The agent is endowed with some biased states that he wishes to justify. After receiving a signal, the agent updates his belief according to the model that maximally supports the bias. This biased updating rule can accommodate interesting phenomena which are inconsistent with the Bayesian framework. For instance, the agent can exhibit the "good-news effect"; that is, he processes good news and bad news asymmetrically. This essay provides a complete characterization of limit beliefs under the biased updating rule. Using the characterization, the paper describes several effects of ambiguity on learning. First, ambiguity can lead to incomplete learning and polarization. Second, ambiguity can lead to overconfidence, and the overconfidence can persist even asymptotically.

The third essay, "Naïve Social Learning with Heterogeneous Model Perceptions" studies an economy in which individuals are connected with each other through a social network and they can observe a sequence of signals and communicate beliefs with their neighbours repeatedly through some naïve rule. Previous research shows if all individuals understand the data-generating process correctly then complete learning can be achieved. This essay re-examines the problem under the assumption that some individuals may misinterpret their information. The formal results in this paper provide a characterization of limit beliefs. Using these results, I find that instead of achieving the wisdom of the crowds, society can suffer from group irrationality—even for some seemingly innocuous misperceptions, correct learning may not be achieved; moreover, individuals may end up forming a belief which is inconsistent with everyone's information.

BIOGRAPHICAL SKETCH

Yang Chen (Jaden) was born in July 1995 in Jiangxi Province, China. He received a B.A. in Economics from the Guanghua School of Management at Peking University and a B.S. in Mathematics as a double major in July 2016. He is expected to receive his Ph.D. degree in Economics from Cornell University in May 2022. He will join the University of North Carolina at Chapel Hill as an Assistant Professor in Economics in July 2022. His research interest is microeconomic theory. His current research focuses on learning under model uncertainty, which studies how people learn in the presence of ambiguous information.

This document is dedicated to my family, friends, and all other people who supported me in my life.

ACKNOWLEDGEMENTS

I am deeply grateful for the help from my thesis advisors: David Easley, Lawrence Blume and Tommaso Denti. David, thank you for the advising and encouragement throughout the process. I would not have explored my current research area, the study of learning and uncertainty, without your guidance. Your devotion to research, teaching and advising inspires me tremendously. I hope that I can become a student you are proud of. Larry, thank you for encouraging me to pursue the academic excellence as best as I can. Your sharp and insightful feedback pushes me to view my research from a different perspective. I also had the opportunity to work with you as a Teaching Assistant for three years. The experience helps me to develop many skills in teaching. Tommaso, thank you for so many excellent comments on my research and for introducing me to the cutting-edge literature. I greatly appreciate your availability to talk and the continuous support. I also want to thank Seth Sanders and Penny Sanders for their gracious help during the job market.

I thank my parents for bringing me to the world and giving me the best they can provide. I thank my friends and colleagues, especially Siguang Li, for their companion, encouragement and help. I want to thank Cheng Su for the unconditional love and support throughout my graduate life.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER 1

SEQUENTIAL LEARNING UNDER INFORMATIONAL AMBIGUITY

1.1 Introduction

1.1.1 Overview of the Results

In many economic problems, individuals need to make decisions based on some information. One standard assumption is that individuals entertain a specific data-generating process (or model) and interpret information according to it. However, in many interesting situations, individuals are unable to pin down a specific model and may face ***model uncertainty***. One good example is social learning. In social learning, individuals observe information from others, but the observations can be both imperfect (e.g., individuals only observe actions but not signals) and limited (e.g., individuals may not repeatedly observe from the same individual), so individuals usually can not specify a unique model to interpret all information at hand. Moreover, due to the lack of prior knowledge, individuals may even be unable to assign a prior over models. In this paper, I introduce model uncertainty into a classical problem of social learning—the sequential learning model (SLM, henceforth), where every individual takes an action sequentially and can observe a private signal and previous actions. This paper finds that the social learning outcomes under model uncertainty and under model certainty have some interesting differences.

Under model certainty, the learning outcome depends on the details of the true data-generating process and its perception. In the pioneering work of

Banerjee (1992) and Bikhchandani et al. (1992) (BHW, henceforth), individuals receive i.i.d. signals from a finite signal space according to a commonly known distribution. One important result is that an ***information cascade*** will arise with probability 1. That is, after some point, individuals will ignore their own signals and follow others even if the action is sub-optimal. Later findings suggest that the occurrence of a cascade relies on the finiteness of signals. Smith and Sørensen (2000) showed that: (i) when signals are unbounded, the society will settle on the correct action in the limit, so an incorrect herding can not occur; (ii) even if signals are bounded, the occurrence of an information cascade is not guaranteed for continuous signals; they showed a weaker result that herding occurs with probability 1, that is, individuals end up taking the same action but any individual could break the herd if he or she received a different signal. Herrera and Hörner (2012) showed that when the data-generating processes satisfy the increasing hazard ratio property (IHRP), an information cascade will not take place. Furthermore, if we allow individuals to perceive a misspecified model, the learning outcome depends more intricately on the details of the true data-generating process and its relation with the perceived data-generating process.[1]

One possible challenge of previous results is that it is empirically difficult to test the true data-generating process or model perceptions in most cases, so the it remains vague which result to expect in the social learning. [2]In this paper, I

[1]For example, we may have complete learning, information cascades or even action non-convergence as implied by Corollary 1.4. Similar results are also noted by recent works in misspecified learning, e.g., Bohren (2016), Bohren and Hauser (2021) and Frick et al. (2020b) but with setups different from the standard SLM.

[2]For example, in terms of the boundedness condition, "whether bounded or unbounded beliefs provide a better approximation to reality is partly an interpretational and partly an empirical question" (Acemoglu et al., 2011), so it remains vague when complete learning can be achieved and when it cannot. As for the more complicated properties, e.g., IHRP, it is harder to provide an intuitive explanation on why they can induce different results in terms of informa-

re-examine the SLM under the assumption that individuals are uncertain about other people's data-generating processes. The paper finds that under sufficient model uncertainty, there exists a unique learning outcome—information cascade, with a positive probability of an incorrect cascade—for *all* possible true data-generating processes. Perhaps more surprisingly, this paper also shows that previous results featuring the absence of a cascade can only represent a knife-edge case from the perspective of model uncertainty.

The model setup is introduced in Section 1.2 and summarized as follows. This paper adopts the stylized SLM framework where there are binary states and actions. Individuals take actions sequentially to match the unknown state of the world. Every individual can observe all previous actions as well as a private signal. The only deviation from the standard SLM is that individuals are ***ambiguous*** about their predecessors' data-generating processes by perceiving a set of feasible data-generating processes, $\mathcal{F}_0$. The size of $\mathcal{F}_0$ intuitively measures the degree of ambiguity, where model certainty corresponds to the case where $\mathcal{F}_0$ is a singleton. The benchmark model assumes that individuals follow the max-min EU model to make choices. It turns out that for all large $\mathcal{F}_0$s, an information cascade occurs almost surely even when the true data-generating processes imply no cascades in the standard case.

At first glance, a large $\mathcal{F}_0$ can include models with different implications, so it is not straightforward which result to expect. The paper's result relies on the finding that the models encouraging an information cascade and the models discouraging it are *asymmetric*. For an individual who is called upon to take an action, the worst possibility if she were to break a herd happens when the herding individuals have highly precise information, whereas the worst possibility if

---

tion cascades.

she followed the herd happens when these predecessors have very imprecise information. The asymmetry in the worst cases implies that as $\mathcal{F}_0$ becomes larger, the cascading force can increase consistently, but the non-cascading force is always restricted. As a consequence, under sufficient ambiguity, the cascading force dominates, so we have an information cascade. Below is a simple example illustrating this idea.

Example 1. The state space $\Theta = \{0,1\}$, and the prior is $\pi_0 = (1/3, 2/3)$. Every individual i takes action $a_i \in \{0,1\}$ sequentially and receives a signal $s_i \in \{h, l\}$. The DGP $g_i\,(s|\theta)$ satisfies

$$\frac{g_i\,(h|1)}{g_i\,(l|1)} = \frac{g_i\,(l|0)}{g_i\,(h|0)} = \gamma_i \in (1,\infty) = \Gamma,$$

where $\gamma_i \overset{i.i.d.}{\sim} h \in \Delta\,(\Gamma)$ describes the signal precision. Individuals know their own precision but is ambiguous about h, so they are ambiguous about others' precision. Let's consider an extreme case, where $\mathcal{F}_0 = \Delta\,(\Gamma)$, so any distribution is possible for h. Suppose that the first individual (he) chose action 1. Denote by $V_2\,(a)$ the minimum EU of the second individual (she) if she takes action a. We have

$$V_2\,(1) = \begin{cases} \gamma_2/\,(\gamma_2+2) & s_2 = h \\ 2/\,(\gamma_2+2) & s_2 = l \end{cases} \text{ and } V_2\,(0) = 0.$$

If individual 2 followed the first individual and chose action 1, the minimum EU is obtained when she believes that individual 1 can only receive uninformative signals. In this case, individual 1's action contains no information which gives $V_2\,(1)$. In contrast, if individual 2 chose a different action, action 0, the the minimum EU is obtained when she believes that individual 1 can only receive the perfectly revealing signal, so acting against him yields a utility of 0, which gives $V_2(0)$.[3] As $V_2(1) > V_2(0)$, individual 2 will follow individual 1 for all possible private signals, so an information cascade occurs.

Notice that the occurrence of an action-1 cascade does not rely on the specific properties of h nor the true state, so an information cascade occurs for all $h \in \Delta(\Gamma)$, and an incorrect cascade occurs with a strictly positive probability. $\square$

The example shows that an information cascade arises under extreme ambiguity, but this paper further notes that to achieve a cascade, the condition can be much weaker. In many interesting situations, a cascade arises only when individuals face a *slight degree of ambiguity*. Section 1.5 characterizes conditions that ensure the almost sure occurrence of an information cascade when signals are **bounded**. Theorem 1.2 provides two sufficient conditions that guarantee a cascade. Intuitively, as long as individuals find it possible that other individuals may have a highly information data-generating process, an information cascade will occur almost surely, regardless of the details of the true data-generating process and other data-generating processes under consideration. The intuition is similar to that in Example 1. We first note that the perception of an informative data-generating process encourages individuals to follow a herd hence creates a cascading force. Due to the asymmetry, if the informativeness is sufficiently high, the presence of any other model is inadequately to offset the cascading force, so an information cascade always occur irrespective of what other models individuals may consider. Interestingly, the conditions proposed by Theorem 1.2 are very easy to hold in many situations. Suppose that the true model is $\overline{F}$, and there is no cascade when individuals correctly perceive $\overline{F}$. If individuals are ambiguous and consider all F such that $||F - \overline{F}|| \leq \varepsilon$, where $||.||$ is some metric

[3] More rigorously, the minimum EU is actually the infimum EU, since $\gamma \in (1, \infty)$.

(e.g., sup-norm metric), then an information cascade occurs almost surely for all $\varepsilon > 0$. It suggests that the standard non-cascade results are not robust, and the statistical properties relevant for a cascade (i.e., IHRP) only matter in knife-edge cases from the perspective of ambiguity.

Section 1.6 shows that similar results also exist for **unbounded** signals. For unbounded signals, an information cascade is more difficult to happen by definition, but this paper establishes that herding occurs almost surely, where an **incorrect herding** occurs with a strictly positive probability, even with arbitrarily small ambiguity. The result implies that the complete learning result in Smith and Sørensen (2000) can be non-robust. Theorem 1.3 provides conditions under which an incorrect herding occurs. The idea is similar to the bounded signal case—if individuals perceive some model that is adequately informative, the cascading force is so strong that it can not be outweighed by any other models, hence an incorrect herding can occur. In some interesting settings, Theorem 1.3 can be easily satisfied so that complete learning only represents a knife-edge case. In Section 1.7, I further provide a necessary and sufficient condition of complete learning under ambiguity for an important class of models—models with power tails (Theorem 1.4). To achieve complete learning, we must impose restrictions on $\mathcal{F}_0$ from both directions—$\mathcal{F}_0$ cannot be overly informative or overly uninformative, since the former will encourage an incorrect herding whereas the latter can lead to action non-convergence.

Section 1.8 extends the discussion to **general ambiguity preferences**. Under the max-min model, individuals are ambiguity-averse and are extremely ambiguity sensitive in the sense that decisions are only affected by the worst outcomes. Section 1.8 focuses on two common alternative models—the α-max-min model and the smooth ambiguity model—and has the following findings. First, a cascade does not require individuals to be ambiguity-averse, and it can also occur with ambiguity-seeking individuals. In the discussion of the α-max-min model, I show that an information cascade occurs for all $\alpha \in [0, 1]$, where $\alpha = 0$ means the max-max model and $\alpha = 1$ means the max-min EU model. Second, a cascade does not rely on the extreme ambiguity sensitivity as in the max-min EU model. As long as individuals are adequately ambiguity sensitive, an information cascade can occur. An example with smooth ambiguity preference shows that an information cascade arises when the curvature of the second-order utility function (i.e., measures the ambiguity sensitivity) is sufficiently large. The discussion implies that main results under the max-min EU model hold for general ambiguity preferences.

Techniques. The analysis under model uncertainty has the following challenges. First, individuals hold a set of posteriors, so we can not keep track of the posterior likelihood ratio as in the literature. To facilitate the analysis, this paper notices that there is a simple statistic—the average likelihood ratio—that captures all relevant information. Employing this fact, this paper is able to analyze the learning process by simply keeping track of this statistic.

Second, under model uncertainty, posteriors no longer exhibit the martingale property, which makes the analysis more difficult especially when signals are unbounded. This challenge is also present in the misspecified learning, where a common approach is to first analyze the local stability of each state and then to extend it globally. Following a similar idea, I re-define the notion of local stability using the average likelihood ratio. One challenge here is that standard techniques to analyze local stability are not applicable to the SLM.

This paper adopts an alternative approach—**infinite series approach**—to facilitate the discussion. The key idea is that a state θ is locally stable (if) and only if the probability that all individuals take action θ is (uniformly) strictly positive when priors are near δ_θ. Besides, it can be further shown that the probability is (uniformly) strictly positive if and only if some infinite series is convergent. As a result, the discussion on local stability is transferred to a simpler problem of determining the convergence of the series. This approach is applied in Theorems 1.3 and 1.4. The relevant literature is discussed below.

1.1.2 Related Literature

This paper is among the few papers that study social learning under ambiguity, especially under model uncertainty. The most relevant paper is Ford et al. (2013), which investigated a sequential trading model where traders face ambiguity and have neo-additive capacity EU preference.[4] They found that in the presence of ambiguity, informed traders can exhibit herding behavior, which is consistent with this paper in some sense. However, their setup and mechanism are different from this paper. First, in their paper, decisions are made in a specific market structure, where trading decisions are jointly determined by beliefs and ask-bid prices, but in this paper, there is no market and decisions are only determined by the beliefs. Second, in their paper, ambiguity also leads to both herding and contrarian, whereas in this paper, ambiguity only produces herding. [5]

[4]There are other papers related to learning under model uncertainty but not in social learning. Examples include Acemoglu et al. (2016), Battigalli et al. (2015), Battigalli et al. (2019b), Chen (2020), Epstein and Schneider (2007)), Fryer Jr et al. (2019), Marinacci (2002, 2015), Marinacci and Massari (2019)

[5]More technically, the occurrence of herding and contrarian in Ford et al. (2013) comes from that posteriors are always bounded away from 0 and 1 under the neo-additive capacity prefer-

Under model uncertainty, individuals will inevitably perceive some incorrect models, so the paper is closely related to the literature on misspecified social learning. The setup is most similar to that of Bohren and Hauser (2021) (also Bohren (2016)), in which they also investigated a sequential learning problem with binary states. One of their main results is that complete learning is *robust* with respect to small misspecifications, which stands in contrast to this paper's finding. The difference arises from their assumption that the society has a positive fraction of autarkic agents who only act according to their private signals. This assumption plays an important role in establishing the local stability of the true state since it leads to a strict Berk-Nash equilibrium (Esponda and Pouzo (2016)). However, their framework does not nest the standard SLM, where a strict Berk-Nash equilibrium can not be obtained. This paper employs a different approach to analyze the local stability and establishes that complete learning in the SLM is *non-robust* in many cases. Frick et al. (2020a,b) also established that complete learning is not robust but in very different settings. Specifically, Frick et al. (2020a) considered a social learning problem (not sequentially) where the state space is continuous and individuals with different preferences randomly meet with each other. Frick et al. (2020b) proposed a local martingale-based approach to analyze misspecified learning and showed the fragility of sequential learning in an environment. Differently, their approach relies on the assumption is that signals are bounded and the fragility of the sequential learning relies on the assumption that risk preference is adequately hetergeneous.

One common feature of Bohren and Hauser (2021) and Frick et al. (2020b) is that they demanded a strict dominance relation to establish local stability, e.g., strict Berk-Nash equilibrium, strict p-dominance. However, under SLM,

ence. In contrast, this paper mainly works with the max-min EU preference, and the occurrence of herding or cascade does not rely on posteriors being bounded away from certainty.

the strict dominance is not satisfied. Technically, this paper complements the literature by employing an infinite series approach to establish local stability. The most similar literature is Rosenberg and Vieille (2019), where they also employed an infinite series to characterize learning efficiency. The difference is that in their setup, individuals perceive a correct model, so complete learning occurs with unbounded signals; in this paper, individuals perceive multiple models, and the infinite series is useful in establishing incomplete learning. Furthermore, applying this paper's arguments, we can show that the efficient learning in their paper can be non-robust with respect to ambiguity (see Remark 1.2).

Under ambiguity, individuals no longer learn in a Bayesian manner, so this paper also belongs to the literature on non-Bayesian social learning. There are papers studying sequential learning with boundedly rational agents, for example, Eyster and Rabin (2010), Guarino and Jehiel (2013) and Dasaratha and He (2020), where individuals follow some naive rule to aggregate information from predecessors. Non-Bayesian social learning is also studied in general network structures, for example, DeMarzo et al. (2003), Golub and Jackson (2010), Li and Tan (2020), Molavi et al. (2018), where individuals apply a rule of thumb when aggregating information from others.

1.2 The Model

There are two possible states of world, $\Theta = \{0, 1\}$. A countably infinite set of individuals $N = \{1, 2, ...\}$ act sequentially. Each individual makes a binary choice $a \in A = \{0, 1\}$ and can observe the choices taken by all predecessors. Individuals have identical utility functions, which have a payoff of 1 when actions match the actual state and a payoff of 0 otherwise. Without loss of generality, the true state $\theta^* = 0$ and is not known to individuals.

Signal Structure

Individuals share a full-support common prior π_0. For simplicity, I assume that $\pi_0(0) = \pi_0(1) = \frac{1}{2}$. Each individual i, will receive a signal $s_i \in \mathcal{S} \subset \mathbb{R}$. Signals are independently (but not necessarily identically) distributed according to $\left\{\overline{G}_1^\theta, \overline{G}_2^\theta, ...\right\}$, where $\overline{G}_i^\theta : \mathcal{S} \to [0,1]$ denotes the cumulative distribution function of s_i when the actual state is θ. I refer to $\overline{G}_i = \left(\overline{G}_i^0, \overline{G}_i^1\right)$ as individual i's *data-generating process*. No signal perfectly reveals the state; therefore, the probability measures induced by $\overline{G}_i^0$ and $\overline{G}_i^1$ are mutually absolutely continuous. Following the convention, I introduce the normalized signal, $\lambda_i(s)$, where $\lambda_i(s) = \frac{d\overline{G}_i^1(s)}{d\overline{G}_i^0(s)}$ denotes the likelihood ratio induced by each signal. The distribution of the likelihood ratio λ_i is denoted by $\overline{F}_i^\theta$, so $\overline{F}_i^\theta$ must satisfy $\lambda = \frac{d\overline{F}_i^1(\lambda)}{d\overline{F}_i^0(\lambda)}$ almost everywhere, which means that receiving a signal λ leads to a likelihood ratio equal to λ. For the rest of this paper, I focus on the normalized signal, λ, and the normalized data-generating processes, $\overline{F}_i^\theta$. For simplicity, I assume that: (i) all signals are continuous, that is, $\overline{F}_i^\theta$ is continuous for all i and θ, and (ii) signals are symmetric in the sense that $\overline{F}_i^1(\lambda) = 1 - \overline{F}_i^0(1/\lambda)$ for all i and λ. All results can be extended to cases where signals are discontinuous and asymmetric. All individuals' data-generating processes are assumed to have a common support, $\Lambda \equiv co(\text{supp}(F_i)) = \left[\frac{1}{\gamma}, \gamma\right] \subset [0, \infty]$, where $\gamma > 1$, meaning that signals are informative. Signals are *bounded* if $\gamma < \infty$ and signals are *unbounded* if $\gamma = \infty$. Let $\mathbb{P}^*$ denote the true probability measure, that is, the measure induced by data-generating processes, $\left\{\overline{F}_1^\theta, \overline{F}_2^\theta, ...\right\}$, conditional on the true state θ^*.

Belief Structure

Individuals are assumed to be *ambiguous* about their predecessors' data-generating processes, but they know that signals are symmetric and independently distributed. As a consequence of the symmetry, every data-generating process, $F = \left(F^0, F^1\right)$, is uniquely determined by one coordinate, so this paper keeps track of F^1 when characterizing every individual's belief set. Denote by $\mathcal{F}$ the set of all possible F^1, both discrete and continuous, with support in $\left[\frac{1}{\gamma}, \gamma\right]$. To be more precise, $\mathcal{F}$ consists of the set of F^1s that constitute some symmetric data-generating process with support in $\left[\frac{1}{\gamma}, \gamma\right]$. Further denote by $\mathcal{F}_{ij}$ the set of individual j's data-generating processes considered possible by individual i, where $\mathcal{F}_{ij} \subset \mathcal{F}$ and $j \neq i$, and refer to $B_i \equiv \left\{\mathcal{F}_{ij} : j \neq i\right\}$ as individual i's *belief structure*. With this, I make the following assumptions throughout this paper.

Assumption 1.1. *[Homogeneous Belief] There exists some $\mathcal{F}_0 \subset \mathcal{F}$ such that $\mathcal{F}_{ij} = \mathcal{F}_0$ for all $i, j \in N$ with $j \neq i$.*

Assumption 1.2. *[Common Knowledge] Individuals' belief structures are common knowledge.*

Assumption 1.1 is that individuals have homogeneous belief structures. The homogeneity is reflected in the following two aspects. First, for any given individual i, this individual is identically ambiguous about every other individual's data-generating process. Second, all individuals have identical belief structures. The homogeneous belief assumption is a simplifying assumption and can be extended to the heterogeneous belief setting. Assumption 1.2 is a standard assumption, with which individuals are able to make iterated inferences based on other individuals' actions.[6] These two assumptions hold in situations where the

[6] Assumption 1.2 is commonly adopted in the literature of SLM. In the standard SLM, it

set of data-generating processes are public information, but individuals lack the information to determine the specific data-generating processes. For simplicity, I would also use the notation $F \in \mathcal{F}_0$, and here $F = \left(F^0, F^1\right)$ is a data-generating process. It actually means that the second coordinate $F^1 \in \mathcal{F}_0$.

Belief-updating Process

Let $h_i = (a_1, ..., a_{i-1})$ be the history observed by individual i. Denote $I_i = \{\lambda_i, h_i\}$, where I_i represents the information available to individual i, which consists of her private signal λ_i and history h_i. Let $\mathcal{I}_i$ be the set of all possible information available to individual i, where $\mathcal{I}_i = \Lambda \times \{0,1\}^{i-1}$. A (pure) strategy for individual i is a mapping $\sigma_i : \mathcal{I}_i \to \{0,1\}$, which maps individual i's information set, I_i, to an action $a_i \in \{0,1\}$. When the actual state is θ, for any individual i, given the predecessors' strategy profile $\sigma_{-i} = (\sigma_1, ..., \sigma_{i-1})$, and their data-generating processes profile $F_{-i} = (F_1, ..., F_{i-1})$, the observed history $h_i = (a_1, ..., a_{i-1})$ is a stochastic process with a probability measure $\mathbb{P}_{F_{-i}}(.|\theta;\sigma_{-i})$. Let $\Pi(h_i, \sigma_{-i})$ denote the set of beliefs over the state space Θ, given history h_i, and strategy profile σ_{-i}, which I refer to as a *public belief set*. It is easy to see that

$$\Pi(h_i, \sigma_{-i}) = \left\{\pi \in \Delta(\Theta) : \pi(\theta) = \mathbb{P}_{F_{-i}}(\theta|h_i;\sigma_{-i}), F_{-i} \in \mathcal{F}_0^{i-1}\right\}$$

where $\mathbb{P}_{F_{-i}}(\theta|h_i;\sigma_{-i})$ is the conditional probability on θ derived from $\mathbb{P}_{F_{-i}}(.|\theta;\sigma_{-i})$, and $\mathcal{F}_0^{i-1}$ is $i-1$ copies of $\mathcal{F}_0$. The public belief set consists of

is assumed that individuals correctly understand the true DGP and the true DGP is common knowledge, for example, Banerjee (1992), Bikhchandani et al. (1992), Smith and Sørensen (2000). In the SLM with model misspecification, individuals may perceive an incorrect DGP, but how individuals perceive the model is commonly known, for example, Bohren (2016), Bohren and Hauser (2021). Similarly, this paper assumes that individuals are uncertain about the DGP and the feasible model set $\mathcal{F}_0$ is commonly known.

conditional probabilities generated by all possible $F_{-i} \in \mathcal{F}_0^{i-1}$ for which the conditional probabilities are well-defined. Based on the public beliefs and private signal λ_i, individual i will form a belief set, $\Pi_i(I_i, \sigma_{-i})$, which I refer to as a *posterior set*. Assuming that individuals use the full Bayesian updating rule axiomatized by Pires (2002) to update beliefs, thus:

$$\Pi_i(I_i, \sigma_{-i}) = \left\{\pi \in \Delta(\Theta) : \pi = BU\left(\pi'; \lambda_i\right), \pi' \in \Pi(h_i, \sigma_{-i})\right\}$$

where $BU(\pi'; \lambda_i)$ denotes the Bayesian update of belief π' based on signal λ_i. In other words, individuals update the public belief set prior-by-prior using Bayes' rule.

This updating rule has the advantage of being straightforward and it has been adopted in many applications (e.g., Bose and Renou (2014)). Two major criticisms of it, however, are: first, the size of the belief set remains unchanged even after learning new information; second, it can lead to dynamic inconsistency (see Machina and Siniscalchi (2014)'s survey). My responses are: the results in this paper are still robust to other updating rules, for example individuals can update the set of data-generating processes based on their observations; besides, in the setting of this paper, individuals only need to make a once-in-a-lifetime decision; therefore, dynamic inconsistency is not relevant here[7].

[7]It remains a question whether dynamic consistency should be maintained in the presence of ambiguity. Notice that Bayes' updating rule comes from Savage's sure-thing principle, which implies both consequentialism and dynamic consistency. Because most ambiguity preferences assume a violation of the sure-thing principle, it becomes hard to maintain both properties. Many papers hence drop dynamic consistency to retain consequentialism (e.g., Pires (2002) and Eichberger et al. (2007)).

Equilibrium Concept

Assume that individuals are ambiguity averse and have max-min expected utility (MEU) preferences as in Gilboa and Schmeidler (1989). An equilibrium concept is defined as the following:

Definition 1. (Equilibrium) A strategy profile $\sigma^* = \left(\sigma_i^*\right)_{i \in N}$ constitutes an *equilibrium* if for all $i \in N$ and all information sets $I_i \in \mathcal{I}_i$, we have:

$$\sigma_i^* \left(I_i\right) \in \arg\max_{a \in \{0,1\}} \inf_{\pi \in \Pi_i\left(I_i, \sigma_{-i}^*\right)} \mathbb{E}_\pi U\left(a, \theta\right) \tag{1.1}$$

, where $U\left(a, \theta\right)$ is the utility function which equals 1 if $a = \theta$ and equals 0 if $a \neq \theta$.

Where no confusion would exist, I omit the equilibrium strategy notation σ^* and denote $\Pi\left(h_i\right)$ and $\Pi_i\left(I_i\right)$ as the equilibrium public belief set and posterior set. To address the tie case, I assume the following "tie-breaking rule": when indifferent, individual i chooses action 1 if $\lambda_i > 1$ and action 0 if $\lambda_i \leq 1$. With the rule, Definition 1 provides a unique pure-strategy equilibrium.

It remains questionable whether the pure-strategy equilibrium is a reasonable concept to consider when mixed strategies are also allowed. The answer is straightforward in the standard model because individuals with expected utility preferences are indifferent to randomization over choices; thus, restricting attention to pure-strategy equilibriums is without loss of generality. However, with ambiguity, there needs to be slightly more justifications. It seems that ambiguity-averse individuals can make themselves better off by using randomizations, as suggested by the Uncertainty Aversion axiom in Gilboa and Schmeidler (1989). However, the Uncertainty Aversion axiom only assumes that individuals have incentives to engage in *ex-post randomization* instead of *ex-ante randomization*, as in the mixed-strategy case.[8] Although there still remains a question of whether individuals can benefit from ex-ante randomization, a growing body of literature suggests that indifference to ex-ante randomization seems a more reasonable assumption.[9] Hence, in this paper, I also assume that individuals are indifferent to ex-ante randomization. Under this assumption, individuals have no incentive to play mixed strategies, therefore Definition 1 accommodates the case where mixed strategies are allowed.

1.3 Equilibrium Strategies and Concepts

This section characterizes individuals' equilibrium strategies. Similar to the standard model, ambiguous individuals' equilibrium strategies can be decomposed into two parts representing information from private signals and public history. This section then characterizes cascade sets and formally defines an information cascade, which will be used in later discussions.

[8]The Uncertainty Aversion axiom is: for all acts f, g and $\alpha \in (0, 1)$, so we have:

$$f \simeq g \Rightarrow \alpha f + (1-\alpha) g \succeq f$$

, where $[\alpha f + (1-\alpha) g](s) \equiv \alpha f(s) + (1-\alpha) g(s)$ for all possible states, $s \in S$. That is, for all states, acts f and g are mixed with a fixed proportion, α and $1-\alpha$, which is the ex-post randomization. Whereas the ex-ante randomization means that individuals first take a lottery with probability α on f and $1-\alpha$ on g, individuals' payoffs are only generated by f or g depending on the lottery outcome.

[9]For example, Saito (2012) suggests that individuals have no incentive to engage in ex-ante randomization when the Certainty Strategic Rationality axiom is assumed. Eichberger et al. (2016) shows that dynamic consistency implies that individuals are indifferent to ex-ante randomizations. Besides, indifference to ex-ante randomizations is also implicitly assumed in the smooth ambiguity model axiomatized by Klibanoff et al. (2005) (in the assumption that individuals have expected utilities on second-order acts).

1.3.1 Characterizations of Equilibrium Strategies

In the standard model, after observing history h_i and private signal λ_i, individual i's posterior belief has a likelihood ratio equal to $\lambda_i \cdot l_i$, where l_i denotes the likelihood ratio of the public belief after observing history h_i. In the equilibrium, individual i will choose action 1 if the product, $\lambda_i \cdot l_i$, is greater than 1 and choose action 0 otherwise. As a result, individuals' decision rules can be decomposed into two parts: the private information part, λ_i, and public information part, l_i.

When individuals are ambiguous, it seems difficult to have such a neat decomposition because the public belief is represented by a set. However, it turns out that we can extend the concept of likelihood ratio and represent the public belief set using the average likelihood ratio for the beliefs featured in in it. Based on this unique fact, we have a parallel characterization of individuals' equilibrium strategies under ambiguity. The idea of the average likelihood ratio is introduced below:

Definition 2. (Average Public Likelihood Ratio) Denote $L_i = \left\{ \frac{\pi(1)}{\pi(0)} : \pi \in \Pi\left(h_i\right) \right\}$, where $\underline{l}_i = \inf L_i$ and $\overline{l_i} = \sup L_i$. Denote $r_i = \sqrt{\overline{l_i} \cdot \underline{l}_i}$, called the *average public likelihood ratio*, based on history h_i.

The average public likelihood ratio r_i is the geometric average of the highest and lowest likelihood ratios in the public belief set, which reflects how likely the public thinks state 1 is (relative to state 0) on average. The following characterizes individuals' equilibrium strategies employing the average public likelihood ratios.

Proposition 1.1. (Characterizations of Equilibrium Strategies) *In the equilibrium,* *for any individual, $i \in N$, and information set, $I_i \in \mathcal{I}_i$, we have:*

(1) When $\lambda_i > 1$: $\sigma_i^*(I_i) = 1$ if and only if $\lambda_i \cdot r_i \geq 1$; $\sigma_i^*(I_i) = 0$ if and only if $\lambda_i \cdot r_i < 1$.

(2) When $\lambda_i \leq 1$: $\sigma_i^*(I_i) = 1$ if and only if $\lambda_i \cdot r_i > 1$; $\sigma_i^*(I_i) = 0$ if and only if $\lambda_i \cdot r_i \leq 1$.

Proof. Denote $\underline{\pi}_i(\theta) = \inf\{\pi(\theta) : \pi \in \Pi_i(I_i)\}$. Suppose that $\lambda_i > 1$, then $a_i = 1$ if and only if $\underline{\pi}_i(1) \geq \underline{\pi}_i(0)$. Note that:

$$\underline{\pi}_i(1) = \frac{\lambda_i \underline{l}_i}{1 + \lambda_i \underline{l}_i} \qquad \underline{\pi}_i(0) = \frac{1}{1 + \lambda_i \overline{l}_i}$$

by solving $\underline{\pi}_i(1) \geq \underline{\pi}_i(0)$, we have: $\lambda_i \geq \frac{1}{\sqrt{\overline{l}_i \cdot \underline{l}_i}} = \frac{1}{r_i}$. Other cases follow symmetrically. □

The average likelihood ratio is an extension of the likelihood ratio in the standard model. It acts as a sufficient statistic for the public history in cases where there are multiple beliefs. Proposition 1.1 shows that individuals' equilibrium strategies can also be represented as the product of two parts. The private information component is still the private signal, λ_i, whereas the public information is captured by the average public likelihood ratio, r_i. When the product, $\lambda_i \cdot r_i$, is greater than 1, reflecting that state 1 is more likely, individuals will choose action 1 and vice versa. For simplicity, "average public likelihood ratio" is sometimes referred to as "public belief" when there is no confusion.

Note that the representation of the average public likelihood ratio relies on individuals' ambiguity preferences. The representation in Definition 2 relies on the assumption that individuals have MEU preferences. When individuals have other ambiguity preferences (e.g., smooth ambiguity preferences), we may have different representations.

1.3.2 Herding, Cascades and Learning

Definition 3. [Herding and Information Cascades] In the equilibrium, we say that

(i) a ***herding*** occurs if there exists some $I < \infty$ and $a \in A$ such that for all $i \geq I$, $a_i = a$;

(ii) an ***information cascade*** occurs if there exists some $I < \infty$ and $a \in A$ such that for all $i \geq I$, we have $\mathbb{P}^* (a_i = a | h_i) = 1$;

A herding occurs if the society ends up taking the same action, and an information cascade occurs if after some point, individuals will only choose one action regardless of their private signals. An information cascade is stronger than a herding. During a herding, individuals would have acted differently if they received different signals, whereas during an information cascade, any realizations of private signals are unable to overturn the herd.[10] When signals are bounded and individuals correctly specify the model, a herding occurs almost surely and the herding can be incorrect, but an information cascade may or may not happen depending on the details of the data-generating processes.

Definition 4. ***Complete learning*** occurs if there exists some $I < \infty$ such that for all $i \geq I$, $a_i = \theta^*$ $\mathbb{P}^*$-almost surely.

[10] These two concepts was distinguished by Smith and Sørensen (2000) and experimentally distinguished by Çelen and Kariv (2004). Anderson and Holt (1997) also provided laboratory evidence of information cascades.

In other words, complete learning occurs if the society eventually settles on the optimal action with probability 1. In the standard framework, complete learning occurs if and only if signals are unbounded.

1.4 Benchmark Case: Cascades under Extreme Ambiguity

To build intuition, it is better to discuss the benchmark case, where individuals are extremely ambiguous.

Assumption 1.3. $\mathcal{F}_0 = \mathcal{F}$.

This assumption says that individuals consider all feasible models with support in $[\gamma, 1/\gamma]$. It describes a situation where individuals only know the range of signals without further knowledge about the true data-generating processes. Under this assumption, we have the following theorem.

Theorem 1.1. *Under Assumption 1.3, an information cascade occurs $\mathbb{P}^*$-almost surely for* ***all*** *possible $\overline{F}_i$s in $\mathcal{F}$.*

One significance of Theorem 1.1 is that an information cascade occurs almost surely for all possible true data-generating processes, regardless of whether signals are bounded or unbounded, discrete or continuous. Besides, it is also true that an incorrect cascade occurs with a strictly positive probability, so the society can settle on the incorrect action for all possible signals (even if signals are unbounded).[11] The intuition comes from the following arguments.

[11] In this paper, whenever an information cascade arises, an incorrect cascade must arise with a strictly positive probability, so I will not state it explicitly in theorems.

Intuition. Suppose that the first i individuals take action 1 (i.e., $a_1 = ... = a_i = 1$), and individual $i+1$ receives a signal $\frac{1}{\gamma}$, the strongest signal for state 0. Suppose that an information cascade did not occur when the first i individuals made decisions. Consider the decision problem of individual $i+1$. As she has max-min EU preference, her decision is determined by the worst scenarios:

- If she follows the herd and takes action 1, the worst case happens when the predecessors' data-generating processes are uninformative. In this case, $\lambda_1 = ... = \lambda_i = 1$. By following the herd, she will be acting against one signal, $\frac{1}{\gamma}$ (her own signal);
- If she breaks the herd and takes action 0, the worst case arises when the predecessors' data-generating processes are most precise (only generating signal γ and $1/\gamma$). In this case, the predecessors' actions perfectly reveal that their signals are all γs. Hence, by taking action 0, individual $i+1$ follows her own signal, but she will be acting against i signal γs.

As can be seen, the forces encouraging a cascade and discouraging it are **asymmetric**. As i increases, the cost of breaking the herd increases consistently as individual $i+1$ will be acting against more and more signal γs in the worst case. However, the cost of herding remains the same due to the fact that, in the worst case, individual i is always acting against one signal $\frac{1}{\gamma}$. When i is sufficiently large, the cost of breaking the herd is higher; therefore, all the following individuals will choose action 1 and an information cascade of action 1 occurs.

Sketch of the Proof. The proof of Theorem 1.1 is not difficult. First, from the equilibrium strategy, the average likelihood ratio, r_i, serves as a sufficient statistic for the history, so we only keep track of r_i. Second, it can be verified that whenever a cascade does not occur, the increment of r_i is bounded away from 1 by some constant. To see that, suppose that $a_i = 1$. Then, the highest likelihood ratio increases by a factor of γ, which corresponds to the increment when individual i has the most precise data-generating process, but the lowest likelihood ratio does not decrease, since intuitively, an action 1 appears as positive news for state 1. Combining these two cases, the average public likelihood ratio should at least increase by a factor of $\sqrt{\gamma}$, which is greater than 1. The case where $a_i = 0$ is symmetric. The rest of the analysis is standard. From the second step, we know that finite number of consecutive actions will trigger a cascade, which further implies that at each period, the probability of a cascade is bounded away from 0, so a cascade must occur with probability 1.

1.5 Information Cascades with Bounded Signals

Last section establishes an information cascade when there is extreme ambiguity. It is natural to ask whether the result still holds in less extreme cases. It turns out that when signals are bounded, an information cascade is easy to occur in the presence of ambiguity. Interestingly, in many situations, the non-cascade results only represent knife-edge cases. To see that, let's first look at two conditions that ensure a cascade.

Theorem 1.2. *If there exists some $F \in \mathcal{F}_0$ such that one of the following conditions holds:*

(1) F is discrete at γ;

(2) F is continuously differentiable on $(\gamma - \varepsilon, \gamma)$ for some $\varepsilon > 0$ with $f^1(\gamma) > \frac{2}{\gamma - 1}$,

where $f^1(\gamma) = \lim_{x \nearrow \gamma} \frac{dF^1}{dx}(x)$. *Then, when signals are bounded, an information cascade occurs* $\mathbb{P}^*$*-almost surely.*

Conditions (1) and (2) can be intuitively interpreted as that some DGP under consideration is sufficiently informative. Specifically, the DGP assigns sufficiently large weights to high-precision signals, that is, signals close to γ, and symmetrically, signals close to $1/\gamma$. Theorem 1.2 says that if individuals find it possible that other individuals can have a highly informative DGP, an information cascade emerges almost surely.

Theorem 1.1 imposes the following restrictions. First, it only requires $\mathcal{F}_0$ to contain one one such F but does not impose other restrictions on $\mathcal{F}_0$. The intuition is based on the observation that the forces encouraging a cascade and discouraging it are asymmetric. Due to the asymmetry, if $\mathcal{F}_0$ contains a highly informative F, the cascading force becomes so strong such that it cannot be offset by other models, so a cascade will occur almost surely regardless of what other models $\mathcal{F}_0$ may contain. Second, it only requires the F to place sufficient weights on the tails but does not impose any restrictions in the middle. It comes from the fact that beliefs will approach the boundary after many identical actions, so by restricting the tail properties of F, we can ensure the occurrence of an information cascade.

Remark 1.1. *One conjecture is that F is the "essential model" under the max-min EU, so individuals would act as if the true model was F. This conjecture is incorrect. In fact, every model in* $\mathcal{F}_0$ *may affect the learning process. It is just that the cascading force from F is too strong such that other models cannot alter the occurrence of a cascade.* [12]

[12] Even if other models do not alter the occurrence of a cascade, but they do affect belief dynamics, convergence speed and so on. As a result, learning under a non-degenerate $\mathcal{F}_0$ containing F is not observationally equivalent to learning under $\{F\}$. Moreover, if F is less informative,

1.5.1 Fragility of the Non-cascade Result

As can be seen, the restrictions in Theorem 1.1 can be very moderate in some sense, which implies that the standard result about cascades can be non-robust. Below is an example.

Example 2. Suppose that individuals perceive the following set of models

$$\mathcal{F}_0 = (1-\varepsilon)\,G + \varepsilon\mathcal{F} \equiv \{F_0 : F_0 = (1-\varepsilon)\,G + \varepsilon F, \text{for } F \in \mathcal{F}\}\,.$$

Recall that $\mathcal{F}$ denotes the set of all possible data-generating processes. The model set, $\mathcal{F}_0$, is constructed by making an $\varepsilon-$perturbation to G using set $\mathcal{F}$. Notice that when $\varepsilon > 0$, there exists some F_0 that is discrete at γ, so an information cascade occurs almost surely for all $\varepsilon > 0$ by Theorem 1.2. □

Example 2 shows that any positive perturbation can produce an information cascade, so we can have an information cascade even if individuals are just "slightly" ambiguous. To characterize the degree of ambiguity explicitly, I follow a common approach in the literature and assume that individuals' model sets are generated by some distance function.

Assumption 1.4. *Suppose that $d : \mathcal{F} \times \mathcal{F} \to \mathbb{R}_+$, and that individuals have the following belief set:*

$$\mathcal{F}_0 = \{F \in \mathcal{F} : d\,(F, G) \leq K\}$$

for some $K \geq 0$ and some $G \in \mathcal{F}$ with a strictly positive density on $[1/\gamma, \gamma]$.

Assumption 1.4 states individuals only consider the set of data-generating processes within distance K to the benchmark model, G. The requirement that

it is possible to construct an example where a cascade may or may not occur depending on other models in $\mathcal{F}_0$.

G has a positive density function is to simplify the discussion by ruling out some irregular Gs. Here, K is referred to as individuals' *ambiguity level*. When $K = 0$, meaning that individuals are not ambiguous, the belief set only contains the benchmark model, G; when $K = \infty$, individuals are extremely ambiguous, as in Assumption 1.3. I focus on the following class of distribution function.

Definition 5. d is *consistent with weak convergence* if for any data-generating process, F, and sequence, $(F_n) \in \mathcal{F}$, satisfying $F_n \Rightarrow F$, we have $d(F_n, F) \to 0$, where "$\Rightarrow$" represents weak convergence.

Consistency with weak convergence has an intuitive interpretation: when the distribution functions of two data-generating processes are close to each other, individuals tend to think that they are similar. Many commonly used distances belong to this class, including the sup-norm metric, total-variation, and Levy-Prokhorov metric. With this class of distance metrics, we have a similar result as in Corollary 1.1.

Corollary 1.1. *Under Assumption 1.4 and that d is consistent with weak convergence, when signals are bounded, an information cascade occurs $\mathbb{P}^*$-almost surely for all $K > 0$.*

The idea is similar to Example 2. Under a metric consistent with weak convergence, any continuous data-generating process can be approximated by a discrete one. Therefore, for all $K > 0$, there exists a discrete model that is within benchmark K to the benchmark model, so a cascade occurs almost surely.

Other Distance Concepts. There are some also distance concepts not consistent with weak convergence. With these distances, the distance between a discrete and a continuous model can be infinity. An interesting example is relative entropy, which was adopted by Hansen and Sargent (2001). The following corollary shows that with relative entropy distance, an information cascade emerges almost surely when individuals are sufficiently ambiguous (but still with a finite degree of ambiguity).

Corollary 1.2. *Under Assumption 1.4 and that d is the relative entropy, where*

$$d(F,G) = \int_{\frac{1}{\gamma}}^{\gamma} \log\left(\frac{dF(\lambda)}{dG(\lambda)}\right) dF(\lambda),$$

then there exists some finite number, $\overline{K}$, such that when signals are bounded, an information cascade occurs $\mathbb{P}^*$-almost surely as long as $K > \overline{K}$.

The proof makes use of condition (2) in Theorem 1.2. One can show that there exists some continuous data-generating process, F, satisfying condition (2) and $d(F,G) < \infty$. Then, we simply need to set $\overline{K} = d(F,G)$. The following example depicts one such data-generating process (detailed analysis is provided in the Appendix).

Example 3. [A continuous DGP that induces a cascade] For simplicity in exposition, I deal with the nominal signal space $S = [0,1]$. Consider the following h:

$$h^1(s) = \begin{cases} 1 + 2\varepsilon(1+\gamma)\cdot s & s \in \left[0, \frac{1}{1+\gamma}\right] \\ 0 & s \in \left(\frac{1}{1+\gamma}, \frac{\gamma}{1+\gamma}\right), \\ 2\varepsilon(1+\gamma)\cdot s + \gamma - 2\varepsilon(1+\gamma) & s \in \left[\frac{\gamma}{1+\gamma}, 1\right] \end{cases} \qquad h^0(s) = h^1(1-s)$$

where $0 < \varepsilon < \frac{\gamma-1}{2\gamma+2}$. After making the transformation, $\lambda = \frac{h^1(s)}{h^0(s)}$, we can express each signal s in terms of likelihood ratios, λ. It can be seen that likelihood ratio $\lambda \in \left[\frac{1}{\gamma}, \frac{1+2\varepsilon}{\gamma-2\varepsilon}\right] \cup \left[\frac{\gamma-2\varepsilon}{1+2\varepsilon}, \gamma\right]$. When ε is smaller, the signals (λ's) are more concentrated around the two tails, meaning that the data-generating process is more "precise". When ε is sufficiently small, condition (2) is satisfied. *As long as the belief set, $\mathcal{F}_0$, contains one such data-generating process, an information cascade occurs almost surely*. Under the assumptions of benchmark G, it is easy to see that a model set with a finite radius includes such a data-generating process. $\square$

1.6 Incorrect Herding with Unbounded Signals

The last section shows that when signals are bounded, the non-cascade results in the literature can be fragile with respect to ambiguity. Similar non-robustness also exists when signals are unbounded. Recall that Smith and Sørensen (2000) showed that complete learning occurs whenever signals are unbounded. However, this section shows that complete learning may collapse even if there is a slight degree of ambiguity.

1.6.1 Incomplete Learning under Ambiguity

To be more specific, under ambiguity, incomplete learning will arise in the form that individuals can herd on the incorrect action with a strictly positive probability. The following theorem provides conditions for an incorrect herding to arise.

Theorem 1.3. *Suppose that for all i, $\overline{F}_i^0(x) \leq ax^\alpha$ with $a, \alpha > 0$ as $x \to 0$. If there exists some $F \in \mathcal{F}_0$ such that $x^p = o\left(F^0(x)\right)$ as $x \to 0$ for some $p \in (0, \alpha)$, herding occurs $\mathbb{P}^*$-almost surely, and an incorrect herding occurs with a $\mathbb{P}^*$-strictly positive probability.*

The condition $x^p = o\left(F^0(x)\right)$ means that the tail of $F^0(x)$ is sufficiently fat, or intuitively, it can be interpreted as that $F^0(x)$ is sufficiently informative. Therefore, Theorem 1.3 conveys two messages: first, when individuals perceive some adequately informative DGP, herding occurs with probability 1, so it is not possible for actions to oscillate; second, an incorrect herding occurs with a strictly positive probability, so complete learning does not hold. Theorem 1.3 can be viewed as a parallel statement of Theorem 1.2 when signals are unbounded. Theorem 1.2 implies that non-cascade is not robust in many cases, and similarly, Theorem 1.3 also implies that complete learning is also not robust in some interesting cases. Corollary 1.3 presents one possibility.

Corollary 1.3. *Suppose that signals are i.i.d. with* $\overline{F}^0(x) = O(x^\alpha)$ *with* $\alpha > 0$ *as* $x \to 0$*. If there exists some* $F^0 \in \mathcal{F}_0$ *such that* $F^0 = O\left(x^{\alpha-\varepsilon}\right)$ *with* $\varepsilon \in (0, \alpha)$ *as* $x \to 0$*, then herding occurs* $\mathbb{P}^*$*-almost surely, and an incorrect herding occurs with a strictly positive probability.*

Corollary 1.3 says that any small ambiguity in the order of F^0 will destroy complete learning learning. Notice that Corollary 1.3 does not impose other restrictions on the belief set, so individuals can also perceive data-generating processes with an order greater than α. But as long as any data-generating process with an order less than α is considered possible, the society will not achieve complete learning. Below is a concrete example.

Example 4. [Ambiguity in Model Parameter] Consider the signal space $\mathcal{S} = (0,1)$, signals are i.i.d. and the data-generating process takes the form of $g_m = \left(g_m^0, g_m^1\right)$, where

$$g_m^0(s) = (m+1)(1-s)^m \text{ and } g_m^1(s) = (m+1)s^m \qquad \text{for } s \in (0,1).$$

It is easy to see that signals are unbounded under g_m (in terms of likelihood ratios). Suppose that the true data-generating process is g_{m_0}. Individuals are ambiguous about the true parameter m_0 and perceive a set $M_\varepsilon = [m_0 - \varepsilon, m_0 + \varepsilon] \subset \mathbb{R}_+$. When $\varepsilon = 0$, there is no ambiguity, so complete learning occurs. However, for *all* $\varepsilon > 0$, complete learning does not occur, and the society will settle on an incorrect action with a strictly positive probability. This implies that any small vagueness around the true parameter can lead to incomplete learning. $\square$

Remark 1.2. *The collapse of complete learning can imply a discontinuity in learning efficiency. For instance, suppose that $m_0 > 1$. When there is no ambiguity, the expected number of incorrect actions is finite as shown by Rosenberg and Vieille (2019), so learning is efficient according to their definition. In contrast, for all $\varepsilon > 0$, an incorrect herding occurs with a strictly positive probability, so the expected number of incorrect actions becomes infinite.*

1.6.2 Sketch of the Proof.

For simplicity in illustration, I sketch the proof under the assumptions of Corollary 1.3. The complete learning in Smith and Sørensen (2000) is a result of the martingale convergence theorem. However, under model ambiguity, posteriors are no longer martingales, so we cannot apply the same martingale technique. This paper adopts a new approach to analyze learning with unbounded signals. The proof consists of the following three steps.

(i) Incorrect herding $p > 0$. The most important step is to show that an incorrect herding occurs with a strictly positive probability. We note that the probability of an incorrect herding is

$$\lim_{i\to\infty} \mathbb{P}^*\left(a_1 = ... = a_i = 1\right) = \prod_{i=1}^{\infty}\left(1 - \overline{F}^0\left(1/r_i\right)\right),$$

where r_i denotes the average public likelihood after $h_i = (1, ..., 1)$. The probability is positive if and only if the infinite series $\sum \overline{F}^0\left(1/r_i\right) < \infty$, which is equivalent to $\sum\left(\frac{1}{r_i}\right)^{\alpha} < \infty$, where α is the order of $\overline{F}^0$ near the neighborhood of 0. Under the assumption of Corollary 1.3, we can find an $\varepsilon \in (0, \alpha)$ such that the growing speed of r_i satisfies $r_{I+t} \geq (t+1)^{\frac{1}{\alpha-\varepsilon}}$ for some I sufficiently large and for $t \geq 1$. It implies that

$$\sum_t \left(\frac{1}{r_{I+t}}\right)^{\alpha} \leq \sum_t \frac{1}{(t+1)^{\frac{\alpha}{\alpha-\varepsilon}}} < \infty,$$

which establishes that an incorrect herding occurs with a strictly positive probability. To get a sense of how the non-robustness arises, notice that if there is no ambiguity, $\varepsilon = 0$, so the infinite series on the RHS becomes $\sum \frac{1}{t+1}$, which diverges and corresponds to the absence of an incorrect herding. However, when $\varepsilon > 0$ even if it is arbitrarily small, the infinite series becomes convergent, so an incorrect herding arises.

(ii) Herding occurs w.p.1. A symmetric argument implies that a correct herding occurs with $p > 0$. It can be further shown that the probabilities of both herding has a *uniform* lower bound when average likelihood ratio r_i is sufficiently large or small, i.e., both states are locally stable. In other words, for all possible history h_i and for all i, the probability that herding eventually occurs is uniformly bounded from below, which implies that herding must occur almost surely.

1.7 Conditions for Complete Learning under Ambiguity

The previous discussion shows that complete learning can easily collapse under ambiguity. One natural question is that: when does complete learning hold under ambiguity (if it is possible)? This section presents a **necessary and sufficient condition** for complete learning within the class of data-generating processes that have power tails.

Definition 6. A data-generating process $F \in \mathcal{F}$ has a *power tail* if there exists some $\alpha > 0$ such that $F^0(x) = O(x^\alpha)$ as $x \to 0$. The power of F, denoted by $\mathcal{P}(F)$, is defined to be α.

A data-generating process has a power tail if it can be approximated by a power function around its tail. It is easy to see that a power-tail data-generating process is unbounded. The power provides an intuitive measure of the signal informativeness, where a larger power means that signals are less informative (around the tails). To see this, note that if F has a larger power, it means that its tails are thinner, so it generates high-precision signals with a lower probability. This section focuses on the power-tail models and imposes the following assumptions.

Assumption 1.5. *$\overline{F}$ has a power tail, and $\mathcal{F}_0$ only contains models with power tails.*

Assumption 1.6. *$\mathcal{F}_0$ is finite where each $F \in \mathcal{F}_0$ has a different power and is differentiable.*

Assumption 1.5 says that the true DGP has a power tail, and individuals only perceive DGPs with power tails. Assumption 1.6 is imposed for simplicity and the discussion can be extended to situations where $\mathcal{F}_0$ is infinite and some models may have identical powers. Theorem 1.4 provides a necessary and sufficient condition for complete learning under these two assumptions.

Theorem 1.4. *Under Assumptions 1.5 and 1.6, complete learning occurs **if and only if** $\mathcal{F}_0$ satisfies*

(i) for all $F \in \mathcal{F}_0$, we have $\mathcal{P}(F) \geq \mathcal{P}(\overline{F})$, and

(ii) there exists some $F \in \mathcal{F}_0$ such that $\mathcal{P}(F) < \mathcal{P}(\overline{F}) + 1$.

Theorem 1.4 says that to establish complete learning, we need to impose restrictions from two directions. On one hand, all perceived models can not be too informative. Specifically, they must be less informative than the true model in the sense they must have higher powers. On the other hand, some perceived model has to be sufficiently informative such that its power does not exceed that of the true model by 1. Before explaining the intuition, it helps to see what will happen if the conditions in Theorem 1.4 are violated.

Corollary 1.4. *Under Assumptions 1.5 and 1.6, (i) if there exists some $F \in \mathcal{F}_0$ such that $\mathcal{P}(F) < \mathcal{P}(\overline{F})$, an incorrect herding occurs with a strictly positive probability; (ii) if for all $F \in \mathcal{F}_0$, $\mathcal{P}(F) \geq \mathcal{P}(\overline{F}) + 1$, actions do not converge with probability 1.*

First, when individuals perceive some highly informative model, learning is incomplete and takes the form of incorrect herding. Second, when all models considered by individuals are insufficiently informative, incomplete learning takes the form of action non-convergence. The first case has been explained. The second case comes from an intuitive argument that if individuals underestimate the informativeness of their predecessors, any herd will be overturned frequently, and as a consequence, the actions will not converge.

Intuition of Theorem 1.4. In order to achieve complete learning, we must exclude two sources of incomplete learning—incorrect herding and action non-convergence. Correspondingly, we also need to restrict $\mathcal{F}_0$ from two directions. To prevent incorrect herding, $\mathcal{F}_0$ must not contain highly informative data-generating processes, which correspond to Theorem 1.4 (i). To prevent action non-convergence, $\mathcal{F}_0$ must not only contain data-generating processes that are too uninformative, which Theorem 1.4 (ii).

1.8 Other Ambiguity Preferences

Many results of this paper can be extended to a wider class of ambiguity preferences. This section focuses on two important examples, the α-max-min EU preference and the smooth ambiguity preference. The discussion shows that results under the max-min EU preference are not as extreme as it appears.

1.8.1 $\alpha-$Max-Min EU Model

Consider first the case where individuals hold the α-maxmin expected utility ($\alpha-$MEU) preferences (Hurwicz, 1951; Ghirardato et al., 2004). With this class of preferences, individual i's utility is

$$V_i(a) = \alpha \cdot \inf_{\pi \in \Pi_i} \mathbb{E}_\pi U(a, \theta) + (1 - \alpha) \cdot \sup_{\pi \in \Pi_i} \mathbb{E}_\pi U(a, \theta)$$

where $\alpha \in [0,1]$. Here α represents the degree of an individual's pessimism, where $\alpha = 1$ corresponds to the MEU model, and $\alpha = 0$ corresponds to the max-max expected utility model.[13] Within this class of models, we have a surprising result as follows.

Proposition 1.2. *When individuals have α-MEU preferences,* ***all*** *previous results hold for all $\alpha \in [0,1]$.*

Proof. It can be verified that the equilibrium strategy takes the identical form as in the MEU case. □

Proposition 1.2 shows that this paper's results hold for all ambiguity attitudes captured by the $\alpha-$MEU model. The intuition is less surprising than it appears. Notice that the key force is the asymmetry between the models encouraging and discouraging a cascade, but the asymmetry does not rely on ambiguity aversion, for example, we can still apply the similar arguments under Theorem 1.1 when individuals have max-max EU preference. It is worth noting that Proposition 1.2 relies on the binary choice structure, and if we allow for multiple actions, ambiguity attitudes affect which action the society settles on, e.g., individuals will herd on a safer action when they are ambiguity-averse but not when they are ambiguity-loving.

[13]Notice that $\alpha = \frac{1}{2}$ does not represent the expected-utility model—it is only that individuals attach the same weight to the best and worst scenario.

1.8.2 Smooth Ambiguity Model

Consider next that individuals hold the smooth ambiguity model as axiomatized by Klibanoff et al. (2005).

$$V_i(a) = \phi^{-1}\left(\int_{\Pi_i} \phi\left[\mathbb{E}_\pi U(a,\theta)\right] d\mu(\pi)\right).$$

where μ stands for the second-order belief, and the curvature of the ϕ function describes the ambiguity attitude. The analysis under the smooth model becomes more difficult as the equilibrium strategy no longer has a simple characterization. However, many results can still hold qualitatively when individuals are sufficiently **ambiguity sensitive**. Below is an example.

Example 5. Let's consider the setup Example 1. Recall that in Example 1, where $S = \{h, l\}$, and the data-generating process has precision γ. Each individual's signal precision $\gamma \in (1, \overline{\gamma}) = \Gamma$ is further drawn from a full-support distribution, h. Signals are bounded, so $\overline{\gamma} < \infty$. Every individual has the following preference.

$$V_i(a) = \left[\int_{\Gamma^{i-1}} \left[\mathbb{E}_{\gamma_1,\ldots,\gamma_{i-1}} u(a,\theta)\right]^{1-\sigma} dh(\gamma_1, \ldots, \gamma_{i-1})\right]^{\frac{1}{1-\sigma}}.$$

This preference is a reformulation of the preference with constant relative ambiguity aversion (CRAA). Parameter σ represents the ambiguity attitude, where $\sigma = 1$ corresponds to ambiguity neutrality, and as σ grows, individuals are becoming more ambiguity-averse.

(i) *When $\sigma = 0$, whether an information cascade can occur depends on the properties of h.*

Note that $\sigma = 0$ corresponds to the situation where individuals have EU preference and hold a correct model perception, h, so the occurrence of a cascade depends h.

(ii) *When $|\sigma|$ is sufficiently large, an information cascade occurs with a strictly positive probability for all possible h.*

As $\sigma \to +\infty$, the preference approaches the max-min EU model, and as $\sigma \to -\infty$, it approaches the max-max EU model. For any fixed $I < \infty$, the continuity in preference implies that belief dynamics before individual I can be arbitrarily close to to those under the max-min or max-max EU when $|\sigma|$ is sufficiently large, which also implies the occurrence of a cascade. More details can be found in Appendix A.1.8. $\square$

The example suggests that information cascades are not unique to the max-min EU model. It also emerges under the smooth ambiguity preference when individuals are sufficiently ambiguity sensitive. Example 5 only looks at the bounded-signal case, but it can be extended to unbounded signals with some restrictions on h (see Appendix A.1.8).

1.9 Discussion: Bayesian vs Ambiguity

From the previous discussion, we know an information cascade occurs under sufficient model ambiguity. This section summarizes how results change if we switch to the Bayesian case.

1.9.1 Bayesian Model Certainty

When individuals are Bayesian and certain about the data-generating processes, the learning outcome depends on the features of the model specifications. When individuals hold a **correctly specified** model perception, the learning outcomes depend on the details of the true data-generating processes as discussed earlier. When individuals hold a **incorrectly specified** model perception, the learning outcomes depend on the interplay between the true and perceived data-generating processes. For example, if individuals overestimate other people's informativeness, the society may herd on the incorrect action; if they underestimate the informativeness, the society may fail to settle on an action as implied by Corollary 1.4.

1.9.2 Bayesian Model Uncertainty

Another case is that individuals are Bayesian and uncertain about the data-generating processes. In this case, the learning outcome depends on the **priors** over the model space. Below are two examples.

Example 6. The model space is $\mathfrak{F} \equiv \mathcal{F}^{\infty}$, where a typical element $\boldsymbol{F} = (F_1, F_2, ...)$ describes a list of all data-generating processes. Individuals hold an identical prior $Q \in \Delta(\mathfrak{F})$. All signals are i.i.d. and unbounded. The true data-generating process sequence is denoted by $\overline{\boldsymbol{F}}$.

(i) *If* $Q\left(\overline{\boldsymbol{F}}\right) > 0$*, complete learning occurs almost surely.*

In other words, if the prior puts a strictly positive probability on $\overline{\boldsymbol{F}}$, or it contains a "grain of the truth", complete learning occurs. This is because when the true model path is assigned a strictly positive probability, limit beliefs will "merge to" the beliefs induced by the true models (Kalai and Lehrer, 1993). Note that if individuals knew the true model, complete learning would occur. As a consequence, it can be verified that complete learning also occur if the prior contains a "grain of the truth". The more rigorous proof can be found in Appendix A.1.9.

(ii) *If* $Q\left(\overline{\boldsymbol{F}}\right) = 0$*, complete learning may not occur.*

If the "grain-of-the-truth" condition fails, limit learning can be different from that under the true model. One example is that Q features an independent distribution across all individuals. In this case, Q is not updated after observations, so the problem becomes learning under model certainty, where the model perception is $F_Q = \mathbb{E}_Q F$. From the previous discussion, learning outcome depends on F_Q and its relation with $\overline{F}$, so it is possible that complete learning does not occur (Appendix A.1.9 provides an example). □

Remark 1.3. *Example 6 shows that when signals are unbounded, the occurrence of complete learning depends on the priors. Similarly, when signals are bounded, the occurrence of a cascade also depends on the priors. Appendix A.1.9 presents one such example. In the example, there are two models, where the first implies a cascade but the second does not. This example shows that a cascade occurs if the prior assigns a sufficiently large weight to the first model but does not occur if the second model is assigned a large weight.*

1.9.3 Ambiguous Model Uncertainty

In Bayesian learning, the learning outcome differs with individuals' model perceptions or priors, so we do know which outcome will arise without the knowledge of priors or model perceptions. This study complements the literature by showing that when individuals consider several models simultaneously, an information cascade can almost surely occur under sufficient ambiguity. The results are driven by the following differences.

1. Ambiguity. Under ambiguity and with max-min EU preference, individuals are unable to assign probabilities, and their decisions are determined by comparing the worst payoffs. In this case, an information cascade arises because of the following asymmetric effects—the worst case for breaking a herd is when the predecessors have precise data-generating processes, but the worst case for following the herd is only when predecessors have imprecise information. The first possibility encourages a cascade whereas the second discourages it. Under sufficient model uncertainty, the encouraging effect outweighs the discouraging effect, so individuals will follow the herd to avoid the larger ambiguity from acting against it, hence an information cascade arises.

2. Bayesian. When individuals are Bayesian, their priors play an important role in determining the learning outcome. It is true that highly informative data-generating processes encourage an information cascade, and the encouraging effect cannot be offset by other data-generating processes in terms of their induced payoffs. However, individuals make decisions based on the payoffs weighted by the posteriors on models, which depend on how their priors are formed. For some priors, the posteriors on these informative data-generating processes will become very small or even zero in the limit. In these cases, the encouraging effect is so small such that an information cascade or an incorrect herding does not occur. Similarly, for some other priors, the encouraging effect can be very large such that a cascade occurs.

Remark 1.4. *Under model uncertainty, the Bayesian approach assumes a **prior** over models, based on which individuals can pin down a unique posterior, then a decision can be made according to that posterior. One challenge is that the learning outcome depends on the priors, and it is difficult to test the priors individuals hold or to justify which priors should be imposed in the model. Differently, the ambiguous approach allows individuals to hold multiple posteriors but imposes an **ambiguity preference** in the decision making. It turns out that when people are sufficiently ambiguity sensitive, and when there is sufficient model uncertainty, an information cascade will occur.*

1.10 Conclusion

This paper studies sequential learning under the assumption that individuals face ambiguity about other people's data-generating processes. This paper finds that under sufficient ambiguity, an information cascade arises with probability 1. Interestingly, the results that feature non-cascades may only represent a knife-edge case from the perspective of ambiguity. This paper provides a new perspective into the mechanism behind cascades and herding. The paper adopts the most standard setup, where there are binary states, binary actions. The Appendix A.2 shows by examples that qualitative results hold with multi-state, multi-action and general updating rules (i.e., α-maximum likelihood updating). This paper also relies on the linear network, so an interesting direction is to extend the result to general networks.

CHAPTER 2

BIASED LEARNING UNDER AMBIGUOUS INFORMATION

2.1 Introduction

It is standard in the learning literature to assume that individuals are rational and update beliefs using Bayes rule. However, a growing body of evidence suggests that individuals often hold *biased* beliefs and process information in a self-serving manner. For example, individuals tend to overinterpret others' compliments to maintain a good self-image, leading to the famous "better-than-average" effect (e.g. Taylor and Brown, 1988; Alicke and Govorun, 2005). Citizens holding different political views and receiving identical information tend to process it differently to support their positions, which can lead to polarization in beliefs (e.g Taber and Lodge, 2006) . In financial markets, many investors interpret evidence to support their current investment philosophy (e.g Graham et al., 2009; Grinblatt and Keloharju, 2009; Chang and Cheng, 2015). Such biases also exist for analysts, and it is well-documented that analysts tend to overreact or underreact to new information and exhibit biases when issuing reports (e.g. Butler and Lang, 1991; Brous and Kini, 1993; Easterwood and Nutt, 1999).

Beliefs of biased individuals typically exhibit non-Bayesian dynamics, so standard techniques in Bayesian learning cannot be directly applied to this situation.[1] It remains unanswered how beliefs evolve and how learning outcomes differ when we account for individuals' biases. This paper provides a framework that enables analysis of biased learning. In the paper, I investigate a standard learning problem where individuals observe a sequence of signals and

[1] See Blume et al. (1998) for a survey on Bayesian learning in economics.

update beliefs about the true state of the world. Different from the standard framework, (i) individuals are ***ambiguous*** about the true data-generating process and consider a set of models as possible, and (ii) individuals are ***biased*** toward some states of the world and interpret signals according to the model that best supports their biases. The ***biased updating rule*** captures the idea that individuals process information in a self-serving manner. The major deviation from Bayes rule is that under the biased rule, individuals can use different models for different signals. This feature allows the biased rule to exhibit interesting non-Bayesian dynamics. Below is one example.

Example 7 (Good News or Bad News)**.** The state of the world is summarized by $\Theta = \{H, L\}$. Individuals receive signals from the signal space $S = \{S_H, S_L\}$. Signals can either be informative or uninformative in the sense that

$$\frac{P\left(S_H|H\right)}{P\left(S_L|H\right)} = \frac{P\left(S_L|L\right)}{P\left(S_H|L\right)} \in \{1, \alpha\} = \mathcal{A} \quad \text{where } \alpha > 1.$$

If the likelihood ratio is α, signals are informative; if the likelihood ratio is 1, signals are uninformative. Individuals are biased toward one of the two states. After receiving signals, they update beliefs by Bayes rule according to the signal interpretation that maximizes the probability of their biased states. It can be verified that individuals will process good news and bad news in an *asymmetric* way. Consider an individual with bias H; if he received a signal S_H (i.e., the good news), he would view this signal as informative hence adopt model α. On the other hand, if he received a signal S_L (i.e., the bad news), he would treat it as uninformative by adopting model 1. In summary, individuals tend to value good news but overlook bad news. This "good-news effect" has been documented by many experimental findings (e.g. Eil and Rao, 2011; Ertac, 2011; Mobius et al., 2011; Coutts, 2019; Barron, 2021) but is not well reconciled by a standard Bayesian framework that requires *consistent* interpretations for all signal realizations. □

In this paper, I examine a more general learning environment in which individuals face a general state space and signal structures and can be biased toward a weighted set of states. Individuals observe a sequence of i.i.d. signals, but they are ambiguous about the true data-generating process. In other words, individuals perceive a set of feasible models, $\mathcal{A}$, and do not assign probabilities to models. After observing the signals, individuals update beliefs via Bayes rule according to the feasible model that can best justify their endowed biases. This learning framework satisfies Bayes rule model-wise, but it can exhibit non-Bayesian features at the aggregate level. As a result, the model nests Bayes rule as a special case, and it can also explain some behavioral patterns inconsistent with Bayesian framework.

In addition to its new implications, this rule preserves tractability. As information accumulates, posteriors approach a limit that can be characterized in a simple form. This paper provides characterizations of limit beliefs under the biased rule (in Theorem 2.1 and Theorem 2.2). The characterizations are based on the seminal work of Berk (1966), and they can be summarized as follows. As the number of signals approaches infinity, (i) individuals will eventually adopt models that minimize a normalized version of relative entropy of their biased states, and (ii) limit beliefs will concentrate on the states that minimize the relative entropy under these models.

Based on these characterizations, I then discuss some effects of ambiguity on learning. First, ambiguity can lead to incomplete learning and polarization. Under sufficient ambiguity, individuals have enough flexibility in interpreting information to justify their biases or settle on states far from the true state in the limit. On the contrary, as ambiguity diminishes, correct learning can be restored in a weak sense, that is, when the state space is finite, correct learning can be achieved under sufficiently small ambiguity. Second, ambiguity can lead to overconfidence. Under ambiguity, perceived models can complement each other during the learning, so a biased individual can become strictly more confident toward their biases than a Bayesian individual with any feasible model perception. The intuition resembles that of Example 7. Model ambiguity enables individuals to exploit good news and hedge against the bad news by choosing to interpret signals asymmetrically, so individuals could justify their biases to a larger extent than if they only perceived a specific model. Moreover, this paper shows that the overconfidence can persist even with arbitrarily many signals, which highlights that the biased rule can be different from Bayesian rule even in the limit.

2.1.1 Related Literature and Contributions

The key contribution of the paper is to propose an updating rule under ambiguity that accounts for people's biases. First, in terms of topics, this paper belongs to the literature on learning under behavioral biases. The most relevant bias is self-serving bias, which means that individuals interpret information in a self-serving manner. This type of bias has been investigated in many experimental studies (e.g. Babcock et al., 1995, 1996; Haisley and Weber, 2010; Deffains et al., 2016), but its effects on learning have not been systematically studied in theoretical works. This paper contributes to the literature by proposing a framework to study the biased learning. Methodologically, this paper employs techniques from misspecified learning literature, along with the idea of model uncertainty, to model biased learning. This approach allows us to derive complete characterizations of limit beliefs for a variety of biases, which are both applicable and new to the literature. In addition to self-serving bias, another relevant bias is confirmatory bias, which is more widely studied along with learning (e.g. Rabin and Schrag, 1999; Fryer Jr et al., 2019). This thread of literature is similar to this paper in that information is processed in a biased manner, but it is different in how the bias affects learning. Under confirmatory bias, individuals interpret information toward the most likely state according to their *current* beliefs, and they are not intrinsically attached to any state. Under the bias in this paper, individuals attach intrinsic value to some state, and they update beliefs to justify it regardless of their current beliefs. The former bias is more motivated by the fact that people exhibit some inertia to their prior judgments, whereas the latter is more motivated by the fact that people seek to maintain self-esteem. Due to these differences, they often lead to different belief dynamics and have different asymptotic properties.

Second, this paper also adds to the literature on learning under ambiguity and model uncertainty, by suggesting a new approach to belief updating.[2] The biased updating rule differs from two well-studied updating rules under ambiguity, the full Bayesian rule and the maximum likelihood rule, in the following aspects. The full Bayesian rule updates all models indiscriminately and leads to a set of posteriors, whereas both the maximum likelihood rule and the biased updating rule only update models that satisfy some criterion, and they often lead to a unique posterior. Their main difference is that the maximum likelihood rule follows an objective criterion and selects models according to

[2] Gilboa and Marinacci (2016) and Machina and Siniscalchi (2014) provide exhaustive surveys on ambiguity-related topics. Below is an incomplete list of research on model uncertainty (Marinacci, 2002, 2015; Marinacci and Massari, 2019; Battigalli et al., 2015, 2019a,b; Chen, 2019).

the probability of generating observed information, whereas the biased rule follows a subjective criterion and selects models according to whether they can maximally support the endowed bias. Under appropriate decision rules, these differences can also lead to different actions as shown in Section 2.7.

Finally, under model uncertainty, individuals will inevitably perceive some incorrect models, so this paper also contributes to the growing literature on learning with misspecified models.[3] This thread of literature mostly adopts Bayes rule, so this paper differs from that literature in the same way as it differs from Bayes rule. The most significant difference is that under the biased updating rule, individuals can select models to accommodate their biases, whereas in the misspecification literature, individuals mostly adhere to a fixed model perception. Due to this difference, this paper's framework can produce phenomena that are inconsistent with misspecified Bayesian learning.

This paper is organized as follows. Section 2.2 presents the benchmark model. Section 2.3 presents some examples that illustrate applications of the biased rule. Section 2.4 discusses major modeling assumptions. Section 2.5 and 2.6 characterize limit beliefs under the biased rule. Section 2.7 discusses some examples where actions are involved. Section 2.8 discusses some extensions. The Appendix includes the omitted proofs and supplementary materials.

[3]An incomplete list of learning under incorrect model include Blume and Easley (1982), Nyarko (1991), Bohren (2016), Fudenberg et al. (2017, 2021), Heidhues et al. (2018), Bohren and Hauser (2021), and Frick et al. (2020a,b).

2.2 Benchmark Model

The state space Θ is a compact subset of a Polish space which is endowed with Borel σ-algebra $\mathcal{B}_\Theta$ and a finite measure m. The true state is denoted by $\theta^* \in \Theta$. Individuals do not know the true state and hold a prior that is specified by some density function $\mu_0 : \Theta \to \mathbb{R}_+$ with respect to m, where μ_0 is continuous and has full support. Individuals receive a sequence of i.i.d. signals $\{s_t\}$, where s_t is a random variable taking values in the signal space S. The signal space S is a Polish space associated with Borel σ-algebra $\mathcal{B}_S$ and a σ-finite measure v. Signals are generated by a ***model*** that belongs to the model space $\mathbb{A}$, where $\mathbb{A}$ is also a Polish space. Conditional on state θ, each model $\alpha \in \mathbb{A}$ induces a signal distribution $f(s|\theta, \alpha)$, which is a density function on S with respect to v. Denote by α^* the true model, so $f(s|\theta^*, \alpha^*)$ represents the true distribution of s_t. Denote by $\mathbb{P}^*$ and $\mathbb{E}^*$as the true probability measure and expectation operator induced by the true signal distribution. Below are some technical assumptions on f:

Assumption 2.1. *The mapping $f : S \times \Theta \times \mathbb{A} \to \mathbb{R}_+$ is jointly continuous, and for all $(\theta, \alpha) \in \Theta \times \mathbb{A}$, the support of $f(s|\theta, \alpha)$ is S.*

Assumption 2.2. *For all $\theta \in \Theta$ and $\alpha \in \mathbb{A}$, there exists an open set U containing (θ, α), such that $\mathbb{E}^* \sup_{(\theta,\alpha)\in U} \log^2 f(s|\theta, \alpha) < \infty$.*

Assumption 2.3. *For all $\theta, \theta' \in \Theta$ and $\alpha \in \mathbb{A}$, $\mathbb{P}^* (\{s \in S : f(s|\theta, \alpha) \neq f(s|\theta', \alpha)\}) > 0$.*

These assumptions are standard and can accommodate a wide class of signal distributions. The purpose of Assumption 2.2 is to impose some boundedness conditions, which allow us to apply the dominated convergence theorem to establish continuity. Assumption 2.3 requires that every model induces different signal distributions under different states, so signals are informative under every perceived model.[4]

Under the standard learning framework, individuals are certain about the true model, and they update beliefs according to their perceived true models, where the perceived models can be either correctly or incorrectly specified. Different from the standard model, this paper assumes that individuals are ***ambiguous*** about the true model. More precisely, individuals consider a compact set of models $\mathcal{A} \subset \mathbb{A}$ as possible and are unable to assign probabilities on $\mathcal{A}$. Intuitively, this means that individuals are uncertain about how to interpret signals, and they only know that the correct interpretation belongs to $\mathcal{A}$. The model set is *correctly specified* if $\alpha^* \in \mathcal{A}$, and is *misspecified* if $\alpha^* \notin \mathcal{A}$.

Every individual is endowed with some bias over Θ, where the ***bias*** is represented by a payoff function $\tau : \Theta \to \mathbb{R}_+$. The support of τ, denoted by $\text{supp}\,(\tau) = \overline{\{\theta : \tau(\theta) > 0\}}$, is referred to as the set of ***biased states***, which are states with strictly positive payoff.[5] Denote by $\mathcal{T}$ the set of all possible biases. I assume that every $\tau \in \mathcal{T}$ is a continuous function on its domain. Every individual derives utility from his belief, called ***belief utility***, which is equal to the expected payoff from τ based on his belief. If biases are confirmed by beliefs to a larger extent, individuals can derive higher utility. More precisely, for all $\alpha \in \mathcal{A}$, denote by $U_t(\tau|\alpha)$ the time-t belief utility if individuals update according to

[4]Without this assumption, we may encounter a trivial situation where learning cannot occur on a subset of states, in which case how the prior μ_0 is formed on this set becomes vital. It needs to be noted that Assumption 2.3 is not required for the characterizations of limit beliefs in Theorem 2.1 and 2.2 but is needed to establish the overconfidence effect in Theorem B.1 in the Appendix.

[5]For simplicity, this paper focuses on positive-valued τs , but the discussion can be extended to negative-valued τ's, where individuals can also deflate the probability of some states.

model α by Bayes rule, that is,

$$U_t\left(\tau|\alpha\right) \equiv \int_{\Theta} \tau\left(\theta\right) \mu_t\left(\theta|\alpha\right) dm\left(\theta\right),$$

where $\mu_t\left(\theta|\alpha\right)$ denotes the Bayes update of $\mu_0\left(\theta\right)$ according to model α at time t. When updating beliefs, individuals adopt the ***biased updating rule*** in the sense that they will update according to the model which generates the highest belief utility, denoted by $U_t\left(\tau\right)$. More formally, denoting by μ_t^τ the time t belief held by an individual with bias τ, we have:

$$\forall \theta \in \Theta: \quad \mu_t^\tau\left(\theta\right) = \mu_t\left(\theta|\alpha_t^\tau\right) \quad \text{where } \alpha_t^\tau \in \arg\max_{\alpha \in \mathcal{A}} U_t\left(\tau|\alpha\right) \tag{2.1}$$

Here, α_t^τ is referred to as the ***model perception*** of an individual with bias τ at time t. Notice that it is possible that multiple models can maximize the belief utility. In this case, I assume that individuals adopt some rule to break the tie, so our biased updating rule always generates a unique belief. Consider again the case in Example 7, the bias-H individual's bias can be described by $\tau\left(\theta\right) = \delta_G\left(\theta\right)$, so he will update his beliefs to maximize the probability of state H. By appropriate definition of τ, this model can accommodate a variety of biases. Section 2.3 discusses some examples.

2.3 Illustrative Examples

In this section, I discuss some examples which illustrate various applications of the biased rule in economics. Some basic properties of the biased rule are also discussed in these examples.

Example 8 (Overconfidence and Excess Entry)**.** A group of candidates $N = \{1, ..., n\}$ decide whether to apply for a prize that will be awarded to the top candidate. Candidates do not know their rankings and they start with a common prior. An ambiguous feedback is then generated by a model in $\mathcal{A}$ and is observed by all candidates. Every candidate is biased and trying to justify himself as the best. Denote by $\mu^i(i)$ as the probability that candidate i thinks that he is the best candidate. In non-trivial cases, we have:

$$\sum_{i \in N} \mu^i(i) = \sum_{i \in N} \max_{\alpha \in \mathcal{A}} \mu(i|\alpha) > \max_{\alpha \in \mathcal{A}} \left[\sum_{i \in N} \mu(i|\alpha) \right] = 1.$$

The sum of the self-perceived probability of being the top candidate is *strictly* greater than 1, which features overconfidence or the "better-than-average" effect. In the case of Bayesian updating with a common model perception, this effect cannot be naturally reconciled due to the fact that the beliefs always sum up to 1.[6] As a result, ambiguous feedback may contribute to a larger volume of applications than non-ambiguous case. This effect can be used to explain excess entry into competitive markets triggered by overconfidence, especially when the performance feedback is ambiguous. For related experimental or empirical evidence, see Heath and Tversky (1991); Camerer and Lovallo (1999); Karelaia and Hogarth (2010). □

Example 9 (Fairness and Preferences for Redistribution-I)**.** A worker has an ability level $\theta \in \{H, L\}$ and is working in some profession. The fairness of the profession is represented by $\alpha \in \{fair, unfair\}$. The worker receives feedback $x_t \in \{G, B\}$ generated by $f(x_t|\theta, \alpha)$, where

$$\begin{cases} f(G|H, fair) = p & f(G|L, fair) = 1 - p \\ f(G|H, unfair) = 1 - p & f(G|L, unfair) = p \end{cases}, \quad \text{where } p > 1/2.$$

[6] Even with heterogeneous model perceptions, the effect is not always true. It is easy to construct an example such that for all possible model-perception combinations $(\alpha_1, ..., \alpha_n)$ in $\mathcal{A}$, the sum of self-perceived probability is less than 1 for some signals, that is, the group may exhibit under-confidence in some cases.

In other words, a high-ability worker is more (less) likely to receive the good feedback G than a low-ability worker when the profession is fair (unfair). Suppose that the true fairness is $\alpha^* = fair$. Assume that $\tau = \delta_H$, so the worker is trying to defend the conjecture that he has high ability. This worker updates beliefs according to the biased updating rule. The following results are easy to verify:

(i) At each period t, if the worker received more good feedback than bad feedback, he would perceive the profession as fair (i.e., $\alpha_t^\tau = fair$), otherwise, he would perceive the profession as unfair (i.e., $\alpha_t^\tau = unfair$).

(ii) As $t \to \infty$, the worker with low ability will almost surely perceive the profession as unfair; the worker with high ability will almost surely perceive the profession as fair.

In this example, high income groups believe that their wealth is due to their ability and low income groups attribute their lower wealth to unfairness. Both groups are biased and may disagree about redistribution plans due to their biases. Therefore, this simple example provides a possible explanation of the phenomenon that low-income groups tend to support redistribution more than high-income groups. Similar effects are supported by recent experimental evidence (e.g. Deffains et al., 2016; Cassar and Klein, 2019). □

Example 10 (Optimistic and Pessimistic Investors-I)**.** The true state of the market can either be good G or bad B. The market consists of two types of investors $\mathcal{T} = \{o, p\}$, where the o-type (i.e., optimists) are biased toward the good state, and the p-type (i.e., pessimists) are biased toward the bad state. Both types of investors hold full-support priors. Investors receive i.i.d. signals overtime.

Signals take values in $S = \{g, m, b\}$ with the true data-generating process being

$$\begin{array}{c|ccc} P(s|\theta) & g & m & b \\ \hline G & \frac{\lambda}{1+\lambda}(1-\varepsilon) & \varepsilon & \frac{1}{1+\lambda}(1-\varepsilon) \\ B & \frac{1}{1+\lambda}(1-\varepsilon) & \varepsilon & \frac{\lambda}{1+\lambda}(1-\varepsilon) \end{array},$$

where $\lambda > 1$ and $\varepsilon \in (0, 1)$. Investors know how to interpret signals g and b in the sense that they know that

$$\frac{P(g|G)}{P(g|B)} = \frac{P(b|B)}{P(b|G)} = \lambda,$$

but they are ambiguous in the interpretation of signal m in the sense that they perceive a set of likelihood ratios induced by m,

$$\frac{P(m|G)}{P(m|B)} \in \left[\frac{1}{1+\delta}, 1+\delta\right],$$

where $\delta \geq 0$ describes the degree of ambiguity. When $\delta = 0$, investors are certain about how to interpret signal m. As δ grows, the uncertainty in interpretation also expands. When δ is sufficiently large, *polarization* of opinions arises almost surely, and we have $\mu_t^o \to \delta_G$ and $\mu_t^p \to \delta_B$ almost surely. In this case, investors become confident in their biased states in the limit, so presenting them with the same information increases instead of decreasing their disagreements. On the contrary, when δ is sufficiently small, polarization disappears and both types of investors will learn the true state correctly.

The intuition is straightforward. When there is sufficient uncertainty, individuals have adequate flexibility to interpret the signals in their most preferred way. Therefore, they can successfully convince themselves that their biased states are true states, which leads to belief polarization. On the contrary, when the degree of uncertainty is low, the room for signal interpretations is also restricted, which leads to complete learning. This example suggests that informational uncertainty can exacerbate biased beliefs, which is also supported by empirical findings (e.g. Ackert and Athanassakos, 1997; Das et al., 1998; Athanassakos and Kalimipalli, 2003; Chang and Choi, 2017). □

Example 11 (Inferring the market fundamental-I)**.** The market price of a commodity is determined by the following demand-and-supply system:

$$\begin{cases} \log Q_t^D = -\alpha^* \log P_t + \varepsilon_t \\ \log Q_t^S = \log P_t + \eta_t \\ Q_t^D = Q_t^S \end{cases} \qquad \text{with } \varepsilon_t \sim \mathcal{N}\left(\theta^*, \sigma^2\left(\theta^*\right)\right) \text{ and } \eta_t \sim \mathcal{N}\left(0, 1\right),$$

where $\alpha^* > 0$ represents the *demand elasticity*, and $\theta^* \in \Theta = \left[\underline{\theta}, \overline{\theta}\right]$ represents the *market fundamental*— a higher θ^* means that the market demand is stronger. An agent can observe the full price history $\{P_1, ..., P_t\}$ up to some period t, and he knows the demand-and-supply relation. Solving for the equilibrium prices, he knows that:

$$\log P_t \sim \mathcal{N}\left(\frac{\theta^*}{\alpha^* + 1}, \frac{1 + \sigma^2\left(\theta^*\right)}{\left(\alpha^* + 1\right)^2}\right).$$

The agent lacks sufficient information to figure out the demand elasticity α^*, and he only knows that $\alpha^* \in \mathcal{A}$, so the true state can not be perfectly detected. If the agent observes a high price, it may that the demand is strong, but it may also be that the high price comes from a low demand elasticity. The agent is endowed with some bias $\tau : \left[\underline{\theta}, \overline{\theta}\right] \to \mathbb{R}_+$ and adopts the biased updating rule to update beliefs. At time t, the agent takes an action to match the true market fundamental based. [7] For example, he selects an r_t to maximize

$$V_t\left(r\right) = -\int_{\underline{\theta}}^{\overline{\theta}} \left(r - \theta\right)^2 \times \mu_t\left(\theta\right) d\theta.$$

The optimal action is $r_t = E_t\theta$, where the expectation is taken under the subjective belief at time t. For instance, an analyst wants to issue a report r_t that can

[7] A brief discussion of the decision rule is Section 2.4.

best reflect the current economic condition, but how he evaluates the economy may be unconsciously influenced by his bias, so he may end up issuing a report that is also biased. The biased reporting of analysts has been documented in many research works, for example, Brous and Kini (1993) and Easterwood and Nutt (1999). By flexibly defining the τ function, we can accommodate a variety of biases. For example, (i) if $\tau = \delta_\theta$, where δ denotes the Dirac delta function, it means that the agent seeks to confirm a specific state θ. (ii) if τ is a constant function, it implies that there is no bias, so the agent adopts the Bayes rule to update his beliefs. (iii) if τ is an increasing (or decreasing) function, the agent is upward (or downward) biased and updates his beliefs to justify a state as higher (or lower) as possible. Unlike previous examples, it is not immediately clear how different biases will lead to different model perceptions and limit beliefs. This example will revisited in Section 2.6. □

2.4 Discussion of the Model

In this section, I discuss some the critical assumptions about the biased learning given in Section 2.2. This section focuses on the two most important elements, the model ambiguity and the bias, and discusses their motivations and interpretations. This section also provides a brief discussion of other modeling assumptions.

2.4.1 Discussion of the Ambiguity

One key element is that individuals are ***ambiguous*** about the true signal-generating process. The presence of model ambiguity can come from multiple sources. From a frequentist perspective, ambiguity can be a result of identification problems. If every model in $\mathcal{A}$ can match the true signal distribution under some state, that is,

$$\forall \alpha \in \mathcal{A}, \quad \exists \theta \in \Theta \text{ s.t. } f\left(s|\alpha,\theta\right) = f\left(s|\alpha^*,\theta^*\right) \quad \text{for all } s \in S,$$

then it is impossible for individuals to figure out the correct model based on the long-run signal frequency.[8] Below is a simple example.

Example 12 (Identification Problem)**.** In Example 9, models $fair$ and $unfair$ induce the following signal distributions

$$\begin{array}{c|cc} fair & G & B \\ \hline H & p & 1-p \\ L & 1-p & p \end{array} \quad \begin{array}{c|cc} unfair & G & B \\ \hline H & 1-p & p \\ L & p & 1-p \end{array}, \text{ where } p > \frac{1}{2}.$$

Suppose that the true signal distribution is $f^* = (p, 1-p)$. Then both models are consistent with the true distribution. For model $fair$, the distribution under state H matches the true distribution; for model $unfair$, the distribution under state L also matches. Therefore, a frequentist cannot identify the true model based on the long-run signal frequency. □

From a Bayesian perspective, even if $\mathcal{A}$ is incorrectly specified, but every model in $\mathcal{A}$ can match the signal distribution equally well in terms of relative

[8] Similar justifications are seen in Battigalli et al. (2015). They introduced a self-confirming equilibrium under model uncertainty, in which players remain ambiguous about strategies that match the true long-run frequency.

entropy, a Bayesian individual with prior on $\mathcal{A}$ is still unable to eliminate any model in the limit.[9] In the previous example, suppose instead that the true model is $f^* = (1/2, 1/2)$, so neither model can match the true distribution. We notice that both models are symmetric with respect to f^*, in other words, they are "equally wrong", so a Bayesian individual is still unable to exclude either model. Another justification is that individuals remain ambiguous simply because they are unable to pin down a specific prior, hence every feasible model can be correct as long as it generates the observed history with a strictly positive probability. In this paper, I am agnostic about how the model set $\mathcal{A}$ is formed. Instead of imposing specific restrictions on the model set $\mathcal{A}$, this paper treats $\mathcal{A}$ as part of the learning environment and aims to derive a general characterization for *all* possible forms of $\mathcal{A}$ (e.g., correctly or incorrectly specified, matching the long-run frequency or not). It is worth noting that in addition to just being general, this flexibility accommodates situations where individuals will adopt a less-likely (but still possible) explanation. These situations may be less interesting in environments with "rational" individuals, but they seem more prevalent for biased individuals who have incentives to distort information.

2.4.2 Discussion of the Bias

Another key element is that individuals are endowed with some ***bias*** when processing new information. First, in terms of **motivation**, the bias is motivated by the idea that people attach intrinsic values to some states of the world and hence will process information to justify those states. For example, individuals enjoy

[9] In generic cases, beliefs on the model space will not converge, and *any* belief on $\mathcal{A}$ can be an accumulation point. Bunke and Milhaud (1998) provide one example, and the idea of Berk-Nash equilibrium proposed by Esponda and Pouzo (2016) is built on a similar spirit.

feeling good about themselves, so they tend to listen to good news and overlook bad news to maintain positive self-esteem. Second, in terms of its **impact on updating**, the bias pins down a posterior. Under model uncertainty, there are two approaches to pin down a posterior: (i) individuals can assign a prior over models, and (ii) individuals can employ a mechanism to determine which model to update. This paper follows the second approach, and the bias works as a model-selection mechanism. To draw a parallel, under maximum likelihood updating, individuals select a model to maximize the likelihood function; similarly, under biased updating, individuals select a model to maximize the bias, which is represented by a belief-utility function.[10]

This paper focuses on the dynamics of beliefs, but given that individuals hold biased beliefs, it is also natural to ask about **actions**. One possibility is that individuals are *naive* and evaluate all choices according to their current beliefs and model perceptions as in Example 11. The naivety is reflected in the following aspects: (i) when forming beliefs, individuals only seek to justify their bias, but do not consider the impacts on actions; (ii) when making decisions, individuals inherit beliefs from the learning stage and make decisions accordingly, but do not manipulate their biases or act in a strategic manner. In Example 11, it implies that the analyst's biased reporting only reflects his interpretation of the information and does not embed any strategic concerns. One natural question is that what do we gain from use of the naivety assumption? First, it has the advantage of being highly tractable, since we can determine all relevant actions by keeping track of the agent's beliefs. Second, it seems more in line with how the bias normally works compared with a sophisticated assumption. Many experimental findings suggest that people do not get biased in a conscious or a

[10]For references on the belief-utility function, see Caplin and Leahy (2001), Bénabou and Tirole (2002), Brunnermeier and Parker (2005) and Köszegi (2006).

strategic manner, so it is difficult to reconcile the idea that individuals are biased but are also sufficiently sophisticated to manipulate their own biases. [11] Due to these reasons, it seems natural to use the naive rule as a benchmark if we want to talk about actions. Under this assumption, Section 2.7 analyzes some examples to illustrate how biased individuals make decisions.

2.4.3 Discussion of Other Assumptions

Some ***other assumptions*** also merit discussion. First, individuals face the constraint that beliefs are updated by **Bayes rule model-wise**. Notice that when maximizing their belief utilities, individuals must face some constraint, since otherwise they can perfectly self-deceive, which makes the problem trivial and unrealistic. The model-wise Bayes constraint can be regarded as a "rationality constraint", which requires that individual would still update in a "rational" way if they knew how signals are generated. Admittedly, it remains debatable whether this is a sensible constraint to impose, but it serves as a natural first-step. Second, individuals have **perfect recall** so that they can go back and revise how to interpret all previous signals. It also remains controversial whether perfect or bounded memory (usually with one-period) seems most appropriate, but the perfect recall assumption has the advantage of enabling a simple characterization of limit beliefs. Third, individuals may **change their world views** (i.e., model perceptions) each time they receive a new signal. This as-

[11] Pronin et al. (2002) showed that people often fail to recognize their own biases (e.g., the self-serving bias, the "better-than-average" effect) even if they can recognize that the bias exists for others. Babcock et al. (1995) showed that the self-serving bias still exists even if bargainers have incentives to process the information objectively, which implies that "the bias does not appear to be deliberate or strategic" (Babcock and Loewenstein, 1997). Loewenstein and Adler (1995) implied that people do not have good awareness of the endowment effect, that is, they tend to underestimate that they would become biased toward an object once they received it.

sumption implies some "discontinuity" of their world views, which needs justifications. First, in many situations, the perceived models will converge, hence a sudden shift of the world view is less frequent in the long run. Second, the shifting of world views is not new to the literature. One prominent example is the maximum likelihood updating, where individuals can modify their model perceptions after observing new information. [12] Third, we can re-interpret the model as that individuals observe all signals first, and then decide their model perceptions and beliefs (assuming that there is no later information). This re-interpretation circumvents the problem and results in the same belief as in the benchmark model. Some other assumptions are worth discussing, e.g., individuals hold a fixed model set and a fixed bias, and Section 2.8 will discuss some extensions of the benchmark model by relaxing these assumptions.

Remark 2.1. *Another concern is the robustness of the updating rule. It is worth noting that the biased updating rule relies on the assumption that the individual always chooses the maximizing model even if many models are very close in their induced belief utility in the limit. This can be a strong assumption since we can construct examples such that a small perturbation leads to very different limit beliefs. This paper intends to provide a benchmark, but it would be of interest to relax it in applications.*

2.5 Relative Entropy and Related Concepts

This section introduces concepts that will be used in later sections. One preliminary characterization of limit beliefs is also presented.

[12] Here are some other examples. Ortoleva (2012) provided a axiomatic foundation for a non-Bayesian updating rule, where individuals can change their paradigms (i.e., the priors) if they observe some event to which their original paradigms assign a small probability. Galperti (2019) studied a persuasion problem under a similar assumption that individuals can shift their paradigms abruptly.

Definition 7. Define the ***relative entropy*** of state θ under model α as

$$\mathcal{R}(\alpha,\theta) \equiv \int_S \log\left[\frac{f(s|\alpha^*,\theta^*)}{f(s|\alpha,\theta)}\right] f(s|\alpha^*,\theta^*)\,dv(s)$$

Further denote $r(\alpha,\theta) \equiv \mathcal{R}(\alpha,\theta) - \min_{\theta'\in\Theta}\mathcal{R}(\alpha,\theta')$, called ***information potential*** of state θ under model α.[13]

The *relative entropy* of state θ under model α provides a measure of closeness between distribution $f(s|\alpha,\theta)$ and the true signal distribution, $f(s|\alpha^*,\theta^*)$. Fixing a model α, if the relative entropy of θ is lower, it means that state θ induces a distribution that is closer to the true distribution. The left graph of Figure 2.1 provides an illustration. The solid curve represents the true distribution, and the dashed and dotted curves represent the distributions of states θ and θ_1 under an alternative model, model α_1. In this graph, the true distribution is "closer" to $f(s|\alpha_1,\theta_1)$ than to $f(s|\alpha_1,\theta)$, and this corresponds to the relation $\mathcal{R}(\alpha_1,\theta_1) < \mathcal{R}(\alpha_1,\theta)$. [14] *Information potential* is a normalized version of relative entropy, where every entropy-minimizing state is normalized to have zero potential. The right graph of Figure 2.1 provides an illustration. This graph plots the relative entropy under α_1 over a continuum of states, where state θ_1 has the minimum relative entropy. The information potential of state θ is equal to its relative entropy minus a normalized term, which is the minimum relative entropy under α_1. As will be seen later, such normalization is useful since it enables comparison *across* models. In the rest of this paper, I will use the term "*zero-potential state*" and "(relative) *entropy-minimizing state*" interchangeably.

Below is a simple example showing how to find zero-potential states.

[13] This concept is well-defined since it is easy to verify that $\mathcal{R}(\alpha,\theta)$ is a continuous function, so the minimum can be obtained.

[14] A more concrete example is the class of normal distribution, $f(s|\alpha,\theta) = \mathcal{N}(\theta,\alpha^2)$. To generate the relation in Figure 2.1, we simply set θ_1 very close to θ^* and θ very far from θ^*.

Example 13 (Fairness and Preferences for Redistribution-II)**.** Consider the case in Example 9. Suppose that the true signal distribution $f^* = (q, 1-q)$ with $q > 1/2$, where the first coordinate represents the probability of feedback G, and the second represents the probability of feedback B. In this example, it is possible that q is different from p, so there may exist model misspecification. Considering $\alpha = fair$, by Definition 7, we have

$$\begin{aligned}\mathcal{R}\left(fair, H\right) &= f^*\left(G\right) \times \log\left(\frac{f^*\left(G\right)}{f\left(G|fair, H\right)}\right) + f^*\left(B\right) \times \log\left(\frac{f^*\left(B\right)}{f\left(B|fair, H\right)}\right) \\ &= q \times \log\left(\frac{q}{p}\right) + (1-q) \times \log\left(\frac{1-q}{1-p}\right).\end{aligned}$$

Analogously, $\mathcal{R}\left(fair, L\right) = q \times \log\left(\frac{q}{1-p}\right) + (1-q) \times \log\left(\frac{1-q}{p}\right)$. Recalling that $p, q > 1/2$, it follows that

$$\mathcal{R}\left(fair, H\right) - \mathcal{R}\left(fair, L\right) = (2q-1) \times \log\left(\frac{1-p}{p}\right) < 0,$$

so state H is the entropy-minimizing state, hence zero-potential state, under model $\alpha = fair$. This fact can be explained more intuitively. Under $\alpha = fair$, it is more likely to receive good feedback (G) than bad feedback (B) if the worker has high ability (H), and the opposite if the worker has low ability (L). Recall that the true distribution is $(q, 1-q)$ with $q > 1/2$, so it is better approximated by the distribution under state H under $\alpha = fair$. Therefore, state H is the zero-potential state under $\alpha = fair$. It can be further shown that if the worker updates according to $\alpha = fair$, he will hold a degenerate belief on the zero-potential state, state H, in the limit. □

Previous examples have simple structures such that limit beliefs can be calculated explicitly. To characterize limit beliefs for more general cases, I introduce the following definition.

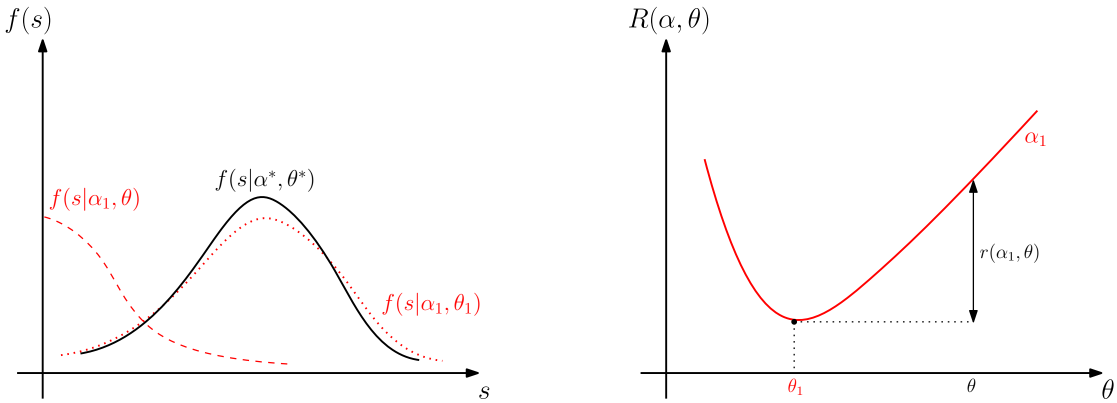

Figure 2.1: Relative Entropy & Information Potential

Definition 8. A sequence of probability measures $\{\mu_t\}$ is ***asymptotically carried*** on set $\mathcal{U}$ if $\lim \mu_t(U) = 1$ for all open sets $U \supset \mathcal{U}$.

This concept was also used by Berk (1966) to characterize limit beliefs under incorrect models. Roughly speaking, if a sequence of probability measures are asymptotically carried on $\mathcal{U}$, then only states in $\mathcal{U}$ will have non-zero weights in the limit, and states outside of $\mathcal{U}$ will be attached zero weight. Lemma 2.1 says that limit beliefs are asymptotically carried on zero-potential states.[15]

Lemma 2.1. *Let $\mathcal{U}_{\mathcal{A}}$ be the zero-potential states under models in $\mathcal{A}$, that is,*

$$\mathcal{U}_{\mathcal{A}} = \{\theta \in \Theta : r(\alpha, \theta) = 0 \textit{ for some } \alpha \in \mathcal{A}\}.$$

Then all individuals' beliefs are asymptotically carried on set $\mathcal{U}_{\mathcal{A}}$ $\mathbb{P}^*$-almost surely.

To understand the intuition, it is useful to review the case in Berk (1966). Berk showed that if individuals update beliefs according to a possibly incorrect model, limit beliefs will settle on zero-potential states under that model. If indi-

[15] It is conceivable that all results still hold in a stronger sense of convergence. Bunke and Milhaud (1998) used a stronger notion of concentration which is expressed in terms of the expected distance to set $\mathcal{U}$ under measures μ_ts. They proved a stronger version of the result in Berk (1966) using this notion.

viduals perceive a model α, and if limit beliefs attach positive weights to a state θ_0, we must have $r(\alpha, \theta_0) = 0$, or equivalently,

$$\theta_0 \in \arg\min_{\theta \in \Theta} \mathcal{R}(\alpha, \theta) \quad \mathbb{P}^* - a.s..$$

Intuitively, relative entropy measures the distance between state θ's induced distribution under α and the true distribution. Therefore, Berk's result captures the idea that limit beliefs will settle on states that generate the closest distributions to the true signal distribution. As in Example 13, under $\alpha = fair$, state H's induced distribution best approximates the true distribution, so the worker will settle on state H in the limit. Back to our model, when individuals can perceive *a set of* models, $\mathcal{A}$, the intuition still holds, and limit beliefs will settle on the set of zero-potential states under models in $\mathcal{A}$. When $\mathcal{A}$ consists of finite number of models, Lemma 2.1 comes from a simple union. When $\mathcal{A}$ consists of infinitely many models, the main difficulty is to establish that beliefs will converge *uniformly* for all models in $\mathcal{A}$, which is discussed in the Appendix.

2.6 A Complete Characterization of Limit Beliefs

Lemma 2.1 provides a coarse characterization of beliefs, which applies to all possible biases but does not address how biases influence model perceptions and beliefs. This section provides a fuller characterization of limit beliefs and model perceptions. This section first discusses a baseline case where individuals have one biased state. The discussion is then extended to the case with general biases.

2.6.1 Baseline Case: Single Biased State

Consider first the case in which individuals have only one biased state, that is, $\tau = \delta_\theta$ for some state $\theta \in \Theta$. Individuals whose biased state is θ are referred to as bias-θ individuals. Their updating rule is:

$$\forall \theta' \in \Theta : \quad \mu_t^\theta \left(\theta'\right) = \mu_t \left(\theta' | \alpha_t^\theta\right) \quad \text{where } \alpha_t^\theta \in \arg\max_{\alpha \in \mathcal{A}} \mu_t \left(\theta | \alpha\right),$$

where μ_t^θ denotes the time-t belief of a bias-θ individual, and α_t^θ denotes the model adopted by the individual at time t. As discussed previously, under this updating rule, individuals seek to confirm one state, so they update beliefs according to the model that delivers the highest possible likelihood for the biased state. The lemma below characterizes limit model perceptions.

Lemma 2.2. *For all $\theta \in \Theta$, let $\mathcal{A}_\infty^\theta$ denote the set of limit points of $\left\{\alpha_t^\theta\right\}$, we have:*

$$\mathcal{A}_\infty^\theta \subset \arg\min_{\alpha \in \mathcal{A}} r\left(\alpha, \theta\right) \quad \mathbb{P}^* - a.s.$$

that is, individuals will asymptotically update according to the models that minimize the information potential of their biased states.

The intuition behind Lemma 2.2 can be summarized as follows. When maximizing the likelihood of some bias θ, individuals will eventually select a model α that minimizes the "distance" between θ and the states that limit beliefs concentrate on, where the "distance" is measured by the difference in relative entropy. From Lemma 2.1, limit beliefs will concentrate on zero-potential states. By Definition 7, minimizing the "distance" to zero-potential states is equivalent to minimizing the information potential. As such, the bias-justifying behavior will lead individuals to select a model that minimizes the information potential of the biased state θ, which implies Lemma 2.2.

If $\mathcal{A}$ satisfies the condition that every model has the same minimum relative entropy,[16] we can further replace the $r(\alpha, \theta)$ in Lemma 2.2 with $\mathcal{R}(\theta, \alpha)$. In this case, individuals will adopt models that minimize the *relative entropy* of their biases in the limit. One simple example is when every model in $\mathcal{A}$ can match the true signal distribution. However, for a general $\mathcal{A}$, the limit model may not be the one that minimizes the relative entropy (e.g., when individuals are not ambiguous in a frequentist or a Bayesian manner). Below is one simple example.

Example 14. Suppose that $\Theta = \{\theta_1, \theta_2\}$, $S = \{s_1, s_2\}$ and $\mathcal{A} = \{\alpha_1, \alpha_2\}$, where

$$\begin{array}{c|cc} \alpha_1 & s_1 & s_2 \\ \hline \theta_1 & \frac{1}{2} & \frac{1}{2} \\ \theta_2 & \frac{2}{5} & \frac{3}{5} \end{array} \qquad \begin{array}{c|cc} \alpha_2 & s_1 & s_2 \\ \hline \theta_1 & \frac{1}{4} & \frac{3}{4} \\ \theta_2 & \frac{3}{4} & \frac{1}{4} \end{array},$$

The true signal distribution is $f^* = (1/3, 2/3)$, where the first coordinate denotes the probability of s_1 and the second coordinate denotes the probability of s_2. In this example, individuals are endowed with a *misspecified* model set (i.e., neither model matches the long-run frequency). It can be seen that θ_2 and θ_1 minimize the relative entropy under α_1 and α_2 respectively. Suppose that an individual has a biased state θ_1. If he chooses to believe in model α_1, he will form a belief δ_{θ_2} in the limit. If he chooses to believe in model α_2, he will form a belief δ_{θ_1} in the limit, which leads him to justify his bias perfectly. As a result, the bias-θ_1 individual will almost surely interpret signals according to model α_2 in the limit. However, it can be verified that model α_1 induces a lower relative entropy than model α_2 at state θ_1. $\square$

Based on Lemma 2.1 and Lemma 2.2, limit beliefs can be characterized by the theorem below.

[16] Formally, it means that $\forall \alpha, \alpha' \in \mathcal{A}$ we have $\min_{\theta \in \Theta} \mathcal{R}(\theta, \alpha) = \min_{\theta \in \Theta} \mathcal{R}(\theta, \alpha')$.

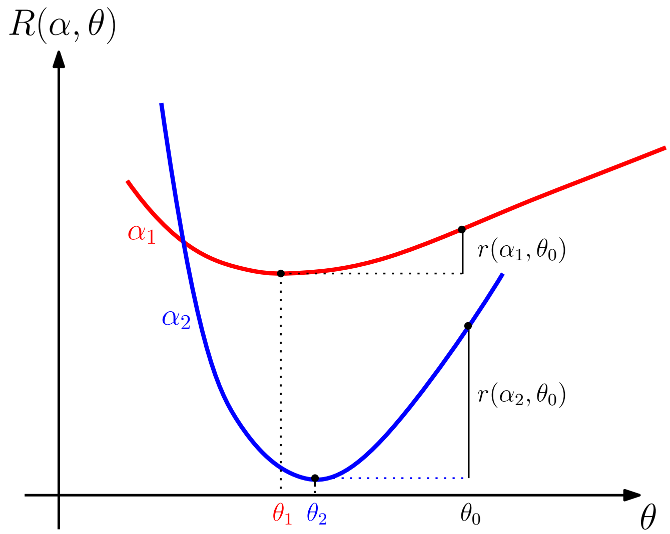

Figure 2.2: Limit beliefs

Theorem 2.1. *For all $\theta \in \Theta$, define a set $\mathcal{U}^{\theta}_{\mathcal{A}}$ as follows*

$$\mathcal{U}^{\theta}_{\mathcal{A}} = \left\{ \theta' \in \Theta : r\left(\alpha', \theta'\right) = 0 \text{ where } \alpha' \in \arg\min_{\alpha \in \mathcal{A}} r\left(\alpha, \theta\right) \right\}.$$

Then bias-θ individuals' beliefs are asymptotically carried on $\mathcal{U}^{\theta}_{\mathcal{A}}$ $\mathbb{P}^*$-almost surely.

To summarize, limit beliefs will concentrate on zero-potential states under the potential-minimizing models of the biased state. Figure 2.2 provides an illustration. Suppose that individuals only consider two models as possible $\mathcal{A} = \{\alpha_1, \alpha_2\}$. From Lemma 2.2, a bias-θ_0 individual will asymptotically adopt the model that minimizes the potential for state θ_0, which is model α_1. From Theorem 2.1, the bias-θ_0 individual will eventually settle on the zero-potential state under α_1, that is, state θ_1. In other words, individuals with bias θ_0 will hold a degenerate belief on state θ_1 almost surely.

Remark 2.2. *If $\mathcal{U}^{\theta}_{\mathcal{A}}$ contains a unique state, bias-θ individuals will almost surely hold a degenerate belief in the limit. When $\mathcal{U}^{\theta}_{\mathcal{A}}$ contains multiple states, there is no guarantee that beliefs will converge to a well-defined limit. In most interesting cases, beliefs of bias-θ individuals will oscillate on set $\mathcal{U}^{\theta}_{\mathcal{A}}$ in the limit. This comes from the fact that the* *log likelihood ratio between any two entropy-minimizing states constitutes a zero-mean random walk, so it oscillates between $-\infty$ and $+\infty$ with probability 1 (see Berk (1966) for a concrete example).*

Below is a numerical example based on Theorem 2.1.

Example 15 (Inferring the market fundamental-II)**.** Consider the case described by Example 11. I assume that $\Theta = [-10, 10]$ with the true state $\theta^* = 2$, the true demand elasticity $\alpha^* = 1$, and the market volatility $\sigma(\theta) = 1$. Therefore, the true distribution of market price is:

$$\log P_t \sim \mathcal{N}\left(1, \frac{1}{2}\right)$$

The perceived set of models is $\mathcal{A} = \left[\frac{1}{1+\delta}, 1+\delta\right]$ where a higher δ corresponds to a higher degree of ambiguity. Suppose that every agent has one biased state $\theta \in \Theta$. For example, investors try to persuade themselves of some market condition in order to support some investment tendency; policy makers want to justify some economic policy by persuading people (or even themselves) of a specific market condition. For simplicity in exposition, I assume that the degree of ambiguity δ is not too large (to avoid a kink in the expression of relative entropy). It can be verified that

$$r(\alpha, \theta) = \left(\frac{\theta}{\alpha+1} - 1\right)^2.$$

Letting $\theta(\alpha)$ denote the *zero-potential state* under model α, we have:

$$\theta(\alpha) = \alpha + 1 \quad \text{for all } \alpha \in \mathcal{A}$$

Further defining $\alpha(\theta)$ to be the *potential-minimizing model* of state θ, we have:

$$\alpha(\theta) = \begin{cases} \frac{1}{1+\delta} & \theta \in \left[-10, \frac{2+\delta}{1+\delta}\right] \\ \theta - 1 & \theta \in \left[\frac{2+\delta}{1+\delta}, 2+\delta\right] \\ 1+\delta & \theta \in [2+\delta, 10] \end{cases},$$

which gives the limit model perception for each possible bias from Lemma 2.2. Using results from Theorem 2.1, the limit belief carrier for each bias is:

$$\mathcal{U}_{\mathcal{A}}^{\theta} = \begin{cases} \left\{\frac{2+\delta}{1+\delta}\right\} & \theta \in \left[-10, \frac{2+\delta}{1+\delta}\right] \\ \{\theta\} & \theta \in \left[\frac{2+\delta}{1+\delta}, 2+\delta\right] \\ \{\delta+2\} & \theta \in [2+\delta, 10] \end{cases}.$$

For all individuals with bias $\theta \in \left[\frac{2+\delta}{1+\delta}, 2+\delta\right]$, they will successfully justify their biased states and hold a Dirac belief δ_θ in the limit. For individuals with bias too high (i.e., $\theta \geq 2+\delta$) or too low (i.e., $\theta \leq \frac{2+\delta}{1+\delta}$), they are unable to justify their biased states. Individuals with a very high bias will interpret price information according to the highest possible demand elasticity, $1+\delta$, and believe in the corresponding zero-potential state, $2+\delta$. Here is the intuition. Recall that a higher demand elasticity implies a lower price, and a higher market fundamental implies a higher price. When faced with the same price sequence, a higher demand elasticity enables individuals to justify a higher market fundamental, so individuals with high (or low) bias will update beliefs according to the model with high (or low) demand elasticity. $\square$

Remark 2.3. *As the degree of ambiguity $\delta \to 0$, the limit belief carrier $\mathcal{U}_{\mathcal{A}} = \left[\frac{2+\delta}{1+\delta}, 2+\delta\right] \to \{\theta^*\}$,which means that limit beliefs become more concentrated around the true state. However, for any degree of ambiguity $\delta > 0$, polarization arises on states in $\mathcal{U}_{\mathcal{A}}$. Depending on their biases, individuals can hold any limit posterior δ_θ with $\theta \in \mathcal{U}_{\mathcal{A}}$, so they can totally disagree on the correct state. Unlike Example 10, in this example, by reducing ambiguity, we only decrease the range of disagreement but can never eradicate the disagreement itself.*

2.6.2 Comparative Statics: Asymptotic Learning and Polarization

Based on the remark under Example 15, this subsection discusses what happens as ambiguity varies. The degree of ambiguity $d(\mathcal{A})$ is defined to be the diameter of $\mathcal{A}$. [17]For simplicity in exposition, this subsection assumes that the perceived model set $\mathcal{A}$ contains the true model α^*. First, as the degree of ambiguity approaches zero, all individuals will learn in a weak sense.

Proposition 2.1. *As ambiguity vanishes, limit beliefs converge* ***weakly*** *to the correct belief δ_{θ^*}, that is,*

$$\lim_{d(\mathcal{A})\to 0}\left[\lim_{t\to\infty}|\int_{\Theta} h\times\mu_t^{\theta}dm-\int_{\Theta} h\times\delta_{\theta^*}dm|\right]\to 0\quad \mathbb{P}^*-a.s.$$

for all bounded and continuous function $h:\Theta\to\mathbb{R}$ and for all $\theta\in\Theta$.

The concept of convergence is an adapted version of convergence in weak topology. Roughly speaking, it says that as ambiguity vanishes, the limit belief carrier approaches the true state. We can have a stronger version of learning if the true state, θ^*, is a "dominant" state for all models within a small neighborhood of the true model, α^*. More precisely, θ^* is *locally dominant* if there exists some non-degenerate closed neighborhood $C\ni\alpha^*$ such that θ^* is the unique zero-potential state for all $\alpha\in C$.

Proposition 2.2. *If θ^* is locally dominant, limit beliefs converge* ***strongly*** *to the correct belief when ambiguity is sufficiently small, that is,*

$$\sup\left\{|\mu_t^{\theta}(E)-\delta_{\theta^*}(E)|:\theta\in\Theta,E\in\mathcal{B}_{\Theta}\right\}\to 0\quad \mathbb{P}^*-a.s.$$

[17]More precisely, $d(\mathcal{A})=\max_{\alpha,\alpha'\in\mathcal{A}}||\alpha-\alpha'||$, where $||.||$ denotes the relevant metric on the model space $\mathbb{A}$.

for all $\mathcal{A}$ such that $d(\mathcal{A}) < \varepsilon$ for some $\varepsilon > 0$.

The concept of convergence in Proposition 2.2 is strong convergence, which requires convergence on all measurable sets. It is a direct implication of Lemma 2.1. If the true state θ^* is locally dominant, $\mathcal{U}_{\mathcal{A}}$ becomes a singleton $\{\theta^*\}$ when the degree of ambiguity is sufficiently small. Therefore, all beliefs converge to the correct belief δ_{θ^*} almost surely. The intuition is that if the true state is locally dominant, then it can be perfectly identified within a small neighborhood around the true model. One important case where the local dominance condition holds is when the state space is finite.

Corollary 2.1. *If $|\Theta| < \infty$, beliefs converge strongly to the correct belief when ambiguity is sufficiently small.*

This corresponds to the case in Example 10. Recall that in this example, $\Theta = \{G, B\}$, and both optimists and pessimists can learn the true state when ambiguity is sufficiently small. To illustrate the idea, a special case of Example 10 is discussed as follows.

Example 16 (Optimistic and Pessimistic Investors-II)**.** Suppose that $\lambda = 2$ and $\varepsilon = 1/2$, and that the true state $\theta^* = B$. Let $\mathcal{A} = \left[\frac{1}{1+\delta}, 1+\delta\right]$ denote the set of likelihood ratios induced by signal m. Every model $\alpha \in \mathcal{A}$ represents a likelihood ratio. It is easy to verify that as $t \to \infty$,

$$\frac{1}{t}\log\left[\frac{\mu_t(G|\alpha)}{\mu_t(B|\alpha)}\right] \to \frac{1}{2}\log\alpha - \frac{1}{6}\log 2 \quad \mathbb{P}^*\text{−}a.s..$$

For pessimistic investors, we have:

$$\frac{1}{t}\log\left[\frac{\mu_t^p(G)}{\mu_t^p(B)}\right] = \min_{\alpha\in\mathcal{A}}\frac{1}{t}\log\left[\frac{\mu_t(G|\alpha)}{\mu_t(B|\alpha)}\right] \to \frac{1}{2}\log\left(\frac{1}{1+\delta}\right) - \frac{1}{6}\log 2 < 0 \quad \mathbb{P}^*\text{−}a.s.,$$

which implies that $\mu_t^p(B) \to 1$ almost surely, so pessimists will almost surely learn the true state, state B. For optimistic investors, we have:

$$\frac{1}{t}\log\left[\frac{\mu_t^o(G)}{\mu_t^o(B)}\right] = \max_{\alpha\in\mathcal{A}}\frac{1}{t}\log\left[\frac{\mu_t(G|\alpha)}{\mu_t(B|\alpha)}\right] \to \frac{1}{2}\log(1+\delta) - \frac{1}{6}\log 2 \quad \mathbb{P}^*-a.s.. \tag{2.2}$$

When the degree of ambiguity is small such that $\delta < 2^{1/3} - 1$, the RHS of (A.12) is strictly negative, which implies that $\mu_t^o(B) \to 1$ almost surely. In this case, correct learning arises for both optimists and pessimists. □

The intuition for Corollary 2.1 is straightforward. When the state space is finite, the distance between the true signal distributions under any two states is bounded away from 0. Therefore, the true state can be identified when all models are sufficiently close to the true model, that is, when the degree of ambiguity is sufficiently small. However, when the state space is rich enough (e.g., uncountably infinite), correct learning may **not** occur. Moreover, polarization can occur for any positive degree of ambiguity as in Example 15. I refer to state θ^* as *singular* if θ^* is the unique zero-potential state under all model $\alpha \in \mathbb{A}$ such that $r(\alpha, \theta^*) = 0$.[18] One simple example satisfying this property is that every model has a unique zero-potential state. We have:

Proposition 2.3. *If θ^* is singular and not locally dominant, **polarization** occurs for any positive degree of ambiguity, that is,*

$$\sup\left\{|\mu_t^\theta(E) - \mu_t^{\theta'}(E)| : \theta, \theta' \in \Theta, E \in \mathcal{B}_\Theta\right\} \to 1 \quad \mathbb{P}^* - a.s.$$

for all $\mathcal{A}$ such that $d(\mathcal{A}) > 0$.

Proposition 2.3 says that with probability 1, there exist individuals with different biases who *totally disagree* on some event. More precisely, there exists

[18]In mathematical language, $\{\theta \in \Theta : r(\alpha,\theta) = 0 \text{ where } \alpha \text{ solves } r(\alpha,\theta^*) = 0\} = \{\theta^*\}$

some event where one individual assigns probability 0 whereas another individual assigns probability 1; that is, individuals' limit beliefs can be mutually singular. Proposition 2.3 shows that learning and belief consensus can be fragile with respect to ambiguity. When individuals are trying to defend their biases, it is possible that a slight degree of ambiguity suffices to destroy agreement. The proof of Proposition 2.3 is not difficult. The singular assumption can be thought of as a regularity assumption. The main assumption driving the result is the assumption that θ^* is not locally dominant, so there exists an identification problem in any small neighborhood around the correct model. As a result, limit posteriors of differently biased individuals can disagree—— they may or may not learn the true state, depending on their biases.

2.6.3 General Case: Multiple Biased States

This section characterizes limit beliefs in a more general framework where individuals have more than one biased state. The characterizations are parallel to the benchmark case, so readers may choose to skip this part. Recall that when each individual only has one biased state, the limit models are those that minimize the potential of *the* biased state. The following lemma extends this result to the case with multiple biased states.

Lemma 2.3. *For all $\tau \in \mathcal{T}$, let $\mathcal{A}_\infty^\tau$ denote the set of limit points of $\{\alpha_t^\tau\}$, we have:*

$$\mathcal{A}_\infty^\tau \subset \arg\min_{\alpha \in \mathcal{A}} \left[\min_{\theta \in supp(\tau)} r(\alpha, \theta) \right] \quad \mathbb{P}^* - a.s.$$

that is, bias-τ individuals will asymptotically update according to the models that minimize the minimum information potential of their biased states.

Lemma 2.2 says that in the case with multiple biased states, individuals seek to minimize the *minimum* potential of their biased states. This comes from the fact that the biased states with the minimum potential will eventually dominate all other biased states, so only the minimum-potential biased states are relevant in the limit. Following the same logic as in Lemma 2.2, individuals will seek to minimize the potential of these relevant biased states, or equivalently, they will minimize the minimum potential of biased states in the limit. Based on this lemma, limit beliefs can be characterized as below.

Theorem 2.2. *For all $\tau \in \mathcal{T}$, define a set $\mathcal{U}_{\mathcal{A}}^{\tau}$ as follows,*

$$\mathcal{U}_{\mathcal{A}}^{\tau} = \left\{ \theta' \in \Theta : r\left(\alpha', \theta'\right) = 0 \textit{ where } \alpha' \in \arg\min_{\alpha \in \mathcal{A}} \left[\min_{\theta \in supp(\tau)} r\left(\alpha, \theta\right) \right] \right\}.$$

Then bias-τ individuals beliefs are asymptotically carried on $\mathcal{U}_{\mathcal{A}}^{\tau}$ $\mathbb{P}^*$-almost surely.

Theorem 2.2 is a parallel statement of Theorem 2.1. It says that limit beliefs will settle on zero-potential states under models that minimize the minimum potential of biased states. Below is an example example that explains how to find limit beliefs.

Example 17. Suppose that $\Theta = \{\theta_1, \theta_2, \theta_3, \theta_4\}$, $S = \{s_1, s_2\}$ and $\mathcal{A} = \{\alpha_1, \alpha_2\}$, where

α_1	s_1	s_2
θ_1	$1/10$	$9/10$
θ_2	$7/8$	$1/8$
θ_3	$2/5$	$3/5$
θ_4	$3/4$	$1/4$

α_2	s_1	s_2
θ_1	$7/9$	$2/9$
θ_2	$1/10$	$9/10$
θ_3	$3/4$	$1/4$
θ_4	$3/5$	$2/5$

,

and the true signal distribution is $f^* = (1/2, 1/2)$. An individual is endowed with bias $\tau = \frac{1}{2}1_{\theta_1} + \frac{1}{2}1_{\theta_2}$, which means that this individual seeks to confirm either state θ_1 or state θ_2, and he assigns equal weights to these two biased states. To calculate his limit beliefs, we follow the steps below.

First, *find the zero-potential state under each model.* It can be verified that θ_3 is the zero-potential state under α_1, and θ_4 is the zero-potential state under α_2.[19] In this case, neither model can justify the individual's biased states in the limit.

Second, *find the minimum-potential biased state under each model.* Under model α_1, the information potential of θ_1 and θ_2 is

$$\begin{aligned} r(\alpha_1, \theta_1) &= \mathcal{R}(\alpha_1, \theta_1) - \mathcal{R}(\alpha_1, \theta_3) = \frac{1}{2} \log\left(\frac{2/5}{1/10}\right) + \frac{1}{2} \log\left(\frac{3/5}{9/10}\right) \\ r(\alpha_1, \theta_2) &= \mathcal{R}(\alpha_1, \theta_2) - \mathcal{R}(\alpha_1, \theta_3) = \frac{1}{2} \log\left(\frac{2/5}{7/8}\right) + \frac{1}{2} \log\left(\frac{3/5}{1/8}\right). \end{aligned}$$

It is easy to see $r(\alpha_1, \theta_1) > r(\alpha_1, \theta_2)$, so state θ_2 is the minimum-potential biased state under α_1. Intuitively, it means that θ_2 is "closer" to the zero-potential state (i.e., state θ_3) than θ_1 in terms of their induced distributions. Similarly, we have

$$\begin{aligned} r(\alpha_2, \theta_1) &= \frac{1}{2} \log\left(\frac{3/5}{7/9}\right) + \frac{1}{2} \log\left(\frac{2/5}{2/9}\right) \\ r(\alpha_2, \theta_2) &= \frac{1}{2} \log\left(\frac{3/5}{1/10}\right) + \frac{1}{2} \log\left(\frac{2/5}{9/10}\right). \end{aligned}$$

Since $r(\alpha_2, \theta_1) < r(\alpha_2, \theta_2)$, state θ_1 is the minimum-potential biased state under α_2. The intuition follows analogously.

Third, *compare the information potential of the minimum-potential biased state under each model.* That is, we need to compare $r(\alpha_1, \theta_2)$ and $r(\alpha_2, \theta_1)$. It is easy

[19] Intuitively, under model α_1, θ_3 induces a distribution $(2/5, 3/5)$ that is closer to the true distribution, $(1/2, 1/2)$, than any other state, so θ_3 is the zero-potential state under α_3. Analogously, we can see that θ_4 is the zero-potential state under α_4.

to see that

$$
\begin{aligned}
r(\alpha_1, \theta_2) &= \frac{1}{2} \log\left(\frac{2/5}{7/8}\right) + \frac{1}{2} \log\left(\frac{3/5}{1/8}\right) \\
&> \frac{1}{2} \log\left(\frac{3/5}{7/9}\right) + \frac{1}{2} \log\left(\frac{2/5}{2/9}\right) = r(\alpha_2, \theta_1),
\end{aligned}
$$

so α_2 minimizes the minimum potential of the biased states. Lemma 2.3 and Theorem 2.2 further imply that this individual will almost surely adopt model α_2 and hold a degenerate belief on the zero-potential state under α_2 in the limit. As a result, his limit belief is δ_{θ_4}. □

In some situations, we can further refine the characterization in Theorem 2.2 by comparing the belief utility $\tau(\theta)$ among all states in $\mathcal{U}_{\mathcal{A}}^{\tau}$.

Corollary 2.2 (Refinement of Theorem 2.2)**.** *Suppose that every $\alpha \in \mathcal{A}$ has a unique zero-potential state and denote*

$$
\mathcal{V}_{\mathcal{A}}^{\tau} = \arg\max \left\{\tau(\theta) : \theta \in \mathcal{U}_{\mathcal{A}}^{\tau}\right\}.
$$

For all possible bias $\tau \in \mathcal{T}$, bias-τ individuals beliefs are asymptotically carried on $\mathcal{V}_{\mathcal{A}}^{\tau}$ $\mathbb{P}^*$-almost surely.

In other words, if the uniqueness of zero-potential state is assumed, limit beliefs will only accumulate states in $\mathcal{U}_{\mathcal{A}}^{\tau}$ with the highest belief utility.[20] Without the uniqueness assumption, it is possible that beliefs may also accumulate on states that do not have the highest weight (see Example 18).

[20]Actually, we only need a weaker assumption that for each τ, there exists a model that has a unique zero-potential state with the highest $\tau(\theta)$ in $\mathcal{U}_{\mathcal{A}}^{\tau}$

2.6.4 The Overconfidence Effect under Ambiguity

In Examples 7 and 8, we can see that the biased rule is different from Bayes rule (i.e., both correctly and incorrectly specified). In these examples, individuals only observe finite number of signals, so it is natural to ask: will the biased rule be identical to Bayes rule in the limit? The answer is no. This section discusses a novel overconfidence effect that can persist even in the limit. Even with arbitrarily many signals, biased individuals can exhibit strictly greater confidence than *any* Bayesian agent (i.e., with arbitrary feasible model perception) in the limit.

Example 18 (Overconfidence Effect under Ambiguity)**.** Suppose that state space $\Theta = \{1, 0, -1\}$ with the true state $\theta^* = 0$. Signals take values in $S = \{g, m, b\}$. Individuals only consider two possible models α_1 and α_2, where

α_1	g	m	b
1	$1-\varepsilon$	$\frac{\varepsilon}{2}$	$\frac{\varepsilon}{2}$
0	$1/4$	$1/2$	$1/4$
-1	$\frac{\varepsilon}{2}$	$\frac{\varepsilon}{2}$	$1-\varepsilon$

α_2	g	m	b
1	$1/2$	$1/4$	$1/4$
0	$\frac{\varepsilon}{2}$	$1-\varepsilon$	$\frac{\varepsilon}{2}$
-1	$1/4$	$1/4$	$1/2$

,

where $\varepsilon > 0$ is sufficiently small (e.g., $\varepsilon = 0.01$). For simplicity, I assume a symmetric true model $f^* = (1/3, 1/3, 1/3)$. Individuals have an upward bias such that $\tau(1) > \tau(0) > \tau(-1)$. At each time t, each individual reports his estimation of the state $r_t \in \{-1, 0, 1\}$ which corresponds to the state he assigns the highest probability to. [21]

Fact 1: If the individual is a Bayesian, the expected report $\mathbb{E}^* r_t$ converges to

[21] If two or more states have the highest probability, individuals adopt some tie-breaking rule (or randomize) when issuing their reports.

0 almost surely for every possible model perception.

If the individual adopts model α_1, his beliefs will almost surely settle on state 0, the unique entropy-minimizing state under α_1, which implies $r_t \to 0$ almost surely (hence $\mathbb{E}^* r_t \to 0$ almost surely). If the individual adopts model α_2, beliefs will oscillate between state 1 and -1, which are both entropy-minimizing states under α_2. Furthermore, it can be shown that limit beliefs will concentrate around 1 and -1 each with probability $1/2$. Therefore, $\mathbb{E}^* r_t$ converges 0 almost surely (the details are shown in the Appendix).[22]

Consider a third party who collects reports from a large group and wants to infer the true state based on the average report. When individuals are Bayesian, the third-party will asymptotically conclude that the true state is 0 for *all* possible model perceptions within the group. From this perspective, the model uncertainty in this example seems harmless. However, the following fact contradicts this conjecture.

Fact 2: If the individual is ambiguous about the true model, the expected report $\mathbb{E}^* r_t$ converges to a *strictly* positive number almost surely.

This fact comes from the following arguments. In the limit, (1) if the belief under α_2 concentrates around state -1, this individual will update according to model α_1 (since he "prefers" state 0 to state 1), which leads to a report equal to 0; (2) if the belief under α_2 concentrates around state 1, this individual will adopt model α_2, which leads to a report equal to 1. Each case occurs with probability

[22] Here is an intuitive way to understand the result. Under model α_2, the distribution in state 1, $(1/2, 1/4, 1/4)$, and the distribution in state -1, $(1/4, 1/4, 1/2)$, are symmetric with respect to the true distribution $f^* = (1/3, 1/3, 1/3)$. Intuitively, these two states are "equally" likely to be the true state, so limit beliefs will accumulate around each state with probability $1/2$. Here, I only use the word "accumulate" since beliefs will not converge (see the remark 2.2).

$1/2$, so the expected report $\mathbb{E}^* r_t$ converges to $1/2$. Consequently, this individual becomes strictly more optimistic than the case where he is not ambiguous. $\square$

The argument above shows an interesting result of overconfidence that cannot be accommodated by the Bayesian framework. It illustrates how ambiguity and bias-defending work together to produce overconfidence. The basic intuition is that different models may have **complementary effects**, which in a similar spirit to the "good-news effect" as in Example 7. Whenever the "good news" occurs such that model α_2 leads to a high state, state 1, individuals can *exploit* this good news by interpreting signals according to α_2. Whenever the "bad news" occurs such that when model α_2 leads to a low state, state -1, individuals can *hedge against* this bad news by interpreting signals according to α_1. As can be seen, ambiguity accommodates the asymmetric treatment of signals under different signal paths, enabling individuals to exhibit greater confidence than any Bayesian individual.

Remark 2.4. *Combining the results in this example and examples in previous sections, one observes that model ambiguity produces two related effects that can lead to overconfidence (relative to Bayesian case). The first is what I call "flexible effect". The ambiguity accommodates the flexibility in signal interpretations, so individuals with different biases can choose to interpret signals differently, which leads to greater confidence than the case with a common model perception (see Example 8 and Example 10). The second effect is the "complementary effect" in this section. This effect utilizes another feature of ambiguity: ambiguity allows every individual to adopt different interpretations under different signal realizations. Under appropriate conditions, models can complement each other, which enables a biased individual to exhibit strictly greater confidence than any Bayesian individual. In Appendix B.2, I provide detailed conditions under which the complementary effect occurs.*

Previous sections focus on the dynamics of beliefs, but it is also interesting to see how actions are different under the biased rule. To provide a fully satisfactory answer, we may need to build a decision theoretical foundation, which is beyond the scope of this paper. To shed light on how actions might be chosen, this section discusses some examples under the assumption that individuals are naive, that is, they choose what is optimal according to their current belief and model perception. This naivety assumption was discussed in Section 2.4 and seems to be a natural benchmark. Under this assumption, I present two examples—one static and one dynamic. The static example compares the biased rule with the full Bayesian rule and the maximum likelihood rule in their action implications. The dynamic example discusses the issue of dynamic (in)consistency.

2.7 Examples: Decisions under the Biased Updating Rule

2.7.1 Differences from the Full Bayesian and Maximum Likelihood Rule

One may wonder how the biased rule is different from two common updating rules under ambiguity—the full Bayesian updating (or FB) and maximum likelihood updating (or ML). Under FB, individuals keep the posteriors updated from all possible models. Under ML, individuals only keep the posteriors updated from the models that maximize the probability of generating the observed information. To facilitate the comparison, I consider a decision environment and examine how decisions differ under each of these updating rules. For the rest of this subsection, I assume that individuals select an action to maximize the minimum expected utility when there are multiple posteriors.

Example 19 (Difference among FB, ML & the Biased Rule)**.** Suppose that $\Theta = \{L, M, R\}$ and $S = \{a, b\}$. An agent holds a flat prior over Θ and can observe a sequence of signals. After observing the signals, the agent chooses among three actions $A = \{l, m, r\}$ with the following payoffs

$$u(l, \theta) = 1_l(\theta), \quad u(r, \theta) = 1_r(\theta), \quad \text{and } u(m, \theta) = \frac{1}{2}.$$

Actions l and r are risky actions that only generate payoffs in specific states, and action m is a safe action that generates a constant payoff of $1/2$. The agent is ambiguous over a set of models $\mathcal{A} = \{\alpha_1, \alpha_2\}$, where

$$\begin{array}{ccc} [\alpha_1] & a & b \\ L & 1/3 & 2/3 \\ M & 1/2 & 1/2 \\ R & 2/3 & 1/3 \end{array} \qquad \begin{array}{ccc} [\alpha_2] & a & b \\ L & 3/4 & 1/4 \\ M & 1/2 & 1/2 \\ R & 1/4 & 3/4 \end{array}.$$

Suppose that the agent observed 8 signals, which consist of 2 signal as and 6 signal bs. The observed signal frequency is hence denoted by $(1/4, 3/4)$. In the following discussion, we compare three different updating rules: the full Bayesian rule, the maximum likelihood rule, and the biased rule. To make the discussion not straightforward, for the biased rule, I assume that the agent is biased toward state M. This is because if the agent is biased toward state L or R, the optimal choice is just action l or r. Under this specification, it turns out that *each* updating rule induces a different optimal action:

updating rule	optimal action
Full Bayesian	$a^{FB} = m$
Maximum Likelihood	$a^{ML} = r$
Biased Rule	$a^{BS} = l$

.

(i) If the agent updates via the ***full Bayesian rule***, his optimal choice will be $a^{FB} = m$. To see that, we first notice that under the given signal structures, each signal can be flexibly interpreted as good news and bad news for any state. For example, model α_1 interprets signal a as bad news for state L, so receiving a signal a will decrease the likelihood of state L. On the contrary, model α_2 interprets signal a as "good news" for state L. If the agent adopts the full Bayesian rule, he is unwilling to take any risky action because he could always interpret the majority signals as bad news for the payoff relevant state. As such, the lowest probability on each state is less than $1/2$, so the agent will opt for the safe action, action m.

(ii) If the agent updates via the ***maximum likelihood rule***, his optimal choice will be $a^{ML} = r$. To see that, first note that under model α_2, state R induces a signal distribution $(1/4, 3/4)$, which perfectly matches the observed signal frequency; under model α_1, the best-matching distribution is $(1/3, 2/3)$, which is induced by state L and does not perfectly match the frequency. When there are sufficiently many signals (i.e., in this example, $n = 8$ suffices), posteriors will put sufficiently large probability on the state that induces the best match, so only the best-matching distribution matters for evaluating the likelihood of each model. As a consequence, model α_2 is more likely to generate the observed signals than model α_1, so the agent will update according to model α_2.

It then follows that the agent will choose action r since state R induces the best-matching distribution.

(iii) If the agent updates via the ***biased rule*** and is biased toward state M, his optimal choice will be $a^{BS} = l$. The intuition is similar to the intuition behind Theorem 2.1. When trying to defend his bias, the agent will update according to the model which minimizes the "distance" between the signal distributions in his biased state and in the best-matching state of that model. In this example, the signal distribution in state M is $(1/2, 1/2)$ under both models. Intuitively, $(1/2, 1/2)$ is "closer" to $(1/3, 2/3)$, the best-matching distribution under α_1, than to $(1/4, 3/4)$, the best-matching distribution under α_2. Therefore, the biased agent will update according to model α_1. Under model α_1, the best-fitting state is state L, so the agent will choose action l.[23] □

Remark 2.5. *To sum up, the biased updating rule is different from the full Bayesian rule and maximum likelihood rule in the following aspects: (i) ML adopts a purely objective criterion and evaluates models according to their probability of generating the observed event, but BS adopts a purely subjective criterion and evaluates models according to their consistency with the bias; (ii) FB updates all models indiscriminately and leads to a set of posteriors, but BS only updates the bias-maximizing model and leads to a unique posterior. The difference between BS and FB is more evident if we compare the relevant posteriors for decisions. Under FB, individuals can use different posteriors to evaluate different actions, so the relevant posterior is choice-dependent; in contrast, under BS, individuals form a biased posterior first and then use that posterior to evaluate all choices, so the relevant posterior is choice-independent. These differences in beliefs can also lead to different actions as in Example 19. Though the example is built on a heuristic decision rule, it suggests that the biased rule has the potential of delivering*

[23] The detailed verification of these arguments are in the Appendix.

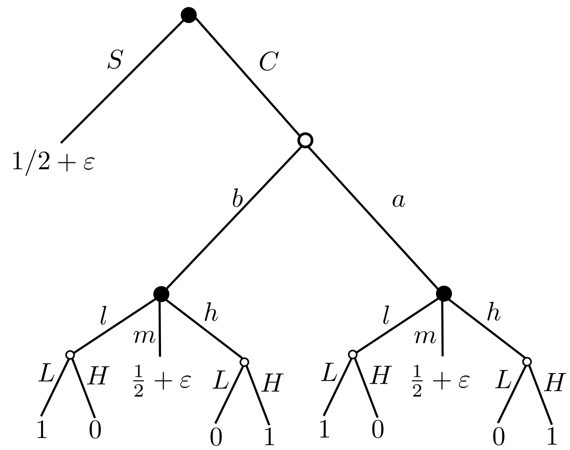

Figure 2.3: Dynamic Inconsistency

new action implications.

2.7.2 Dynamic (In)consistency of the Biased Agents

Another relevant question is what individuals would do if they had to choose at different points in time. This subsection looks at a dynamic decision environment, where individuals use the biased updating rule and make decisions according to the naive decision rule. One implication is that individuals may suffer from dynamic inconsistency. Below is a simple example.

Example 20 (Dynamic Inconsistency)**.** The payoff relevant state space is $\Theta = \{H, L\}$. The signal space is $S = \{a, b\}$. An individual is faced with a simple problem of information acquisition as in Figure 2.3. Before the start of the game, the true state is drawn by the nature and is not known to the individual. At the initial decision node, the individual decides whether to observe a free signal—if he chooses not to observe (S), he gets a constant payoff of $1/2 + \varepsilon$, where $\varepsilon > 0$ and is sufficiently small; if he chooses to observe (C), the signal is then realized, and this individual needs to choose an action $x \in \{h, l, m\}$ based on the signal realization, where action h and l only generate a payoff of 1 in state H and L respectively, and action m generates a constant payoff of $1/2 + \varepsilon$. Finally, the true state is revealed, and the individual receives the payoff.

The individual is ambiguous about how to interpret the signal such that he perceives a set of likelihood ratios

$$\frac{P(a|H)}{P(b|H)} = \frac{P(b|L)}{P(a|L)} = \alpha \quad \text{where } \alpha \in \{1/4, 4\} \equiv \mathcal{A}.$$

Here, $\alpha = 4$ interprets signal a as more indicative of state H, whereas $\alpha = 1/4$ interprets signal a as more indicative of state L. The true model α^* belongs to $\mathcal{A}$. The individual starts with a prior $\mu_0 = (1/2, 1/2)$ and a model perception $\alpha_0 = 4$. The individual is biased toward state H and updates his belief according to the biased updating rule. The individual is risk neutral, so he evaluates every *plan* according to its expected payoff. Here, a plan is a complete description of the individual's action at each decision node. For example, plan *Chl* says that the individual will choose C at the initial node and then choose action h upon signal a and action l upon signal b.

(i) Suppose that the individual has the commitment power to be ***dynamically consistent***. That is, the individual is able to commit to a plan at the beginning and carry it out throughout the game. It is easy to see that the optimal plan is *Chl*, that is, the individual will accept the free signal and then choose what is optimal according to his initial signal interpretation.[24]

(ii) Suppose that the individual has no commitment power, then his choice will be ***dynamically inconsistent***. At the initial decision node, the individual

[24] Recall that the initial signal interpretation is $\alpha_0 = 4$, so the individual will choose h after signal a and l after signal b.

naively evaluates all plans according to his current model perception, so the *ex-ante* optimal plan is still Chl, which requires him to observe the signal. Once the signal is realized, the individual will interpret the signal as good news for his biased state, state H. For example, if he observed a signal b, he would view b as more indicative of state H by changing his model perception to $\alpha_1 = 1/4$, which leads to a posterior $\mu_1 = \left(\frac{4}{5}, \frac{1}{5}\right)$. Given μ_1, action h generates the highest expected payoff, so the individual will choose action h instead of the planned action l. Therefore, even though Chl is *ex-ante* optimal, the individual is not willing to carry it out *ex-post*, which gives rise to dynamic inconsistency.

(iii) Suppose that the individual has no commitment power and follows a ***sophisticated*** decision rule. That is, the individual can correctly anticipate his future behavior, and he only chooses the best plan from the set of plans that can be actually carried out (see, Strotz (1955), Siniscalchi (2011)). We can think of the individual as consisting of two selves, the *ex-ante* self and the *ex-post* self. The *ex-ante* self evaluates a plan according to the ex-ante belief and model perception, but he understands how the *ex-post* self will react to information. The goal of the *ex-ante* self is to choose a plan that maximizes his own benefits. In this example, if he chooses to observe the signal, the only implementable plan is Chh. It is easy to see that $V(Chh) = 1/2 < 1/2 + \varepsilon = V(S)$, where $V(p)$ denotes the ex-ante expected payoff of plan p, so he will choose *not* to observe the signal at the very beginning. This leads to the seemingly paradoxical phenomenon of **information avoidance**. To explain in words, when the individual anticipates that he will handle information in a highly biased manner, he may find it optimal to reject the information in the first place. In this example, information avoidance can be viewed as the individual's commitment to "debias" himself. $\square$

2.8 Extensions

Some aspects of the model can be relaxed to incorporate more realistic concerns. For example, in the biased updating, individuals are trying to justify some *fixed* bias throughout the learning process. It would also be interesting to consider a **belief-dependent biased updating rule**, in which individuals can modify their bias to tailor their current beliefs. Below is an example.

Example 21 (Belief-Dependent Bias I)**.** Suppose that $\Theta = \{G, B\}$. Consider an individual who is initially biased toward state θ. His bias process $\left\{b_t^\theta\right\}$ evolves according to the following rule

$$b_t^\theta = \begin{cases} G & \text{if } \mu_t\left(G\right) \geq 1/2 \\ B & \text{if } \mu_t\left(G\right) < 1/2 \end{cases}.$$

In other words, if state G is most likely, individuals will become biased toward G; if state B is most likely, individuals will become biased toward B. Given the time-t bias b_t^θ, the next-period belief μ_{t+1}^θ is updated according to the model that can best support the current bias, b_t^θ. This example in a similar spirit to the model of confirmatory bias in Rabin and Schrag (1999), where individuals are biased toward the most likely state according to their current beliefs.[25] □

Remark 2.6. *In Appendix B.2.2, I present a more general framework that accommodates multiple states and a large class of bias-evolving rules. It is worth mentioning that the belief-dependent biased rule is not identical to the maximum likelihood rule. Even though the bias is updated overtime, individuals are still biased at every period, so information is always processed in a non-objective manner.*

[25] In Rabin and Schrag (1999), if an agent received a signal that supports his current belief, he would correctly read the signal; if he received a signal that contradicts his current belief, he would misread the signal with a strictly positive probability.

Another interesting extension is that individuals also care about being correct when justifying their bias. In the benchmark model, individuals keep all models without narrowing them down, so it would be interesting to consider a setup where individuals can also discard some models during the learning. One possibility is to consider the **ρ-maximum likelihood biased updating rule**, where individuals apply the biased rule to a subset of models that pass some likelihood test as follows.

Example 22 (ρ-maximum likelihood biased updating rule)**.** For all $\rho \in [0,1]$, denote by

$$\mathcal{A}_{\rho}^{t} \equiv \left\{ \alpha \in \mathcal{A} : \mu\left(s_1, s_2, ..., s_t | \alpha\right) \geq (1-\rho) \times \max_{\alpha' \in \mathcal{A}} \mu\left(s_1, s_2, ..., s_t | \alpha'\right) \right\}.$$

At time t, individuals only consider the models in $\mathcal{A}_{\rho}^{t}$ and apply the biased rule with these models to obtain the next-period belief. The ρ-maximum likelihood biased rule can be viewed as a combination of the biased updating rule and the ρ-maximum likelihood rule in Epstein and Schneider (2007). Specifically, when $\rho = 0$, it refines the maximum updating rule; when $\rho = 1$, it corresponds to the benchmark biased rule. To characterize limit beliefs, we only need to replace $\mathcal{A}$ with $\mathcal{A}_{\rho}^{\infty}$ in the statement of relevant theorems, where $\mathcal{A}_{\rho}^{\infty}$ represents the set of models that will survive the likelihood test in the limit. □

Some other extensions are worth pursuing. For example, we can consider a richer learning environment. This paper only deals with a simple learning environment where individuals receive exogenous i.i.d. signals. A natural extension is to introduce endogenous signals which could also depend on actions taken by individuals. It is conceivable that similar characterizations should still hold in the limit.

2.9 Concluding Remarks

This paper develops a framework to study biased learning and provides a comprehensive discussion of limit beliefs under this rule. The paper highlights the fact that informational ambiguity can contaminate learning and lead to overconfidence when individuals have bias-confirming incentives. This result mainly comes from two forces. First, ambiguity accommodates multiple interpretations of signals, so individuals have the flexibility to distort evidence toward their favorite directions. Second, ambiguous models have complementary effects, with which individuals can exploit favorable news and hedge against bad news, so biased individuals can become more confident than any Bayesian individual. Due to these forces, when there is sufficient ambiguity, individuals may end up learning incorrectly and more confidently. Some topics are not covered in this paper and are worth pursuing. First, it would be of interest to provide an axiomatic foundation for the biased rule suggested by this paper. Second, this paper focuses on the case where individuals are only driven by bias-justifying incentives, so it would be more realistic to consider other constraints (e.g., preferences for accuracy, consistency). Third, this paper only studies a passive learning problem and it would be interesting to investigate a problem where signals are endogenously driven. Lastly, the learning processes of individuals are independent of each other in this paper, so allowing for dependence across individuals seems a natural next-step (e.g., strategic interactions, social learning).

CHAPTER 3

NAIVE SOCIAL LEARNING WITH HETEROGENEOUS MODEL PERCEPTIONS

3.1 Introduction

A standard setup in social learning is that individuals are located in a network, receive a sequence of informative signals and communicate with their neighbors. In practice, it is usually difficult for individuals to make Bayes inferences from their neighbors when the network is large and complex. As such, one popular approach is to assume that individuals are naive and follow a rule of thumb when aggregating the information from their neighbors.[1] Many previous works have shown that the society can successfully aggregate information even if individuals are naive, which provide support for the wisdom of the crowds. However, one implicit assumption is that *all* individuals are able to interpret their signals precisely. The assumption can be restrictive especially when the size of the society is very large. In reality, the society is more likely to be populated with individuals with heterogeneous perceptions of the true data-generating process, some are correct whereas others may be incorrect. It remains unknown how the learning outcome will change and whether communications improve the learning efficiency.

This paper investigates a social learning problem, where individuals are naive when learning from neighbors and are possibly incorrect when interpret-

[1] Below are a few examples of Bayesian social learning Banerjee (1992) and Bikhchandani et al. (1992) studied learning in a linear network and showed that herding and information cascades can arise. The discussion was extended to general networks by Acemoglu et al. (2011). Gale and Kariv (2003) and Rosenberg et al. (2009) studied a problem where individuals repeatedly observe neighbors' actions and update beliefs by Bayes rule.

ing private information. All individuals are connected through a fixed network and attempt to learn the true state of the world. At the beginning of each period, individuals communicate their current beliefs with their neighbors via some naive learning rule, which is assumed to be the DeGroot's rule in the benchmark model. The communication lasts for multiple periods. After the communication, every individual receives a private signal drawn from a data-generating process conditional on the true state. The key feature of this paper is that individuals can misspecify the true data-generating process, and moreover, they can hold heterogeneous model perceptions. After observing their signals, individuals update the beliefs according to the perceived data-generating processes. The updated beliefs give the next-period beliefs, and the process repeats itself.

The presence of misspecifications brings the following changes to the original problem. First, there dos not exist an obvious state for beliefs to settle on. When all individuals correctly specify the data-generating process, the natural candidate for limit beliefs to settle on is the true state, since every individual's private information must be consistent with the true state. In contrast, with misspecifications, it is not clear which state beliefs will settle on due to the heterogeneity in signal interpretations. Second, some nice statistical properties no longer hold after the introduction of misspecifications, e.g., the martingale property, which makes it difficult to apply the standard techniques in social learning.

In the paper, I present a set of tractable characterizations of limit beliefs for the learning problem. The key idea is that the limit beliefs are closely related to the society's *weighted relative entropy* between all perceived and true data-generating processes, where the weight vector measures the centrality of the network. First, whenever beliefs converge, they must settle on a state that minimizes the society's weighted relative entropy (Theorem 3.1). Intuitively, it describes the idea that the social consensus must minimize the weighted conflicts of all members such that it minimizes the "weighted distance" between individuals' perceived and the true data-generating processes. Second, it is possible that beliefs do not converge, and they will almost surely not converge if (or only if) there dos not exist a state that (strictly) minimizes the society's inner-weighted relative entropy, which is a stronger version of the weighted relative entropy (Theorem 3.2). One implication based on these characterizations is that social learning can lead to *group irrationality*. The society may end up believing in state that is contradictory to *every* member's private information. Even if every individual are able to learn optimally in an isolated environment, the society may not be able to achieve the optimal learning after communications. Examples are presented in Section 3.2.

Related Literature. The paper belongs to the literature on social learning, especially the naive social learning. DeMarzo et al. (2003) investigated a social learning problem where individuals aggregate neighbors' opinions repeatedly and showed that individuals may exhibit persuasion bias. Golub and Jackson (2010) looked at a situation where individuals average their neighbors' beliefs according to the DeGroot's rule, and they found that naive rule can successfully aggregate information. Jadbabaie et al. (2012) studied a naive social learning problem, and their setup is most similar to this paper. In their paper, individuals receive infinitely many signals and exchange beliefs with neighbors repeatedly. They showed that complete learning can be achieved under some regular conditions. Molavi et al. (2018) adopts an axiomatic approach and presents a comprehensive analysis of non-Bayesian social learning. They provided conditions for learning and non-learning. In the paper, they also adopted the weighted relative entropy to characterize the learning speed. The difference is that this paper uses this concept and its variants to characterize the state that limit beliefs settle on. Other examples of naive social learning include Li and Tan (2020), in which individuals treat the local network as the whole network and apply Bayesian updating locally; Eyster and Rabin (2010) and Dasaratha and He (2020), which investigated naive social learning in the linear network.

This paper also belongs to the literature on misspecified learning. Bohren (2016) and Bohren and Hauser (2021) investigated a sequential learning problem where individuals hold misspecified beliefs and make Bayes inferences. Frick et al. (2020a) showed that social learning can be fragile with respect to model misspecification. Frick et al. (2020b) proposed an original order based on the Kullback-Leibler divergence to analyze misspecified learning. This paper employs a similar order when characterizing limit beliefs. The notion of the p-inner-weighted relative entropy plays a similar role as their p-dominance relation, but this paper's characterization also combines the network structure and induces a different ordering of states. Bowen et al. (2021) studied a social learning problem where individuals are learning from shared news. They showed that misperceptions can lead to polarizations. This paper is also related to the literature on learning under model uncertainty. In the paper, we can think of the society as facing an aggregate model uncertainty in the sense that each individual can perceive a different version of the true data-generating processes. Chen (2019) showed that sufficient model uncertainty can result in inefficient learning in the form of information cascades and incorrect herding. This paper has a similar idea that correct learning can break down if some individual is sufficiently misspecified (see Example 24 and 28).

3.2 Examples

Consider the simplest network consisting of two individuals, $N = \{1, 2\}$. The state space is Θ, and the true state $\theta^* \in \Theta$ is unknown to both individuals. Each individual holds a full-support prior μ_{i0}. In each period $t \in \{1, 2, ...\}$, they first communicate their last-period beliefs and apply the DeGroot's rule to aggregate information. The interim belief $v_{i,t}$ is given by

$$v_{i,t} = \frac{1}{2}\mu_{1,t-1} + \frac{1}{2}\mu_{2,t-1}.$$

After the communication, a signal is then realized, but individuals may not correctly understand the data-generating process. They then apply Bayes rule to the interim belief to obtain the posterior, that is,

$$\forall \theta \in \Theta : \quad \mu_{i,t}(\theta) = \frac{v_{i,t}(\theta)\, \hat{l}_i(s_{i,t}|\theta)}{\sum_{\theta' \in \Theta} v_{i,t}(\theta')\, \hat{l}_i(s_{i,t}|\theta')},$$

where $\hat{l}_i$ denotes the data-generating process perceived by individual i. If both individuals correctly specify the true data-generating processes, it is known that they will learn the true state almost surely under moderate conditions, hence the wisdom of the crowds can achieve.[2] However, correct learning may not be achievable with the following "harmless" specifications.

Example 23. [Coarse Thinking] There are three possible states $\Theta = \{\alpha, \beta, \gamma\}$. Suppose that the model perceptions $\left(\hat{l}_1, \hat{l}_2\right)$ are

$\hat{l}_1(s\|\theta)$	H	L
α	9/10	1/10
β	1/2	1/2
γ	2/3	1/3

$\hat{l}_2(s\|\theta)$	H	L
α	1/2	1/2
β	9/10	1/10
γ	2/3	1/3

.

[2] Jadbabaie et al. (2012) and Molavi et al. (2018) showed that all individuals will almost surely learn the true state if the society does not face the aggregate identification problem.

Signals are i.i.d. according to $\hat{l}_1$. At time t, individuals need to take an action $a_{i,t} \in \{h, l\}$ to maximize the one-period expected payoff of $u(a, \theta)$, where

$$u(h,\theta) = \begin{cases} 1 & \theta \in \{\alpha, \beta\} \\ 0 & \theta = \gamma \end{cases} \text{ and } u(l,\theta) = 1 - u(h,\theta),$$

so the optimal action in states α and β is h, and the optimal action in state γ is l. We say that learning is optimal if $a_{i,t} \to a(\theta^*)$ as $t \to \infty$ for both i, where $a(\theta^*)$ denotes the optimal action. In this problem, individuals are only interested in learning whether $\theta^* \in \{\alpha, \beta\}$. Although individual 2's model is not precisely specified, he only rearranges the distributions between α and β. At first glance, this rearrangement seems harmless and should not affect the optimal learning. Besides, it is also easy to see that with these specifications, both individuals could achieve optimal learning if they were to learn independently. Perhaps surprisingly, learning may no longer be optimal when individuals can communicate with each other. For instance, suppose that $\theta^* = \beta$, hence the optimal action is h. It can be verified that in social learning, both individuals will assign probability 1 to state γ almost surely in the limit. As a consequence, the society agrees on the sub-optimal action, action l, even if no one would take that action in independent learning. $\square$

Remark 3.1. *This example has a similar spirit to the "coarse thinking". We can think of that individuals partition the state space into two categories, $\mathcal{C}_1 = \{\alpha, \beta\}$ and $\mathcal{C}_2 = \{\gamma\}$. They only care about identifying the correct category and are not interested in learning the exact state in each category. In the example, individuals are "coarsely" correct in the sense that they correctly identify the set of distributions within each category. However, this example shows that even with coarsely correct specifications, individuals may still settle on the incorrect category in the end.*

Example 24. [Qualitative Thinking] Suppose that the state space $\Theta = \{G, B\}$, and the signal space is $S = \{S_G, S_B\}$. The perceived data-generating processes are

$$\begin{array}{c|cc} \hat{l}_i\,(s|\theta) & S_G & S_B \\ \hline G & p_i & 1-p_i \\ B & 1-p_i & p_i \end{array}, \qquad \text{where } p_i > 1/2.$$

The true data-generating process takes the same form but with parameter p^*, where $p^* > 1/2$. Although p_i may not be equal to p^*, both individuals correctly understand the "direction" of each signal, that is, they know that signal S_θ is more indicative of state θ. This type of misspecification also seems innocuous. This is because as long as individuals correctly understand the directions of both signals, they can infer the true state by looking at which signal is the majority—if in the long run, there are more S_G than S_B, they know that the true state is G; otherwise, the true state is B. As a consequence, both individuals are able to identify the true state independently. However, in social learning, the true state may not be learned, and it is possible that beliefs do not converge with probability 1. It can be verified that if there exists some i such that p_i is sufficiently close to 1, beliefs will oscillate with probability 1. $\square$

Remark 3.2. *In this example, individuals are correct qualitatively but may not quantitatively. That is to say, they understand the implication of each signal but may be incorrect in specifying the exact probability. The qualitative learning is sufficient for the correctness in independent learning but not sufficient for that in social learning. In the example, beliefs almost surely do not converge to the truth in the sense that they oscillate infinitely often. In the paper, we will see that beliefs can even converge to the incorrect state with a strictly positive probability (Example 28).*

3.3 Model Setup

Network Structure. A society consists of a finite set of individuals, $N = \{1, 2, 3, ..., n\}$. Individuals are located in a deterministic social network G, where $g_{ij} \in [0, 1]$ reflects the weight that individual i assigns to individual j. Let $N_i \subset N$ denote the neighbors of individual i, which are the set of individuals that i assigns strictly positive weights to. The network is *self-influential* if $g_{ii} > 0$ for all $i \in N$, which means that individuals attach strictly positive weights to themselves. The network is *strongly connected* if for all $i, j \in N$, there exists a path $i_1, i_2, ..., i_k$ connecting them with $i_1 = i$ and $i_k = j$ and $g_{i_m i_{m+1}} > 0$ for all $m \in \{1, 2, ..k - 1\}$. Throughout this paper, I follow the convention and assume that the network G is self-influential and strongly connected. This assumption ensures that G is irreducible and aperiodic.

Signal Structure. The state space Θ is finite, and the true state $\theta^* \in \Theta$ is fixed. Every individual $i \in N$ holds a full-support prior $\mu_{i,0} \in \Delta(\Theta)$ and is trying to learn the true state. At each period $t \in \{1, 2, ...\}$, individual i receives a signal $s_{i,t} \in S_i$ generated by data-generating process $l_i(s|\theta)$ conditional on $\theta = \theta^*$, where the signal space S_i is also finite. Let $L = (l_1, l_2, ..., l_n)$ denote the *true signal structure* of the society. Individuals may misunderstand their true signal structure. Let $\hat{l}_i(s|\theta)$ denote the data-generating process perceived by individual i, and let $\hat{L} = \left(\hat{l}_1, \hat{l}_2, ..., \hat{l}_n\right)$ denote *perceived signal structures* of the society. I assume that: (i) every individual's perceived signals are not perfectly informative, so $\hat{l}_i(s|\theta) \in (0, 1)$ for all s and θ; and (ii) every state is identifiable at the aggregate level, that is, for all $i \in N$ and $\theta, \theta' \in \Theta$, there exists some $s_i \in S_i$ such that $\hat{l}_i(s_i|\theta) / \hat{l}_i(s_i|\theta') \neq 1$.

Belief Updating Process. The learning process consists of two separate stages: learning from neighbors and learning from private signals. This paper follows the literature and assumes that individuals apply the naive rule when learning from neighbors and apply the Bayes rule to incorporate private signals. First, at the beginning of each period t, individuals communicate beliefs with their neighbors back and forth via the DeGroot's rule for $k_t \in \mathbb{Z}_{++}$rounds. For simplicity, I assume that k_ts are bounded by some constant $K < \infty$. This phase is referred to as *social learning*, and the resulting beliefs are denoted by vector $\boldsymbol{v_t} = (v_{1,t}, ..., v_{n,t})$, where

$$\boldsymbol{v_t}(\theta) = \underbrace{G \cdot ... \cdot G}_{k_t} \cdot \boldsymbol{\mu_{t-1}}(\theta),$$

for all $\theta \in \Theta$ and $t \geq 1$. Second, individuals apply the Bayes rule to update the beliefs from the communication stage. This phase is referred to as *private learning*, and the resulting beliefs give the current-period posterior $\boldsymbol{\mu_t} = (\mu_{1,t}, ..., \mu_{n,t})$, so

$$\boldsymbol{\mu_t}(\theta) = BU\left(\boldsymbol{v_t}(\theta), \boldsymbol{s_t} | \hat{L}\right),$$

where (i) $\boldsymbol{\mu_t}$ represents the vector of time-t beliefs, (ii) $BU\left(\boldsymbol{\mu_t}, \boldsymbol{s_t} | \hat{L}\right)$ represents the Bayes update of $\boldsymbol{\mu_t}$ conditional on signals $\boldsymbol{s_t}$ and according to perceived models $\hat{L}$. To get a better idea of how the updating works, suppose that $k_t = 1$. The updating rule can be written as follows,

$$\mu_{i,t}(\theta) = BU\left(v_{i,t}(\theta), s_{i,t} | \hat{l}_i\right), \qquad \text{where } v_{i,t}(\theta) = \sum_{j \in N(i)} g_{i,j} \times \mu_{j,t-1}(\theta)$$

for all $i \in N$, $\theta \in \Theta$ and $t \geq 1$. As can be seen, individuals first take a simple average of their neighbors' beliefs and then apply the Bayes rule to the updated belief.

Remark 3.3. *The benchmark model assumes that individuals apply the DeGroot's* *model to communicate beliefs. The results in this paper can be extended to a wider class of naive rules, which will be discussed in Section 3.4.4.*

3.4 Main Results: Characterizations of Limit Beliefs

3.4.1 Relative Entropy

To characterize the asymptotic dynamics of beliefs, I first introduce the concept of relative entropy as follows.

Definition 9. For all $i \in N$ and $\theta \in \Theta$, define the ***relative entropy*** of state θ for individual i as

$$D_{KL}^{i}(\theta) \equiv \mathbb{E} \log \left(\frac{l_i(s|\theta^*)}{\hat{l}_i(s|\theta)} \right).$$

[3]Further let Θ_i be the set of states that minimize $D_{KL}^{i}(\theta)$ for individual i.

The relative entropy provides a notion of distance between individual i's perceived and true data-generating processes. The set Θ_i represents the states that best fit the true distribution under individual i's perceived model in the sense that they induce the smallest relative entropy. In the case of independent learning, Θ_i is equal to the set of states that limit beliefs will settle on.

Lemma 3.1. *If individual i were to learn independently, we have $\mu_{i,t}(\Theta_i) \rightarrow 1$ with probability 1.*

In other words, all states outside of Θ_i will be assigned zero probability in the limit. The result is based on Berk (1966). Intuitively, it means that individuals

[3]More precisely, $D_{KL}^{i}(\theta) = \sum_{s \in S_i} l_i(s|\theta^*) \times \log \left(\frac{l_i(s|\theta^*)}{\hat{l}_i(s|\theta)} \right)$.

will believe in the best-fitting state in the limit. Further suppose that individuals participate in the social learning as described in the setup.

Lemma 3.2. *If individuals participate in the naive social learning, and if they correctly specify their models, we have $\mu_{i,t}(\theta^*) \to 1$ with probability 1 for all i.*

The lemma corresponds to the work on naive social learning with correct model perceptions as in Jadbabaie et al. (2012) and Molavi et al. (2018). The intuition is that if every individual correctly specifies the true data-generating process, we have $\theta^* \in \Theta_i$ for all i. By assumption, there is no identification problem at the aggregate level, $\{\theta^*\} = \cap_{i=1}^{n} \Theta_i$, hence θ^* is the only state that best fits everyone's information. Conceivably, beliefs can only settle on θ^* whenever they converge, and the convergence can be proved using some martingale arguments.

3.4.2 Characterization of Convergent Beliefs

However, when some individuals misspecify the true model, it is no longer true that $\theta^* \in \Theta_i$ for all i, so there does not exist a obvious state for limit beliefs to settle on. The following theorem provides a simple characterization for limit beliefs in this case.

Theorem 3.1. *Whenever beliefs converge, they converge to a point-mass belief on θ_0 satisfying*

$$\theta_0 \in \arg\min_{\theta \in \Theta} \sum_{i=1}^{n} w_i \times D_{KL}^{i}(\theta) \tag{3.1}$$

except for null events, where $\boldsymbol{w} = (w_1, ..., w_n)$ is the stationary distribution of G.[4]

[4]In other words, $\boldsymbol{w}$ satisfies: (i) $\boldsymbol{w}^T = \boldsymbol{w}^T G$, and (ii) $\sum w_i = 1$, $w_i \geq 0$.

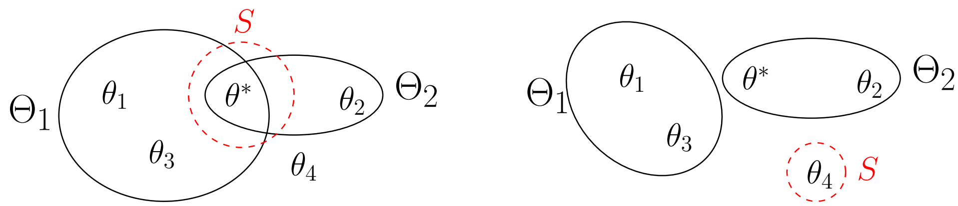

Figure 3.1: Illustration of Limit Beliefs

Theorem 3.1 says that beliefs can only settle on a state that minimizes the ***weighted relative entropy*** of the society in the limit. The weight $\boldsymbol{w}$ corresponds to the stationary distribution of the Markov process with a transition matrix equal to G. More importantly, $\boldsymbol{w}$ describes the eigenvector *centrality* of the communication network, where a higher w_i means that individual i is better connected and more influential in the network. Intuitively, the characterization says that the society can only agree on a state that minimizes the average distance between each member's perceived and true data-generating processes. In other words, social learning will select a state that is the least controversial.

If all individuals correctly specify the data-generating processes, Theorem 3.1 implies that the state that minimizes the society's weighted relative entropy must belong to the intersection of all Θ_is. In other words, only the states in the intersection of all Θ_is can "survive". Due to the fact the intersection leads to a smaller set, there are less states confusing the true state in social learning, which corresponds to the wisdom of the crowds. In the left graph of Figure 3.1, individual 1 can only identify $\Theta_1 = \{\theta_1, \theta^*, \theta_3\}$, and individual 2 can only identify $\Theta_2 = \{\theta^*, \theta_2\}$. In social learning, they can identify $\Theta_1 \cap \Theta_2 = \{\theta^*\}$, and correctly learning can be achieved.

If some individuals incorrectly specify the data-generating processes, the wisdom of the crowd does not have a straightforward parallel. In the left graph of Figure 3.1, we have $\Theta_1 \cap \Theta_2 = \emptyset$, so state satisfies both individuals' information. A natural conjecture would be that limit beliefs will accumulate on $\Theta_1 \cup \Theta_2$ so that both individuals' information can be satisfied to some extent. Perhaps surprisingly, Theorem 3.1 implies that beliefs may settle on a state outside of $\Theta_1 \cup \Theta_2$, for example state θ_3 in the right graph of Figure 3.1. This is comes from the simple fact that the state that minimizes the weighted relative entropy may not minimize the relative entropy for all individuals. As a consequence, the society can settle on a state that all individuals would eliminate in independent learning. This explains the group irrationality as in in Example 23.

Example 25. (Coarse Thinking, continued) In Example 23, it is easy to see that

$$G = \begin{bmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{bmatrix} \qquad \text{and} \qquad w^T = (1/2, 1/2)\,,$$

which comes from the symmetric network structure. The perceived data-generating processes are

$\hat{l}_1(s\|\theta)$	H	L
α	$9/10$	$1/10$
β	$1/2$	$1/2$
γ	$2/3$	$1/3$

$\hat{l}_2(s\|\theta)$	H	L
α	$1/2$	$1/2$
β	$9/10$	$1/10$
γ	$2/3$	$1/3$

,

and the true signal distribution $l^* = (1/2, 1/2)$. Let $I(\theta)$ denote the weighted relative entropy of state θ. It is easy to verify that

$$I(\alpha) = w_1 \times D^1_{KL}(\alpha) + w_2 \times D^2_{KL}(\alpha) = \frac{1}{4} \times \log\left(\frac{5}{9}\right) + \frac{1}{4} \times \log(5)\,.$$

By symmetry, we have $I(\beta) = I(\alpha)$. Similarly, we have $I(\gamma) = \frac{1}{2}\log\left(\frac{3}{4}\right) + \frac{1}{2}\log\left(\frac{3}{2}\right)$. It can be seen that $I(\gamma) < I(\alpha) = I(\beta)$, so state γ minimizes the society's weighted relative entropy. Theorem 3.1 implies that beliefs can only converge to δ_γ. It can be further verified that beliefs will actually converge. The verification can be found in the Appendix. $\square$

A Regular Assumption

It is well known that in the Bayesian learning, beliefs do not converge when there are multiple states that minimize the relative entropy.[5] Similar results are also true in the naive social learning. To facilitate the discussion, I impose the following assumption throughout the paper.

Assumption 3.1. *There is a unique state θ_0 that minimizes the society's weighted relative entropy.*

Under this assumption, there is a unique state that minimizes the "weighted distance" between the perceived and true data-generating processes. This excludes an obvious source of non-convergence, which comes from the multiplicity of entropy-minimizing states.

3.4.3 Characterization of Non-convergent Beliefs

Theorem 3.1 says that whenever beliefs converge, they must settle on a state that minimizes the society's weighted relative entropy. One natural question is that if there is a unique state that minimizes the society's weighted relative entropy, will beliefs converge and settle on the state in the limit? Unfortunately, this is not true. Recall that in Example 24, both individuals correctly specify the "direction" of each signal, so the true state uniquely minimizes the society's weighted relative entropy. However, beliefs do not converge in Example 24. Therefore, the next question will be under which conditions will beliefs fail to converge?

[5] This comes from the fact that the log likelihood ratio between any two entropy-minimizing states constitutes a zero-mean random walk, so it oscillates between $-\infty$ and $+\infty$ with probability 1 (see Berk (1966) for a concrete example).

To answer the question, I introduce a stronger version of the weighted relative entropy below.

Definition 10. For all $\theta, \theta' \in \Theta$, define an order $\succ^*$ such that

$$\theta \succ^* \theta' \qquad \text{if} \qquad \mathbb{E} \log \left(\sum_{i=1}^{n} w_i \times \frac{\hat{l}_i \left(s_i | \theta' \right)}{\hat{l}_i \left(s_i | \theta \right)} \right) < 0,$$

where $\boldsymbol{w} = (w_1, ..., w_n)$ denotes the stationary distribution of G. The weak order $\succeq^*$ is defined by replacing the strict inequality with a weak one.

Here, $\theta \succeq^* \theta'$ is called that θ has a lower ***inner-weighted relative entropy*** than state θ'. To see how the name is motivated, notice that if θ has a lower weighted relative entropy than θ', or

$$\sum w_i \times D^i_{KL} \left(\theta \right) \leq \sum w_i \times D^i_{KL} \left(\theta' \right),$$

we will have

$$\sum_{i=1}^{n} w_i \times \mathbb{E} \log \left(\frac{\hat{l}_i \left(s_i | \theta' \right)}{\hat{l}_i \left(s_i | \theta \right)} \right) \leq 0.$$

As can be seen, the inner-weighted relative entropy moves the weight vector $\boldsymbol{w}$ inside the log function.[6]

Remark 3.4. *The inner-weighted relative entropy induces a stronger order than the weighted relative entropy. Let $\succeq_*$ denote the order induced by the weighted relative entropy. From the concavity of the log function, we have if $\theta \succeq^* \theta'$, then $\theta \succeq_* \theta'$.*

[6] The order induced by the inner-weighted relative entropy is related to the $p-$dominance order in Frick et al. (2020b)). The order in their paper is defined as $\theta \succeq^p \theta'$ if $\mathbb{E} \left(\frac{\hat{l}(s|\theta')}{\hat{l}(s|\theta)} \right)^p \leq 1$. The order in this paper differs from theirs in the following aspects. First, this paper is about naive learning in social network, and the order embeds a network centrality parameter $\boldsymbol{w}$. Second, this paper's order features "double average": (i) the likelihood ratio is averaged (using the centrality vector); (ii) a log function is applied to the averaged likelihood ratio and is then averaged (using the expectation).

Notice that if the network has one person, these two notions become equivalent, which is the simple relative entropy.

I then define the state that minimizes the (inner-)weighted relative entropy as follows.

Definition 11. State θ minimizes the inner-weighted relative entropy if $\theta \succeq^* \theta'$ for all $\theta' \neq \theta$. State θ strictly minimizes the inner-weighted relative entropy if the relation is strict.

Unlike the weighted relative entropy, it is possible that there is no state that minimizes the inner-weighted relative entropy. In other words, the order induced by the inner-weighted relative entropy can be ***incomplete***. Recall that the inner-weighted relative entropy induces a stronger order, which also implies that it is more difficult to rank two states under this notion. As a consequence, it is possible that two states can not be compared, which leads to the incompleteness. Interestingly, it turns out the incompleteness is the key for the non-convergence.

Theorem 3.2. *There exists some $k_0 < \infty$ such that for all $\{k_t\}$ with $\liminf k_t \geq k_0$, beliefs almost surely do not converge*

(i) if there is no state that minimizes the inner-weighted relative entropy;

(ii) only if there is no state that strictly minimizes the inner-weighted relative entropy.

Let's first ignore the restriction on $\{k_t\}$. Theorem 3.2 says that the existence of the a state that minimizes the inner-weighted relative entropy serves as an *almost necessary and sufficient condition* for beliefs not to converge. The only gap between the necessary and sufficient conditions is whether the order is strict, so the characterization is very tight. The restriction on $\{k_t\}$ says that individuals communicate for sufficiently many rounds in the limit. With this restriction, the characterization in Theorem 3.2 holds in *all but finite* possible rounds of communications. The only possibility that Theorem 3.2 might break down is when the communication occurs for small number of rounds. To save notation, I refer to the condition in Theorem 3.2 as "*under adequate communications*" henceforth. The goal of the restriction is for technical simplicity, and its necessity is an interesting question to explore. [7]

Remark 3.5. *The threshold value k_0 depends on the network structure, signal structure and the perceptions. For some simple networks, e.g., complete networks, it is straightforward to see that the characterization holds for all possible k_t.*

The following example illustrates the difference between convergence and non-convergence.

Example 26. Consider the same setup as in Example 23. Now suppose that the model perceptions $\left(\hat{l}_1, \hat{l}_2\right)$ are

$\hat{l}_1(s\|\theta)$	H	L
α	x	$1-x$
β	$1/2$	$1/2$
γ	$2/3$	$1/3$

$\hat{l}_2(s\|\theta)$	H	L
α	$1/2$	$1/2$
β	x	$1-x$
γ	$2/3$	$1/3$

,

where $x \in (0,1)$. Signals are i.i.d. according to $\hat{l}_1$, and that $\theta^* = \beta$, so the true distribution $l^* = (1/2, 1/2)$. Let $I(\theta)$ denote the weighted relative entropy of

[7] If communications occur for sufficiently many rounds at each period, individuals' beliefs will become sufficiently close after the communication. More importantly, their belief ratios are sufficiently close to 1, so the proof does not need to worry that belief ratios go to infinity near the Dirac beliefs. It is possible that these restrictions can be relaxed or even dropped.

state θ. It can be verified that

$$I(\alpha) = I(\beta) = \frac{1}{4}\log\frac{1}{2x} + \frac{1}{4}\log\frac{1}{2-2x} \quad \text{and} \quad I(\gamma) = \frac{1}{2}\log\left(\frac{3}{4}\right) + \frac{1}{2}\log\left(\frac{3}{2}\right)$$

When x is sufficiently small ($x \to 0$) or sufficiently large ($x \to 1$), we have $I(\gamma) < I(\alpha) = I(\beta)$, so state γ is the unique state that minimizes the weighted relative entropy. Theorem 3.1 implies that γ is a candidate state for limit beliefs to settle on. Example 23 and 25 show that beliefs will almost surely converge to δ_γ when x is large, i.e., $x = 9/10$. However, convergence does *not* occur for small xs.

$$\begin{aligned} H(\alpha,\gamma) \equiv & \mathbb{E}\log\left(w_1 \times \frac{\hat{l}_1(s_1|\alpha)}{\hat{l}_1(s_1|\gamma)} + w_2 \times \frac{\hat{l}_2(s_2|\alpha)}{\hat{l}_2(s_2|\gamma)}\right) \\ = & \frac{1}{4} \times \log\left[\left(\frac{3}{4}x + \frac{3}{8}\right)\left(\frac{3}{4}x + \frac{3}{4}\right)\left(\frac{15}{8} - \frac{3}{2}x\right)\left(\frac{9}{4} - \frac{3}{2}x\right)\right]. \end{aligned}$$

It can be verified that $H(\alpha,\gamma) > 0$ when x is sufficiently small, so we have $\gamma \not\succeq^* \alpha$. One the other hand, since $\succeq^*$ is stronger than the order induced by the weighted relative entropy, and γ minimizes the weighted relative entropy, so we also have $\alpha \not\succeq^* \gamma$. Symmetrically, we have $\gamma \not\succeq^* \beta$ and $\beta \not\succeq^* \gamma$. As a consequence, there is no state that minimizes the inner-weighted relative entropy when x is sufficiently small. Theorem 3.2 implies that beliefs almost surely do not converge. Figure 3.2 provides an illustration. The left graph corresponds to $x = 9/10$, where beliefs converge; the right graph corresponds to $x = 1/10$, where beliefs oscillate. □

One implication of Theorem 3.2 is that the society can face group irrationality in the sense that beliefs may fail to settle on the true state even if it minimizes the relative entropy for each individual. Let's re-examine Example 24.

Example 27. (Qualitative Thinking, continued) Consider the case in Example

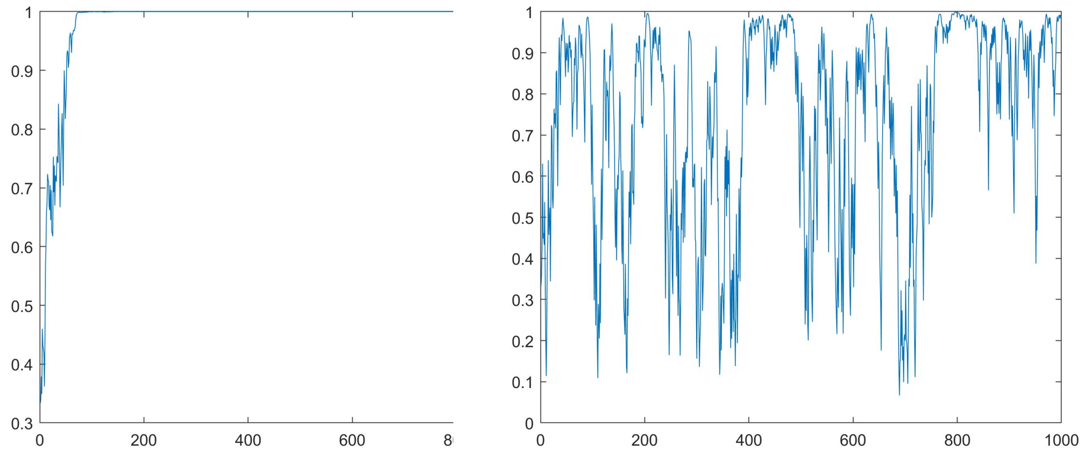

Note: The horizontal axis is t, and the vertical axis is $v_t(\gamma)$. In the left graph, $x = 1/10$. In the right graph, $x = 9/10$.

Figure 3.2: (Non) Convergence of beliefs

24, where the perceived data-generating processes are

$\hat{l}_i(s\vert\theta)$	S_G	S_B
G	p_i	$1-p_i$
B	$1-p_i$	p_i

, where $p_i > 1/2$,

and the true model is $p > 1/2$. Suppose that $\theta^* = G$, then it is easy to see that G minimizes the relative entropy for each individual hence minimizes the weighted relative entropy. Therefore, we also have $B \not\succeq^* G$. To show the other direction, we notice that

$$
\begin{aligned}
&H(G,B) \equiv \mathbb{E}\log\left(w_1 \times \frac{\hat{l}_1(s_1|G)}{\hat{l}_1(s_1|B)} + w_2 \times \frac{\hat{l}_2(s_2|G)}{\hat{l}_2(s_2|B)}\right) \\
&= p^2 \log\left(\frac{p_1}{2-2p_1} + \frac{p_2}{2-2p_2}\right) + p(1-p)\log\left(\frac{p_1}{2-2p_1} + \frac{1-p_2}{2p_2}\right) \\
&+ (1-p)p\log\left(\frac{1-p_1}{2p_1} + \frac{p_2}{2-2p_2}\right) + (1-p)^2\log\left(\frac{1-p_1}{2p_1} + \frac{1-p_2}{2p_2}\right).
\end{aligned}
$$

Fixing other terms, as $p_i \to 1$, we have $H(G,B) \to \infty$. Therefore, if one individual has a sufficiently large p_i, we have $G \not\succeq^* B$ as well. In this case, beliefs almost surely do not converge. □

3.4.4 Generalized DeGroot's Style Learning

To get a better idea of the mechanism, it would be helpful to place the DeGroot's style learning to a general class of learning rules and investigate how results change.

3.4.5 Generalized DeGroot's Style Learning

Recall that under the DeGroot's rule, individuals average their neighbors' beliefs. This section focuses on a generalized version of the DeGroot's rule, where individuals average the p-th order of their neighbors' beliefs. Below is the description of the learning rule.

The p-DeGroot's Rule. Let Ψ denote a social learning rule, which is defined as a self-mapping on $\Delta^n(\Theta)$. For any belief $\boldsymbol{\mu} = (\mu_1, ..., \mu_n)$, the vector $\Psi(\boldsymbol{\mu}) = (\Psi_1(\boldsymbol{\mu}), ..., \Psi_n(\boldsymbol{\mu}))$ represents the beliefs after one-round of communication. This section focuses on a updating rule, where individuals average the p-th order of their neighbors' beliefs. More precisely,

$$(\Psi_i(\boldsymbol{\mu})(\theta))^p = \left[\sum_{j \in N(i)} g_{ij} \times \left(\mu_j(\theta)\right)^p\right] \times C,$$

where C is a normalizing constant to ensure that all probabilities sum up to 1. In other words, the p-th order of the posterior is equal to the averaged p-th order of the prior. This paper calls the updating rule the ***p*-DeGroot's rule**.[8]

[8] To be more precise, $\Psi_i(\boldsymbol{\mu})(\theta) = \frac{\left[\sum_{j \in N(i)} g_{ij} \times \left(\mu_j(\theta)\right)^p\right]^{1/p}}{\sum_{\theta' \in \Theta}\left[\sum_{j \in N(i)} g_{ij} \times \left(\mu_j(\theta')\right)^p\right]^{1/p}}$. In Molavi et al. (2018), this updating rule belongs to the class of the weakly separable learning rule, which features homogeneous of degree 1 and takes a CES functional form (see their Example 1).

Remark 3.6. *By varying p, we can describe a variety of belief aggregation rules. Two important special cases are $p = 1$ and p close to 0. When $p = 1$, it becomes the standard DeGroot's rule, so*

$$\Psi_i(\boldsymbol{\mu})(\theta) = \sum_{j \in N(i)} g_{ij} \times \mu_j(\theta).$$

When $p \to 0$, it approaches the log-linear rule axiomatized by Molavi et al. (2018), where

$$\log \frac{\Psi_i(\boldsymbol{\mu})(\theta)}{\Psi_i(\boldsymbol{\mu})(\theta')} = \sum_{j \in N_i} g_{ij} \times \log \frac{\mu_j(\theta)}{\mu_j(\theta')}.$$

Under this rule, individuals take the average of the log-likelihood ratios of their neighbors. Henceforth, I use $p = 0$ to refer to the log-linear rule.

Belief Updating Process. The belief updating process is similar to the benchmark model. At the beginning of each period t, individuals communicate beliefs with their neighbors according to the p-DeGroot's rule for $k_t \in \mathbb{Z}_{++}$rounds. The resulting beliefs vector $\boldsymbol{v_t} = (v_{1,t}, ..., v_{n,t})$ is given by

$$\boldsymbol{v_t}(\theta) = \underbrace{\Psi \circ ... \circ \Psi}_{k_t}[\boldsymbol{\mu_{t-1}}(\theta)],$$

which is given by the k-th iteration of function Ψ. After the communication state, individuals then apply the Bayes rule to the aggregated beliefs, which gives the current period beliefs. The process repeats itself infinitely many times.

3.4.6 Characterizations of Limit Beliefs

Interestingly, when individuals adopt this general class of updating rules, we can establish similar results as in the previous section. I first characterize the convergent beliefs as follows.

Lemma 3.3. *Suppose that individuals adopt the p-DeGroot's rule. Let θ_0 be the state that minimizes the weighted relative entropy,*

$$\theta_0 \in \arg\min_{\theta \in \Theta} \sum_{i=1}^{n} w_i \times D_{KL}^{i}(\theta), \tag{3.2}$$

(i) When $p \geq 0$, whenever beliefs converge, they can only converge to δ_{θ_0} except for null events.

(ii) When $p \leq 0$, beliefs will converge to δ_{θ_0} with a strictly positive probability under adequate communications.

Lemma 3.3 shows that the characterization in Theorem 3.1 still holds in some forms, that is, the state that minimizes the society's weighted relative entropy provides a benchmark for limit beliefs. We also notice that the statements for $p > 0$ and $p < 0$ take slightly different forms. Their differences are discussed below.

$p > 0$: Belief Non-convergence

The discussion of $p > 0$ is very similar to the benchmark case. I first define a stronger version of the weighted relative entropy as follows.

Definition 12. For all $\theta, \theta' \in \Theta$, define an order $\succeq_p^*$ such that

$$\theta \succeq_p^* \theta' \qquad \text{if} \qquad \frac{1}{p}\mathbb{E}\log\left(\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)}\right)^p\right) \leq 0,$$

where $\boldsymbol{w} = (w_1, ..., w_n)$ denotes the stationary distribution of G, and $\theta \succeq_p^* \theta''$ is called that θ has a lower ***p-inner-weighted relative entropy*** than state θ'.

Definition 13. State θ minimizes the p-inner-weighted relative entropy if $\theta \succeq_p^* \theta'$ for all $\theta' \neq \theta$. State θ strictly minimizes the p-inner-weighted relative entropy if the relation is strict.

Notice that the inner-weighted relative entropy in the last section corresponds to the special case where $p = 1$. Similarly, when $p > 0$, the p-inner-weighted relative entropy induces a *stronger* order than the weighted relative entropy, so $\succeq_p^*$ can be ***incomplete***, which can lead to belief non-convergence.

Theorem 3.3. *Suppose that individuals apply the p-DeGroot's rule with $p > 0$. Under adequate communications, beliefs almost surely do not converge*

(i) if there is no state that minimizes the p-inner-weighted relative entropy;

(ii) only if there is no state that strictly minimizes the p-inner-weighted relative entropy.

The statement of Theorem 3.3 is almost identical to Theorem 3.2, so the discussion is skipped here. Let's then look at the case where $p < 0$.

$p < 0$: Multiplicity of Limits

We first notice that when $p < 0$, the p-inner-weighted relative entropy induces a *weaker* order than the weighted relative entropy. In other words, whenever $\theta \succeq_* \theta'$, we must have $\theta \succeq_p^* \theta'$. In this case, the order induced by the p-inner-weighted relative entropy is complete (since $\succeq_*$) is complete, so the existence of a weighted-entropy-minimizing state is guaranteed. Though the order is no longer incomplete, a new issue arises, which is that the induced order can be ***intransitive***. A direct implication is that there may exist *multiple* states that strictly minimize the p-inner-weighted relative entropy, so we may have multiple limit beliefs.

Theorem 3.4. *Suppose that individuals apply the p-DeGroot's rule with $p < 0$. Under adequate communications, beliefs will converge with a strictly positive probability. Specifically, beliefs will converge to δ_θ with a strictly positive probability for all θs that strictly minimize the p-inner-weighted relative entropy.*

Notice that even if there is a unique state that minimizes the weighted relative entropy, beliefs can settle on a different state in the limit. In this situation, the group irrationality arises because the society can settle on an incorrect state with a strictly positive probability, even if all individuals are able to identify the true state.

Example 28. Consider the case in Example 24, where the perceived data-generating processes are

$$\begin{array}{c|cc} \hat{l}_i\left(s|\theta\right) & S_G & S_B \\ \hline G & p_i & 1-p_i \\ B & 1-p_i & p_i \end{array}, \qquad \text{where } p_i > 1/2,$$

and the true model is $p > 1/2$. Suppose that $\theta^* = G$. But individuals adopt the following naive rule to aggregate beliefs,

$$v_{i,t}\left(\theta\right) = \frac{\left(\frac{1}{2}\frac{1}{\mu_{1,t-1}(\theta)} + \frac{1}{2}\frac{1}{\mu_{2,t-1}(\theta)}\right)^{-1}}{\sum_{\theta' \in \Theta}\left(\frac{1}{2}\frac{1}{\mu_{1,t-1}(\theta')} + \frac{1}{2}\frac{1}{\mu_{2,t-1}(\theta')}\right)^{-1}},$$

which corresponds to $p = -1$. Under this rule,

$$\frac{1}{v_{i,t}\left(\theta\right)} \sim \frac{1}{2}\frac{1}{\mu_{1,t-1}\left(\theta\right)} + \frac{1}{2}\frac{1}{\mu_{2,t-1}\left(\theta\right)},$$

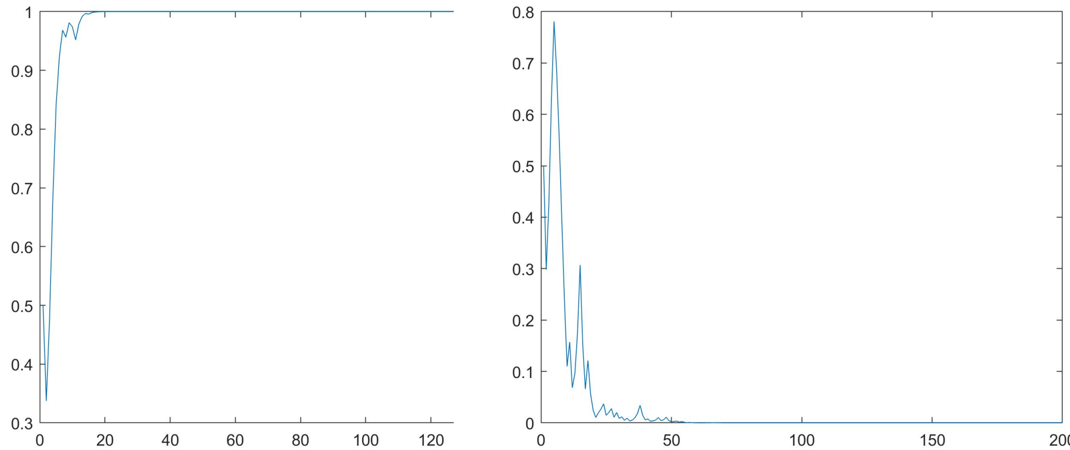

Note: The horizontal axis is t, and the vertical axis is $v_t(G)$. The prior $\mu_{1,0} = \mu_{2,0} = (0.5, 0.5)$, the true parameter $p = 0.6$, and the perceived parameters are $p_1 = 0.55, p_2 = 0.8$.

Figure 3.3: Multiple Belief Limits

so individuals are averaging the reciprocal of neighbors' beliefs. Since G minimizes the relative entropy for each individual, so we have $B \succ_{-1} G$. Notice that

$$H_{-1}(B, G) \equiv -\mathbb{E} \log \left(w_1 \times \left(\frac{\hat{l}_1(s_1|B)}{\hat{l}_1(s_1|G)} \right)^{-1} + w_2 \times \left(\frac{\hat{l}_2(s_2|B)}{\hat{l}_2(s_2|G)} \right)^{-1} \right)$$
$$= -p^2 \log \left(\frac{p_1}{2 - 2p_1} + \frac{p_2}{2 - 2p_2} \right) - p(1-p) \log \left(\frac{p_1}{2 - 2p_1} + \frac{1 - p_2}{2p_2} \right)$$
$$- (1-p) p \log \left(\frac{1 - p_1}{2p_1} + \frac{p_2}{2 - 2p_2} \right) - (1-p)^2 \log \left(\frac{1 - p_1}{2p_1} + \frac{1 - p_2}{2p_2} \right).$$

Fixing other terms, as $p_i \to 1$, we have $H_{-1}(B, G) \to -\infty$, so $G \succ_{-1} B$. To sum up, both B and G strictly minimize the -1-inner-weighted relative entropy, hence beliefs converge to δ_G and δ_B with a strictly positive probability. Figure 3.3 provides an illustration. □

Remark 3.7. *As $p \to 0$, the induced order $\succeq_p^*$ is becoming more and more complete and transitive. When $p = 0$, which corresponds to the log-linear rule, $\succeq_p^*$ degenerates to the weighted relative entropy, which is both complete and transitive. It is easy to* *show that beliefs will converge to the state that minimizes the weighted relative entropy with probability 1, so Theorem 3.1 tightly characterizes limit beliefs.*

3.5 Conclusion

This paper investigates a social learning problem where individuals are naive in learning from others and may be incorrect in interpreting their own information. This paper derives a set of characterizations of limit beliefs for this problem. One key feature of the characterizations is that the society has a tendency to settle on a state that minimizes the weighted distance between the true and the perceived data-generating processes, and the weight describes the network's centrality. This condition holds for a variety of naive updating rules, so it serves as a good benchmark for limit beliefs. This paper further notes that it is possible that beliefs fail to converge or converge to multiple limits. To describe these properties, this paper employs a variant of the weighted relative entropy. One implication is that it is likely that the society can face group irrationality unless all members precisely interpret the signals. For many updating rules, the society may fail to learn optimally even if each member can achieve optimal learning independently.

APPENDIX A

APPENDIX OF CHAPTER 1

A.1 Omitted Proofs of Chapter 1

A.1.1 Proof of Theorem 1.1

Lemma A.1. *For all data-generating process $F \equiv \left(F^0, F^1\right)$, we have:*

(1) $F^0\left(r\right) > F^1\left(r\right)$ except when both are equal to 0 or 1;

(2) $\frac{F^0(r)}{F^1(r)} \geq \frac{1}{r}$ and $\frac{1-F^1\left(\frac{1}{r}\right)}{1-F^0\left(\frac{1}{r}\right)} \geq \frac{1}{r}$ for $r \in (0, \infty)$ (strictly when $F^1\left(r\right) > 0$ and $F^0\left(\frac{1}{r}\right) < 1$);

(3) $\frac{F^0(r)}{F^1(r)}$ and $\frac{1-F^1\left(\frac{1}{r}\right)}{1-F^0\left(\frac{1}{r}\right)}$ are weakly decreasing (strictly on $supp\left(F\right)$).

Proof. See Smith and Sørensen (2000) Lemma A.1. □

Lemma A.2. *Define $C_0 = \left[0, \frac{1}{\gamma}\right]$ and $C_1 = [\gamma, \infty]$. Whenever $r_i \in C_\theta$, an information cascade of action θ occurs.*

Proof. It follows directly from the definition of the information cascade. □

Step 1: The increment (or the decrement) of r_i is bounded by some constant.

Lemma A.3. *Under Assumption 1.3, for all $r_i \in \left(\frac{1}{\gamma}, \gamma\right)$, we have*

$$\frac{r_{i+1}}{r_i} \begin{cases} \geq \sqrt{\gamma} & \text{if } a_i = 1 \\ \leq \frac{1}{\sqrt{\gamma}} & \text{if } a_i = 0 \end{cases}.$$

Proof. Suppose that $r_i \in \left(\frac{1}{\gamma}, \gamma\right)$, hence a cascade does not occur. If $a_i = 1$, we have

$$r_{i+1} = \sqrt{\sup_{F_i \in \mathcal{F}_0} \frac{1 - F_i^1\left(\frac{1}{r_i}\right)}{1 - F_i^0\left(\frac{1}{r_i}\right)} \times \inf_{F_i \in \mathcal{F}_0} \frac{1 - F_i^1\left(\frac{1}{r_i}\right)}{1 - F_i^0\left(\frac{1}{r_i}\right)}} \times r_i.$$

Let F_γ be the data-generating process such that $\text{supp}\left(F_\gamma\right) = \left\{\gamma, \frac{1}{\gamma}\right\}$. Intuitively, F_γ is the "most informative" data-generating process that only generates signals with the highest and the lowest likelihood ratios. For all $r_i \in \left(\frac{1}{\gamma}, \gamma\right)$, we have

$$\sup_{F_i \in \mathcal{F}_0} \frac{1 - F_i^1\left(\frac{1}{r_i}\right)}{1 - F_i^0\left(\frac{1}{r_i}\right)} \geq \frac{1 - F_\gamma^1\left(\frac{1}{r_i}\right)}{1 - F_\gamma^0\left(\frac{1}{r_i}\right)} = \frac{\mathbb{P}_{F_\gamma}^1\left(\gamma\right)}{\mathbb{P}_{F_\gamma}^0\left(\gamma\right)} = \gamma, \tag{A.1}$$

where $\mathbb{P}_{F_\gamma}^\theta\left(\gamma\right)$ denotes the probability of observing a signal γ in state θ, the first equality comes from $\text{supp}\left(F_\gamma\right) = \left\{\gamma, \frac{1}{\gamma}\right\}$, and the last equality comes from the definition of normalized signals. From Lemma A.1, we know that

$$\inf_{F_i \in \mathcal{F}_0} \frac{1 - F_i^1\left(\frac{1}{r_i}\right)}{1 - F_i^0\left(\frac{1}{r_i}\right)} \geq 1. \tag{A.2}$$

Combining (C.3) and (C.4), we obtain $r_{i+1} \geq \sqrt{\gamma} \times r_i$ when $a_i = 1$. The discussion for $a_i = 0$ is symmetric. □

Step 2: An information cascade occurs almost surely $\mathbb{P}^*$.

Unbounded Signals. When signals are unbounded, the occurrence of a cascade is easy to see. This is because when $\gamma = \infty$, which Lemma A.3 implies that

$$r_1 = \begin{cases} \infty & \text{if } a_1 = 1 \\ 0 & \text{if } a_1 = 0 \end{cases},$$

so a cascade occurs immediately after the first action.

Bounded Signals. When signals are bounded, Lemma A.3 also implies that we can find some $K < \infty$ such that for all $r_i \in \left(\frac{1}{\gamma}, \gamma\right)$, K consecutive action θs will lead r_i to enter cascade set θ, hence triggering an information cascade of action θ. Specifically, whenever $r_i \geq 1$, K consecutive signals $\lambda_i, \lambda_{i+1}, ..., \lambda_{i+K-1} > 1$ lead to $a_i = a_{i+1} = ... = a_{i+K-1} = 1$ and result in a cascade of action 1 afterwards. Further note that the probability of receiving a signal $\lambda_i > 1$ satisfies

$$\frac{\mathbb{P}^*(\lambda_i > 1)}{1 - \mathbb{P}^*(\lambda_i > 1)} = \frac{1 - \overline{F}_i^0(1)}{\overline{F}_i^0(1)} = \frac{\overline{F}_i^1(1)}{\overline{F}_i^0(1)} \geq \lim_{r \to \min(\text{supp}\overline{F}_i)} \frac{\overline{F}_i^1(r)}{\overline{F}_i^0(r)} \geq \frac{1}{\gamma},$$

where the second equality comes from the symmetry of signals, and the inequality comes from the Lemma A.1 (iii). As a result, we have $\mathbb{P}^*(\lambda_i > 1) \geq \frac{1}{1+\gamma}$, and

$$\mathbb{P}^*(\text{Cascade}|r_i \geq 1) \geq \mathbb{P}^*(\lambda_i, \lambda_{i+1}, ..., \lambda_{i+K-1} > 1|r_i \geq 1) \geq \left(\frac{1}{1+\gamma}\right)^K > 0.$$

Symmetrically, we also have $\mathbb{P}^*(\text{Cascade}|r_i < 1) \geq \left(\frac{\gamma}{1+\gamma}\right)^K > 0$. Therefore, for all possible history h_i, we have $\mathbb{P}^*(\text{Cascade}|h_i) \geq \varepsilon$ for some $\varepsilon > 0$. Levy's 0-1 Law shows that as $i \to \infty$, we $\mathbb{P}^*$-almost surely have

$$\mathbb{P}^*(\text{Cascade}|h_i) \to \mathbb{P}^*(\text{Cascade}|h_\infty) = 1_{\text{Cascade}} \in \{0, 1\}.$$

Recall that $\mathbb{P}^*(\text{Cascade}|h_i) > \varepsilon > 0$ for all i, so $1_{\text{Cascade}} = 1$ $\mathbb{P}^*$-almost surely, in other words, an information cascade almost surely happens.

A.1.2 Proof of Theorem 1.2

Condition (1): Suppose that there exists some $F \in \mathcal{F}_0$ that is discrete at γ. Due to the symmetry, F^0 is discrete at $\frac{1}{\gamma}$. Denote $p = \mathbb{P}_{F^0}\left(\frac{1}{\gamma}\right) > 0$, which is the probability that F^0 puts on $\frac{1}{\gamma}$. Suppose that $a_i = 1$, for $r_i \in \left(\frac{1}{\gamma}, \gamma\right)$, we have:

$$\bar{l}_{i+1} = \bar{l}_i \times \sup_{F_i \in \mathcal{F}_0} \frac{1 - F_i^1\left(\frac{1}{r_i}\right)}{1 - F_i^0\left(\frac{1}{r_i}\right)} \geq \bar{l}_i \times \frac{1 - F^1\left(\frac{1}{r_i}\right)}{1 - F^0\left(\frac{1}{r_i}\right)} \tag{A.3}$$

$$\geq \bar{l}_i \cdot \left[\lim_{r \to \gamma} \frac{1 - F^1\left(\frac{1}{r}\right)}{1 - F^0\left(\frac{1}{r}\right)}\right] = \bar{l}_i \cdot \frac{1 - \frac{1}{\gamma} \cdot p}{1 - p}, \tag{A.4}$$

where the inequality line comes from Property (3) in Lemma A.1, and the last equality comes from the discreteness of signals. Besides, we have $\underline{l}_{i+1} \geq \underline{l}_i$, so

$$r_{i+1} \geq \sqrt{\frac{1 - \frac{1}{\gamma} \cdot p}{1 - p}} r_i \equiv \beta \times r_i$$

Symmetrically, when $a_i = 0$, we have $r_{i+1} \leq \frac{1}{\beta} \times r_i$. From the proof of Theorem 1.1, an information cascade occurs $\mathbb{P}^*$-almost surely.

Condition (2): Suppose that there exists some $F^1 \in \mathcal{F}_0$ such that F^1 is continuously differentiable on $(\gamma - \varepsilon, \gamma)$ with $F^{1'}\left(\gamma^-\right) > \frac{2}{\gamma - 1}$. When F^1 is discrete at γ, an information cascade occurs almost surely as implied by condition (1). I thus only focus on the case where F^1 is continuous at γ. Suppose that $a_i = 1$, we have:

$$r_{i+1} = r_i \cdot \sqrt{\sup_{F_i \in \mathcal{F}_0} \frac{1 - F_i^1\left(\frac{1}{r_i}\right)}{1 - F_i^0\left(\frac{1}{r_i}\right)} \cdot \inf_{F_i \in \mathcal{F}_0} \frac{1 - F_i^1\left(\frac{1}{r_i}\right)}{1 - F_i^0\left(\frac{1}{r_i}\right)}} \geq r_i \cdot \sqrt{\frac{1 - F^1\left(\frac{1}{r_i}\right)}{1 - F^0\left(\frac{1}{r_i}\right)}} \equiv I\left(r_i\right)$$

Let $I'\left(\gamma\right) \equiv \lim_{\delta \to 0} I'\left(\gamma - \delta\right)$ and $f^\theta\left(\gamma\right) \equiv F^{\theta'}\left(\gamma^-\right)$. It is easy to verify

$$I'\left(\gamma\right) = \gamma \cdot \left[\frac{1}{\gamma} + \frac{1}{2}\left(f^0\left(\gamma\right) - f^1\left(\gamma\right)\right)\right] = 1 - \left(\frac{\gamma - 1}{2}\right) f^1\left(\gamma\right) < 0,$$

where the last equality comes from $f^0(\gamma) = \frac{1}{\gamma} f^1(\gamma)$. Because F^1 is continuously differentiable on $(\gamma - \varepsilon, \gamma)$, there exists some $\varepsilon_0 > 0$ such that for all $r \in [\gamma - \varepsilon_0, \gamma)$, $I'(r) < 0$. Since $I(\gamma) = \gamma$, we have $I(r) \geq \gamma$ for all $r \in [\gamma - \varepsilon_0, \gamma]$. For all $r_i \in \left(\frac{1}{\gamma - \varepsilon_0}, \gamma - \varepsilon_0\right)$, if $a_i = 1$, we have:

$$\frac{r_{i+1}}{r_i} \geq \sqrt{\frac{1 - F^1\left(\frac{1}{r_i}\right)}{1 - F^0\left(\frac{1}{r_i}\right)}} \geq \sqrt{\frac{1 - F^1\left(\frac{1}{\gamma - \varepsilon_0}\right)}{1 - F^0\left(\frac{1}{\gamma - \varepsilon_0}\right)}} > 1.$$

So, for all r_i, there exists a $K < \infty$ such that after K action 1s, we have $r_i \geq \gamma - \varepsilon_0$. Also note that if $r_i \in [\gamma - \varepsilon_0, \gamma]$ and $a_i = 1$, we have $r_{i+1} \geq I(r_i) \geq \gamma$, so $K + 1$ consecutive action 1s will trigger a cascade of action 1. Similarly, $K + 1$ consecutive action 0s will trigger a cascade of action 0. Applying the proof of Theorem 1.1 again, we can show that r_i will enter the cascade set almost surely.

A.1.3 Proof of Corollary 1.1

The idea of the proof is to make use of condition (1) in Theorem 1.2. As mentioned in the main text: weak convergence implies that we can construct a F that is discrete at γ and is sufficiently close to G (under d). Below is the explicit construction.

Construction of a Discrete Approximation of G. Let $x_0^n \equiv 1$, $\Delta_n \equiv \frac{\gamma - 1}{n}$, $x_i^n \equiv x_0^n + \Delta_n \cdot i$, note that $x_n^n = \gamma$. Consider the following partition:

$$\tau^n = \left\{ [\frac{1}{x_n^n}, \frac{1}{x_{n-1}^n}), ..., [\frac{1}{x_2^n}, \frac{1}{x_1^n}), [\frac{1}{x_1^n}, 1], (1, x_1^n], ..., (x_{n-1}^n, x_n^n] \right\}$$

Since the benchmark distribution $G \in \mathcal{F}$ (recall that $G(x) = \mathbb{P}_G(\lambda \leq x|1)$, that is G is the data-generating process in state 1 by definition). We have:

$$G\left(x_i^n\right) - G\left(x_{i-1}^n\right) = G^0\left(\frac{1}{x_{i-1}^n}\right) - G^0\left(\frac{1}{x_i^n}\right) = \int_{\frac{1}{x_i^n}}^{\frac{1}{x_{i-1}^n}} dG^0\left(\lambda\right)$$
$$= \int_{\frac{1}{x_i^n}}^{\frac{1}{x_{i-1}^n}} \frac{1}{\lambda} dG\left(\lambda\right) = \frac{1}{v_i^n} \cdot \left[G\left(\frac{1}{x_{i-1}^n}\right) - G\left(\frac{1}{x_i^n}\right)\right],$$

for some $v_i^n \in [x_{i-1}^n, x_i^n)$, where G^0 is the corresponding data-generating process in state 0. So we have:

$$v_i^n = \frac{G\left(\frac{1}{x_{i-1}^n}\right) - G\left(\frac{1}{x_i^n}\right)}{G\left(x_i^n\right) - G\left(x_{i-1}^n\right)}$$

Define the following cutoff points:

$$\varrho^n = \left\{\frac{1}{\gamma}, \frac{1}{v_{n-1}^n}, \frac{1}{v_{n-2}^n}, ..., v_{n-2}^n, v_{n-1}^n, \gamma\right\}$$

Construct a discrete distribution $\mathbb{P}_n\left(.|1\right)$ that puts all the mass on the elements of ϱ^n:

(1) For all $i \leq n-1$, let

$$\mathbb{P}_n\left(v_i^n|1\right) = G\left(\frac{1}{x_{i-1}^n}\right) - G\left(\frac{1}{x_i^n}\right)$$
$$\mathbb{P}_n\left(\frac{1}{v_i^n}|1\right) = G\left(x_i^n\right) - G\left(x_{i-1}^n\right)$$

Let $\mathbb{P}_n\left(v_i^n|0\right) = \frac{1}{v_i^n} \cdot \mathbb{P}_n\left(v_i^n|1\right)$.

(2) For $i = n$, let

$$\mathbb{P}_n\left(\frac{1}{\gamma}|1\right) = \frac{\gamma}{1+\gamma}\left(1 - G\left(x_{n-1}^n\right) + G\left(\frac{1}{x_{n-1}^n}\right)\right)$$
$$\mathbb{P}_n\left(\gamma|1\right) = \frac{1}{1+\gamma}\left(1 - G\left(x_{n-1}^n\right) + G\left(\frac{1}{x_{n-1}^n}\right)\right)$$

Let $\mathbb{P}_n\left(\gamma|0\right) = \frac{1}{\gamma}\mathbb{P}_n\left(\gamma|1\right)$ and $\mathbb{P}_n\left(\frac{1}{\gamma}|0\right) = \mathbb{P}_n\left(\frac{1}{\gamma}|1\right)\gamma$.

The idea of this construction is: we assign all the weights of G in the interval $[x_{i-1}^n, x_i^n)$ (and $[\frac{1}{x_i^n}, \frac{1}{x_{i-1}^n})$) on the cutoff point v_i^n (or $1/v_i^n$). Let F_n be the c.d.f. of $\mathbb{P}_n(.|1)$. It can be verified that: (1) $F_n \in \mathcal{F}$, since by construction, $\mathbb{P}_n$ is a symmetric signal-generating process (with the data-generating process on state 0 given by $\mathbb{P}_n(.|0)$;

(2) $F_n \Rightarrow G$, since as $n \to \infty$, the division becomes finer and finer, the distance between $F_n(.)$ and $G(.)$ shrinks to 0 (i.e., $d(F_n, G) \to 0$).

A.1.4 Proof of Corollary 1.2

Proof. The idea of this proof makes use of condition (2) in Theorem 1.2. I prove this corollary by constructing a $F \in \mathcal{F}$ satisfying condition (2) and satisfy $d(F,G) < \infty$. Then we just need to set $\overline{K} \equiv d(F,G)$. As in the Example 3, I deal with a signal space $S \equiv [0,1]$ and consider the following h:

$$h^1(s) = \begin{cases} 1 + 2\varepsilon(1+\gamma) \cdot s & s \in \left[0, \frac{1}{1+\gamma}\right] \\ 2\varepsilon(1+\gamma) \cdot s + (1-2\varepsilon)\gamma - 2\varepsilon & s \in \left[\frac{\gamma}{1+\gamma}, 1\right] \end{cases}$$

$$h^0(s) = \begin{cases} 2\varepsilon(1+\gamma) \cdot (1-s) + (1-2\varepsilon)\gamma - 2\varepsilon & s \in \left[0, \frac{1}{1+\gamma}\right] \\ 1 + 2\varepsilon(1+\gamma) \cdot (1-s) & s \in \left[\frac{\gamma}{1+\gamma}, 1\right] \end{cases}$$

where $h^\theta(s)$ is the p.d.f. of the data-generating process in state θ and $\varepsilon > 0$. For a signal s, the likelihood ratio induced by it is $\lambda(s) = \frac{h^1(s)}{h^0(s)}$. By changing the variable (from s to λ), we can equivalently express $h^1(s)$ as a p.d.f. $f^1(\lambda)$, where f has support $\left[\frac{1}{\gamma}, \frac{1+2\varepsilon}{\gamma-2\varepsilon}\right] \cup \left[\frac{\gamma-2\varepsilon}{1+2\varepsilon}, \gamma\right]$. Besides $F \in \mathcal{F}$ because it represents a symmetric data-generating process. Due to the full supportness of G, we have: $d(F,G) < \infty$ for all $\varepsilon > 0$.

For $s \in \left[\frac{\gamma}{1+\gamma}, 1\right]$, the normalized signal is

$$\lambda = \frac{h^1(s)}{h^0(s)} = \frac{2\varepsilon(1+\gamma) \cdot s + \gamma - 2\varepsilon(1+\gamma)}{1 + 2\varepsilon(1+\gamma) \cdot (1-s)}$$

so

$$s = \frac{[1 + 2\varepsilon(1+\gamma)]\lambda - \gamma + 2\varepsilon(1+\gamma)}{2\varepsilon(1+\gamma)(\lambda+1)} = \frac{(1+\rho)\lambda - \gamma + \rho}{\rho(\lambda+1)} \text{ where } \rho = 2\varepsilon(1+\gamma)$$
$$\frac{ds}{d\lambda} = \frac{(1+\rho)\rho(\lambda+1) - \rho[(1+\rho)\lambda - \gamma + \rho]}{[\rho(\lambda+1)]^2} = \frac{1+\gamma}{\rho(1+\lambda)^2}$$

the transformed PDF $f^1(\lambda)$ becomes:

$$f^1(\lambda) = \left[\rho \times \frac{(1+\rho)\lambda - \gamma + \rho}{\rho(\lambda+1)} + \gamma - \rho\right] \times \frac{1+\gamma}{\rho(1+\lambda)^2}$$
$$F'(\gamma-) = \lim_{\lambda \to \gamma} f^1(\lambda) = \frac{\gamma}{1+\gamma}\frac{1}{\rho} = \frac{\gamma}{2(1+\gamma)^2 \varepsilon}$$

It is easy to see that there exists some $\bar{\varepsilon} > 0$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$, $F'(\gamma-) > \frac{2}{\gamma-1}$, so condition (2) of Theorem 1.2 is satisfied. Let's just set $\varepsilon = \bar{\varepsilon}/2$ and F is the corresponding distribution function. Let $\overline{K} = d(F, G)$, it is easy to verify that $d(F, G) < \infty$ under the assumptions of G. When $K \geq \overline{K}$, the belief set $\mathcal{F}_0$ satisfies the condition (2) thus an information cascade occurs almost surely. $\square$

A.1.5 Proof of Theorem 1.3

Auxiliary Results: Local Stability under Ambiguity

I first introduce the notion of local stability under ambiguity. Following concepts and results are parallel to those in Bayesian learning, especially learning with misspecified model.

Definition 14. [Local (un)stability under Ambiguity]

(i) State 0 (or state 1) is *locally stable* if there exists some $r \in \mathbb{R}_{++}$ (or $R \in \mathbb{R}_{++}$) and $\varepsilon > 0$ such that $\mathbb{P}_{r_0}(r_i \to 0) > \varepsilon$ (or $\mathbb{P}_{r_0}(r_i \to \infty) > \varepsilon$) for all prior sets Π_0 with average likelihood ratio $r_0 < r$ (or $r_0 > R$).

(ii) State 0 (or state 1) is *locally unstable* if there exists some $r \in \mathbb{R}_{++}$ (or $R \in \mathbb{R}_{++}$) such that $\mathbb{P}_{r_0}(r_i > r) = 1$ (or $\mathbb{P}_{r_0}(r_i < R) = 1$) for all prior sets Π_0 with average likelihood ratio $r_0 < r$ (or $r_0 > R$).

Intuitively speaking, state θ is locally stable if beliefs will converge to δ_θ with a strictly positive probability for all priors within a neighborhood of δ_θ. Symmetrically, priors are locally unstable if beliefs will escape from a small neighborhood almost surely. We have the following results.

Lemma A.4. *Under the assumptions of Theorem 1.3, a herding of action* 0 *(or* 1*) occurs if and only if $r_i \to 0$ (or $r_i \to \infty$).*

Proof. Due to the symmetry, I only prove the result for a herding of action 1.

"If" part. Suppose that $r_i \to \infty$, we must have a herding of action 1, since if an action 0 is observed, we have

$$r_{i+1} = r_i \times \sqrt{\sup_{F_i \in \mathcal{F}_0} \frac{F_i^1(1/r_i)}{F_i^0(1/r_i)} \times \inf_{F_i \in \mathcal{F}_0} \frac{F_i^1(1/r_i)}{F_i^0(1/r_i)}} \leq r_i \times \sqrt{\frac{1}{r_i} \times \frac{1}{r_i}} = 1,$$

where the inequality comes from Lemma A.1 (2), contradicting $r_i \to \infty$.

"Only if" part. Suppose that a herding of action 1 occurs. Lemma A.1 (1) implies that

$$r_{i+1} = r_i \times \sqrt{\sup_{F \in \mathcal{F}_0} \frac{1 - F_i^1(1/r_i)}{1 - F_i^0(1/r_i)} \times \inf_{F \in \mathcal{F}_0} \frac{1 - F_i^1(1/r_i)}{1 - F_i^0(1/r_i)}} \geq r_i,$$

so $\{r_i\}$ is an increasing sequence, hence it converges in $\mathbb{R} \cup \{+\infty\}$. If r_i does not converge to infinity, it must converge to some $R < \infty$. Let F be the crucial DGP that $\mathcal{F}_0$ contains. We then have

$$r_{i+1} \geq r_i \times \sqrt{\frac{1 - F^1(1/r_i)}{1 - F^0(1/r_i)}}. \tag{A.5}$$

Take limit on both sides of (C.1), we obtain $R \geq \sqrt{\frac{1-F^1(1/R)}{1-F^0(1/R)}} \times R$, so $\sqrt{\frac{1-F^1(1/R)}{1-F^0(1/R)}} \leq 1$. However, since F has unbounded signals, Lemma A.1 (1) implies that $\sqrt{\frac{1-F^1(1/R)}{1-F^0(1/R)}} > 1$ when $R < \infty$, which is a contradiction. As a consequence, $r_i \to \infty$. □

Lemma A.5. *If both* 0 *and* 1 *are locally stable, then (i) herding occurs almost surely, and (ii) an incorrect herding occurs with a strictly positive probability.*

Proof. (ii) follows directly from Lemma A.4 and the definition of local stability. It remains to prove that herding occurs almost surely. Since state 1 and 0 are locally stable, once beliefs enter $C = \{r_i < r\} \cup \{r_i > R\}$, it will remain in C with a strictly positive probability. Denote by $H = \{r_i \to 0\} \cup \{r_i \to \infty\}$, which represents the event of herding by Lemma A.4. Notice that whenever r_i is not in C, we know that $r_i \in [r, R]$ is bounded, so K consecutive actions lead beliefs to enter C, which is positive-probability event.[1] After beliefs enter C, with a strictly positive probability, we either have $r_i \to 0$ or $r_i \to \infty$ depending on which neighborhood r_i enters. In other words, a herding will occur with a strictly positive probability. As a result, we can find a constant $\varepsilon' > 0$ such that for all

[1]On $\{r_i \leq R\}$, we have $r_{i+1} \leq r_i \times \sqrt{\frac{F^1(1/r_i)}{F^0(1/r_i)}} \leq r_i \times \sqrt{\frac{F^1(1/R)}{F^0(1/R)}}$ for any $F \in \mathcal{F}_0$ after an action 0. Hence, $r_{i+1}/r_i \leq \sqrt{\frac{F^1(1/R)}{F^0(1/R)}} < 1$, so the decrement is bounded by some constant less than 1. Since $r_i < R < \infty$, finite steps will make $r_i < r$. The case for $\{r_i \geq r\}$ is symmetric.

possible history h_i, $\mathbb{P}^*(H|h_i) > \varepsilon'$. Applying the Levy's 0-1 Law, $\mathbb{P}^*(H|h_i) \to \mathbb{P}^*(H|h_\infty) = 1_H \in \{0,1\}$, so H is a probability-1 event. □

From Lemma A.5, we know that we only need to establish the local stability of both states to prove Theorem 1.3.

Step 1: Establish the Local Stability of State 1

To show that state 1 is locally stable, we need to show that there exists some $R < \infty$ such that for all $r_0 \geq R$, the probability of an action-1 herding is greater than some $\varepsilon > 0$. Recall that

$$\mathbb{P}^0_{r_0}(Herd_1) = \lim_{i\to\infty} \mathbb{P}^0_{r_0}(a_1 = a_2 = ...a_i = 1) = \prod_{i=1}^{\infty}\left[1 - F_i^0\left(\frac{1}{r_i}\right)\right] \geq \prod_{i=1}^{\infty}\left[1 - a\left(\frac{1}{r_i}\right)^{\alpha}\right], \tag{A.6}$$

where r_i represents the average public likelihood ratio after $h_i = (1,1,...,1)$. Recall that

$$r_{i+1} = r_i \times \sqrt{\sup_{F\in\mathcal{F}_0} \frac{1-F^1(1/r_i)}{1-F^0(1/r_i)} \times \inf_{F\in\mathcal{F}_0} \frac{1-F^1(1/r_i)}{1-F^0(1/r_i)}} \geq r_i \times \sqrt{\frac{1-F^1(1/r_i)}{1-F^0(1/r_i)}},$$

where F denotes the model in $\mathcal{F}_0$ such that $x^p = o\left(F^0(x)\right)$. We first state the following lemma.

Lemma A.6. $\sqrt{G_F(1/x)} = \sqrt{\frac{1-F^1(x)}{1-F^0(x)}} \sim 1 + \frac{1}{2}F^0(x)$ *as* $x \to 0$.

Proof. In Rosenberg and Vieille (2019), they showed that

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

or equivalently, $\frac{1-F^1(x)}{1-F^0(x)} \sim 1 + F^0(x)$, so $\sqrt{\frac{1-F^1(x)}{1-F^0(x)}} \sim \sqrt{1+F^0(x)} = 1 + \frac{1}{2}F^0(x) + o\left(F^0(x)\right)$, which proves the lemma. □

Let $q \in (p, \alpha)$ and consider the following limit.

$$\begin{aligned}
\lim_{r\to\infty} \frac{\sqrt{\frac{1-F^1(1/r)}{1-F^0(1/r)}}-1}{\left(1+\frac{1}{r^q}\right)^{1/q}-1} &= \lim_{r\to\infty} \frac{\sqrt{\frac{1-F^1(1/r)}{1-F^0(1/r)}}-1}{\frac{1}{r^q}} \times \lim_{r\to\infty} \frac{\frac{1}{r^q}}{\left(1+\frac{1}{r^q}\right)^{1/q}-1} \\
&= \lim_{r\to\infty} \frac{\frac{1}{2}F^0\left(1/r\right)}{\frac{1}{r^q}} \times \lim_{r\to\infty} \frac{\frac{1}{r^q}}{\left(1+\frac{1}{r^q}\right)^{1/p}-1} \qquad \text{(A.7)} \\
&> \lim_{r\to\infty} \frac{\frac{1}{2}\left(1/r\right)^p}{\frac{1}{r^q}} \times q = \infty,
\end{aligned}$$

where (C.8) follows from Lemma A.6. From the proof of Lemma A.4, we know that $\{r_i\}$ is increasing during an action-1 herd, so $r_i \geq R$ for all i. Therefore, we can choose R to be sufficiently large such that for all $i \geq 0$,

$$\sqrt{\frac{1-F^1\left(1/r_i\right)}{1-F^0\left(1/r_i\right)}} \geq \left(1+\frac{1}{r_i^q}\right)^{1/q},$$

which further implies that

$$r_{i+1} \geq r_i \times \sqrt{\frac{1-F^1\left(1/r_i\right)}{1-F^0\left(1/r_i\right)}} \geq r_i \times \left(1+\frac{1}{r_i^q}\right)^{1/q} = \left(r_i^q+1\right)^{1/q}.$$

After iterations, we can obtain

$$r_i \geq \left(r_0^q+i\right)^{1/q}, \qquad \forall i \geq 1. \qquad \text{(A.8)}$$

After substituting (A.8) into (C.11), we know that for all $r_0 \geq R$,

$$\begin{aligned}
\mathbb{P}^0_{r_0}\left(Herd_1\right) &\geq \prod_{i=1}^{\infty}\left[1-a\times\left(\frac{1}{r_i}\right)^{\alpha}\right] \geq \prod_{i=1}^{\infty}\left[1-a\times\frac{1}{\left(r_0^q+i\right)^{\alpha/q}}\right] \\
&\geq \prod_{i=1}^{\infty}\left[1-a\times\frac{1}{\left(R^q+i\right)^{\alpha/q}}\right].
\end{aligned}$$

Here, we also choose the R to be sufficiently large such that $1 - a \times \frac{1}{R^\alpha} > 0$, so $1 - a \times \frac{1}{(R^q+i)^{\alpha/q}} \in (0,1)$ for all $i \geq 1$. Notice that the infinite product

$\prod_{i=1}^{\infty}\left[1-a\times\frac{1}{(R^q+i)^{\alpha/q}}\right]>0$ if and only if the infinite series $\sum a\times\frac{1}{(R^q+i)^{\alpha/q}}<\infty$. Since $q<\alpha$, we know that $\sum a\times\frac{1}{(R^q+i)^{\alpha/q}}<\infty$, hence convergent, so

$$\mathbb{P}^0_{r_0}\left(Herd_1\right)\geq\prod_{i=1}^{\infty}\left[1-a\times\frac{1}{(R^q+i)^{\alpha/q}}\right]\equiv\varepsilon>0,$$

which establishes the local stability of state 1.

Step 2: Establish the Local Stability of State 0

The case for state 0 is symmetric to Step 1. Let r_i denotes the average likelihood ratio after $h_i=(0,...,0)$. From symmetry, we have

$$\mathbb{P}^0_{r_0}\left(Herd_0\right)=\prod_{i=1}^{\infty}F^0\left(\frac{1}{r_i}\right)=\prod_{i=1}^{\infty}\left[1-F^1\left(r_i\right)\right]\geq\prod_{i=1}^{\infty}\left[1-F^0\left(r_i\right)\right]=\mathbb{P}^0_{1/r_0}\left(Herd_1\right).$$

Roughly speaking, this relation says that the probability of a correct herding is higher than that of an incorrect herding. The intuition is straightforward, as the society is receiving some information, so the action is more likely to be correct than incorrect. From Step 1, there exists R such that $\mathbb{P}^0_{1/r_0}\left(Herd_1\right)\geq\varepsilon>0$ for all $1/r_0>R$. Let $r=1/R$, so we also have $\mathbb{P}^0_{r_0}\left(Herd_0\right)\geq\varepsilon>0$ for all $r_0<r$, which establishes the local stability of state 0.

A.1.6 Proof of Theorem 1.4

Local stability and complete learning

Lemma A.7. *Complete learning occurs if and only if $r_i\to 0$ with probability 1.*

Proof. It follows by the definition of complete learning. First, during complete learning, there must be a herding of action 0 after some point, so $r_i\to 0$ with probability 1. Second, if $r_i \to 0$ with probability 1, a herding of action 0 will eventually occur, since an action 1 will lead to $r_i \geq 1$. $\square$

Lemma A.8. *Complete learning occurs if* 0 *is locally stable and state* 1 *is locally unstable, only if state* 0 *is not locally unstable and state 1 is not locally stable.*

Proof. The proof is similar to the proof of Lemma A.5. (i) "if" part. Since state 1 is locally unstable, beliefs will enter $\{r_i < R\}$ infinitely many often. Whenever $r_i < R$, we can find a finite K such that K consecutive action 0s lead to $r_i < r$, and this probability is greater than some positive constant. From the facts that state 0 is locally stable and that the process $\{r_i\}$ is a Markov process, whenever $\{r_I < r\}$ for some I, $r_i \to 0$ with a probability greater than ε. As a consequence, we can find a constant $\varepsilon' > 0$ such that for all possible history h_i, $\mathbb{P}(r_t \to 0|h_i) > \varepsilon'$. Applying the Levy's 0-1 Law as in the proof of Lemma A.5, we know that complete learning occurs. (ii) "only if" part. If state 0 is locally unstable, beliefs will escape from the neighborhood around $r = 0$ with probability 1, which is inconsistent with complete learning. If state 1 is locally stable, we have $r_i \to \infty$ with a positive probability, which also contradicts complete learning. $\square$

A.1.7 Proof of Theorem 1.4

I first state the following proposition.

Proposition A.1. *Under Assumptions 1.5 and 1.6, we have:*

(a) if for all $F \in \mathcal{F}_0$, $\mathcal{P}(F) \geq \mathcal{P}(\overline{F})$, state 1 *is locally unstable;*

(b) if there exists some $F \in \mathcal{F}_0$ such that $\mathcal{P}(F) < \mathcal{P}(\overline{F})$, state 1 is locally stable;

(c) if for all $F \in \mathcal{F}_0$, $\mathcal{P}(F) \geq \mathcal{P}(\overline{F}) + 1$, state 0 is locally unstable;

(d) if there exists some $F \in \mathcal{F}_0$ such that $\mathcal{P}(F) < \mathcal{P}(\overline{F}) + 1$, state 0 is locally stable.

From Lemma A.8, we know that Proposition A.1 implies Theorem 1.4, so I now prove Proposition A.1 as follows. For simplicity in notation, I define $\overline{\alpha} \equiv \mathcal{P}(\overline{F})$, $\alpha_{max} \equiv \max_{F \in \mathcal{F}_0} \mathcal{P}(F)$ and $\alpha_{min} \equiv \min_{F \in \mathcal{F}_0} \mathcal{P}(F)$. The data-generating processes with the maximum and minimum power are denoted by F_{max} and F_{min}.

Step 1: Proof of Proposition A.1 (a)

To show that state 1 is locally unstable, it suffices to show that a herding of action 1 cannot occur for all priors r_0 sufficiently large.[2] Given r_0, the probability of a herding of action 1 is as follows

$$\lim_{i \to \infty} \mathbb{P}^0_{r_0}(a_1 = a_2 = ...a_i = 1) = \prod_{i=1}^{\infty} \mathbb{P}^0_{r_0}(a_i = 1 | h_i) = \prod_{i=1}^{\infty} \left[1 - F^0\left(\frac{1}{r_i}\right)\right],$$

where r_i represents the average likelihood ratio after $h_i = (1, 1, ..., 1)$. The probability is equal to 0 if and only if $\sum F^0\left(\frac{1}{r_i}\right) = \infty$, or equivalently, $\sum \frac{1}{r_i^{\overline{\alpha}}} = \infty$. Note that $\{r_i\}$ is determined by the following dynamics

$$r_{i+1} = r_i \times \sqrt{\max_{F \in \mathcal{F}_0} \frac{1 - F^1(1/r_i)}{1 - F^0(1/r_i)} \times \min_{F \in \mathcal{F}_0} \frac{1 - F^1(1/r_i)}{1 - F^0(1/r_i)}}.$$

When r_0 is sufficiently large, we have $\frac{1-F^1(1/r_i)}{1-F^0(1/r_i)} \sim 1 + F^0(1/r_i)$ for all i, since $r_i \geq r_0$ is also sufficiently large. Therefore, we have

$$r_{i+1} = r_i \times \sqrt{\frac{1 - F^1_{min}(1/r_i)}{1 - F^0_{min}(1/r_i)} \times \frac{1 - F^1_{max}(1/r_i)}{1 - F^0_{max}(1/r_i)}} \leq r_i \times \frac{1 - F^1_{min}(1/r_i)}{1 - F^0_{min}(1/r_i)}$$

[2] In other words, action 0 occurs infinitely many often. Recall that after an action 0, we must have $r_i \leq 1$, so beliefs cannot remain in a small neighborhood around δ_1.

for all r_0 sufficiently large. By the definition of F_{min}, we have $\frac{1-F^1_{min}(1/r_i)}{1-F^0_{min}(1/r_i)} \sim 1 + F^0_{min}\left(1/r_i\right) \sim 1 + C_{min} \times \frac{1}{r_i^{\alpha_{min}}}$, for some constant $C_{min} > 0$.

Suppose that all $F \in \mathcal{F}_0$, $\mathcal{P}\left(F\right) \geq \mathcal{P}\left(\overline{F}\right)$, in other words, $\alpha_{min} \geq \overline{\alpha}$. We have

$$
\begin{aligned}
\lim_{r\to\infty} \frac{\frac{1-F^1_{min}(1/r)}{1-F^0_{min}(1/r)} - 1}{\left(1 + \frac{2\overline{\alpha}C_{min}}{r^{\overline{\alpha}}}\right)^{1/\overline{\alpha}} - 1} &= \lim_{r\to\infty} \frac{\frac{1-F^1_{min}(1/r)}{1-F^0_{min}(1/r)} - 1}{\frac{2\overline{\alpha}C_{min}}{r^{\overline{\alpha}}}} \times \frac{\frac{2\overline{\alpha}C_{min}}{r^{\overline{\alpha}}}}{\left(1 + \frac{2\overline{\alpha}C_{min}}{r^{\overline{\alpha}}}\right)^{1/\overline{\alpha}} - 1} \\
&= \lim_{r\to\infty} \frac{C_{min} \times \frac{1}{r^{\alpha_{min}}}}{\frac{2\overline{\alpha}C_{min}}{r^{\overline{\alpha}}}} \times \overline{\alpha} \\
&= \frac{1}{2} \times \lim_{r\to\infty} \frac{1}{r^{\alpha_{min}-\overline{\alpha}}} = \begin{cases} 0 & \alpha_{min} > \overline{\alpha} \\ \frac{1}{2} & \alpha_{min} = \overline{\alpha} \end{cases} < 1,
\end{aligned}
$$

which implies that $\frac{1-F^1_{min}(1/r_i)}{1-F^0_{min}(1/r_i)} < \left(1 + \frac{2\overline{\alpha}C_{min}}{r_i^{\overline{\alpha}}}\right)^{1/\overline{\alpha}}$ when r_0 is sufficiently large. Therefore, for all $i \geq 0$,

$$
\begin{aligned}
r_{i+1} &< \left(1 + \frac{2\overline{\alpha}C_{min}}{r_i^{\overline{\alpha}}}\right)^{1/\overline{\alpha}} \times r_i = \left(r_i^{\overline{\alpha}} + 2\overline{\alpha}C_{min}\right)^{1/\overline{\alpha}} \\
r_{i+1} &< \left(r_{i+1}^{\overline{\alpha}} + 2\overline{\alpha}C\right)^{1/\overline{\alpha}} < \left(r_i^{\overline{\alpha}} + 2\overline{\alpha}C_{min} \times 2\right)^{1/\overline{\alpha}} \\
&\ldots \\
r_{i+t} &< \left(r_i^{\overline{\alpha}} + 2\overline{\alpha}C_{min} \times t\right)^{1/\overline{\alpha}}.
\end{aligned}
$$

As a consequence, when r_0 is sufficiently large,

$$
\sum_{i=1}^{\infty} \frac{1}{r_i^{\overline{\alpha}}} > \sum_{i=1}^{\infty} \frac{1}{r_0^{\overline{\alpha}} + 2\overline{\alpha}C \times i} = \infty,
$$

so an herding of action 1 occurs with probability 0, which implies that state 1 is locally unstable.

Step 2: Proof of Proposition A.1 (b)

To show that state 1 is locally stable, we need to show that the probability of an action-1 herding is greater than some $\varepsilon > 0$ when r_0 is large. Recall that

$$\mathbb{P}_{r_0}^{0}\left(Herd_1\right) = \lim_{i\to\infty} \mathbb{P}_{r_0}^{0}\left(a_1 = a_2 = ...a_i = 1\right) = \prod_{i=1}^{\infty}\left[1 - F^0\left(\frac{1}{r_i}\right)\right],$$

so in order to establish local stability, we need to find a *uniform* lower bound of the probability on the RHS for all large r_0s.

Suppose that $F^0\left(x\right) \sim \overline{C} \times x^{\overline{\alpha}}$ for some constant $\overline{C} > 0$. When $r_0 \geq R$ with R sufficiently large, we have $\frac{F^0(x)}{\overline{C}\times x^{\overline{\alpha}}} \in \left[1-\varepsilon_1, 1+\varepsilon_1\right]$ for some $\varepsilon_1 > 0$, so

$$\mathbb{P}_{r_0}^{0}\left(Herd_1\right) = \prod_{i=1}^{\infty}\left[1 - F^0\left(\frac{1}{r_i}\right)\right] \geq \prod_{i=1}^{\infty}\left[1 - \left(1+\varepsilon_1\right) \times \overline{C} \times \frac{1}{r_i^{\overline{\alpha}}}\right]. \tag{A.9}$$

Here, R is sufficiently large such that the infinite product on the RHS is strictly positive. On the other hand, when R is large, we have

$$r_{i+1} = r_i \times \sqrt{\frac{1 - F_{min}^1\left(1/r_i\right)}{1 - F_{min}^0\left(1/r_i\right)} \times \frac{1 - F_{max}^1\left(1/r_i\right)}{1 - F_{max}^0\left(1/r_i\right)}}.$$

Define $\beta = (1-\varepsilon)\frac{C_{min}\times\alpha_{min}}{2}$ for some small $\varepsilon > 0$. We have

$$\begin{aligned}
&\lim_{r\to\infty} \frac{\sqrt{\frac{1-F^1_{min}(1/r_i)}{1-F^0_{min}(1/r_i)} \times \frac{1-F^1_{max}(1/r_i)}{1-F^0_{max}(1/r_i)}} - 1}{\left(1+\frac{\beta}{r^{\alpha_{min}}}\right)^{1/\alpha_{min}} - 1} \\
&= \lim_{r\to\infty} \frac{\sqrt{\frac{1-F^1_{min}(1/r_i)}{1-F^0_{min}(1/r_i)} \times \frac{1-F^1_{max}(1/r_i)}{1-F^0_{max}(1/r_i)}} - 1}{\sqrt{\frac{1-F^1_{min}(1/r_i)}{1-F^0_{min}(1/r_i)}} - 1} \times \lim_{r\to\infty} \frac{\sqrt{\frac{1-F^1_{min}(1/r_i)}{1-F^0_{min}(1/r_i)}} - 1}{\left(1+\frac{\beta}{r^{\alpha_{min}}}\right)^{1/\alpha_{min}} - 1} \\
&= 1 \times \lim_{r\to\infty} \frac{\sqrt{\frac{1-F^1_{min}(1/r_i)}{1-F^0_{min}(1/r_i)}} - 1}{\left(1+\frac{\beta}{r^{\alpha_{min}}}\right)^{1/\alpha_{min}} - 1} \\
&= \lim_{r\to\infty} \frac{\sqrt{\frac{1-F^1_{min}(1/r_i)}{1-F^0_{min}(1/r_i)}} - 1}{\frac{\beta}{r^{\alpha_{min}}}} \times \lim_{r\to\infty} \frac{\frac{\beta}{r^{\alpha_{min}}}}{\left(1+\frac{\beta}{r^{\alpha_{min}}}\right)^{1/\alpha_{min}} - 1} \\
&= \frac{C_{min}\times\alpha_{min}}{2\beta} = \frac{1}{1-\varepsilon} > 1.
\end{aligned}$$

When R sufficiently large, we have

$$r_{i+1} \geq r_i \times \left(1+\frac{\beta}{r_i^{\alpha_{min}}}\right)^{1/\alpha_{min}} = \left(r_i^{\alpha_{min}}+\beta\right)^{1/\alpha_{min}} \Rightarrow r_i \geq \left(r_0^{\alpha_{min}}+\beta\times i\right)^{1/\alpha_{min}}. \tag{A.10}$$

Similarly, $r_i \geq (r_0 + \beta\times i)^{1/\alpha_{min}}$. Combining (C.7) and (C.14), we obtain

$$\begin{aligned}
\mathbb{P}^0_{r_0}(Herd_1) &\geq \prod_{i=1}^{\infty}\left[1-(1+\varepsilon_1)\times\overline{C}\times\frac{1}{r_i^{\overline{\alpha}}}\right] \\
&\geq \prod_{i=1}^{\infty}\left[1-(1+\varepsilon_1)\times\overline{C}\times\frac{1}{\left(r_0^{\alpha_{min}}+\beta\times i\right)^{\overline{\alpha}/\alpha_{min}}}\right] \\
&\geq \prod_{i=1}^{\infty}\left[1-(1+\varepsilon_1)\times\overline{C}\times\frac{1}{\left(R^{\alpha_{min}}+\beta\times i\right)^{\overline{\alpha}/\alpha_{min}}}\right]
\end{aligned}$$

for all $r_0 \geq R$. Here, R is chosen to be sufficiently large such that each term is strictly positive. Suppose that there exists some $F \in \mathcal{F}_0$ such that $\mathcal{P}(F) < \mathcal{P}(\overline{F})$, which implies that $\alpha_{min} < \overline{\alpha}$, so

$$\sum \frac{1}{\left(R^{\alpha_{min}}+\beta\times i\right)^{\overline{\alpha}/\alpha_{min}}} < \infty,$$

which further implies that

$$\mathbb{P}_{r_0}^0\left(Herd_1\right) \geq \prod_{i=1}^{\infty}\left[1-\left(1+\varepsilon_1\right) \times \overline{C} \times \frac{1}{\left(R^{\alpha_{min}}+\beta \times i\right)^{\overline{\alpha}/\alpha_{min}}}\right] \equiv \delta > 0,$$

for all $r_0 \geq R$. In other words, the probability of an action-1 herding is greater than $\delta > 0$, which proves that state 1 is locally stable.

Step 3: Proof of Proposition A.1 (c) & (d)

The proofs of Proposition A.1 (c) and (d) are almost identical to the proofs of (a) and (b). The only difference is that the cutoff value becomes $\mathcal{P}\left(\overline{F}\right)+1$. To see where the difference arises, we note that the probability of an action-0 herd is as follows.

$$\mathbb{P}_{r_0}^0\left(Herd_0\right) = \lim_{i\to\infty} \mathbb{P}_{r_0}^0\left(a_1 = a_2 = ...a_i = 0\right) = \prod_{i=1}^{\infty} F^0\left(\frac{1}{r_i}\right) = \prod_{i=1}^{\infty}\left[1-F^1\left(r_i\right)\right],$$

where r_i denotes the average likelihood ratio after $h_i = (0, ..., 0)$. An action-0 herd occurs with a strictly positive probability if and only if $\sum F^1\left(r_i\right) < \infty$. During a herd of action 0, we have $r_i \to 0$. Besides, it can be verified that $\overline{F}^1\left(x\right) = O\left(x^{\overline{\alpha}+1}\right)$ as $x \to 0$. [3]Therefore, an action-0 herd occurs with a strictly positive probability if and only if $\sum r_i^{\overline{\alpha}+1} < \infty$. The rest of the proofs are exactly symmetric to those of (a) and (b).

[3]Recall that $\overline{F}^0\left(x\right) \sim \overline{C} \times x^{\overline{\alpha}}$ as $x \to 0$, so

$$\lim_{x\to 0}\frac{\overline{F}^1\left(x\right)}{x^{\overline{\alpha}+1}} = \lim_{x\to 0}\frac{\overline{f}^1\left(x\right)}{\left(\overline{\alpha}+1\right)x^{\overline{\alpha}}} = \frac{1}{\overline{\alpha}+1}\lim_{x\to 0}\frac{\overline{f}^0\left(x\right)}{x^{\overline{\alpha}-1}} = \frac{\overline{\alpha}}{\overline{\alpha}+1}\lim_{x\to 0}\frac{\overline{F}^0\left(x\right)}{x^{\overline{\alpha}}} = \frac{\overline{\alpha}}{\overline{\alpha}+1}\overline{C},$$

hence $\overline{F}^1\left(x\right) = O\left(x^{\overline{\alpha}+1}\right)$ as $x \to 0$.

A.1.8 Discussion of Example 5

To show that a cascade occurs with a strictly positive probability, it suffices to construct an example. For simplicity, I assume $\pi_0 = (2/3, 1/3)$ and $\overline{\gamma} > 2$, but examples can be constructed for general π_0 and $\overline{\gamma}$.

Suppose that $a_1 = a_2 = a_3 = 1$. It can be shown that an information cascade occurs for individual 4 when $|\sigma|$ is sufficiently large. Let $\lambda_i(\sigma)$ be the normalized signal such that individual i is indifferent between two actions. When $i = 1$, $\lambda_1(\sigma) = 2$. When $i = 2$, $\lambda_2(\sigma)$ is the solution to $V_2(1) - V_2(0) = 0$, or equivalently,

$$\left[\int_1^{\overline{\gamma}} \left(\frac{\gamma_1 \lambda_2(\sigma)}{2 + \gamma_1 \lambda_2(\sigma)}\right)^{1-\sigma} h(d\gamma_1)\right]^{\frac{1}{1-\sigma}} - \left[\int_1^{\overline{\gamma}} \left(\frac{2}{2 + \gamma_1 \lambda_2(\sigma)}\right)^{1-\sigma} h(d\gamma_1)\right]^{\frac{1}{1-\sigma}} = 0.$$

When $i = 3$, $\lambda_3(\sigma)$ is the solution to $V_3(1) - V_3(0) = 0$, or equivalently,

$$\left[\int_{\lambda_2(\sigma)}^{\overline{\gamma}} \int_1^{\overline{\gamma}} \left(\frac{\gamma_1 \gamma_2 \lambda_3(\sigma)}{2 + \gamma_1 \gamma_2 \lambda_3(\sigma)}\right)^{1-\sigma} h(d\gamma) + \int_1^{\lambda_2(\sigma)} \int_1^{\overline{\gamma}} \left(\frac{\gamma_1 \lambda_3(\sigma)}{2 + \gamma_1 \lambda_3(\sigma)}\right)^{1-\sigma} h(d\gamma)\right]^{\frac{1}{1-\sigma}} \tag{A.11}$$

$$= \left[\int_{\lambda_2(\sigma)}^{\overline{\gamma}} \int_1^{\overline{\gamma}} \left(\frac{2}{2 + \gamma_1 \gamma_2 \lambda_3(\sigma)}\right)^{1-\sigma} h(d\gamma) + \int_{\frac{1}{r_2(\sigma)}}^{\overline{\gamma}} \int_1^{\overline{\gamma}} \left(\frac{2}{2 + \gamma_1 \lambda_3(\sigma)}\right)^{1-\sigma} h(d\gamma)\right]^{\frac{1}{1-\sigma}},$$

where $\gamma = (\gamma_1, \gamma_2)$ and (C.5) comes from that when $\gamma_2 > \lambda_2(\sigma)$, individual 2 must receive a signal $s_2 = h$, but when $\gamma_2 < \lambda_2(\sigma)$, both signals can justify $a_2 = 1$. The expression for $\lambda_4(\sigma)$ can be written analogously. It is easy to see that both $\lambda_i(\sigma)$ is continuous in σ. Denote by $\lambda_i(\infty) = \lim_{\sigma \to \infty} \lambda_i(\sigma)$, which corresponds to the average likelihood ratio under the max-min EU model. It can be verified that

$$\lambda_2(\infty) = 2/\sqrt{\overline{\gamma}}, \quad \lambda_3(\infty) = 2/\overline{\gamma}, \quad \lambda_4(\infty) = 2/\overline{\gamma}^2,$$

so when σ is sufficiently large, $\lambda_4(\sigma)$ is sufficiently close to $2/\overline{\gamma}^2 < 1/\overline{\gamma}$, so an information cascade occurs.

When signals are unbounded, i.e., $\overline{\gamma} = \infty$.

The occurrence of a cascade also exists for unbounded signals, or $\overline{\gamma} = 0$. We have the following fact.

Fact A.1. *Suppose that $h(\gamma) \sim C \times \frac{1}{\gamma^{\alpha}}$ for some $C, \alpha > 0$ as $\gamma \to \infty$.*

(i) When $\sigma = 0$, complete learning occurs.

(ii) When σ is sufficiently large, an information cascade occurs almost surely.

Proof. Suppose that $a_1 = 1$, and that individual 2 received an opposite signal, signal l, and that her signal precision is γ_2. Her utility of each action is (the prior is assumed to be flat).

$$V_2(0) = \left[\int_1^{\infty} \left[\mathbb{P}_{\gamma_1}(\theta = 0|I_2)\right]^{1-\sigma} h(\gamma_1)\, d\gamma_1\right]^{\frac{1}{1-\sigma}} = \left[\int_1^{\infty} \left[\frac{\gamma_2}{\gamma_1+\gamma_2}\right]^{1-\sigma} h(\gamma_1)\, d\gamma_1\right]^{\frac{1}{1-\sigma}}$$
$$V_2(1) = \left[\int_1^{\infty} \left[\mathbb{P}_{\gamma_1}(\theta = 1|I_2)\right]^{1-\sigma} h(\gamma_1)\, d\gamma_1\right]^{\frac{1}{1-\sigma}} = \left[\int_1^{\infty} \left[\frac{\gamma_1}{\gamma_1+\gamma_2}\right]^{1-\sigma} h(\gamma_1)\, d\gamma_1\right]^{\frac{1}{1-\sigma}}.$$

Individual 2 will break the herd only if her signal precision γ_2 satisfies $V_2(0) \geq V_2(1)$. When σ is sufficiently large, or more specifically, when $\sigma > \alpha + 1$, we have

$$\begin{aligned} V_2(0) &\leq \left[M + \int_R^{\infty} \left[\frac{\gamma_2}{\gamma_1+\gamma_2}\right]^{1-\sigma} \frac{2C}{\gamma_1^{\alpha}} d\gamma_1\right]^{\frac{1}{1-\sigma}} \\ &= \left[M + \int_R^{\infty} \frac{2C}{\gamma_2} \times \frac{(\gamma_1+\gamma_2)^{\sigma-1}}{\gamma_1^{\alpha}} d\gamma_1\right]^{-\frac{1}{\sigma-1}} = 0, \end{aligned}$$

for some $M, R < \infty$. It comes from that when $\sigma > \alpha + 1$, we have $\frac{(\gamma_2\gamma_1+1)^{\sigma-1}}{\gamma_1^{\alpha}} \to \infty$ as $\gamma_1 \to \infty$, so the integral diverges. Further notice that $\frac{\gamma_1\gamma_2}{\gamma_1\gamma_2+1} \geq \frac{1}{2}$, so

$$V_2(1) = \left[\int_1^{\infty} \left[\frac{\gamma_1}{\gamma_1+\gamma_2}\right]^{1-\sigma} h(\gamma_1)\, d\gamma_1\right]^{\frac{1}{1-\sigma}} \geq \frac{1}{1+\gamma_2} > 0.$$

To sum up, $V_2(1) > V_2(0)$ for all $s \in S$ and for all $\gamma_2 \in (1, \infty)$, this implies that individual 2 will choose action 1 regardless of her private signal, so an information cascade occurs. $\square$

A.1.9 Discussion of Example 6: Bayesian Model Uncertainty

Recall that in Example 6, the set of model paths is $\mathfrak{F} = \mathcal{F}^{\infty}$, the prior is $Q \in \Delta(\mathfrak{F})$, all signals are i.i.d. and unbounded, and the true model paths is $\overline{\boldsymbol{F}} = (\overline{F}, \overline{F}, ...)$.

Example 6 (i): *If* $Q(\overline{\boldsymbol{F}}) > 0$, *then complete learning occurs.*

Proof. Denote by $\mathbb{P}_{\overline{F}}$ the belief that individuals would form if they knew the true model, and by $\mathbb{P}_Q$ individuals' subjective beliefs under Q. Denote by $l_i^{\overline{F}} = \frac{\mathbb{P}_{\overline{F}}(\theta=1|h_i)}{\mathbb{P}_{\overline{F}}(\theta=0|h_i)}$ the likelihood ratio based on the true model $\overline{F}$, and by l_i^Q the likelihood ratio based on Q. First note that $\left\{l_i^{\overline{F}}\right\}$ is a martingale under the true measure $\mathbb{P}^*$, where $\mathbb{P}^* = \mathbb{P}_{\overline{F}}^0$. The Martingale Convergence Theorem implies that there exists a random variable $l_\infty^{\overline{F}}$ such that $l_i^{\overline{F}} \to l_\infty^{\overline{F}}$ and $l_\infty^{\overline{F}} < \infty$ $\mathbb{P}^*$-almost surely. Since $Q(\overline{\boldsymbol{F}}) > 0$, Kalai and Lehrer (1993) implies that $\mathbb{P}_Q$ merges to $\mathbb{P}_{\overline{F}}$ almost surely ($\mathbb{P}_{\overline{F}}$), that is, for all $\varepsilon > 0$, there exists some $I < \infty$ such that

$$l_i^{\overline{F}} / l_i^Q \in (1-\varepsilon, 1+\varepsilon) \quad \text{for all } i \geq I \qquad \mathbb{P}_{\overline{F}} - a.s.. \tag{A.12}$$

Note that $\mathbb{P}_{\overline{F}} \gg \mathbb{P}^* = \mathbb{P}_{\overline{F}}^0$, the previous relation also holds $\mathbb{P}^*$-almost surely. It implies that l_i^Q converges to some limit l_∞^Q, and that $l_\infty^Q = l_\infty^{\overline{F}}$ $\mathbb{P}^*$-almost surely.

The dynamics of $l_i^{\overline{F}}$ is given by

$$l_{i+1}^{\overline{F}} = \begin{cases} l_i^{\overline{F}} \times \frac{1-\overline{F}^1\left(1/l_i^Q\right)}{1-\overline{F}^0\left(1/l_i^Q\right)} & \text{if } a_i = 1 \\ l_i^{\overline{F}} \times \frac{\overline{F}^1\left(1/l_i^Q\right)}{\overline{F}^0\left(1/l_i^Q\right)} & \text{if } a_i = 0 \end{cases}.$$

In the limit, we have $l_\infty^Q = l_\infty^{\overline{F}} \equiv l_\infty$, where l_∞ satisfies

$$l_\infty = \begin{cases} l_\infty \times \frac{1-\overline{F}^1(1/l_\infty)}{1-\overline{F}^0(1/l_\infty)} & \text{if } a_\infty = 1 \\ l_\infty \times \frac{\overline{F}^1(1/l_\infty)}{\overline{F}^0(1/l_\infty)} & \text{if } a_\infty = 0 \end{cases}.$$

All these claims hold $\mathbb{P}^*$-almost surely. From Lemma A.1, we know that $l_\infty \in \{0, \infty\}$. Since l_∞ is almost surely finite, so we must have $l_\infty = 0$, which means that complete learning occurs $\mathbb{P}^*$-almost surely. □

Example 6 (ii): *If $Q\left(\overline{F}\right) = 0$, complete learning may not occur.*

Proof. Suppose that Q features an independent distribution across individuals, so $Q\left(F_1, ..., F_n\right) = q\left(F_1\right) \times ...q\left(F_n\right)$ for all possible F_is and all n, where q is a distribution over models. As explained in Example 6, the problem is identical to that individuals perceive $F_Q = \mathbb{E}_Q F = \sum_{F \in \text{supp}(q)} F \times q\left(F\right)$. For convenience, I assume that q has a finite support. Suppose that $\overline{F}^0\left(x\right) \sim x^{\overline{\alpha}}$ as $x \to 0$, and that there is some $F \in \text{supp}\left(q\right)$ such that $F^0\left(x\right) \sim x^\alpha$ as $x \to 0$, where $\alpha < \overline{\alpha}$. In this case, $F_Q^0\left(x\right) \geq x^\alpha > x^{\overline{\alpha}}$ as $x \to 0$. In other words, the perceived model F_Q is more informative that the true model $\overline{F}$, so an incorrect herding occurs with a strictly positive probability as implied by Corollary 1.4. □

Example 6 assumes that signals are unbounded. It is also true that learning outcome depends on the prior when signals are bounded. Below is an example.

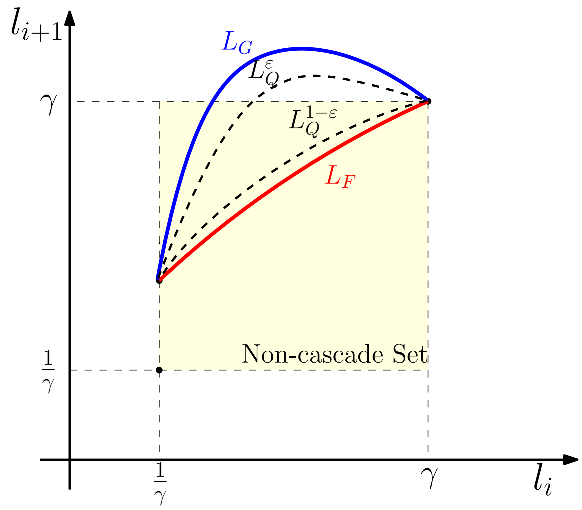

Figure A.1: Information Cascades under Different q

Example 29. Suppose that there are two possible data-generating processes, F and G. In Figure A.1, L_G and L_F represent the public likelihood ratios when individuals perceive the true model as G and F respectively.

(i) A cascade may or may not occur depending on the prior.

For instance, let Q features an independent distribution with a marginal distribution q. If G is assigned a large weight by q, a cascade occurs; if F is assigned a large weight, a cascade does not occur. To see this, suppose that F is assigned a small weight, $q(F) = \varepsilon$, hence G is assigned a large weight. Then, L_Q is very close to L_G, so it enters the cascade set, which implies that an information cascade will occur. On the contrary, suppose that F is assigned a small weight, $q(F) = 1 - \varepsilon$. In this case, L_Q is close to L_F, so it is trapped in the non-cascade set, implying that a cascade will not occur. $\square$

A.2 Supplementary Materials

A.2.1 Multiple States and Actions

Multiple States

The analysis of multiple states becomes more complicated as the equilibrium strategy does not have a simple characterization. It is conceivable that qualitative results still hold. Below is a simple example.

Example 30. [Multi-state Case] Suppose that the state space $\Theta = \{0, 1, ..., K\}$, and the action space $A = \Theta$. Similarly, individuals get a payoff of 1 if the action matches the true state and a payoff of 0 if otherwise, and all priors are flat. Individual i has a data-generating process g_i with the following form

$$\begin{array}{c|cccc} C_i \times g_i\left(s|\theta\right) & s_0 & s_1 & ... & s_K \\ \hline 0 & \gamma_i & 1 & ... & 1 \\ 1 & 1 & \gamma_i & ... & 1 \\ \vdots & \vdots & \vdots & & \vdots \\ K & 1 & 1 & ... & \gamma_i \end{array}, \quad \text{where } \gamma_i \in [1, \overline{\gamma}],$$

and C_i is a normalized term to ensure that all probabilities add up to 1. In this example, s_θ represents the good news for state θ, and all signals are symmetric. Individuals are ambiguous about the γ_is and consider any element in $[1, \overline{\gamma}]$ as possible.

(i) From $a_1 = \theta_1$, we know that $\mathfrak{s}_1 = s_{\theta_1}$, where $\mathfrak{s}_i$ denotes the signal of individual i. It is easy to derive that for all $\mu \in \Pi_1$, we have

$$\mu\left(\theta_1\right) = \frac{\gamma_1}{\gamma_1 + K}, \text{ for some } \gamma_1 \in [1, \overline{\gamma}],$$

and $\mu(\theta) = \frac{1}{\gamma_1 + K}$ for all $\theta \neq \theta_1$.

(ii) From $a_2 = \theta_1$, we know that the data-generating process, g_2 (featured by γ_2), and the signal, $\mathfrak{s}_2$ must satisfy the following inequality:

$$\min_{\pi \in \Pi_1} \frac{\pi(\theta_1) g_2(\mathfrak{s}_2|\theta_1)}{\sum \pi(\theta) g_2(\mathfrak{s}_2|\theta)} \geq \min_{\pi \in \Pi_1} \frac{\pi(\theta') g_2(\mathfrak{s}_2|\theta')}{\sum \pi(\theta) g_2(\mathfrak{s}_2|\theta)} \quad \forall \theta' \in \Theta.$$

If $\mathfrak{s}_2 = s_{\theta_1}$, the signal increases the likelihood of state θ_1 and decreases the likelihood of all other states, so γ_2 can be any element in $[1, \overline{\gamma}]$. If $\mathfrak{s}_2 = s_\theta$ with $\theta \neq \theta_1$, we have

$$\min_{\pi \in \Pi_1} \frac{\pi(\theta_1)}{\pi(\theta) \times \gamma_2 + \sum_{\theta' \neq \theta} \pi(\theta')} \geq \min_{\pi \in \Pi_1} \frac{\pi(\theta) \times \gamma_2}{\pi(\theta) \times \gamma_2 + \sum_{\theta' \neq \theta} \pi(\theta')},$$

which implies that

$$\frac{1}{\gamma_2 + K} \geq \frac{\gamma_2}{\gamma_2 + \overline{\gamma} + K - 1},$$

so $\gamma_2 \leq \overline{\gamma}_2$ for some $\overline{\gamma}_2 \in (1, \overline{\gamma}]$. To sum up, when $\gamma_2 > \overline{\gamma}_2$, a_2 perfectly reveals that $\mathfrak{s}_2 = s_{\theta_1}$, but when $\gamma_2 \leq \overline{\gamma}_2$, a_2 is consistent with all s_θ thus is uninformative. As a consequence, for all $\mu \in \Pi_2$, we have

$$\mu(\theta_1) = \begin{cases} \frac{\gamma_1 \gamma_2}{\gamma_1 \gamma_2 + K} & \gamma_1 \in [1, \overline{\gamma}], \gamma_2 \in [\overline{\gamma}_2, \overline{\gamma}] \\ \frac{\gamma_1}{\gamma_1 + K} & \gamma_1 \in [1, \overline{\gamma}], \gamma_2 \in [1, \overline{\gamma}_2] \end{cases},$$

and $\mu(\theta) = \frac{1 - \mu(\theta_1)}{K}$ for all $\theta \neq \theta_1$.

(iii) From $a_3 = \theta_1$, we know that $\mathfrak{s}_3 = s_\theta$ with $\theta \neq \theta_1$ is consistent with γ_3 satisfying

$$\min_{\pi \in \Pi_2} \frac{\pi(\theta_1)}{\pi(\theta) \times \gamma_3 + \sum_{\theta' \neq \theta} \pi(\theta')} \geq \min_{\pi \in \Pi_2} \frac{\pi(\theta) \times \gamma_3}{\pi(\theta) \times \gamma_3 + \sum_{\theta' \neq \theta} \pi(\theta')},$$

which implies that

$$\frac{1}{\gamma_3 + K} \geq \frac{\gamma_3}{\gamma_3 + \overline{\gamma}^2 + K - 1},$$

so $\gamma_2 \leq \overline{\gamma}_3$ for some $\overline{\gamma}_3 \in (1, \overline{\gamma}]$, where $\overline{\gamma}_3 > \overline{\gamma}_2$. Therefore, for all $\mu \in \Pi_3$, we have

$$\mu(\theta_1) = \begin{cases} \frac{\gamma_1\gamma_2\gamma_3}{\gamma_1\gamma_2+K} & \gamma_1 \in [1, \overline{\gamma}], \gamma_2 \in [\overline{\gamma}_2, \overline{\gamma}], \gamma_3 \in [\overline{\gamma_3}, \overline{\gamma}] \\ \frac{\gamma_1\gamma_2}{\gamma_1\gamma_2+K} & \gamma_1 \in [1, \overline{\gamma}], \gamma_2 \in [\overline{\gamma}_2, \overline{\gamma}], \gamma_3 \in [1, \overline{\gamma_3}] \\ \frac{\gamma_1\gamma_3}{\gamma_1\gamma_3+K} & \gamma_1 \in [1, \overline{\gamma}], \gamma_2 \in [1, \overline{\gamma}_2], \gamma_3 \in [\overline{\gamma_3}, \overline{\gamma}] \\ \frac{\gamma_1}{\gamma_1+K} & \gamma_1 \in [1, \overline{\gamma}], \gamma_2 \in [1, \overline{\gamma}_2], \gamma_3 \in [1, \overline{\gamma_3}] \end{cases}.$$

By induction, for all $i \geq 3$, $\mathfrak{s}_i = s_\theta$ with $\theta \neq \theta_1$ is consistent with γ_i satisfying

$$\frac{1}{\gamma_i + K} \geq \frac{\gamma_i}{\gamma_i + \overline{\gamma}^{i-1} + K - 1}.$$

When i is sufficiently large, the RHS is smaller than the LHS for all $\gamma_i \in [1, \overline{\gamma}]$. That is, for all possible data-generating processes and for all signals, individuals will find it optimal to follow the herd and choose action θ_1. In other words, an information cascade can arise after finite number of individuals for all possible combinations of data-generating processes. $\square$

Multiple Actions

When there are multiple actions, we will still have an information cascade in situations where a cascade is absent. However, with multiple actions, the learning outcomes depend more intricately on the ambiguity attitudes. As shown in an earlier version of this paper, if there exists a safe action and if individuals are ambiguity averse, there will be an information cascade on the safe action. In contrast, if individuals are ambiguity loving, they will only settle on the uncertain actions.

A.2.2 Other Updating Rules

Suppose that individuals hold follow the α-**maximum likelihood rule** as in Epstein and Schneider (2007) and update the model set $\mathcal{F}_0$ over time. That is,

$$\mathcal{F}_{-i} \mid h_i = \left\{ F_{-i} : \mathbb{P}_{F_{-i}} \left(h_i | \sigma_{-i} \right) \geq \alpha \cdot \sup_{F_{-i} \in \mathcal{F}_{-i}} \mathbb{P}_{F_{-i}} \left(h_i | \sigma_{-i} \right) \right\}$$

where $\alpha \in [0, 1]$. Notice that $\alpha = 1$ corresponds to the maximum likelihood updating, and $\alpha = 0$ corresponds to the full Bayesian updating.

Proposition A.2. *Suppose that individuals use $\alpha - MLU$ to update their beliefs. Under Assumption 1.3, for all $\alpha \in [0, 1)$, an information cascade occurs with strictly positive probability.*

Proof. Notice that

$$\mathbb{P}_{F_{-i}} \left(h_i \right) = \mathbb{P}_{F_{-i}} \left(a_1 \right) \mathbb{P}_{F_{-i}} \left(a_2 | a_1 \right) ... \mathbb{P}_{F_{-i}} \left(a_{i-1} | a_1, a_2, ..., a_{i-2} \right)$$

Consider the action profile where $a_1 = a_2 = ... = a_{i-1} = 1$, which is a positive probability event for any finite i. Suppose $F^*_{-i} = \left(F^*_1, ..., F^*_{i-1} \right) \in \arg\max \mathbb{P}_{F_{-i}} \left(h_i \right)$. Notice that the maximum can be obtained. Since $\mathbb{P}_{F^*_1} \left(a_1 \right) = \frac{1}{2}$, which holds for all F_1 continuous at 1. We can just let $F^*_2 = ... = F^*_{i-1}$ be uninformative data-generating process. In this case $\mathbb{P}_{F_2} \left(a_2 | a_1 \right) = ... = \mathbb{P}_{F_{-i}} \left(a_{i-1} | a_1, a_2, ..., a_{i-2} \right) = 1$. The maximum is obtained. I then define $F_{-i} \equiv \left(F^*_1, ..., F^*_{i-2}, F_{i-1} \right)$, then $F_{-i} \in \mathcal{F}_{-i} \mid h_i$ only if $\mathbb{P}_{F_{-i}} \left(h_i \right) \geq \alpha \cdot \mathbb{P}_{F^*_{-i}} \left(h_i \right)$ or

$$\begin{aligned} &\mathbb{P}_{F_{-i}} \left(a_{i-1} | h_{i-1} \right) \geq \alpha \mathbb{P}_{F^*_{-i}} \left(a_{i-1} | h_{i-1} \right) = \alpha \\ &\mathbb{P}^0_{F_{-i}} \left(a_{i-1} | h_{i-1} \right) \mathbb{P}_{F_{-i}} \left(\theta = 0 | h_{i-1} \right) + \mathbb{P}^1_{F_{-i}} \left(a_{i-1} | h_{i-1} \right) \mathbb{P}_{F_{-i}} \left(\theta = 1 | h_{i-1} \right) \geq \alpha \end{aligned} \tag{A.13}$$

Since $a_1 = ... = a_{i-2} = 1$ or $h_{i-1} = \{1, ..., 1\}$, it is easy to verify that: $\mathbb{P}_{F_{-i}}(\theta = 1|h_{i-1}) \geq \mathbb{P}_{F_{-i}}(\theta = 0|h_{i-1})$ for all $F_{-i} \in \mathcal{F}_{-i}$, which means that a sequence of action 1 reveals that state 1 is more likely. Since $a_{i-1} = 1$,we also have $\mathbb{P}_{F_{-i}}(a_{i-1}|h_{i-1}; \theta = 1) \geq \mathbb{P}_{F_{-i}}(a_{i-1}|h_{i-1}; \theta = 0)$. So

$$\mathbb{P}_{F_{-i}}(a_{i-1}|h_{i-1}; \theta = 0)\, \mathbb{P}_{F_{-i}}(\theta = 0|h_{i-1}) + \mathbb{P}_{F_{-i}}(a_{i-1}|h_{i-1}; \theta = 1)\, \mathbb{P}_{F_{-i}}(\theta = 1|h_{i-1})$$
$$\geq \mathbb{P}_{F_{-i}}(a_{i-1}|h_{i-1}; \theta = 0)\, \frac{1}{2} + \mathbb{P}_{F_{-i}}(a_{i-1}|h_{i-1}; \theta = 1)\, \frac{1}{2}$$

So inequality (A.13) is true when

$$\mathbb{P}_{F_{-i}}(a_{i-1}|h_{i-1}; \theta = 0)\, \frac{1}{2} + \mathbb{P}_{F_{-i}}(a_{i-1}|h_{i-1}; \theta = 1)\, \frac{1}{2} \geq \alpha \tag{A.14}$$

Denote r_i as the average public likelihood ratio after observing h_i. Assume that there is no information cascade yet, suppose that $i \geq 2$. So $r_i \in (1, \gamma)$. From individuals' equilibrium strategies, we have: $\mathbb{P}_{F_i}(a_i|h_i; \theta) = 1 - F_i^{\theta}\left(\frac{1}{r_i}\right)$. Consider the following F_i where $\text{supp}(F_i) = \left\{\frac{1}{\gamma}, 1, \gamma\right\}$. Let f_i^{θ} be the p.m.f. of F_i^{θ}. Suppose that $f_i^0(\gamma) = f_i^1\left(\frac{1}{\gamma}\right) = p$ thus $f_i^0\left(\frac{1}{\gamma}\right) = f_i^1(\gamma) = p\gamma$, where $p \in \left[0, \frac{1}{\gamma+1}\right]$. We have:

$$\mathbb{P}_{F_i}(a_i|h_i; 0) = 1 - p\gamma$$
$$\mathbb{P}_{F_i}(a_i|h_i; 1) = 1 - p$$

Then (A.14) gives: $p \leq \frac{2-2\alpha}{1+\gamma}$. Then I just take $p = \frac{2-2\alpha}{1+\gamma}$, the F_i constructed with this p belongs to $\mathcal{F}_{-i} \mid h_i$. We have as long as $r_i \in (1, \gamma)$,

$$\frac{r_{i+1}}{r_i} = \frac{1 - p\gamma}{1 - p} > 1 \text{ when } \alpha < 1$$

So an information cascade occurs after finite steps thus with strictly positive probability. $\square$

Remark A.1. *Notice that a cascade does not occur at $\alpha = 1$, the maximum likelihood updating. This is because under MLU, there exists an "over-fitting problem". A* *herding can be best justified when all followers have uninformative data-generating processes. As such, under MLU, individuals will only keep very uninformative models in* $\mathcal{F}_0$*, so beliefs stop updating after the first person in the herd.*

APPENDIX B

APPENDIX OF CHAPTER 2

B.1 Omitted Proofs of Chapter 2

B.1.1 Some Auxiliary Lemmas

I first define several concepts, some of which are also seen in Berk (1966). Let v be an finite measure on $(\Theta, \mathfrak{B}_\Theta)$ that admits a density function $f_v \equiv \frac{dv}{dm}$ which is continuous on its support, $\text{supp}(v)$.[1] The support of v refers to the closure of points f_v taking strictly positive values on. For all $\mathcal{U} \subset \Theta$, all continuous function $G(\alpha, \theta) : \mathcal{A} \times \Theta \to \mathbb{R}$, and all $t \in \{1, 2, ...\} \cup \{\infty\}$, I define

$$\begin{aligned}
\mathcal{U}||G(\alpha,\theta)||_t^{\theta,v} &\equiv \left[\int_{\mathcal{U}} |G(\alpha,\theta)|^t dv(\theta)\right]^{1/t} \\
\mathcal{U}||G(\alpha,\theta)||_\infty^{\theta,v} &\equiv \sup\{|G(\alpha,\theta)| : \theta \in \mathcal{U} \cap \text{supp}(v)\}.
\end{aligned}$$

Further define

$$H_t(\alpha,\theta) \equiv \frac{1}{t}\sum_{i=1}^{t} \log f(s_i|\alpha,\theta) \text{ and } H(\alpha,\theta) \equiv \mathbb{E}^* \log f(s_i|\alpha,\theta).$$

We have the following lemmas.

Lemma B.1 (Continuity)**.** *(1) $H(\alpha,\theta) : \mathcal{A} \times \Theta \to \mathbb{R}$ is continuous (2) $\mathcal{U}||\exp H(\alpha,\theta)||_\infty^{\theta,v} : \mathcal{A} \to \mathbb{R}_{++}$ is continuous for all compact $\mathcal{U} \subset \Theta$.*

Proof. (1) is a direct result from Assumption 2.1 and 2.2 combined with the dominated convergence theorem. The continuity of $\mathcal{U}||\exp H(\alpha,\theta)||_\infty^{\theta,v}$ follows from (1) and Berge's maximum theorem. □

[1] In the main text, I also use v to denote the benchmark measure on the signal space $(S, \mathcal{B}_S)$, but it is easy to infer from the context which one I am referring to.

Lemma B.2 (Uniform Convergence)**.** *For all compact $\mathcal{U} \subset \Theta$,*

$$\max_{\alpha\in\mathcal{A}} \left| \mathcal{U} || \exp H_t(\alpha,\theta) ||_t^{\theta,v} - \mathcal{U} || \exp H(\alpha,\theta) ||_\infty^{\theta,v} \right| \to 0$$

$\mathbb{P}^*$-almost surely.

Proof. On one hand, we can show that

$$\begin{aligned} &\max_{\alpha\in\mathcal{A}} \left(\mathcal{U} || \exp H_t(\alpha,\theta) ||_t^{\theta,v} - \mathcal{U} || \exp H(\alpha,\theta) ||_\infty^{\theta,v} \right) \\ \le &\max_{\alpha\in\mathcal{A}} \mathcal{U} || \exp H_t(\alpha,\theta) - \exp H(\alpha,\theta) ||_\infty^{\theta,v} \end{aligned}$$

where the inequality is implied by Holder's inequality and the triangle inequality. From the uniform law of large numbers (ULLN), we have:

$$\max_{(\alpha,\theta)\in\mathcal{A}\times\mathcal{U}\cap\mathrm{supp}(v)} |H_t(\alpha,\theta) - H(\alpha,\theta)| \to 0 \quad \mathbb{P}^* - a.s.,$$

which further implies that

$$\max_{\alpha\in\mathcal{A}} \mathcal{U} || \exp H_t(\alpha,\theta) - \exp H(\alpha,\theta) ||_\infty^{\theta,v} \tag{B.1}$$

$$= \max_{(\alpha,\theta)\in\mathcal{A}\times\mathcal{U}\cap\mathrm{supp}(v)} |\exp H_t(\alpha,\theta) - \exp H(\alpha,\theta)| \to 0 \quad \mathbb{P}^* - a.s., \tag{B.2}$$

so we have

$$\limsup \max_{\alpha\in\mathcal{A}} \left(\mathcal{U} || \exp H_t(\alpha,\theta) ||_t^{\theta,v} - \mathcal{U} || \exp H(\alpha,\theta) ||_\infty^{\theta,v} \right) \le 0 \quad \mathbb{P}^* - a.s.. \tag{B.3}$$

On the other hand, we have

$$\begin{aligned} &\min_{\alpha\in\mathcal{A}} \left(\mathcal{U} || \exp H_t(\alpha,\theta) ||_t^{\theta,v} - \mathcal{U} || \exp H(\alpha,\theta) ||_\infty^{\theta,v} \right) \\ \ge &- \underbrace{\max_{\alpha\in\mathcal{A}} \mathcal{U} || \exp H_t(\alpha,\theta) - \exp H(\alpha,\theta) ||_\infty^{\theta,v}}_{(a)} \\ &+ \underbrace{\min_{\alpha\in\mathcal{A}} \left(\mathcal{U} || \exp H(\alpha,\theta) ||_t^{\theta} - \mathcal{U} || \exp H(\alpha,\theta) ||_\infty^{\theta,v} \right).}_{(b)} \end{aligned}$$

I next show that both (a) and (b) converge to 0 almost surely. The convergence of (a) follows from (B.2). Notice that

$$\begin{aligned}\mathcal{U}||\exp H(\alpha,\theta)||_t^{\theta,v} &= \left[\int_{\mathcal{U}\cap\text{supp}(v)} |\exp H(\alpha,\theta)|^t dv(\theta)\right]^{1/t} \\ &\to \sup_{\theta\in\mathcal{U}\cap\text{supp}(v)} |\exp H(\alpha,\theta)| = \mathcal{U}||\exp H(\alpha,\theta)||_\infty^{\theta,v}.\end{aligned}$$

Since $\exp H(\alpha,\theta)$ is continuous, $h_t(\alpha) \equiv \mathcal{U}||\exp H(\alpha,\theta)||_t^{\theta,v}$ is also continuous in α. We know that $\{h_t(\alpha)\}$ converges uniformly to $h_\infty(\alpha)$ from Dini's theorem (since $\{h_t(\alpha)\}$ is an increasing function sequence), so (b) also converges to 0. Therefore,

$$\liminf \min_{\alpha\in\mathcal{A}} \left(\mathcal{U}||\exp H_t(\alpha,\theta)||_t^{\theta,v} - \mathcal{U}||\exp H(\alpha,\theta)||_\infty^{\theta,v}\right) \geq 0 \quad \mathbb{P}^* - a.s.. \tag{B.4}$$

Combining (B.3) and (B.4), we have

$$\max_{\alpha\in\mathcal{A}} \left|\mathcal{U}||\exp H_t(\alpha,\theta)||_t^{\theta,v} - \mathcal{U}||\exp H(\alpha,\theta)||_\infty^{\theta,v}\right| \to 0 \quad \mathbb{P}^* - a.s..$$

□

Corollary B.1 (Composite Uniform Convergence)**.** *For all continuous function $\omega : \mathbb{R}_{++} \to \mathbb{R}$ and all compact $\mathcal{U} \subset \Theta$, we must have*

$$\max_{\alpha\in\mathcal{A}} \left|\omega\left(\mathcal{U}||\exp H_t(\alpha,\theta)||_t^{\theta,v}\right) - \omega\left(\mathcal{U}||\exp H(\alpha,\theta)||_\infty^{\theta,v}\right)\right| \to 0 \quad \mathbb{P}^* - a.s.$$

Proof. Denote set $C \subset \mathbb{R}$ as the range of function $\mathcal{U}||\exp H(\alpha,\theta)||_\infty^{\theta,v}$ (as a function of α defined on $\mathcal{A}$). Due to the fact that $\mathcal{U}||\exp H(\alpha,\theta)||_\infty^{\theta,v}$ is continuous and $\mathcal{A}$ is compact, C must be a compact set as well. Further denote $C^\varepsilon = \{x \in \mathbb{R} : \min_{y\in C}|x-y| \leq \varepsilon\}$, which is the set of points within distance ε to set C, so C^ε is also compact. Since ω is continuous and C^ε is compact,

ω is a uniformly continuous function on C^{ε}. As a result, for all $\xi > 0$, there exists some $\delta > 0$ such that $|\omega(x) - \omega(y)| < \xi$ whenever $|x - y| < \delta$ and $x, y \in C^{\varepsilon}$. From Lemma B.2, for almost all signal paths and for all α, we make the distance between $\mathcal{U}||\exp H_t(\alpha,\theta)||_t^{\theta,v}$ and $\mathcal{U}||\exp H(\alpha,\theta)||_{\infty}^{\theta,v}$ sufficiently small by making t sufficiently large. As such, for all $\xi > 0$, the difference between $\omega\left(\mathcal{U}||\exp H_t(\alpha,\theta)||_t^{\theta,v}\right)$ and $\omega\left(\mathcal{U}||\exp H_t(\alpha,\theta)||_{\infty}^{\theta,v}\right)$ is a.s. uniformly bounded by ξ for sufficiently large t, which establishes the corollary. $\square$

B.1.2 Proof of Lemma 2.1

The idea of the proof resembles Berk (1966), but the main difference is to prove that beliefs converge *uniformly* for all models, which can be established from the previous lemmas.

Proof. For all $x > 0$, I define a sequence of sets $\{\mathcal{U}_{\mathcal{A}}^x\}$ where:

$$\mathcal{U}_{\mathcal{A}}^x \equiv \left\{\theta \in \Theta : \min_{\alpha \in \mathcal{A}} r(\alpha,\theta) \leq \frac{1}{x}\right\}. \tag{B.5}$$

It is easy to verify that $\{\mathcal{U}_{\mathcal{A}}^x\}$ is decreasing in x (with respect to set inclusion) with the limit being $\mathcal{U}_{\mathcal{A}}$. It is easy to verify that $\mathcal{U}_{\mathcal{A}}^x$ is a closed subset of Θ, hence compact. For all open set U containing $\mathcal{U}_{\mathcal{A}}$, there must exist some $x > 0$ such that $\mathcal{U}_{\mathcal{A}}^x \subset U$.[2] If $\mathcal{U}_{\mathcal{A}}^x = \Theta$, the claim is trivially correct, so the rest of the proof focuses on the case where $\mathcal{U}_{\mathcal{A}}^x \subsetneqq \Theta$. I also define a set $\mathcal{U}_{\mathcal{A}}^{-x}$ by flipping the direction of inequality in (C.14), so $\mathcal{U}_{\mathcal{A}}^{-x}$ is also a compact set.

With some abuse of notation, I use μ to denote the measure corresponding

[2]From Berk (1966), suppose not, $\mathcal{U}_{\mathcal{A}}^x \cap U^c$ is a nested system of closed non-empty set, which must have non-empty intersection, i.e., $\cap \mathcal{U}_{\mathcal{A}}^x \cap U^c = \mathcal{U}_{\mathcal{A}} \cap U^c \neq \emptyset$, which contradicts $U \supset \mathcal{U}_{\mathcal{A}}$.

to the density function μ. For all model $\alpha \in \mathcal{A}$, we have:

$$
\begin{aligned}
&\min_{\alpha\in\mathcal{A}} \frac{1}{t}\log\left[\frac{\mu_t\left(U|\alpha\right)}{\mu_t\left(U^c|\alpha\right)}\right] \\
\geq &\min_{\alpha\in\mathcal{A}} \frac{1}{t}\log\left[\frac{\mu_t\left(\mathcal{U}_{\mathcal{A}}^{x}|\alpha\right)}{\mu_t\left(\mathcal{U}_{\mathcal{A}}^{-x}|\alpha\right)}\right] \\
= &\min_{\alpha\in\mathcal{A}} \frac{1}{t}\log\left[\frac{\int_{\mathcal{U}_{\mathcal{A}}^{x}}\prod_{i=1}^{t} f\left(s_i|\alpha,\theta\right)\mu\left(\theta\right)dm\left(\theta\right)}{\int_{\mathcal{U}_{\mathcal{A}}^{-x}}\prod_{i=1}^{t} f\left(s_i|\alpha,\theta\right)\mu\left(\theta\right)dm\left(\theta\right)}\right] \\
= &\min_{\alpha\in\mathcal{A}} \log\frac{\left(\int_{\mathcal{U}_{\mathcal{A}}^{x}}\left(\exp\left(\frac{1}{t}\sum_{i=1}^{t}\log f\left(s_i|\alpha,\theta\right)\right)\right)^{t}\mu\left(\theta\right)dm\left(\theta\right)\right)^{1/t}}{\left(\int_{\mathcal{U}_{\mathcal{A}}^{-x}}\left(\exp\left(\frac{1}{t}\sum_{i=1}^{t}\log f\left(s_i|\alpha,\theta\right)\right)\right)^{t}\mu\left(\theta\right)dm\left(\theta\right)\right)^{1/t}} \\
= &\min_{\alpha\in\mathcal{A}} \log\left[\frac{\mathcal{U}_{\mathcal{A}}^{x}||\exp H_t\left(\alpha,\theta\right)||_t^{\theta,\mu}}{\mathcal{U}_{\mathcal{A}}^{-x}||\exp H_t\left(\alpha,\theta\right)||_t^{\theta,\mu}}\right] \\
\geq &\underbrace{\min_{\alpha\in\mathcal{A}}\left(\log\left[\frac{\mathcal{U}_{\mathcal{A}}^{x}||\exp H_t\left(\alpha,\theta\right)||_t^{\theta,\mu}}{\mathcal{U}_{\mathcal{A}}^{-x}||\exp H_t\left(\alpha,\theta\right)||_t^{\theta,\mu}}\right]-\log\left[\frac{\mathcal{U}_{\mathcal{A}}^{x}||\exp H\left(\alpha,\theta\right)||_\infty^{\theta,\mu}}{\mathcal{U}_{\mathcal{A}}^{-x}||\exp H\left(\alpha,\theta\right)||_\infty^{\theta,\mu}}\right]\right)}_{(a)} \\
&+\min_{\alpha\in\mathcal{A}} \log\left[\frac{\mathcal{U}_{\mathcal{A}}^{x}||\exp H\left(\alpha,\theta\right)||_\infty^{\theta,\mu}}{\mathcal{U}_{\mathcal{A}}^{-x}||\exp H\left(\alpha,\theta\right)||_\infty^{\theta,\mu}}\right]
\end{aligned}
$$

Rearranging the terms of (a), we get

$$
\begin{aligned}
(a) \geq &\min_{\alpha\in\mathcal{A}}\left[\log\left(\mathcal{U}_{\mathcal{A}}^{x}||\exp H_t\left(\alpha,\theta\right)||_t^{\theta,\mu}\right)-\log\left(\mathcal{U}_{\mathcal{A}}^{x}||\exp H\left(\alpha,\theta\right)||_\infty^{\theta,\mu}\right)\right] \\
&+\min_{\alpha\in\mathcal{A}}\left[\log\left(\mathcal{U}_{\mathcal{A}}^{-x}||\exp H_t\left(\alpha,\theta\right)||_\infty^{\theta,\mu}\right)-\log\left(\mathcal{U}_{\mathcal{A}}^{-x}||\exp H\left(\alpha,\theta\right)||_t^{\theta,\mu}\right)\right]
\end{aligned}
\tag{B.6}
$$

Corollary B.1 implies that

$$
\min_{\alpha\in\mathcal{A}}\left[\log\left(\mathcal{U}||\exp H_t\left(\alpha,\theta\right)||_t^{\theta,\mu}\right)-\log\left(\mathcal{U}||\exp H_t\left(\alpha,\theta\right)||_\infty^{\theta,\mu}\right)\right]\to 0 \quad \mathbb{P}^*-a.s.,
$$

for all compact set $\mathcal{U}\subset\Theta$. Recall that $\mathcal{U}_{\mathcal{A}}^{x}$ and $\mathcal{U}_{\mathcal{A}}^{-x}$ are both compact, so the RHS of (C.8) must converge to 0 almost surely, which implies that the lim inf of (a) is non-negative, so we have

$$\begin{aligned}
&\liminf \left(\min_{\alpha \in \mathcal{A}} \frac{1}{t} \log \left[\frac{\mu_t (U|\alpha)}{\mu_t (U^c|\alpha)} \right] \right) \\
\geq & \min_{\alpha \in \mathcal{A}} \log \left[\frac{\mathcal{U}_{\mathcal{A}}^{x} || \exp H(\alpha, \theta) ||_{\infty}^{\theta, \mu}}{\mathcal{U}_{\mathcal{A}}^{-x} || \exp H(\alpha, \theta) ||_{\infty}^{\theta, \mu}} \right] \quad \mathbb{P}^* - a.s. \\
= & \min_{\alpha \in \mathcal{A}} \left(\max_{\theta \in \mathcal{U}_{\mathcal{A}}^{x}} \mathbb{E}^* \log f (s|\alpha, \theta) - \max_{\theta \in \mathcal{U}_{\mathcal{A}}^{-x}} \mathbb{E}^* \log f (s|\alpha, \theta) \right) \\
= & \min_{\alpha \in \mathcal{A}} \left(\min_{\theta \in \mathcal{U}_{\mathcal{A}}^{-x}} \mathbb{E}^* \log \frac{f (s|\alpha^*, \theta^*)}{f (s|\alpha, \theta)} - \max_{\theta \in \mathcal{U}_{\mathcal{A}}^{x}} \mathbb{E}^* \log \frac{f (s|\alpha^*, \theta^*)}{f (s|\alpha, \theta)} \right) \\
= & \min_{\alpha \in \mathcal{A}} \left(\min_{\theta \in \mathcal{U}_{\mathcal{A}}^{-x}} \mathcal{R} (\alpha, \theta) - \min_{\theta \in \mathcal{U}_{\mathcal{A}}^{x}} \mathcal{R} (\alpha, \theta) \right) \\
= & \min_{\alpha \in \mathcal{A}} \left(\min_{\theta \in \mathcal{U}_{\mathcal{A}}^{-x}} r (\alpha, \theta) - \min_{\theta \in \mathcal{U}_{\mathcal{A}}^{x}} r (\alpha, \theta) \right) \geq \frac{1}{x} > 0,
\end{aligned}$$

where the last weak inequality comes from the facts that: (i) $\min_{\theta \in \mathcal{U}_{\mathcal{A}}^{-x}} r (\alpha, \theta) \geq 1/x$ from the definition of $\mathcal{U}_{\mathcal{A}}^{-x}$, and (ii) $r (\alpha, \theta) \geq 0$ from the definition of r. As a result, $\min_{\alpha \in \mathcal{A}} \mu_t (U|\alpha) \to 1$ $\mathbb{P}^*$-almost surely. For all possible bias τ, we must have $\mu_t^{\tau} (U|\alpha) \geq \min_{\alpha \in \mathcal{A}} \mu_t (U|\alpha)$, so the result is proved. □

B.1.3 Proof of Lemma 2.2 and Theorem 2.1

Proof. Denote by $U_t (\theta)$ the time-t utility of an individual with bias θ. We have

$$U_t (\theta) = \max_{\alpha \in \mathcal{A}} \mu_t (\theta|\alpha) = \max_{\alpha \in \mathcal{A}} \frac{\prod_{i=1}^{t} f (s_i|\alpha, \theta) \mu (\theta)}{\int \prod_{i=1}^{t} f (s_i|\alpha, \theta') \mu (\theta') dm (\theta')},$$

or equivalently,

$$\begin{aligned}
\frac{1}{t} \log U_t (\theta) = & \max_{\alpha \in \mathcal{A}} \left(\frac{1}{t} \log \left[\prod_{i=1}^{t} f (s_i|\alpha, \theta) \mu (\theta) \right] \right. \\
& \left. - \frac{1}{t} \log \left[\int_{\Theta} \prod_{i=1}^{t} f (s_i|\alpha, \theta') \mu (\theta') dm (\theta') \right] \right) \\
= & \max_{\alpha \in \mathcal{A}} \left[H_t (\alpha, \theta) - \log \left(\Theta || \exp H_t (\alpha, \theta) ||_{t}^{\theta, \mu} \right) \right] + \frac{1}{t} \log \mu (\theta)
\end{aligned}$$

Uniform law of large number (ULLN) implies that

$$\max_{\alpha\in\mathcal{A}} \left| H_t(\alpha,\theta) - H(\alpha,\theta) \right| \to 0 \quad \mathbb{P}^* - a.s.. \tag{B.7}$$

Corollary B.1 implies that

$$\max_{\alpha\in\mathcal{A}} \left| \log\left(\Theta || \exp H_t(\alpha,\theta) ||_t^{\theta,\mu}\right) - \log\left(\Theta || \exp H(\alpha,\theta) ||_\infty^{\theta,\mu}\right) \right| \to 0 \quad \mathbb{P}^* - a.s.. \tag{B.8}$$

From (B.7) and (C.9), we can find a set E with $\mathbb{P}^*(E) = 1$ such that for all signal path $s^\infty \in E$, for all $\varepsilon > 0$, there exists some T such that for all $t \geq T$, we have

$$\begin{aligned}
\forall \alpha \in \mathcal{A}: \quad & \left| H_t(\alpha,\theta) - H(\alpha,\theta) \right| \leq \varepsilon/2 \\
& \left| \log\left(\Theta || \exp H_t(\alpha,\theta) ||_t^{\theta,\mu}\right) - \log\left(\Theta || \exp H(\alpha,\theta) ||_\infty^{\theta,\mu}\right) \right| \leq \varepsilon/2,
\end{aligned}$$

which implies that for all $t \geq T$,

$$\begin{aligned}
& \max_{\alpha\in\mathcal{A}} \left[H_t(\alpha,\theta) - \log\left(\Theta || \exp H_t(\alpha,\theta) ||_t^{\theta,\mu}\right) \right] \\
\in & \left[\max_{\alpha\in\mathcal{A}} \left[H(\alpha,\theta) - \log\left(\Theta || \exp H(\alpha,\theta) ||_\infty^{\theta,\mu}\right) \right] \pm \varepsilon \right].
\end{aligned}$$

Therefore,

$$\begin{aligned}
& \max_{\alpha\in\mathcal{A}} \left[H_t(\alpha,\theta) - \log\left(\Theta || \exp H_t(\alpha,\theta) ||_t^{\theta,\mu}\right) \right] \\
\to & \max_{\alpha\in\mathcal{A}} \left(H(\alpha,\theta) - \log\left(\Theta || \exp H(\alpha,\theta) ||_\infty^{\theta,\mu}\right) \right) \quad \mathbb{P}^* - a.s..
\end{aligned}$$

Consequently, we get

$$\begin{aligned}
\frac{1}{t} \log U_t(\theta) \to & \max_{\alpha\in\mathcal{A}} \left(H(\alpha,\theta) - \log\left(\Theta || \exp H(\alpha,\theta) ||_\infty^{\theta,\mu}\right) \right) \quad \mathbb{P}^* - a.s. \\
& = \max_{\alpha\in\mathcal{A}} \left(\mathbb{E}^* \log f(s|\alpha,\theta) - \log\left(\max_{\theta\in\Theta} \exp \mathbb{E}^* \log f(s|\alpha,\theta) \right) \right) \\
& = \max_{\alpha\in\mathcal{A}} \left(\mathbb{E}^* \log f(s|\alpha,\theta) - \max_{\theta\in\Theta} \mathbb{E}^* \log f(s|\alpha,\theta) \right) \\
& = \max_{\alpha\in\mathcal{A}} \left(-\mathbb{E}^* \log \frac{f(s|\alpha^*,\theta^*)}{f(s|\alpha,\theta)} + \min_{\theta\in\Theta} \mathbb{E}^* \log \frac{f(s|\alpha^*,\theta^*)}{f(s|\alpha,\theta)} \right) \\
& = -\min_{\alpha\in\mathcal{A}} r(\alpha,\theta)
\end{aligned}$$

Since α_t^θ maximizes $U_t(\theta)$ for each t (hence maximizes $\frac{1}{t}\log U_t(\theta)$), it follows immediately that

$$\mathcal{A}_\infty^\theta \subset \min_{\alpha \in \mathcal{A}} r(\alpha, \theta) = \mathcal{A}^\theta \quad \mathbb{P}^* - a.s.,$$

so individuals will only adopt models that minimize the information potential of state θ. Since only models in $\mathcal{A}^\theta$ will be adopted in the limit, the learning problem is essentially the same as if bias-θ individuals only perceive the model set $\mathcal{A}^\theta$. Using Lemma 2.1, it is easy to verify that limit beliefs of bias-θ individuals will settle on zero-potential states under $\mathcal{A}^\theta$, which gives Theorem 2.1. □

B.1.4 Proof of Proposition 2.1

Proof. Define $\mathcal{A}(\delta) \equiv \{\alpha \in \mathbb{A} : ||\alpha - \alpha^*|| \leq \delta\}$, where $||.||$ denotes the relevant metric in $\mathbb{A}$. Further define

$$\mathcal{U}(\delta) \equiv \left\{\theta \in \Theta : \min_{\alpha \in \mathcal{A}(\delta)} r(\alpha, \theta) = 0\right\}$$

which is the set of zero-potential states when the model set is $\mathcal{A}(\delta)$. From Berge's maximum theorem, we know that $V(\theta, \delta) = \min_{\alpha \in \mathcal{A}(\delta)} r(\alpha, \theta)$ is a continuous function, so $\mathcal{U}(\delta)$ is a compact-valued upper semi-continuous correspondence. Further denote

$$D(\delta) = \max_{\theta \in \mathcal{U}(\delta)} ||\theta - \theta^*||,$$

which describes the "size" of $\mathcal{U}(\delta)$. Applying Berge's theorem again, we know that $D(\delta)$ is also upper semi-continuous functions of δ. Notice that when $\delta = 0$, we have $\mathcal{U}(\delta) = \{\theta^*\}$, so $D(0) = 0$. Since $D(\delta) \geq 0$, the upper semi-continuity of D implies that $D(\delta) \to 0$ as $\delta \to 0$. Note that for all compact set $\mathcal{A}$, we can always bound it using some $\mathcal{A}(\delta)$. Therefore, we must have:

$$\max_{\theta \in \mathcal{U}_{\mathcal{A}}} ||\theta - \theta^*|| \to 0 \quad \text{as } d(\mathcal{A}) \to 0 \tag{B.9}$$

Further denote $\mathcal{U}_{\mathcal{A}}^{\varepsilon} = \left\{\theta \in \Theta : \min_{\theta' \in \mathcal{U}_{\mathcal{A}}} ||\theta - \theta'|| \leq \varepsilon\right\}$, which is the set of states that are within distance ε to $\mathcal{U}_{\mathcal{A}}$. For all bounded and continuous function $h : \Theta \to \mathbb{R}$ and for all $\varepsilon > 0$, we have:

$$\begin{aligned}\lim_{t\to\infty} \left| \int_{\Theta} h \times \mu_t^{\theta} dm - \int_{\Theta} h \times \delta_{\theta^*} dm \right| &= \lim_{t\to\infty} \left| \int_{\mathcal{U}_{\mathcal{A}}^{\varepsilon}} h \times \mu_t^{\theta} dm - \int_{\mathcal{U}_{\mathcal{A}}^{\varepsilon}} h \times \delta_{\theta^*} dm \right| \\ &\leq \max_{\theta \in \mathcal{U}_{\mathcal{A}}^{\varepsilon}} |h(\theta) - h(\theta^*)| \quad \mathbb{P}^* - a.s.\end{aligned}$$

where the first equality comes from the fact that $\lim \mu_t^{\theta}\left(\mathcal{U}_{\mathcal{A}}^{\varepsilon}\right) = 1$ for all $\varepsilon > 0$ (implied by Lemma 2.1). From (C.3), as $d(\mathcal{A}) \to 0$, we have $\mathcal{U}_{\mathcal{A}}^{\varepsilon} \to B_{\varepsilon}(\theta^*)$ under the Hausdorff metric, where $B_{\varepsilon}(\theta^*)$ denotes the ε-closed neighborhood of θ^*. Therefore, for all $\varepsilon > 0$, we have:

$$\lim_{d(\mathcal{A})\to 0} \lim_{t\to\infty} \left| \int_{\Theta} h \times \mu_t^{\theta} dm - \int_{\Theta} h \times \delta_{\theta^*} dm \right| \leq \max_{\theta \in B_{\varepsilon}(\theta^*)} |h(\theta) - h(\theta^*)| \quad \mathbb{P}^* - a.s. \tag{B.10}$$

Letting $\varepsilon \to 0$, the RHS of (C.4) converges to 0 from the continuity of h, so the claim is proved. □

B.1.5 Proof of Proposition 2.2 and Corollary 2.1

Proof. Proposition 2.2 follows immediately from Lemma 2.1. Let's then prove Corollary 2.1. For all $\theta \neq \theta^*$, we have

$$\mathcal{R}(\alpha^*, \theta) = \mathbb{E}^* \log \frac{f(s|\theta^*, \alpha^*)}{f(s|\theta, \alpha)} > -\log \mathbb{E}^* \frac{f(s|\theta, \alpha^*)}{f(s|\theta^*, \alpha^*)} = 0,$$

where the strict inequality comes from Jensen's inequality and Assumption 2.3. Since Θ is finite, there exists some $\delta > 0$ such that $\mathcal{R}(\alpha^*, \theta) > \delta$ for all $\theta \neq \theta^*$ (note that $\mathcal{R}(\alpha^*, \theta^*) = 0$ by definition). From the continuity of $\mathcal{R}(\alpha, \theta)$ (from Lemma B.1), when $||\alpha - \alpha^*||$ is sufficiently small, we have $\mathcal{R}(\alpha, \theta) > \delta/2$ and $\mathcal{R}(\alpha, \theta^*) < \delta/2$ for all $\theta \neq \theta^*$. Therefore, when $\mathcal{A}$ is sufficiently small (around the true model), θ^* is the unique zero-potential state for all models in $\mathcal{A}$, which implies that beliefs will converge to δ_{θ^*} for all possible bias. $\square$

B.1.6 Proof of Proposition 2.3

Proof. Since θ^* is not locally dominant, for all $\mathcal{A}$ such that $d(\mathcal{A}) > 0$, there exists some model $\alpha_0 \in \mathcal{A}$ such that $r(\alpha_0, \theta_0) = 0$ for some $\theta_0 \neq \theta^*$. Therefore, we have $\min_{\alpha \in \mathcal{A}} r(\alpha, \theta_0) = 0$. From Theorem 2.1, the limit belief carrier for individuals with bias θ_0 is:

$$\mathcal{U}_{\mathcal{A}}^{\theta_0} = \left\{\theta' \in \Theta : r\left(\alpha', \theta'\right) = 0 \text{ where } r\left(\alpha', \theta_0\right) = 0\right\}$$

Since θ^* is singular, we have: (i) $\theta^* \notin \mathcal{U}_{\mathcal{A}}^{\theta_0}$, and (ii) $\mathcal{U}_{\mathcal{A}}^{\theta^*} = \{\theta^*\}$. Besides, it is easy to verify that $\mathcal{U}_{\mathcal{A}}^{\theta_0}$ is a closed set from the continuity of $r(\alpha, \theta)$. From the property of metrizable space, we can find two disjoint open sets U_1 and U_2 to separate $\mathcal{U}_{\mathcal{A}}^{\theta_0}$ and $\mathcal{U}_{\mathcal{A}}^{\theta^*}$ in the sense that $U_1 \supset \mathcal{U}_{\mathcal{A}}^{\theta_0}$ and $U_2 \supset \mathcal{U}_{\mathcal{A}}^{\theta^*}$. From Theorem 2.1, $\mu_t^{\theta_0}(U_1) \to 1$ and $\mu_t^{\theta^*}(U_1) \to 0$ almost surely, so we have:

$$\lim_{t \to \infty} |\mu_t^{\theta_0}(U_1) - \mu_t^{\theta^*}(U_1)| = 1 \quad \mathbb{P}^* - a.s.,$$

which proves the claim. $\square$

B.1.7 Proof of Lemma 2.3 and Theorem 2.2

We prove the benchmark case and the general case separately for better clarity. This is because when Θ is continuous, the benchmark case requires the bias τ to be a Dirac delta function δ_θ, which cannot be directly accommodated in the general case (i.e., we assume that the every bias τ is a well-defined function, but the Dirac delta function is not a function).

Proof. The proof resembles the proof of the single-biased state. For all possible bias $\tau \in \mathcal{T}$,we have

$$U_t(\tau) = \max_{\alpha\in\mathcal{A}} U_t(\tau|\alpha) = \max_{\alpha\in\mathcal{A}} \frac{\int \prod_{i=1}^t f(s_i|\alpha,\theta)\,\mu(\theta)\,\tau(\theta)\,dm(\theta)}{\int \prod_{i=1}^t f(s_i|\alpha,\theta)\,\mu(\theta)\,dm(\theta)}.$$

I define a measure v such that $v(E) = \int_E \mu(\theta)\,\tau(\theta)\,dm(\theta)$ for all measurable set E. In other words, $\mu(\theta)\,\tau(\theta)$ is the density function of measure v with respect to m. Denote by Θ_τ the support of τ, which are the set of biased states. As such,

$$\begin{aligned}\frac{1}{t}\log U_t(\tau) &= \max_{\alpha\in\mathcal{A}}\Bigg(\frac{1}{t}\log\left[\int_{\Theta_\tau}\prod_{i=1}^t f(s_i|\alpha,\theta)\,dv(\theta)\right] \\ &\quad -\frac{1}{t}\log\left[\int_{\Theta}\prod_{i=1}^t f(s_i|\alpha,\theta)\,\mu(\theta)\,dm(\theta)\right]\Bigg) \\ &= \max_{\alpha\in\mathcal{A}}\left[\log\left(\Theta_\tau||\exp H_t(\alpha,\theta)\,||_t^{\theta,v}\right) - \log\left(\Theta||\exp H_t(\alpha,\theta)\,||_t^{\theta,\mu}\right)\right].\end{aligned}$$

Corollary B.1 implies that

$$\max_{\alpha\in\mathcal{A}} |\log\left(\Theta||\exp H_t(\alpha,\theta)\,||_t^{\theta,\mu}\right) - \log\left(\Theta||\exp H(\alpha,\theta)\,||_\infty^{\theta,\mu}\right)| \to 0 \quad \mathbb{P}^* - a.s., \tag{B.11}$$

and

$$\max_{\alpha\in\mathcal{A}} |\log\left(\Theta_\tau||\exp H_t(\alpha,\theta)\,||_t^{\theta,v}\right) - \log\left(\Theta_\tau||\exp H(\alpha,\theta)\,||_\infty^{\theta,v}\right)| \to 0 \quad \mathbb{P}^* - a.s.. \tag{B.12}$$

Following the exact same argument as in the proof of Lemma 2.2 and Theorem

2.1, we have

$$
\begin{aligned}
&\frac{1}{t}\log U_t(\tau) \to \max_{\alpha\in\mathcal{A}}\left[\log\left(\Theta_\tau\|\exp H(\alpha,\theta)\|_\infty^{\theta,v}\right) - \log\left(\Theta\|\exp H(\alpha,\theta)\|_\infty^{\theta,\mu}\right)\right] \\
&= \max_{\alpha\in\mathcal{A}}\left(\log\left(\max_{\theta\in\Theta_\tau}\exp\mathbb{E}^*\log f(s|\alpha,\theta)\right) - \log\left(\max_{\theta\in\Theta}\exp\mathbb{E}^*\log f(s|\alpha,\theta)\right)\right) \\
&= \max_{\alpha\in\mathcal{A}}\left(-\min_{\theta\in\Theta_\tau}\mathbb{E}^*\log\frac{f(s|\alpha^*,\theta^*)}{f(s|\alpha,\theta)} + \min_{\theta\in\Theta}\mathbb{E}^*\log\frac{f(s|\alpha^*,\theta^*)}{f(s|\alpha,\theta)}\right) \\
&= -\min_{\alpha\in\mathcal{A}}\left[\min_{\theta\in\Theta_\tau} r(\alpha,\theta)\right] \quad \mathbb{P}^* - a.s.
\end{aligned}
$$

Since α_t^τ maximizes $U_t(\tau)$ for each t (hence maximizes $\frac{1}{t}\log U_t(\tau)$), it follows immediately that

$$
\mathcal{A}_\infty^\theta \subset \min_{\alpha\in\mathcal{A}}\left[\min_{\theta\in\mathrm{supp}(\tau)} r(\alpha,\theta)\right] \quad \mathbb{P}^* - a.s..
$$

Similarly, limit beliefs of bias-τ individuals will settle on zero-potential states under models on the RHS. □

B.1.8 Proof of Corollary 2.2

Proof. If $\theta_1 \in \mathcal{U}_\mathcal{A}^\tau \backslash \mathcal{V}_\mathcal{A}^\tau$, there must exist some θ_2 such that $\tau(\theta_2) > \tau(\theta_1)$. Denote by $\mathcal{A}(\theta) = \{\alpha \in \mathcal{A} : r(\alpha,\theta) = 0\}$, the set of models under which state θ is a zero-potential state. For all $\alpha_1 \in \mathcal{A}(\theta_1)$ and all $\alpha_2 \in \mathcal{A}(\theta_2)$, we have:

$$
\lim U_t(\tau|\alpha_1) \to \tau(\theta_1) < \tau(\theta_2) = \lim U_t(\tau|\alpha_2) \quad \mathbb{P}^* - a.s.
$$

where the convergence comes from the fact that every model has a unique zero-potential state, so $\mu(\theta|\alpha)$ converges to the Dirac measure on the zero-potential state. Since all models in $\mathcal{A}(\theta_1)$ are strictly dominated in terms of utilities, so $\mathcal{A}(\theta_1)$ will not be chosen in the limit, which implies that θ_1 is not asymptotically carried on. □

B.1.9 Verification of $\mathbb{E}a_t \to 0$ in Example 18

I am going to verify that if individuals update according to model α_2, the expected report approaches 0 in the limit.

Proof. I assume a tie-breaking rule that when indifferent between two actions, individuals choose the action with a larger index.[3] Denote by $E_t = \left\{\frac{\mu_t(1|\alpha_2)}{\mu_t(-1|\alpha_2)} \geq 1\right\}$ and $A_t = \left\{\frac{\mu_t(1|\alpha_2)}{\mu_t(0|\alpha_2)} \geq 1\right\}$ and $B_t = \left\{\frac{\mu_t(-1|\alpha_2)}{\mu_t(0|\alpha_2)} > 1\right\}$. By definition, we have:

$$
\begin{aligned}
\mathbb{E}^* r_t &= \mathbb{P}^*\left(r_t = 1\right) \times 1 + \mathbb{P}^*\left(r_t = -1\right) \times (-1) \\
&= \mathbb{P}^*\left(E_t \cap A_t\right) \times 1 + \mathbb{P}^*\left(E_t^c \cap B_t\right) \times (-1) && \text{(B.13)} \\
&= \left[\mathbb{P}^*\left(E_t\right) - \mathbb{P}^*\left(E_t \cap A_t^c\right)\right] - \left[\mathbb{P}^*\left(E_t^c\right) - \mathbb{P}^*\left(E_t^c \cap B_t^c\right)\right]. && \text{(B.14)}
\end{aligned}
$$

Denoting by $S_t = \log\left[\frac{\mu_t(1|\alpha_2)}{\mu_t(-1|\alpha_2)}\right]$ and applying Bayes rule, we get

$$
S_{t+1} = S_t + \log\left[\frac{f\left(s_t|1,\alpha_2\right)}{f\left(s_t|-1,\alpha_2\right)}\right] \equiv S_t + X_t,
$$

where $X_t = \log\left[\frac{f(s_t|1,\alpha_2)}{f(s_t|-1,\alpha_2)}\right]$. It is easy to verify that state 1 and -1 have the same relative entropy under model α_2, so $\mathbb{E}^* X_t = 0$. Denoting by $\sigma^2 = \mathbb{E}^* X_t^2$, the Central Limit Theorem (CLT) implies that:

$$
\frac{S_t}{\sigma\sqrt{t}} \Rightarrow \mathcal{N}\left(0,1\right),
$$

where "$\Rightarrow$" means convergence in distribution. As a result,

$$
\lim_{t\to\infty} \mathbb{P}^*\left(E_t\right) = \lim_{t\to\infty} \mathbb{P}^*\left(S_t \geq 0\right) = \lim_{t\to\infty} \mathbb{P}^*\left(\frac{S_t}{\sigma\sqrt{t}} \geq 0\right) = \frac{1}{2}.
$$

Similarly, we have $\lim_{t\to\infty} \mathbb{P}^*\left(E_t^c\right) = \frac{1}{2}$. Since the relative entropy of state 1 and -1 under α_2 is strictly than the relative entropy of state 0 under α_2, we have:

$$
\lim_{t\to\infty} \mathbb{P}^*\left(A_t\right) = \lim_{t\to\infty} \mathbb{P}^*\left(B_t\right) = 1,
$$

[3] This is only for convenience purpose. The result also holds for other tie-breaking rules.

which directly implies that

$$\lim_{t\to\infty} \mathbb{P}^* \left(E_t \cap A_t^c\right) = \lim_{t\to\infty} \mathbb{P}^* \left(E_t^c \cap B_t^c\right) = 0.$$

Taking limits on both sides of (C.5), we get $\lim_{t\to\infty} \mathbb{E}^* r_t = 0$, so the expected action under model α_2 approaches 0. $\square$

B.1.10 Verification of the Claims in Example 19

(I) Suppose that the agent updates via the *full Bayesian rule*. Given our payoff structures, the agent's max-min expected utility of choosing action l is $V^{FB}(l) = \underline{\mu}(L)$, where $\underline{\mu}(L)$ denotes the lowest probability on state L. Similarly, we have $V^{FB}(r) = \underline{\mu}(R)$, and $V^{FB}(m) = 1/2$. From the previous discussion, the lowest probability on state L arises when the agent updates according to model α_2, where signal b (i.e., the majority signal) is interpreted as "bad news" for state L. Symmetrically, the lowest probability on state R arises when he updates according to model α_1. Simple calculations show that

$$\underline{\mu}(L) = \mu(L|\alpha_2) = \frac{\left(\frac{3}{4}\right)^2 \cdot \left(\frac{1}{4}\right)^6}{\left(\frac{3}{4}\right)^2 \cdot \left(\frac{1}{4}\right)^6 + \left(\frac{1}{2}\right)^2 \cdot \left(\frac{1}{2}\right)^6 + \left(\frac{1}{4}\right)^2 \cdot \left(\frac{3}{4}\right)^6} \approx 0.009$$

$$\underline{\mu}(R) = \mu(R|\alpha_1) = \frac{\left(\frac{2}{3}\right)^2 \cdot \left(\frac{1}{3}\right)^6}{\left(\frac{1}{3}\right)^2 \cdot \left(\frac{2}{3}\right)^6 + \left(\frac{1}{2}\right)^2 \cdot \left(\frac{1}{2}\right)^6 + \left(\frac{2}{3}\right)^2 \cdot \left(\frac{1}{3}\right)^6} \approx 0.04,$$

which are less than $V^{FB}(m) = 1/2$, so the optimal choice is $a^{FB} = m$.

(II) Suppose that the agent updates via the *maximum likelihood rule*. We can calculate the probabilities of the observed signals under each model,

$$\mathbb{P}\left(\boldsymbol{s}|\alpha_1\right) = \frac{1}{3} \times \left(\frac{1}{3}\right)^2 \cdot \left(\frac{2}{3}\right)^6 + \frac{1}{3} \times \left(\frac{1}{2}\right)^2 \cdot \left(\frac{1}{2}\right)^6 + \frac{1}{3} \times \left(\frac{2}{3}\right)^2 \cdot \left(\frac{1}{3}\right)^6 \approx 0.0048$$
$$\mathbb{P}\left(\boldsymbol{s}|\alpha_2\right) = \frac{1}{3} \times \left(\frac{3}{4}\right)^2 \cdot \left(\frac{1}{4}\right)^6 + \frac{1}{3} \times \left(\frac{1}{2}\right)^2 \cdot \left(\frac{1}{2}\right)^6 + \frac{1}{3} \times \left(\frac{1}{4}\right)^2 \cdot \left(\frac{3}{3}\right)^6 \approx 0.02,$$

so model α_2 is more likely to generate the observed signals. The agent then updates according to α_2, and his posterior on state L is

$$\mu^{ML}\left(R\right) = \mu\left(R|\alpha_2\right) = \frac{\left(\frac{1}{4}\right)^2 \cdot \left(\frac{3}{4}\right)^4}{\left(\frac{3}{4}\right)^2 \cdot \left(\frac{1}{4}\right)^4 + \left(\frac{1}{2}\right)^2 \cdot \left(\frac{1}{2}\right)^4 + \left(\frac{1}{4}\right)^2 \cdot \left(\frac{3}{4}\right)^4} \approx 0.53 > \frac{1}{2}.$$

The expected utility from action r is $V^{ML}\left(r\right) = \mu^{ML}\left(R\right) > \frac{1}{2} \geq \max\left\{V^{ML}\left(m\right), V^{ML}\left(l\right)\right\}$, so the optimal choice is $a^{ML} = r$.

(III) Suppose that the agent updates via the *biased updating rule* and is biased toward state M. Simple application of Bayes rule shows that

$$\mu\left(M|\alpha_1\right) = \frac{\frac{1}{3} \times \left(\frac{1}{2}\right)^2 \cdot \left(\frac{1}{2}\right)^6}{\mathbb{P}\left(\boldsymbol{s}|\alpha_1\right)} > \frac{\frac{1}{3} \times \left(\frac{1}{2}\right)^2 \cdot \left(\frac{1}{2}\right)^6}{\mathbb{P}\left(\boldsymbol{s}|\alpha_2\right)} = \mu\left(M|\alpha_2\right),$$

where the inequality comes from the fact that $\mathbb{P}\left(\boldsymbol{s}|\alpha_1\right) < \mathbb{P}\left(\boldsymbol{s}|\alpha_2\right)$. Therefore, the biased agent will update according to model α_1. Under model α_1, the probability of state L is

$$\mu^{BS}\left(L\right) = \mu\left(L|\alpha_1\right) = \frac{\left(\frac{1}{3}\right)^2 \cdot \left(\frac{2}{3}\right)^6}{\left(\frac{1}{3}\right)^2 \cdot \left(\frac{2}{3}\right)^6 + \left(\frac{1}{2}\right)^2 \cdot \left(\frac{1}{2}\right)^6 + \left(\frac{2}{3}\right)^2 \cdot \left(\frac{1}{3}\right)^6} \approx 0.68 > \frac{1}{2}.$$

The expected utility from action l is $V^{BS}\left(l\right) = \mu^{BS}\left(L\right) > \frac{1}{2} \geq \max\left\{V^{BS}\left(m\right), V^{BS}\left(r\right)\right\}$, so the optimal choice should be $a^{BS} = l$.

B.2 Supplementary Materials

B.2.1 Overconfidence under Ambiguity

This subsection expands on the intuition developed in Example 18 and provides conditions under which the overconfidence exists in the limit. For all model $\alpha \in \mathbb{A}$, denote by Θ_α the set of zero-potential states under model α. Let $\overline{\tau}(\alpha)$ and $\underline{\tau}(\alpha)$ denote the highest and the lowest $\tau(\theta)$ for all $\theta \in \Theta_\alpha$ respectively.

Definition 15. Model α *τ-dominates* α', if $\underline{\tau}(\alpha) > \overline{\tau}(\alpha')$. The *$\tau$-core of* $\mathcal{A}$, denoted by $C^\tau(\mathcal{A})$, consists of all models in $\mathcal{A}$ which are not τ-dominated by any other model in $\mathcal{A}$.

The τ-core of $\mathcal{A}$ consists of all models whose zero-potential states are not strictly dominated by any other model in terms of τ. Following the same logic as in Corollary 2.2, in the limit, only models in the τ-core of $\mathcal{A}$ matter for an individual with bias τ. As a result, when analyzing individuals with bias τ, we can restrict our attention to the τ-core models without loss of generality. The following concept is essential in characterizing the overconfidence effect.

Definition 16. For all $\alpha, \alpha' \in \mathbb{A}$, models α and α' are ***τ-complementary*** if

$$\left[\overline{\tau}(\alpha) - \overline{\tau}(\alpha')\right] \cdot \left[\underline{\tau}(\alpha) - \underline{\tau}(\alpha')\right] < 0.$$

Model set $\mathcal{A}$ is called ***τ-complementary*** if for all $\alpha \in C^\tau(\mathcal{A})$, there exists some $\alpha' \in \mathcal{A}$ such that α and α' are τ-complementary.

The notion of τ-complementarity describes some "crossing" feature of utilities generated by zero-potential states under α and α'. Roughly speaking, if

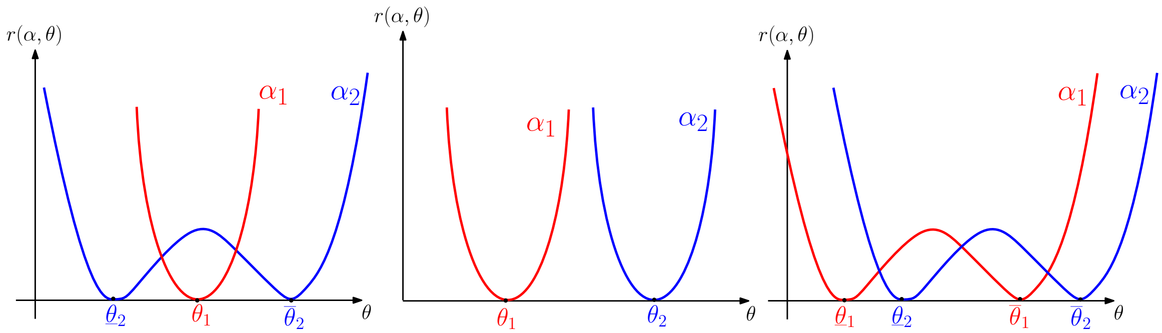

Figure B.1: Illustration of τ-complementary

the highest utility under α is higher than the highest utility under α', then the lowest utility under α must be lower than the lowest utility under α'. Figure B.1 illustrates the concept, where the vertical axis denotes the information potential $r(\alpha, \theta)$. Assuming that τ is strictly monotone, the model set $\{\alpha_1, \alpha_2\}$ is τ-complementary in the left graph, but is not τ-complementary in the middle and right graph.

Theorem B.1 (The Overconfidence Effect)**.** *Assume that $|\Theta| < \infty$. For all bias $\tau \in \mathcal{T}$, if $\mathcal{A}$ is τ-complementary, then for all $\alpha \in \mathcal{A}$, there is some $\varepsilon > 0$ such that:*

$$\liminf_{t\to\infty} \mathbb{P}^* \left[U_t(\tau) - U_t(\tau|\alpha) \geq \varepsilon\right] > 0$$

Recall that $U_t(\tau)$ denotes the time-t belief utility of an individual with bias τ, and $U_t(\tau|\alpha)$ denotes the time-t belief utility if this individual updates according to model α. By definition, $U_t(\tau) \geq U_t(\tau|\alpha)$ for all $\alpha \in \mathcal{A}$, so Theorem B.1 implies that the expected belief utility of a biased individual is *strictly* higher than that of a Bayesian individual with any feasible model perception. This result is interesting since it shows that the effect of ambiguity cannot be replaced by any specific model, so biased individuals can exhibit behavioral patterns that are absent in the Bayesian framework.

The intuition has been discussed in Example 18. To illustrative it more directly, let us look at the left graph in Figure B.1. Suppose that τ is strictly increasing and $\mathcal{A} = \{\alpha_1, \alpha_2\}$. Following the arguments in Example 18, in the limit, the τ-individual will adopt model α_2 when the beliefs under α_2 accumulate around the higher zero-potential state $\overline{\theta}_2$, and adopt model α_1 when the beliefs under α_2 accumulate around the lower zero-potential state $\underline{\theta}_2$. We can think of α_2 as a "riskier" model than α_1 in the sense that it can generate a higher utility gain but can also incur a larger utility loss in the limit (graphically, its zero-potential states are more "spread-out" than model α_1). The presence of the "safe model" α_1 allows individuals to exploit the utility gain as well as hedge against the utility loss from α_2. As a result, individuals are strictly better-off with the presence of model ambiguity.

Proof of Theorem B.1

Proof. Consider any model α and its τ-complementary model α'. Without loss of generality, I assume that $\overline{\tau}(\alpha') > \overline{\tau}(\alpha)$ (the case $\overline{\tau}(\alpha') < \overline{\tau}(\alpha)$ follows an analogous argument). By definition, $\overline{\tau}(\alpha') = \tau(\theta_1)$ for some $\theta_1 \in \Theta_{\alpha'}$, where $\Theta_{\alpha'} = \{\theta_1, \theta_2, ..., \theta_{K+1}\}$. Notice that we must have $K \geq 1$ since $K = 0$ (i.e., $\Theta_{\alpha'} = \{\theta_1\}$) implies that $\overline{\tau}(\alpha') = \underline{\tau}(\alpha') > \overline{\tau}(\alpha) \geq \underline{\tau}(\alpha)$, which contradicts the fact that α and α' are τ-complementary. Therefore, it is meaningful to define two K-dimensional vectors:

$$\boldsymbol{S}_t = \left(\log \frac{\mu_t(\theta_2|\alpha')}{\mu_t(\theta_1|\alpha')}, ..., \log \frac{\mu_t(\theta_{K+1}|\alpha')}{\mu_t(\theta_1|\alpha')}\right),$$
$$\boldsymbol{X}_t = \left(\log \frac{f(s_t|\theta_2, \alpha')}{f(s_t|\theta_1, \alpha')}, ..., \log \frac{f(s_t|\theta_{K+1}, \alpha')}{f(s_t|\theta_1, \alpha')}\right).$$

By Bayes rule, for all $k \in \{1, ..., K\}$, we have

$$\frac{\mu_{t+1}(\theta_{k+1}|\alpha')}{\mu_{t+1}(\theta_1|\alpha')} = \frac{\mu_t(\theta_{k+1}|\alpha')}{\mu_t(\theta_1|\alpha')} \times \frac{f(s_t|\theta_{k+1}, \alpha')}{f(s_t|\theta_1, \alpha')},$$

which implies that

$$\boldsymbol{S}_{t+1} = \boldsymbol{S}_t + \boldsymbol{X}_t \text{ or } \boldsymbol{S}_t = \sum_{i=1}^{t} \boldsymbol{X}_i. \tag{B.15}$$

For all $k \in \{1, ..., K\}$, we have $\theta_{k+1} \in \Theta_{\alpha'}$, so the relative entropy of state $k+1$ and state 1 must be equal, which implies $\mathbb{E}^* \boldsymbol{X}_t = 0$. Applying the multidimensional Central Limit Theorem (CLT) , we get

$$\frac{\boldsymbol{S}_t}{\sqrt{t}} \Rightarrow \mathcal{N}_K(0, \Sigma) \equiv Z \text{ with } \Sigma = \begin{pmatrix} \mathbb{E}^* \boldsymbol{X}_{t1}^2 & \mathbb{E}^* \boldsymbol{X}_{t1} \boldsymbol{X}_{t2} & \cdots & \mathbb{E}^* \boldsymbol{X}_{t1} \boldsymbol{X}_{tK} \\ \mathbb{E}^* \boldsymbol{X}_{t2} \boldsymbol{X}_{t1} & \mathbb{E}^* \boldsymbol{X}_{t2}^2 & & \mathbb{E}^* \boldsymbol{X}_{t2} \boldsymbol{X}_{tK} \\ \vdots & \vdots & & \vdots \\ \mathbb{E}^* \boldsymbol{X}_{tK} \boldsymbol{X}_{t1} & \mathbb{E}^* \boldsymbol{X}_{tK} \boldsymbol{X}_{t2} & \cdots & \mathbb{E}^* \boldsymbol{X}_{t1} \boldsymbol{X}_{tK} \end{pmatrix}, \tag{B.16}$$

where X_{tk} denotes the k-th component of X_t, and $\mathcal{N}_K(0, \Sigma)$ denotes the K-dimensional multivariate normal distribution with covariance matrix Σ. Notice that it can be verified that Σ is a well-defined since every entry of Σ is finite from Assumption 2.2. Based on Assumption 2.3, we can further show that Σ is positive definite, so Z admits a strictly positive density function on $\mathbb{R}^K$.[4] As a result, from (B.16), for all $M \in \mathbb{R}_{++}$, we have

$$\lim_{t \to \infty} \mathbb{P}^* \left(\boldsymbol{S}_t \leq -\log M \cdot \boldsymbol{e} \right) = \lim_{t \to \infty} \mathbb{P}^* \left(\frac{\boldsymbol{S}_t}{\sqrt{t}} \leq -\frac{\log M}{\sqrt{t}} \boldsymbol{e} \right) = \mathbb{P}(Z \leq 0) \equiv \xi > 0, \tag{B.17}$$

where $\boldsymbol{e} = (1, 1, ..., 1) \in \mathbb{R}^K$. Defining $E_t = \{ \boldsymbol{S}_t \leq -\log M \cdot \boldsymbol{e} \}$, for all signal paths $s^\infty \in E_t$, we have

$$\mu_t(\theta_{k+1} | \alpha') \leq \frac{1}{M} \cdot \mu_t(\theta_1 | \alpha') \quad \text{for all } k \in \{1, ..., K\}.$$

Summing up all $k \in \{1, ..., K\}$, we get

$$\mu_t(\Theta_{\alpha'} | \alpha') \leq \left(1 + \frac{K}{M} \right) \times \mu_t(\theta_1 | \alpha'),$$

[4]see Fudenberg et al. (2021) for a detailed argument on the positive definiteness.

so (B.17) implies that

$$\liminf_{t\to\infty} \mathbb{P}^* \left(\mu_t\left(\theta_1|\alpha'\right) \geq \frac{M}{M+K} \cdot \mu_t\left(\Theta_{\alpha'}|\alpha'\right)\right) \geq \xi,$$

or equivalently,

$$\liminf_{t\to\infty} \mathbb{P}^* \left(\left\{\mu_t\left(\theta_1|\alpha'\right) \geq \frac{M}{M+K} \cdot \mu_t\left(\Theta_{\alpha'}|\alpha'\right)\right\} \cap \left\{\mu_t\left(\Theta_{\alpha'}|\alpha'\right) \geq \frac{M+K}{M+K+1}\right\}\right)$$
$$+ \mathbb{P}^* \left(\left\{\mu_t\left(\theta_1|\alpha'\right) \geq \frac{M}{M+K} \cdot \mu_t\left(\Theta_{\alpha'}|\alpha'\right)\right\} \cap \left\{\mu_t\left(\Theta_{\alpha'}|\alpha'\right) < \frac{M+K}{M+K+1}\right\}\right) \geq \xi > 0 \tag{B.18}$$

Applying the Berk (1966, Sect. 4, Main Theorem) result, we have $\lim_{t\to\infty} \mathbb{P}^* \left(\mu_t\left(\Theta_{\alpha'}|\alpha'\right) < \frac{M+K}{M+K+1}\right) = 0$, so (B.18) implies that

$$\liminf_{t\to\infty} \mathbb{P}^* \left(\left\{\mu_t\left(\theta_1|\alpha'\right) \geq \frac{M}{M+K} \cdot \mu_t\left(\Theta_{\alpha'}|\alpha'\right)\right\} \cap \left\{\mu_t\left(\Theta_{\alpha'}|\alpha'\right) \geq \frac{M+K}{M+K+1}\right\}\right) \geq \xi,$$

which directly implies that

$$\liminf_{t\to\infty} \mathbb{P}^* \left(\mu_t\left(\theta_1|\alpha'\right) \geq \frac{M}{M+K+1}\right) \geq \xi > 0 \tag{B.19}$$

for all possible $M \in \mathbb{R}_{++}$.

For simplicity in notation, I normalize the τ such that the highest value is bounded by 1 (e.g., by multiplying some constant. It is easy to see that this transformation will not change the problem). Define $\varepsilon = \frac{1}{3}\left[\overline{\tau}\left(\alpha'\right) - \overline{\tau}\left(\alpha\right)\right]$, which is a strictly positive number by definition. From the definition, we know that $U_t\left(\tau\right) = \max_\alpha U_t\left(\tau|\alpha\right) \geq U_t\left(\tau|\alpha'\right)$, so

$$\liminf_{t\to\infty} \mathbb{P}^* \left[U_t\left(\tau\right) - U_t\left(\tau|\alpha\right) \geq \varepsilon\right] \geq \liminf_{t\to\infty} \mathbb{P}^* \left[U_t\left(\tau|\alpha'\right) - U_t\left(\tau|\alpha\right) \geq \varepsilon\right].$$

Define $E_t = \left\{U_t\left(\tau|\alpha'\right) - U_t\left(\tau|\alpha\right) \geq \varepsilon\right\}$ and $E_t^1 = \left\{\mu_t\left(\theta_1|\alpha'\right) \geq \frac{\varepsilon+1}{3\varepsilon+1}\right\}$ and $E_t^2 = \left\{\mu_t\left(\Theta_\alpha|\alpha\right) \geq \frac{\varepsilon+1}{3\varepsilon+1}\right\}$. For all $s^\infty \in E_t^1 \cap E_t^2$, we have:

$$\begin{aligned} U_t\left(\tau|\alpha'\right) - U_t\left(\tau|\alpha\right) &\geq \mu_t\left(\theta_1|\alpha'\right) \cdot \overline{\tau}\left(\alpha'\right) - \left(\mu_t\left(\Theta_\alpha|\alpha\right) \cdot \overline{\tau}\left(\alpha\right) + \left(1 - \mu_t\left(\Theta_\alpha|\alpha\right)\right) \cdot 1\right) \\ &\geq \frac{\varepsilon+1}{3\varepsilon+1}\left[\overline{\tau}\left(\alpha'\right) - \overline{\tau}\left(\alpha\right)\right] - \frac{2\varepsilon}{3\varepsilon+1} = \varepsilon, \end{aligned}$$

so $s^{\infty} \in E_t$, which implies that $E_t^1 \cap E_t^2 \subset E_t$. As a result,

$$\begin{aligned}
\liminf_{t\to\infty} \mathbb{P}^* \left[E_t\right] &\geq \liminf \mathbb{P}^* \left[E_t^1 \cap E_t^2\right] \\
&\geq \liminf \left(\mathbb{P}^* \left[E_t^1\right] - \mathbb{P}^* \left[E_t^1 \cap \left(E_t^2\right)^c\right]\right) \\
&\geq \liminf \left(\mathbb{P}^* \left[E_t^1\right] - \mathbb{P}^* \left[\left(E_t^2\right)^c\right]\right) \\
&\geq \liminf \mathbb{P}^* \left[E_t^1\right] + \liminf \mathbb{P}^* \left[E_t^2\right] - 1 \qquad \text{(B.20)}
\end{aligned}$$

From Berk's result, we know that $\lim \mathbb{P}^* \left(E_t^2\right) = 1$. From (C.10), we can set $M = \frac{\varepsilon+1}{2\varepsilon}(K+1)$, so $\liminf \mathbb{P}^* \left(E_t^1\right) \geq \xi > 0$. Consequently,

$$\liminf_{t\to\infty} \mathbb{P}^* \left[U_t\left(\tau|\alpha'\right) - U_t\left(\tau|\alpha\right) \geq \varepsilon\right] = \liminf_{t\to\infty} \mathbb{P}^* \left[E_t\right] \geq \xi > 0,$$

so the claim is proved. □

B.2.2 Belief-Dependent Bias

In this section, I present a more general version of ***belief-dependent biased updating rule*** than Examples 21 and characterize limit beliefs under this rule. Suppose that the state space is finite, that is, $\Theta = \{\theta_1, ..., \theta_N\}$ for some finite N, and the signal space S is also finite. The belief simplex $\Delta(\Theta)$ is divided into N regions $\Delta_{\theta_1}, ..., \Delta_{\theta_N}$, which can possibly intersect. For all $\theta \in \Theta$, the Dirac belief δ_θ is uniquely contained in the interior of region Δ_θ and not in any other region. At time $t = 0$, every individual is endowed with some initial biased state $\theta \in \Theta$ and a full-support prior $\mu_0^\theta \in \Delta_\theta$. The biased state can change overtime, where b_t^θ denotes the time-t biased state of the individual with initial bias θ. The bias process $\left\{b_t^\theta\right\}$ evolves with beliefs according to the following rule:

$$\forall \theta, \theta_{i_1}, .., \theta_{i_k} \in \Theta: \quad b_t^\theta \in \left\{\theta_{i_1}, \theta_{i_2}, ..\theta_{i_k}\right\} \text{ if } \mu_t^\theta \in \Delta_{\theta_{i_1}} \cap \Delta_{\theta_{i_2}} .. \cap \Delta_{\theta_{i_k}} \qquad \text{(B.21)}$$

In other words, if the time-t belief μ_t^θ enters some region, the time-t bias b_t^θ changes to the state that corresponds to that region. For the case where μ_t^θ belongs to the intersection of multiple regions, b_t^θ is determined by some tie-breaking rule. For realistic purposes, I assume that the tie-breaking rule features inertia on the initial bias. More specifically, if $b_{t-1}^\theta = \theta$ and $\mu_t^\theta \in \Delta_\theta$, then we must have $b_t^\theta = \theta$. In other words, if the last-period bias is the initial bias, and if the current belief stays in the region favoring the initial bias, then the individual will stick to his initial bias. Given the last period bias b_{t-1}^θ, the time-t belief μ_t^θ is updated according to the model that can best support b_{t-1}^θ. More precisely,

$$\forall \theta, \theta' \in \Theta: \quad \mu_t^\theta\left(\theta'\right) = \mu_{t-1}\left(\theta' | \alpha_t^\theta\right) \text{ where } \alpha_t^\theta \in \arg\max_{\alpha \in \mathcal{A}} \mu\left(b_{t-1}^\theta | \alpha\right), \tag{B.22}$$

where the specific value of α_t^θ is determined by some tie-breaking rule when there are multiple maximizers. Notice that the only difference between (B.22) and the biased updating rule is that the bias b_t^θ can change overtime. Below is another example.

Example 31 (Belief-Dependent Bias II)**.** Define $\Delta_G = \{\mu : \mu(G) \geq k\}$ and $\Delta_B = \{\mu : \mu(B) \geq k\}$, for some $k \leq 1/2$. Now suppose that the updating rule of $\left\{b_t^\theta\right\}$ evolves as follows

$$b_t^\theta = \begin{cases} G & \text{if } \mu_t(G) > 1-k \\ B & \text{if } \mu_t(G) < k \\ b_{t-1}^\theta & \text{if } \mu_t(G) \in [k, 1-k] \end{cases}.$$

If the belief uniquely belongs to Δ_G (or Δ_B), the bias is just equal to G (or B). If the belief falls into the intersection of Δ_G and Δ_B, the bias is determined by the last-period bias. This captures the idea that individuals exhibit some inertia to the current bias, that is, they are willing to switch their bias only if they face strong opposing evidence. When $k = 1/2$, it corresponds to confirmation bias model in Example 21; when $k = 0$, it corresponds to the benchmark situation where individuals hold a fixed bias. $\square$

Characterizations of Limit Beliefs

This subsection provides the characterizations of limit beliefs under the belief-dependent biased updating rule. For simplicity of discussion, this section maintains the assumption that every model has a unique zero-potential state.

Assumption B.1. *For all $\alpha \in \mathcal{A}$, there exists a unique $\theta \in \Theta$ such that $r(\alpha, \theta) = 0$.*

In other words, under every model, there is a unique state that delivers the "smallest" distance between the induced distribution and the true distribution. This assumption is adopted throughout this section. The first result is that all beliefs will converge almost surely as stated below.

Proposition B.1. *For all initial bias $\theta \in \Theta$, we have $\mu_t^\theta \to \mu_\infty^\theta$ $\mathbb{P}^*$-almost surely, where μ_∞^θ is a random variable that satisfies*

$$\sum_{\theta' \in \mathcal{U}_{\mathcal{A}}} \mathbb{P}^* \left(\mu_\infty^\theta = \delta_{\theta'} \right) = 1,$$

where $\mathcal{U}_{\mathcal{A}}$ denotes the set of zero-potential states under $\mathcal{A}$.

Proposition B.1 says that: (i) beliefs will almost surely converge, and (ii) the limit belief is a Dirac belief on some zero-potential state. One natural question is that: will every zero-potential state be visited with a strictly positive probability? Unfortunately, it is not always true. Below is one example.

Example 32. Consider the case described in Example 21. There are two individuals, one with initial bias G, the other with initial bias B. Both individuals hold the same prior that $\mu_0(G) = \mu_0(B) = 1/2$. The set of models is $\mathcal{A} = \{\alpha_1, \alpha_2\}$ and the data-generating processes are described below

$$\begin{array}{c|cc} \alpha_1 & a & b \\ \hline G & 3/4 & 1/4 \\ B & 1/4 & 3/4 \end{array} \quad \begin{array}{c|cc} \alpha_2 & a & b \\ \hline G & 1/4 & 3/4 \\ B & 3/4 & 1/4 \end{array},$$

where a and b are two possible signals. Suppose that the true state is G and the correct model α_1. It is easy to see that both states are zero-potential states: if individuals update according to model α_1, they will believe in state G; if they update according to model α_2, they will believe in state B.

Consider an individual with initial bias G, if he received more signal a, he would adopt model α_1; if he received more signal b, he would adopt model α_2. In both cases, the belief on state G is higher than $1/2$, which prevents the individual from modifying his initial bias. As a result, we must have $\mu_t^G(G) \geq 1/2$ for all t. Proposition B.1 further implies that $\mu_t^G \to \delta_G$ almost surely, which means that every individual will perfectly confirm his bias with probability 1, so not every zero-potential state will be visited with a positive probability. □

In this example, all signals are "controversial" in the sense that they can be interpreted as both good news and bad news for every state. Therefore, individuals can always confirm their initial biases by misinterpreting any negative news as positive. The following assumption restricts this kind of misinterpretation.

Assumption B.2 (Existence of non-controversial signals)**.** *For all $\theta \in \Theta$, there exists some signal $s_\theta \in S$ such that:*

$$\min_{\theta' \neq \theta} \frac{f(s|\alpha, \theta)}{f(s|\alpha, \theta')} > 1 \textit{ for all } \alpha \in \mathcal{A}$$

Assumption B.2 says that for every state θ, there a signal s_θ that is unanimously regarded as the good news for state θ by all models. Under this assumption, it is impossible to interpret any signal as good news to an arbitrary state as in Example 32. Notice that Assumption B.2 can be weak since it only requires the existence of one such s_θ without restricting how large the probability is. It is possible that all models only agree on a tiny fraction of signals but disagree on signals that occur with a large probability. Under this assumption, we have the following result.

Proposition B.2. *If Assumption B.2 holds, then for all initial bias $\theta \in \Theta$ and for all state $\theta' \in \mathcal{U}_{\mathcal{A}}$, we have $\mathbb{P}^* \left(\mu_\infty^\theta = \delta_{\theta'}\right) > 0$.*

Proposition B.1 shows a strong result that for *all* possible initial bias, *every* zero-potential state will be visited with a strictly positive probability. One implication is that **complete learning may not occur for all individuals**. Unless the true state θ^* minimizes relative entropy for all possible models (i.e., $\mathcal{U}_{\mathcal{A}} = \{\theta^*\}$), incorrect learning arises with a strictly positive probability for every individual. Even if an individual is initially biased toward the true state, he may arrive at some incorrect state with a positive probability. Recall that in the case with fixed bias, when the model set is correctly specified, individuals with correct bias can almost surely learn the true state. As such, Proposition B.1 suggests that allowing bias to change can even make it harder to achieve correct learning in some sense.

Previous two propositions hold for individuals with all possible initial biases. However, it is conceivable that individuals starting from different regions (i.e., with different initial biases) may visit each state with different probabilities. Therefore, it becomes natural to ask: how does the distribution of μ_∞^θ vary with the initial bias θ? This question is generally very hard to answer since it is a challenging task to solve for the exact distribution of μ_∞^θ. However, it is possible to characterize the distribution in a limit case where the "inertia" of the initial bias is sufficiently large.

Definition 17. For every θ, define $R_\theta = \sup\left\{r \geq 0 : B_{\frac{r}{r+1}}(\delta_\theta) \subset \Delta_\theta\right\}$, where $R_\theta \in [0, \infty]$.

Intuitively, the magnitude of R_θ measures the size of Δ_θ, which captures the degree of reluctance of switching the bias. If $R_\theta \to 0$, the region Δ_θ approaches the Dirac belief δ_θ, meaning that individuals are only willing to hold the bias θ within a small neighborhood of δ_θ. If $R_\theta \to \infty$, the region Δ_θ approaches the whole space, and individuals with initial bias θ will stick to it for a large set of beliefs. We have the following proposition.

Proposition B.3. *For all initial bias $\theta \in \Theta$, as $R_\theta \to \infty$, we have*

$$\sum_{\theta' \in \mathcal{U}_{\mathcal{A}}^\theta} \mathbb{P}^*\left(\mu_\infty^\theta = \delta_{\theta'}\right) \to 1,$$

where $\mathcal{U}_{\mathcal{A}}^\theta$ denotes the set of zero-potential states under models that minimize the potential of state θ.

When the inertia is sufficiently large, with a large probability, individuals will land on zero-potential states under models that minimize the potential of their initial biases, which correspond to the limit belief carrier in the benchmark case (see Theorem 2.1). This proposition builds the connection between the model with fixed bias and the model with belief-dependent bias. It demonstrates that the fixed-bias model approximates the situation where the degree of inertia is sufficiently large.

Proof of Proposition B.1

Let $\theta(\alpha)$ be the zero-potential state under model α. We have

$$\mathbb{P}^*\left(\left\{\omega : \min_{\alpha\in\mathcal{A}} |\mu_t(\theta(\alpha)|\alpha) - 1| \to 0\right\}\right) = 1 \tag{B.23}$$

where the uniform convergence comes from the same reasoning as in the proof of Lemma 1. Define $B_\varepsilon(\delta_\theta) \equiv \{\mu \in \Delta(\Theta) : \mu(\theta) > 1-\varepsilon\}$, so (B.23) further implies that for all (small) $\varepsilon \in (0,1)$ and for all (large) $\delta \in (0,1)$, there exists some $T < \infty$ such that:

$$\mathbb{P}^*(E) \equiv \mathbb{P}^*\left(\left\{\omega : \mu_t(.|\alpha) \in B_\varepsilon\left(\delta_{\theta(\alpha)}\right) \text{ for all } \alpha \in \mathcal{A} \text{ and } t \geq T\right\}\right) > \delta. \tag{B.24}$$

In (C.11), I set ε to be sufficiently small such that $B_\varepsilon\left(\delta_{\theta(\alpha)}\right)$ is uniquely contained in the region containing $\delta_{\theta(\alpha)}$, that is, $B_\varepsilon\left(\delta_{\theta(\alpha)}\right) \subset \Delta_{\theta(\alpha)} \setminus \cup_{\theta\neq\theta(\alpha)} \Delta_\theta$. We then have the following claim

Claim B.1. *For all signal path $\omega \in E$, the belief μ_t^θ is trapped in $B_\varepsilon(\delta_{\theta'})$ for some θ' forever.*

Proof. Let α_T be the time-T model adopted by the individual, where T is as defined in (C.11).

(1) When $t = T$, $\mu_T^\theta = \mu_T(.|\alpha_T)$, by definition, we have $\mu_T^\theta \in B_\varepsilon\left(\delta_{\theta(\alpha_T)}\right)$. For sufficiently small ε, $\mu_T^\theta \in B_\varepsilon\left(\delta_{\theta(\alpha_T)}\right)$ implies that $\mu_T^\theta \in \Delta_{\theta(\alpha)}$ and $\mu_T^\theta \notin \Delta_\theta$ for all $\theta \neq \theta(\alpha)$. Therefore, $b_T = \theta(\alpha_T)$.

(2) When $t = T+1$,

$$\mu_{T+1}^\theta(\theta(\alpha_T)) = \max_{\alpha\in\mathcal{A}} \mu_{T+1}(\theta(\alpha_T)|\alpha) \geq \mu_{T+1}(\theta(\alpha_T)|\alpha_T) > 1-\varepsilon,$$

where the equality comes from $b_T = \theta(\alpha_T)$, and the last weak inequality comes from the definition of E. Consequently, $\mu^{\theta}_{T+1} \in B_\varepsilon\left(\delta_{\theta(\alpha_T)}\right)$ and $b_{T+1} = \theta(\alpha_T)$.

(3) The same reasoning applies to all $t \geq T+2$, so $\mu^{\theta}_t \in B_\varepsilon\left(\delta_{\theta(\alpha_T)}\right)$ and $b_t = \theta(\alpha_T)$ for all $t \geq T$. $\square$

Previous claim implies that

$$E \subset \left\{\omega : \mu^{\theta}_t \in B_\varepsilon(\delta_{\theta'}) \text{ for some } \theta' \in \Theta \text{ for all } t \geq T\right\} \equiv E_1.$$

By definition,

$$\mathbb{P}^*\left(\left\{\omega : |\limsup \mu^{\theta}_t(\theta'') - \liminf \mu^{\theta}_t(\theta'')| \leq \varepsilon \text{ for all } \theta'' \in \Theta\right\}\right) \geq \mathbb{P}^*(E_1) \geq \delta \tag{B.25}$$

Note that (B.25) holds for arbitrary ε and δ between 0 and 1. Letting $\varepsilon \to 0$ and $\delta \to 1$, we have

$$\mathbb{P}^*\left(\left\{\omega : \mu^{\theta}_t \text{ converges}\right\}\right) = 1,$$

so μ^{θ}_t will converge to some limit measure μ^{θ}_∞. Previous discussion shows that μ^{θ}_t will be trapped in arbitrarily small neighborhood around some Dirac belief $\delta_{\theta(\alpha)}$ with arbitrarily large probability. Therefore, whenever μ^{θ}_t converges, it must converge to some Dirac belief on a zero-potential state. In other words,

$$\sum_{\theta' \in \mathcal{U}_{\mathcal{A}}} \mathbb{P}^*\left(\mu^{\theta}_\infty = \delta_{\theta'}\right) = 1,$$

so the proposition is proved.

Proof of Proposition B.2

For state $\theta' \in \mathcal{U}_{\mathcal{A}}$, let α' be any model such that $r(\alpha', \theta') = 0$. We first have the following lemma:

Lemma B.3. *For all $\varepsilon, \delta \in (0,1)$, there exists some $\varepsilon' \in (0, \varepsilon)$ such that for all $\mu \in B_{\varepsilon'}(\delta_{\theta'})$, we have:*

$$\mathbb{P}^*_\mu \left(\omega : \mu_t \left(|\alpha'\right) \in B_\varepsilon \left(\delta_{\theta'}\right) \text{ for all } t \geq 1\right) > \delta$$

where $\mathbb{P}^*_\mu$ denotes the probability measure if the individual has a prior μ. In other words, the probability that the posterior always stays inside a small neighborhood around $\delta_{\theta'}$ can be made arbitrarily large if the prior is sufficiently close to $\delta_{\theta'}$.

Proof. Detailed proofs are omitted for exposition. The idea is to make use of the supermartingale property near the Dirac belief. A similar result can be found in Frick et al. (2020b). $\square$

Proposition B.2 can be proved as below.

Proof. We can pick the ε such that $B_\varepsilon(\delta_{\theta'}) \subset \Delta_{\theta'} \setminus \cup_{\theta \neq \theta'} \Delta_\theta$ and fix any $\delta > 0$. By Assumption B.2, there exists some $K > 1$ such that

$$\forall \theta \neq \theta' : \quad \frac{f\left(s_{\theta'}|\alpha, \theta'\right)}{f\left(s_{\theta'}|\alpha, \theta\right)} > K \text{ for all } \alpha \in \mathcal{A}$$

Let $T = \left\lceil \log_K^{\frac{1-\varepsilon'}{\mu_0(\theta')}} \right\rceil + 1$, where $\mu_0(\theta')$ denotes the individual's initial prior on state θ'. If $s_1 = ... = s_T = s_{\theta'}$, we must have $\mu_T(\theta'|\alpha) > 1 - \varepsilon'$ for all model $\alpha \in \mathcal{A}$, so $\mu_T^\theta \in B_{\varepsilon'}(\delta_{\theta'})$ and $b_T = \theta'$. Due to the fact that data-generating process has full support, it occurs with a strictly positive probability that the first T consecutive signals are $s_{\theta'}$. Denote by E the event that $\mu_t(|\alpha')$ is trapped in $B_\varepsilon(\delta_{\theta'})$ for all $t \geq T$. Combining Lemma B.3 and the fact that $\{s_1 = ... = s_T = s_{\theta'}\}$ is a positive-probability event, it is easy to see that $\mathbb{P}^*(E) > 0$. The rest of the proof resembles the proof of Proposition B.1. For all signal paths in E, we have

$b_T = \theta'$, so

$$\mu_{T+1}^{\theta}\left(\theta'\right) \geq \max_{\alpha \in \mathcal{A}} \mu_{T+1}\left(\theta' | \alpha\right) \geq \mu_{T+1}\left(\theta' | \alpha'\right) > 1 - \varepsilon,$$

where the last inequality comes from the fact that $\mu_{T+1}\left(|\alpha'\right) \in B_{\varepsilon}\left(\delta_{\theta'}\right)$. Therefore, $\mu_{T+1}^{\theta} \in B_{\varepsilon}\left(\delta_{\theta'}\right)$ and $b_{T+1} = \theta'$. The same argument applies for all $t > T$. Consequently, for all signal paths in E, we have $\mu_t^{\theta} \in B_{\varepsilon}\left(\delta_{\theta'}\right)$ for all $t \geq T$. Proposition B.1 implies that μ_t^{θ} converges to some Dirac belief almost surely, so the only possible limit is $\delta_{\theta'}$. As a result, for all signal paths in E, we have $\mu_t^{\theta} \to \delta_{\theta'}$ with exception only on null events. Recall that $\mathbb{P}^*(E) > 0$, so $\mathbb{P}^*\left(\mu_t^{\theta} \to \delta_{\theta'}\right) > 0$. Since θ' is an arbitrary zero-potential state, Proposition B.2 is thus proved. $\square$

Proof of Proposition B.3

Proposition B.3 is intuitively straightforward, since we would expect the fixed bias case should be approximated by the situation where inertia is sufficiently large. To establish this limit-preserving property, we still need some technical arguments. The idea of the proof is illustrated in Figure B.2. Consider an individual whose initial bias is M and suppose that $\mathcal{U}_{\mathcal{A}}^{M} = \{R\}$. Since M is not a zero-potential state,[5] the belief will almost surely escape from Δ_M in finite time. Besides, the larger the Δ_M becomes (i.e., the larger the inertia becomes), the longer the belief remains in Δ_M. Consequently, the escaping belief μ_{τ}^{M} approaches the limit case where individuals hold a fixed bias M. From Theorem 2.1, we know that if the bias is fixed, the belief $\mu_t^{M} \to \delta_R$ almost surely. As a result, if Δ_M is sufficiently large, with a sufficiently large probability, μ_{τ}^{M} will land in a small neighborhood of δ_R and stay within that neighborhood forever,

[5] if it is, M should belong to $\mathcal{U}_{\mathcal{A}}^{M}$ by definition.

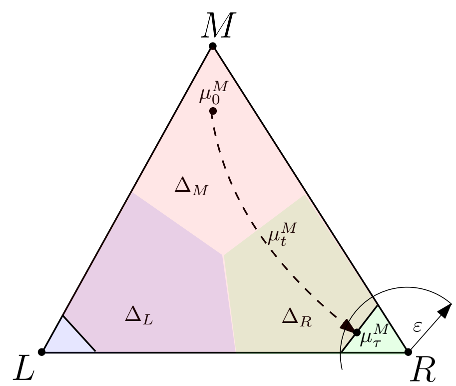

Figure B.2: Illustration of Proposition B.3

in which case we must have $\mu_t^M \to \delta_R$ from Proposition B.1.

Below is the formal proof. Defining a stopping time $\tau = \inf\left\{t : \mu_t^\theta \notin \Delta_\theta\right\}$, we have the following lemmas.

Lemma B.4. *There exists some $\theta_0 \in \Theta$ and some constant $A > 0$ such that*

$$\frac{\mu_\tau(\theta_0)}{\mu_\tau(\theta)} / \frac{\mu_0(\theta_0)}{\mu_0(\theta)} \geq A \times (R_\theta - 1) \tag{B.26}$$

The LHS of (B.26) represents the increment of the likelihood ratio between state θ_0 and θ at the stopping time, τ. Notice that the RHS of (B.26) is increasing in R_θ. Lemma B.4 states a simple fact that when the inertia increases, it becomes harder to escape from Δ_θ in the sense that it requires a larger increment of beliefs for some state.

Proof. From the definition of τ and R_θ, we have $\mu_\tau(\theta) \leq \frac{1}{R_\theta}$, which implies that $\sum_{\tilde{\theta}\neq\theta} \mu_\tau\left(\tilde{\theta}\right) > 1 - \frac{1}{R_\theta}$. Therefore, there must exist some $\theta_0 \in \Theta$ such that $\mu_\tau(\theta_0) \geq \frac{1}{N}\left(1 - \frac{1}{R_\theta}\right)$, which further implies that $\frac{\mu_\tau(\theta_0)}{\mu_\tau(\theta)} \geq \frac{1}{N}(R_\theta - 1)$. □

Lemma B.5. *For all $T < \infty$, as $R_\theta \to \infty$, we have $\mathbb{P}^*(\tau < T) \to 0$.*

This lemma states a simple fact that as the inertia increases, it takes longer for individuals to change their biases.

Proof. From Lemma B.4, we know that:

$$\begin{aligned}
\mathbb{P}^*\left(\tau < T\right) &\leq \sum_{t<T}\sum_{\theta_0\neq\theta}\mathbb{P}^*\left(\frac{\mu_t\left(\theta_0\right)}{\mu_t\left(\theta\right)}/\frac{\mu_0\left(\theta_0\right)}{\mu_0\left(\theta\right)} \geq A\times\left(R_\theta - 1\right)\right) \\
&= \sum_{t<T}\sum_{\theta_0\neq\theta}\mathbb{P}^*\left(\prod_{k\leq t}\frac{f\left(s_k|\alpha,\theta_0\right)}{f\left(s_k|\alpha,\theta\right)} \geq A\times\left(R_\theta - 1\right)\right) \\
&\leq \sum_{t<T}\sum_{\theta_0\neq\theta}\mathbb{P}^*\left(\sum_{k\leq t}\left|\log\frac{f\left(s_k|\alpha,\theta_0\right)}{f\left(s_k|\alpha,\theta\right)}\right| \geq \log A\left(R_\theta - 1\right)\right)
\end{aligned}$$

Applying Markov's inequality, we get:

$$\begin{aligned}
&\mathbb{P}^*\left(\sum_{k\leq t}\left|\log\frac{f\left(s_k|\alpha,\theta_0\right)}{f\left(s_k|\alpha,\theta\right)}\right| \geq \log A\left(R_\theta - 1\right)\right) \leq \frac{t\cdot\mathbb{E}^*\left|\log\frac{f(s|\alpha,\theta_0)}{f(s|\alpha,\theta)}\right|}{\log A\left(R_\theta - 1\right)} \\
\leq &\frac{T\left(\mathbb{E}^*|\log f\left(s|\alpha,\theta_0\right)| + \mathbb{E}^*\log|f\left(s|\alpha,\theta\right)|\right)}{\log A\left(R_\theta - 1\right)} \leq \frac{Tb}{\log A\left(R_\theta - 1\right)},
\end{aligned}$$

for some constant $b \geq 0$, where the last inequality comes from the fact that $\mathbb{E}^*|\log f\left(s|\alpha,\theta_0\right)|$ is uniformly bounded (implied by the Assumption 2.2 and Jensen's inequality). Therefore, we have:

$$\mathbb{P}^*\left(\tau < T\right) \leq \frac{T^2Nb}{\log A\left(R_\theta - 1\right)} \to 0 \quad \text{as } R_\theta \to \infty$$

□

Based on the lemmas above, Proposition B.3 can be proved as follows.

Proof. (i) Consider first the case where $\theta \in \mathcal{U}_{\mathcal{A}}^{\theta}$. Since every model has a unique zero-potential state, we have $\mathcal{U}_{\mathcal{A}}^{\theta} = \{\theta\}$. Let α be the model such that $r\left(\alpha,\theta\right) = 0$. From the fact that $\mu_t\left(.|\alpha\right) \to \delta_\theta$, for all $\varepsilon,\delta \in (0,1)$, there exists some T such that

$$\mathbb{P}^*\left(\mu_t\left(.|\alpha\right) \in B_\varepsilon\left(\delta_\theta\right) \text{ for all } t \geq T\right) \geq 1 - \delta^2,$$

Similar to the previous proof, I choose the ε to be sufficiently small such that $B_\varepsilon\left(\delta_\theta\right)$ is contained in the interior of $\Delta_\theta \backslash \cup_{\theta'\neq\theta}\Delta_{\theta'}$. From Lemma B.5, there exists some $R < \infty$ such that when $R_\theta \geq R$, we have:

$$\mathbb{P}^* \left(\{ \mu_t (.|\alpha) \in B_\varepsilon (\delta_\theta) \text{ for all } t \geq T \} \cap \{ \tau \geq T+1 \} \right) \geq 1 - \delta \tag{B.27}$$

Denote by E_1 the event inside the $\mathbb{P}^*$ operator on the LHS of (B.27). For all $s^\infty \in E_1$, since $\tau(s^\infty) > T$, we must have $\mu_T^\theta \in \Delta_\theta$, so $b_T^\theta = \theta$. Since $\mu_t (.|\alpha) \in B_\varepsilon (\delta_\theta)$ for all $t \geq T$, we have:

$$\mu_{T+1}^\theta (\theta) \geq \max_{\alpha \in \mathcal{A}} \mu_{T+1} (\theta|\alpha) > 1 - \varepsilon,$$

so $\mu_{T+1}^\theta \in B_\varepsilon (\delta_\theta)$, thus $b_{T+1}^\theta = \theta$. The same reasoning applies for all $t \geq T+1$, so we have μ_t^θ remains in the small ball $B_\varepsilon (\delta_\theta)$. Since μ_t^θ converges to a Dirac belief almost surely, we must have for almost every signal path s^∞ in E_1, we have $\mu_t^\theta \to \delta_\theta$. Therefore, $\mathbb{P}^* \left(\mu_t^\theta \to \delta_\theta \right) \geq 1 - \delta$ whenever $R_\theta \geq R$. Since δ can be chosen arbitrarily, we must have $\lim_{R_\theta \to \infty} \mathbb{P}^* \left(\mu_t^\theta \to \delta_\theta \right) = 1$.

(ii) Suppose that $\theta \notin \mathcal{U}_\mathcal{A}^\theta$. The idea of the proof is similar. Denote by $\hat{\mu}_t^\theta$ the belief of the individual with bias θ as in the benchmark model (i.e., this individual has a fixed bias θ). Theorem 2.1 implies that for all $\varepsilon, \delta \in (0,1)$, there exists some T such that

$$\sum_{\theta' \in \mathcal{U}_\mathcal{A}^\theta} \mathbb{P}^* \left(\hat{\mu}_t^\theta \in B_\varepsilon (\delta_{\theta'}) \text{ for all } t \geq T \right) \geq 1 - \delta^2 \tag{B.28}$$

Similarly, I choose the ε to be sufficiently small such that each $B_\varepsilon (\delta_{\theta'})$ is contained in the interior of $\Delta_{\theta'} \setminus \cup_{\theta'' \neq \theta', \theta} \Delta_{\theta''}$. From Lemma B.5, there exists some $R < \infty$ such that when $R_\theta \geq R$, we have:

$$\sum_{\theta' \in \mathcal{U}_\mathcal{A}^\theta} \mathbb{P}^* \left(\left\{ \hat{\mu}_t^\theta \in B_\varepsilon (\delta_{\theta'}) \text{ for all } t \geq T \right\} \cap \{ \tau \geq T+1 \} \right) \geq 1 - \delta$$

Since $\theta \notin \mathcal{U}_\mathcal{A}^\theta$, we must have $\mathbb{P}^* (\tau < \infty) = 1$. This is because if $\{\tau = \infty\}$ occurs with a positive probability, then we have $\hat{\mu}_t^\theta = \mu_t^\theta$ for all t with a positive probability. If that is the case, we must have $\mu_t^\theta = \hat{\mu}_t^\theta \to \delta_{\theta'}$ for some $\theta' \in \mathcal{U}_\mathcal{A}^\theta$ with a positive probability. However, since $\theta \notin \mathcal{U}_{\mathcal{A}}^{\theta}$, $\mu_t^{\theta}(\theta)$ must escape Δ_θ in finite time, which is a contradiction. For each $\theta' \in \mathcal{U}_{\mathcal{A}}^{\theta}$, denote by $E_{\theta'}$ the event inside the $\mathbb{P}^*$ operator above. For all $s^\infty \in E_{\theta'}$ it is a routine to verify that $\mu_{\tau+k}^{\theta} \in B_\varepsilon(\delta_{\theta'})$ for all $k \geq 1$. From the facts that (i) μ_t^{θ} almost surely converges to some Dirac measure, and (ii) $\tau < \infty$, we must have $\mu_t^{\theta} \to \delta_{\theta'}$ for all signal paths in $E_{\theta'}$ (expect for null events), so

$$\sum_{\theta' \in \mathcal{U}_{\mathcal{A}}^{\theta}} \mathbb{P}^* \left(\mu_t^{\theta} \to \delta_{\theta'} \right) \geq 1 - \delta.$$

Since δ is chosen arbitrarily between 0 and 1, we must have $\sum_{\theta' \in \mathcal{U}_{\mathcal{A}}^{\theta}} \mathbb{P}^* \left(\mu_t^{\theta} \to \delta_{\theta'} \right) \to 1$ as $R_\theta \to \infty$. $\square$

APPENDIX C

APPENDIX OF CHAPTER 3

C.1 Omitted Proofs of Chapter 3

C.1.1 Proof of Theorem 3.1

Step 1: *Whenever beliefs converge, they must converge to the same Dirac belief except for null events*

Suppose that $\boldsymbol{\mu_t} = (\mu_{1,t}, ..., \mu_{n,t}) \rightarrow (\mu_{1,\infty}, ..., \mu_{n,\infty}) = \boldsymbol{\mu_\infty}$, and that $\mu_{i,\infty}$ is not a point-mass belief for some $i \in N$. We can find two different states, θ_1 and θ_2, such that $\mu_{i,\infty}(\theta_1), \mu_{i,\infty}(\theta_2) \in (0,1)$. From the interaction rule and the fact that signals are not perfectly informative, we know that

$$\mu_{j,\infty}(\theta_1), \mu_{j,\infty}(\theta_2) \in (0,1) \qquad \forall j \in N^{-1}(i),$$

where $N^{-1}(i)$ denotes the set of individuals who can observe i's beliefs. Since the network is connected, we have

$$\mu_{j,\infty}(\theta_1), \mu_{j,\infty}(\theta_2) \in (0,1) \qquad \forall j \in N.$$

From the assumption, there exists some $j \in N$ and $s_j \in S_j$ such that $\frac{\hat{l}_j(s_j|\theta_1)}{\hat{l}_j(s_j|\theta_2)} \neq 1$. Once a signal s_j is observed, the likelihood ratio between θ_1 and θ_2 will change by some fixed magnitude. To maintain the limit, s_j can only occur for finite periods, which is a probability-0 event. As such, the limit can only be Dirac beliefs except for null events. It is then straightforward to see that all beliefs must converge to the same Dirac belief. This is because if individuals settle on different

Dirac beliefs, some individual will form a mixture belief after communicating with his or her neighbors, which corresponds to the first case, a contradiction. As a consequence, whenever beliefs converge, they must converge to the same Dirac belief except for null events.

Step 2: θ_0 *must minimize the society's weighted relative entropy*

The idea is to show that beliefs cannot accumulate around state θ that does not minimize the weighted relative entropy. For all i, θ and t, the updating rule implies that

$$\mu_{i,t+1}(\theta) = \frac{v_{i,t+1}(\theta) \times \hat{l}_i(s_{i,t+1}|\theta)}{\sum_{\theta' \in \Theta} v_{i,t+1}(\theta') \times \hat{l}_i(s_{i,t+1}|\theta')}, \tag{C.1}$$

where $v_{i,t+1}(\theta)$ denotes the belief of individual i after all communications at period t satisfying

$$\begin{pmatrix} v_{1,t+1}(\theta) \\ v_{2,t+1}(\theta) \\ \vdots \\ v_{n,t+1}(\theta) \end{pmatrix} = G^{k_{t+1}} \times \begin{pmatrix} \mu_{1,t}(\theta) \\ \mu_{1,t}(\theta) \\ \vdots \\ \mu_{n,t}(\theta) \end{pmatrix}.$$

Taking logarithm on both sides of (C.1), we obtain

$$\begin{aligned} \log \mu_{it+1}(\theta) &= \log\left(\sum_{j \in N(i)} g_{ij}^{t+1} \times \mu_{j,t}(\theta)\right) + \log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta)}{\sum_{\theta'} v_{i,t+1}(\theta') \times \hat{l}_i(s_{i,t+1}|\theta')}\right) \\ &\geq \sum_{j \in N(i)} g_{ij}^{t+1} \times \log(\mu_{j,t}(\theta)) + \log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta)}{\sum_{\theta'} v_{i,t+1}(\theta') \times \hat{l}_i(s_{i,t+1}|\theta')}\right), \end{aligned} \tag{C.2}$$

where g_{ij}^{t+1} denotes the $i-j$ th entry of matrix $G^{k_{t+1}}$. Notice that the inequality comes from the concavity of the log function. Multiplying both sides of (C.3) by w_i, and summing over i's, we obtain

$$\sum_{i=1}^{n} w_i \log \mu_{it+1}(\theta) \geq \sum_{i=1}^{n} w_i \log \mu_{it}(\theta) + \sum_{i=1}^{n} w_i \log \left(\frac{\hat{l}_i(s_{i,t+1}|\theta)}{\sum_{\theta'} v_{i,t+1}(\theta') \times \hat{l}_i(s_{i,t+1}|\theta')} \right) \tag{C.3}$$

Let state θ_0 be a state that minimizes the weighted relative entropy, that is,

$$\theta_0 \in \arg\min_{\theta \in \Theta} \sum_{i=1}^{n} w_i \times D_{KL}^{i}(\theta). \tag{C.4}$$

The idea is to show that any state $\theta_1 \in \Theta$ that does not satisfy (C.4) cannot be a limit. To establish the result, I introduce the notion of local stability, which is commonly used in the literature of learning with incorrect models.

Definition 18. A state θ is *locally unstable* if there exists $\varepsilon > 0$ such that for all prior $\mu_0 \in B_\varepsilon(\delta_\theta)$, there almost surely exists $t < \infty$ such that $\mu_t \notin B_\varepsilon(\delta_\theta)$.

Since the belief process $\{\mu_{1.t}, ..., \mu_{n,t}\}$ exhibits the Markov property, it follows immediately that if θ_1 is locally unstable, it is impossible for beliefs to converge to δ_{θ_1}. The goal is to show that θ_1 is locally unstable. Define $m_t(\theta_0) \equiv \sum_{i=1}^{n} w_i \times \log \mu_{i,t}(\theta_0)$. Let $B_{\varepsilon_1}(\delta_{\theta_1})$ be the ε_1-neighborhood of the Dirac belief on state θ_1, that is,

$$B_{\varepsilon_1}(\delta_{\theta_1}) \equiv \{\mu \in \Delta(\Theta) : \mu(\theta_1) \geq 1 - \varepsilon_1\}.$$

Define a stopping time $T = \inf\left\{t : \mu_{i,t} \notin B_{\varepsilon_1}(\delta_{\theta_1}) \text{ for some } i\right\}$, which is the first time that $\mu_{i,t}$ escapes from $B_{\varepsilon_1}(\delta_{\theta_1})$. To show that θ_1 is not locally stable, we must show that $\{T < \infty\}$ occurs with probability 1, or equivalently, $\{T = \infty\}$ occurs with probability 0.

This can be established by contradiction. Suppose that $\{T = \infty\}$ is a positive probability event. On $\{T = \infty\}$, we have

$$
\begin{aligned}
m_{t+1}(\theta_0) &\geq m_t(\theta_0) + \sum_{i=1}^{n} w_i \times \log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta)}{\sum_{\theta' \in \Theta} v_{i,t+1}(\theta') \times \hat{l}_i(s_{i,t+1}|\theta')}\right) \\
&\geq m_t(\theta_0) + \sum_{i=1}^{n} w_i \times \log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta)}{\sum_{\theta' \in \Theta} v_{i,t+1}(\theta') \times \hat{l}_i(s_{i,t+1}|\theta')}\right), \qquad \forall t \geq 0.
\end{aligned}
$$

Summing over t and taking limit, we obtain

$$
\liminf \frac{m_t(\theta_0)}{t} \geq \sum_{i=1}^{n} w_i \times \mathbb{E}_i^* \log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta_0)}{(1-\varepsilon_1)\hat{l}_i(s_{i,t+1}|\theta_1) + \varepsilon_1}\right) \qquad \mathbb{P}^* - a.s, \tag{C.5}
$$

$$
= \sum_{i=1}^{n} w_i \times \mathbb{E}_i^* \log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta_0)}{(1-\varepsilon_1)\hat{l}_i(s_{i,t+1}|\theta_1) + \varepsilon_1}\right), \tag{C.6}
$$

where (C.5) comes from the Strong Law of Large Numbers. By making ε_1 sufficiently small, we can ensure that the RHS of (C.5) is positive. But it implies that $m_t(\theta_0) \to +\infty$, which is a contradiction. Therefore, $\{T = \infty\}$ is a probability-0 event. As a result, limit beliefs can only settle on θ_0 that satisfies (C.4).

C.1.2 Verification of Example 25

I consider a little more general version of Example 25 and suppose that the model perceptions $\left(\hat{l}_1, \hat{l}_2\right)$ are

$\hat{l}_1(s\|\theta)$	H	L
α	x	$1-x$
β	$1/2$	$1/2$
γ	$2/3$	$1/3$

$\hat{l}_2(s\|\theta)$	H	L
α	$1/2$	$1/2$
β	x	$1-x$
γ	$2/3$	$1/3$

.

Note that Example 25 corresponds to the case where $x = 9/10$. We can then show that beliefs will converge when x is sufficiently large ($x = 9/10$ is one example).

Proof of convergence when x is large

Let's then prove the convergence of the beliefs. Denoting by $\mathcal{F}_t \equiv \sigma(s_0, ..., s_t)$, we notice that

$$\begin{aligned}
&\mathbb{E}\left[\log v_{t+1}(\gamma) \mid \mathcal{F}_t\right] \\
&= \log v_t(\gamma) + \mathbb{E} \log\left(\frac{1}{2}\frac{\hat{l}_1(s_{1,t+1}|\gamma)}{\sum_{\hat{\theta}\in\Theta} v_t(\hat{\theta})\,\hat{l}_1(s_{1,t+1}|\hat{\theta})} + \frac{1}{2}\frac{\hat{l}_2(s_{2,t+1}|\gamma)}{\sum_{\hat{\theta}\in\Theta} v_t(\hat{\theta})\,\hat{l}_2(s_{2,t+1}|\hat{\theta})}\right) \\
&\geq \log v_t(\gamma) + \frac{1}{2}\mathbb{E}\log\left(\frac{\hat{l}_1(s|\gamma)}{\sum_{\hat{\theta}\in\Theta} v_t(\hat{\theta})\,\hat{l}_1(s|\hat{\theta})}\right) + \frac{1}{2}\mathbb{E}\log\left(\frac{\hat{l}_2(s|\gamma)}{\sum_{\hat{\theta}\in\Theta} v_t(\hat{\theta})\,\hat{l}_2(s|\hat{\theta})}\right) \\
&\geq \log v_t(\gamma) + \min_{v\in\Delta(\Theta)}\left[\frac{1}{2}\mathbb{E}\log\left(\frac{\hat{l}_1(s|\gamma)}{\sum_{\hat{\theta}\in\Theta} v(\hat{\theta})\,\hat{l}_1(s|\hat{\theta})}\right) + \frac{1}{2}\mathbb{E}\log\left(\frac{\hat{l}_2(s|\gamma)}{\sum_{\hat{\theta}\in\Theta} v(\hat{\theta})\,\hat{l}_2(s|\hat{\theta})}\right)\right].
\end{aligned}$$

Denote by $g(v) \equiv \frac{1}{2}\mathbb{E}\log\left(\frac{\hat{l}_1(s|\gamma)}{\sum_{\hat{\theta}\in\Theta} v(\hat{\theta})\hat{l}_1(s|\hat{\theta})}\right) + \frac{1}{2}\mathbb{E}\log\left(\frac{\hat{l}_2(s|\gamma)}{\sum_{\hat{\theta}\in\Theta} v(\hat{\theta})\hat{l}_2(s|\hat{\theta})}\right)$. Expanding the expression of g, we get

$$\begin{aligned}
g(v) = &\frac{1}{2}\log\frac{v(\alpha)\,x + v(\beta)\,\frac{1}{2} + v(\gamma)\,\frac{2}{3}}{2/3} + \frac{1}{2}\log\frac{v(\alpha)\,(1-x) + v(\beta)\,\frac{1}{2} + v(\gamma)\,\frac{1}{3}}{1/3} \\
&+ \frac{1}{2}\log\frac{v(\alpha)\,\frac{1}{2} + v(\beta)\,x + v(\gamma)\,\frac{2}{3}}{2/3} + \frac{1}{2}\log\frac{v(\alpha)\,\frac{1}{2} + v(\beta)\,(1-x) + v(\gamma)\,\frac{1}{3}}{1/3}.
\end{aligned}$$

It is easy to verify that the minimizing v must satisfy $v(\alpha) = v(\beta) = \frac{1-v(\gamma)}{2}$. Substituting out $v(\alpha)$ and $v(\beta)$, we get

$$\min_{v\in\Delta(\Theta)} g(v) = \min_{v(\gamma)\in[0,1]}\left(\log\frac{\frac{1}{4} + \frac{1}{2}x + \left(\frac{5}{12} - \frac{1}{2}x\right)v(\gamma)}{2/3} + \log\frac{\frac{3}{4} - \frac{1}{2}x + \left(\frac{1}{2}x - \frac{5}{12}\right)v(\gamma)}{1/3}\right).$$

It is easy to show that when x is sufficiently close to 1, the minimizing $v(\gamma) = 1$ (specifically, $x = 9/10$ satisfies the "sufficiently close" condition). In this case, we have $\min_{v\in\Delta(\Theta)} g(v) = 0$. From (A.14), when x is sufficiently close, we have

$$\mathbb{E}\left[\log v_{t+1}(\gamma) \mid \mathcal{F}_t\right] \geq \log v_t(\gamma) + \min_{v\in\Delta(\Theta)} g(v) = \log v_t(\gamma),$$

so $\{\log v_t(\gamma)\}$ constitutes a non-positive submartingale. Applying the martingale convergence theorem, we know that $\log v_t(\gamma)$ converges almost surely. Besides, the only possible limit is $v_\infty(\gamma) = 1$, so $v_t \to \delta_\gamma$, which implies that $\mu_{it} \to \delta_\gamma$.

C.1.3 Proof of Theorem 3.2

I first consider a benchmark case where $k_t = \infty$ and prove that Theorem 3.2 holds in this situation. I then show that the results do not change when k_t is sufficiently large in the limit. We first note that when $k_t = \infty$, we have

$$\boldsymbol{v_t} = \lim_{k\to\infty} G^k \times \boldsymbol{\mu_{t-1}} = \begin{bmatrix} \boldsymbol{w^T} \\ \boldsymbol{w^T} \\ \vdots \\ \boldsymbol{w^T} \end{bmatrix} \times \boldsymbol{\mu_{t-1}},$$

where the convergence comes from the fact that G is irreducible and aperiodic. As such, we have $v_{it} = v_{jt} \equiv v_t$, where v_t denotes the common belief after the communication, called *social belief* at time t. The transition law of social beliefs is as follows

$$v_{t+1}(\theta) = \sum w_i \times BU_i(v_t, s_{i,t})(\theta) = \sum w_i \times \frac{v_t(\theta)\hat{l}_i(s_{i,t}|\theta)}{\sum_{\theta'\in\Theta} v_t(\theta')\hat{l}_i(s_{i,t}|\theta')}. \tag{C.7}$$

Step 1: The Sufficient Part of Theorem 3.2 (when $k_t = \infty$)

I first show that beliefs almost surely do not converge if for all $\theta_0 \in \Theta$, there exists some $\theta \neq \theta_0$ such that

$$\mathbb{E}\log\left(\sum_{i=1}^{n} w_i \times \frac{\hat{l}_i\left(s_i|\theta\right)}{\hat{l}_i\left(s_i|\theta_0\right)}\right) > 0. \tag{C.8}$$

To establish the non-convergence, I am going to show that δ_{θ_0} is locally unstable. Whenever $v_t \in B_\varepsilon\left(\delta_{\theta_0}\right)$,

$$\begin{aligned}
\log\left[\frac{v_{t+1}\left(\theta\right)}{v_{t+1}\left(\theta_0\right)}\right] &= \log\left[\frac{v_t\left(\theta\right)}{v_t\left(\theta_0\right)}\right] + \log\left[\frac{\sum w_i \times \frac{\hat{l}_i\left(s_{i,t}|\theta\right)}{\sum_{\theta'\in\Theta} v_t(\theta')\hat{l}_i\left(s_{i,t}|\theta'\right)}}{\sum w_i \times \frac{\hat{l}_i\left(s_{i,t}|\theta_0\right)}{\sum_{\theta'\in\Theta} v_t(\theta')\hat{l}_i\left(s_{i,t}|\theta'\right)}}\right] \\
&\geq \log\left[\frac{v_t\left(\theta\right)}{v_t\left(\theta_0\right)}\right] + \log\left[\frac{\sum w_i \times \frac{\hat{l}_i\left(s_{i,t}|\theta\right)}{\varepsilon+(1-\varepsilon)\hat{l}_i\left(s_{i,t}|\theta_0\right)}}{\sum w_i \times \frac{\hat{l}_i\left(s_{i,t}|\theta_0\right)}{(1-\varepsilon)\hat{l}_i\left(s_{i,t}|\theta_0\right)}}\right] \\
&\geq \log\left[\frac{v_t\left(\theta\right)}{v_t\left(\theta_0\right)}\right] + \log\left[\sum w_i \times \frac{\left(1-\varepsilon\right)\hat{l}_i\left(s_{i,t}|\theta\right)}{\varepsilon + \left(1-\varepsilon\right)\hat{l}_i\left(s_{i,t}|\theta_0\right)}\right].
\end{aligned}$$

Suppose that $v_1 \in B_\varepsilon\left(\delta_{\theta_0}\right)$. We want to show that v_t will escape from $B_\varepsilon\left(\delta_{\theta_0}\right)$ with probability 1. The idea is similar to the proof of Theorem 3.1. Suppose that $v_t \in B_\varepsilon\left(\delta_{\theta_0}\right)$ for all t, and we have

$$\begin{aligned}
\frac{1}{T}\log\frac{v_{T+1}\left(\theta\right)}{v_{T+1}\left(\theta_0\right)} &\geq \frac{1}{T}\log\frac{v_0\left(\theta\right)}{v_0\left(\theta_0\right)} + \frac{1}{T}\sum_{t=1}^{T}\log\left[\sum w_i \times \frac{\left(1-\varepsilon\right)\hat{l}_i\left(s_{i,t}|\theta\right)}{\varepsilon + \left(1-\varepsilon\right)\hat{l}_i\left(s_{i,t}|\theta_0\right)}\right] \\
&\to \mathbb{E}\log\left(\sum w_i \times \frac{\left(1-\varepsilon\right)\hat{l}_i\left(s_{i,t}|\theta\right)}{\varepsilon + \left(1-\varepsilon\right)\hat{l}_i\left(s_{i,t}|\theta_0\right)}\right) \quad \text{as } T\to\infty
\end{aligned}$$

except for probability-0 events, where the convergence comes from the strong law of large numbers. Since $\mathbb{E}\log\left(\sum w_i \times \frac{\hat{l}_i(s_i|\theta)}{\hat{l}_i(s_i|\theta_0)}\right) > 0$, by making $\varepsilon \to 0$, we have $\frac{1}{T}\log\frac{v_{T+1}(\theta)}{v_{T+1}(\theta_0)}$ converges to a positive constant as $T\to\infty$, so $\frac{v_{t+1}(\theta)}{v_{t+1}(\theta_0)} \to \infty$, which is a contradiction. Therefore, state θ_0 is locally unstable, hence cannot be a limit. Since relation (C.8) holds for all θ_0, we know that every state is locally unstable, so beliefs almost surely do not converge.

Step 2: The Necessary Part of Theorem 3.2 (when $k_t = \infty$)

I adopt the proof by contraposition and show that if there exists a $\theta_0 \in \Theta$, such that $\mathbb{E}\log\left(\sum_{i=1}^n w_i \times \frac{\hat{l}_i(s_i|\theta)}{\hat{l}_i(s_i|\theta_0)}\right) < 0$ for all $\theta \neq \theta_0$, then beliefs converge with a strictly positive probability. I first state the following lemma.

Lemma C.1. *Suppose that X is a random variable satisfying $X \leq C < \infty$. If $\mathbb{E}\log X < 0$, then there exists some $\rho > 0$ such that $\mathbb{E}X^\rho < 1$.*

Proof. Since X is bounded, we can apply the dominated convergence theorem and get

$$\lim_{\rho\to 0+}\mathbb{E}\left(\frac{X^\rho - 1}{\rho}\right) = \mathbb{E}\left(\lim_{\rho\to 0+}\frac{X^\rho - 1}{\rho}\right) = \mathbb{E}\left(\lim_{\rho\to 0+} X^\rho \times \log X\right) = \mathbb{E}\log X < 0.$$

So, there exists $\rho > 0$ such that $\mathbb{E}\left(\frac{X^\rho - 1}{\rho}\right) < 0$, or $\mathbb{E}X^\rho < 1$. □

Based on Lemma C.1, we have the following corollary.

Corollary C.1. *There exists some $\rho > 0$ such that $\mathbb{E}\left(\sum_{i=1}^n w_i \times \frac{\hat{l}_i(s_i|\theta)}{\hat{l}_i(s_i|\theta_0)}\right)^\rho < 1$ for all $\theta \neq \theta_0$.*

Here, I follow the approach in Frick et al. (2020b) and look at the stopped process

$$m_t(\theta) \equiv \left[\frac{v_{t\wedge T}(\theta)}{v_{t\wedge T}(\theta_0)}\right]^\rho,$$

where $T = \inf\left\{t : v_t \notin B_\varepsilon\left(\delta_{\theta_0}\right)\right\}$ and $v_1 \in B_{\varepsilon'}\left(\delta_{\theta_0}\right)$. In other words, T represents the first time that v_t escapes from the ε-neighborhood around δ_{θ_0}. I define $\mathcal{V}_t = \sigma(v_1, ..., v_t)$, the σ-algebra generated by the first t periods social beliefs.

$$\mathbb{E}\left(m_{t+1}(\theta)\,|\mathcal{V}_t\right) = m_t(\theta) \times \mathbb{E}\left[\frac{\sum w_i \times \frac{\hat{l}_i(s_{i,t}|\theta)}{\sum_{\theta'\in\Theta} v_t(\theta')\hat{l}_i(s_{i,t}|\theta')}}{\sum w_i \times \frac{\hat{l}_i(s_{i,t}|\theta_0)}{\sum_{\theta'\in\Theta} v_t(\theta')\hat{l}_i(s_{i,t}|\theta')}}\right]^\rho.$$

Notice that as $\varepsilon \to 0$,

$$\mathbb{E}\left[\frac{\sum w_i \times \frac{\hat{l}_i(s_{i,t}|\theta)}{\sum_{\theta' \in \Theta} v_t(\theta')\hat{l}_i(s_{i,t}|\theta')}}{\sum w_i \times \frac{\hat{l}_i(s_{i,t}|\theta_0)}{\sum_{\theta' \in \Theta} v_t(\theta')\hat{l}_i(s_{i,t}|\theta')}}\right]^\rho \to \mathbb{E}\left(\sum_{i=1}^{n} w_i \times \frac{\hat{l}_i\left(s_i|\theta\right)}{\hat{l}_i\left(s_i|\theta_0\right)}\right)^\rho < 1,$$

so $\{m_t(\theta)\}$ is a (bounded) supermartingale when ε is sufficiently small. Optional stopping theorem implies that

$$\mathbb{E}m_T(\theta) \leq m_1(\theta). \tag{C.9}$$

Suppose that $T < \infty$ with probability 1, which means that v_t escapes from the ε-neighborhood with probability 1. By definition of $B_\varepsilon\left(\delta_{\theta_0}\right)$, we know that $\mu_T(\theta_0) \leq 1 - \varepsilon$.

Besides, we can find $\theta_1 \neq \theta_0$ such that $\mu_T(\theta_1) \geq \frac{\varepsilon}{|\Theta|}$ with probability at least $1/|\Theta|$. In other words, there exists some θ_1 such that $\mathbb{E}m_T(\theta_1) \geq \frac{\varepsilon}{|\Theta|^2(1-\varepsilon)}$. We also note that $m_1(\theta) \leq \frac{\varepsilon'}{1-\varepsilon'}$. By making ε' very small, we can make $m_1(\theta) < m_T(\theta_1)$, which contradicts (C.9). Therefore, v_t will remain in the ε-neighborhood of δ_{θ_0} with a strictly positive probability. In this case, the only possibility is that $v_t \to \delta_{\theta_0}$, so v_t will converge with a strictly positive probability when $v_1 \in B_{\varepsilon'}\left(\delta_{\theta_0}\right)$. From the fact that $\{v_t\}$ is a Markov process, and that it takes finite steps to enter $B_{\varepsilon'}\left(\delta_{\theta_0}\right)$, we know that $v_t \to \delta_{\theta_0}$ with a strictly positive probability for all full-support priors.

Step 3: From $k_t = \infty$ to $\liminf k_t$ sufficiently large

The argument can be easily extended to finite but sufficiently large k_t. As k_t becomes sufficiently large, each individual's belief $v_{i,t}$ will become sufficiently close to the consensus belief v_t in the sense that $\frac{v_{i,t}}{v_t}$ is sufficiently close to 1. It is easy to see that similar martingale arguments still hold.

C.1.4 Proof of Theorem 3.3

The idea is similar to the proof of Theorem 3.1. First, it is easy to see that beliefs can only converge to some Dirac belief whenever they converge. Next, it only remains to show that the Dirac belief concentrates on the state that minimizes the weighted relative entropy. Denote by $\Psi^{(k)}(\boldsymbol{\mu})$ the resulting belief vector after k rounds of communications.

Step 1: Approximation of Ψ near the Dirac beliefs.

The updating rule implies that

$$\mu_{i,t+1}(\theta_0) = \frac{\Psi_i^{(k_{t+1})}(\boldsymbol{\mu_t})(\theta_0) \times \hat{l}_i(s_{i,t+1}|\theta_0)}{\sum_{\theta\in\Theta} \Psi_i^{(k_{t+1})}(\boldsymbol{\mu_t})(\theta) \times \hat{l}_i(s_{i,t+1}|\theta)}, \tag{C.10}$$

where $\Psi^{(k)}$ represents the k-level composition of Ψ. Taking logarithm on both sides, we obtain

$$\log \mu_{i,t+1}(\theta_0) = \log \Psi_i^{(k_{t+1})}(\boldsymbol{\mu}_t) + \log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta_0)}{\sum_{\theta\in\Theta} \Psi_i^{(k_{t+1})}(\boldsymbol{\mu_t}) \times \hat{l}_i(s_{i,t+1}|\theta)}\right). \tag{C.11}$$

Suppose that $\boldsymbol{\mu} \in B_\varepsilon(\delta_{\theta_1})$, where ε is a small number. It is easy to see that $\Psi^{(k)}(\boldsymbol{\mu}) \in B_\varepsilon(\delta_{\theta_1})$, since individuals are averaging their beliefs. I state and prove the following lemma.

Lemma C.2. *Suppose that $\boldsymbol{\mu} \in B_\varepsilon(\delta_{\theta_1})$. Then we have*

$$\Psi_i^{(k)}(\boldsymbol{\mu})(\theta) \in \left[\frac{\left(\sum_j g_{ij}^{(t)} \times \left(\Psi_j^{(k-t)}(\boldsymbol{\mu})(\theta)\right)^p\right)^{1/p}}{[\varepsilon \times (|\Theta|-1)+1]^t}, \frac{\left(\sum_j g_{ij}^{(t)} \times \left(\Psi_j^{(k-t)}(\boldsymbol{\mu})(\theta)\right)^p\right)^{1/p}}{(1-\varepsilon)^t}\right]$$

for all $k \in \mathbb{N}$ and $t \leq k$, where $g_{ij}^{(t)}$ denotes the i-j-th element of G^t.

Proof. Suppose that $t = 1$. By definition, we have

$$\begin{aligned}
\Psi_i^{(k)}(\boldsymbol{\mu})(\theta) &= \frac{\left(\sum_{j \in N(i)} g_{ij} \times \left(\Psi_j^{(k-1)}(\boldsymbol{\mu})(\theta)\right)^p\right)^{1/p}}{\sum_{\theta' \in \Theta}\left(\sum_{j \in N(i)} g_{ij} \times \left(\Psi_j^{(k-1)}(\boldsymbol{\mu})(\theta')\right)^p\right)^{1/p}} \\
&\geq \frac{\left(\sum_{j \in N(i)} g_{ij} \times \left(\Psi_j^{(k-1)}(\boldsymbol{\mu})(\theta)\right)^p\right)^{1/p}}{\sum_{\theta' \neq \theta_1}\left(\sum_{j \in N(i)} g_{ij} \times \varepsilon^p\right)^{1/p} + \left(\sum_{j \in N(i)} g_{ij} \times 1\right)^{1/p}} \\
&= \frac{\left(\sum_{j \in N(i)} g_{ij} \times \left(\Psi_j^{(k-1)}(\boldsymbol{\mu})(\theta)\right)^p\right)^{1/p}}{\varepsilon \times (|\Theta| - 1) + 1}. \qquad \text{(C.12)}
\end{aligned}$$

The inequality comes from that $\boldsymbol{\mu} \in B_\varepsilon\left(\delta_{\theta_1}\right)$, which implies

$$\Psi_j^{(k-1)}(\boldsymbol{\mu})(\theta') \begin{cases} \leq \varepsilon & \text{when } \theta' \neq \theta_1 \\ \leq 1 & \text{when } \theta' = \theta_1 \end{cases}. \qquad \text{(C.13)}$$

We can apply the iterated arguments to (C.12) and obtain

$$\begin{aligned}
\Psi_i^{(k)}(\boldsymbol{\mu})(\theta) &\geq \frac{\sum_j g_{ij} \times \left[\sum_l g_{jl} \times \log \Psi_l^{(k-2)}(\boldsymbol{\mu})(\theta)\right]}{(\varepsilon \times (|\Theta| - 1) + 1)^2} \\
&= \frac{\left(\sum_j g_{ij}^{(2)} \times \left(\Psi_j^{(k-2)}(\boldsymbol{\mu})(\theta)\right)^p\right)^{1/p}}{(\varepsilon \times (|\Theta| - 1) + 1)^2} \\
&\ldots\ldots \\
&\geq \frac{\left(\sum_j g_{ij}^{(t)} \times \left(\Psi_j^{(k-t)}(\boldsymbol{\mu})(\theta)\right)^p\right)^{1/p}}{(\varepsilon \times (|\Theta| - 1) + 1)^t}.
\end{aligned}$$

Similarly, $\boldsymbol{\mu} \in B_\varepsilon\left(\delta_{\theta_1}\right)$ implies that $\Psi_j^{(k-1)}\left(\boldsymbol{\mu}\right)\left(\theta'\right) \begin{cases} \geq 0 & \text{when } \theta' \neq \theta_1 \\ \geq 1-\varepsilon & \text{when } \theta' = \theta_1 \end{cases}$, so

$$\begin{aligned} \Psi_i^{(k)}\left(\boldsymbol{\mu}\right)\left(\theta\right) &\leq \frac{\left(\sum_{j \in N(i)} g_{ij} \times \left(\Psi_j^{(k-1)}\left(\boldsymbol{\mu}\right)\left(\theta\right)\right)^p\right)^{1/p}}{1-\varepsilon} \\ &\leq \ldots \leq \frac{\left(\sum_j g_{ij}^{(t)} \times \left(\Psi_j^{(k-t)}\left(\boldsymbol{\mu}\right)\left(\theta\right)\right)^p\right)^{1/p}}{\left(1-\varepsilon\right)^t}, \end{aligned}$$

which proves the lemma. □

We then have the following corollary.

Corollary C.2. *Suppose that $k_t \leq K$ for all t. Then we have*

$$\Psi_i^{(k_t)}\left(\boldsymbol{\mu}\right)\left(\theta\right) \in \left[\frac{\left(\sum_j g_{ij}^{(k_t)} \times \left(\mu_j\left(\theta\right)\right)^p\right)^{1/p}}{\left[\varepsilon \times \left(|\Theta|-1\right)+1\right]^K}, \frac{\left(\sum_j g_{ij}^{(k_t)} \times \left(\mu_j\left(\theta\right)\right)^p\right)^{1/p}}{\left(1-\varepsilon\right)^K}\right]$$

Proof. It follows directly from Lemma C.2. □

Step 2: Prove that case for $p > 0$.

Let θ_0 be a state that minimizes the weighted relative entropy, and θ_1 be a state that does not. The idea is to show that δ_{θ_1} is locally unstable. Suppose that $\mu_0 \in B_\varepsilon\left(\delta_{\theta_1}\right)$, and we want to show that beliefs escape from $B_\varepsilon\left(\delta_{\theta_1}\right)$ with probability

1. Suppose that it is not the case. From Corollary C.2, we get

$$
\begin{aligned}
&\log \mu_{i,t+1}(\theta_0) \\
=&\log \Psi_i^{(k_{t+1})}(\boldsymbol{\mu_t})(\theta_0) + \log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta_0)}{\sum_{\theta\in\Theta}\Psi_i^{(k_{t+1})}(\boldsymbol{\mu}_t)\times \hat{l}_i(s_{i,t+1}|\theta)}\right) \\
\geq&\log\left[\frac{\left(\sum_j g_{ij}^{(k_{t+1})}\times(\mu_j(\theta))^p\right)^{1/p}}{[\varepsilon\times(|\Theta|-1)+1]^K}\right] + \log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta_0)}{(1-\varepsilon)\,\hat{l}_i(s_{i,t+1}|\theta_1)+\varepsilon}\right) \\
=&\frac{1}{p}\times\log\left(\sum_j g_{ij}^{(k_{t+1})}\times(\mu_{j,t}(\theta_0))^p\right) - K\log(\varepsilon\times(|\Theta|-1)+1) \\
&+\log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta_0)}{(1-\varepsilon)\,\hat{l}_i(s_{i,t+1}|\theta_1)+\varepsilon}\right) \\
&\geq \sum_j g_{ij}^{(k_t)}\log(\mu_{j,t}(\theta_0)) - K\log[\varepsilon\times(|\Theta|-1)+1] + \log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta_0)}{(1-\varepsilon)\,\hat{l}_i(s_{i,t+1}|\theta_1)+\varepsilon}\right).
\end{aligned}
$$

Multiplying both sides by w_i, and summing over i's, we obtain

$$
\begin{aligned}
\sum w_i\times\log\mu_{i,t+1}(\theta_0) \geq \sum w_i\times\log\mu_{i,t}(\theta_0) + \sum w_i\times\log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta_0)}{(1-\varepsilon)\,\hat{l}_i(s_{i,t+1}|\theta)+\varepsilon}\right) \\
-K\log[\varepsilon\times(|\Theta|-1)+1],
\end{aligned}
$$

which implies that

$$
\begin{aligned}
&\frac{1}{t+1}\sum w_i\times\log\mu_{i,t+1}(\theta_0) \\
\geq&\frac{1}{t+1}\sum w_i\times\log\mu_{i,0}(\theta_0) \\
&+\frac{1}{t+1}\sum_{i,t} w_i\times\log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta_0)}{(1-\varepsilon)\,\hat{l}_i(s_{i,t+1}|\theta)+\varepsilon}\right) - K\log[\varepsilon\times(|\Theta|-1)+1] \\
\to&\mathbb{E}\left(\sum w_i\times\log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta_0)}{(1-\varepsilon)\,\hat{l}_i(s_{i,t+1}|\theta)+\varepsilon}\right)\right) - K\log[\varepsilon\times(|\Theta|-1)+1].
\end{aligned}
$$

By letting $\varepsilon\to 0$, we have $\sum w_i\times\log\mu_{i,t+1}(\theta_0)\to\infty$, which establishes the similar contradiction as in the proof of Theorem 3.1.

Step 3: Prove that case for $p < 0$.

This follows from the proof of Theorem 3.4.

C.1.5 Proof of Theorem 3.3

Similarly, when $k_t = \infty$, the consensus is reached at each period, so $v_{it} = v_{jt} \equiv v_t$, where v_t denotes the common belief after the communication, called *social belief* at time t. The key idea is to show that under the p-DeGroot's rule, and when there are sufficiently many communications, we have

$$[v_t(\theta)]^p \approx \sum w_i \times (\mu_{i,t-1}(\theta))^p .$$

Once we establish that, the rest of the proof follows exactly the same as the proof of Theorem 3.2.

Step 1: Show that $[v_t(\theta)]^p \approx \sum w_i \times (\mu_{i,t-1}(\theta))^p$ under adequate communications.

Lemma C.3. *Suppose that $\boldsymbol{\mu} \in B_\varepsilon(\delta_{\theta_0})$, where θ_0 is an arbitrary state in Θ. For all $\varepsilon' > 0$, there exists some $\varepsilon > 0$ and $k_0 \in \mathbb{N}_{++}$ such that*

$$\frac{\left[\Psi_i^{(k)}(\boldsymbol{\mu})(\theta)\right]^p}{\sum w_i \times (\mu_i(\theta))^p} \in \left(\frac{1}{1+\varepsilon'}, 1+\varepsilon'\right)$$

for all $k \geq k_0$, $i \in N$ and $\theta \in \Theta$.

Proof. From Lemma C.2,

$$\left[\Psi_i^{(k)}(\boldsymbol{\mu})(\theta)\right]^p \in \left[\frac{\sum_j g_{ij}^{(k)} \times (\mu_j(\theta))^p}{[\varepsilon \times (|\Theta|-1)+1]^{kp}}, \frac{\sum_j g_{ij}^{(k)} \times (\mu_j(\theta))^p}{(1-\varepsilon)^{kp}}\right].$$

Notice that $g_{ij}^{(k)} \to w_j$ as $k \to \infty$. So, there exists some k_0 such that $\frac{g_{ij}^{(k_0)}}{w_j} \in \left(\frac{1}{1+2\varepsilon'}, 1+2\varepsilon'\right)$. By making ε sufficiently small, we can ensure that

$$\frac{\left[\Psi_i^{(k_0)}(\boldsymbol{\mu})(\theta)\right]^p}{\sum w_i \times (\mu_i(\theta))^p} \in \left(\frac{1}{1+\varepsilon'}, 1+\varepsilon'\right)$$

for all $i \in N$ and $\theta \in \Theta$. Since individuals are only averaging their beliefs, for all $k \geq k_0$, we must have

$$\frac{\left[\Psi_i^{(k)}(\boldsymbol{\mu})(\theta)\right]^p}{\sum w_i \times (\mu_i(\theta))^p} \in \left(\frac{\min_j \Psi_j^{(k_0)}(\boldsymbol{\mu})(\theta)}{\sum w_i \times (\mu_i(\theta))^p}, \frac{\max_j \Psi_j^{(k_0)}(\boldsymbol{\mu})(\theta)}{\sum w_i \times (\mu_i(\theta))^p}\right) \subset \left(\frac{1}{1+\varepsilon'}, 1+\varepsilon'\right),$$

which proves the claim. □

Step 2: Establish the non-convergence results.

The rest of the proof resembles the proof of Theorem 3.2. Lemma C.3 implies that all $v_{i,t}$s are sufficiently close (in terms of ratio) to a consensus belief

$$v_t(\theta) \equiv \frac{\left[\sum w_i \times (\mu_{i,t-1}(\theta))^p\right]^{1/p}}{\sum_{\theta' \in \Theta}\left[\sum w_i \times (\mu_{i,t-1}(\theta'))^p\right]^{1/p}}.$$

The next period consensus belief v_{t+1} can be approximated by the following expression

$$\left[\frac{v_{t+1}(\theta)}{v_{t+1}(\theta_0)}\right]^p = \frac{\sum w_i \times (BU_i(v_t, s_{i,t})(\theta))^p}{\sum w_i \times (BU_i(v_t, s_{i,t})(\theta_0))^p}. \tag{C.14}$$

The expression takes form similar to (C.7). The rest of the proof follows identically as in Theorem 3.2 (we only need to add a power p).

C.1.6 Proof of Theorem 3.4

The proof is similar to the proof Theorem 3.3 and 3.2. First, as in the proof of Theorem 3.3, we can approximate $[v_t(\theta)]^p$ with $\sum w_i \times (\mu_i(\theta))^p$. I therefore focus on the process of v_t, described by (C.14), as it describes the true belief process sufficiently well under adequate communications. Let θ_0 be a state that minimizes the p-inner-weighted relative entropy, so $\frac{1}{p} \times \mathbb{E} \log \left(\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i(s_i|\theta)}{\hat{l}_i(s_i|\theta_0)} \right)^p \right) < 0$,or equivalently, we have

$$\mathbb{E} \log \left(\frac{1}{\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i(s_i|\theta_0)}{\hat{l}_i(s_i|\theta)} \right)^{|p|}} \right) < 0.$$

From Lemma C.1, we know that there exists some $\rho > 0$ such that

$$\mathbb{E} \left(\frac{1}{\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i(s_i|\theta_0)}{\hat{l}_i(s_i|\theta)} \right)^{|p|}} \right)^{\rho} < 1.$$

From (C.14), we have

$$\left[\frac{v_{t+1}(\theta)}{v_{t+1}(\theta_0)} \right]^{|p| \times \rho} = \left[\frac{v_t(\theta)}{v_t(\theta_0)} \right]^{|p| \times \rho} \times \left[\frac{\sum w_i \times \left(\frac{\sum_{\theta' \in \Theta} v_t(\theta') \hat{l}_i\left(s_{i,t}|\theta'\right)}{\hat{l}_i\left(s_{i,t}|\theta_0\right)} \right)^{|p|}}{\sum w_i \times \left(\frac{\sum_{\theta' \in \Theta} v_t(\theta') \hat{l}_i\left(s_{i,t}|\theta'\right)}{\hat{l}_i\left(s_{i,t}|\theta\right)} \right)^{|p|}} \right]^{\rho}.$$

Similarly, we define $m_t(\theta) \equiv \left[\frac{v_{t \wedge T}(\theta)}{v_{t \wedge T}(\theta_0)} \right]^{|p| \times \rho}$, where $T = \inf \left\{ t : v_t \notin B_\varepsilon\left(\delta_{\theta_0}\right) \right\}$ and $v_1 \in B_{\varepsilon'}\left(\delta_{\theta_0}\right)$. As $\varepsilon \to 0$, we have

$$\mathbb{E} \left[\frac{\sum w_i \times \left(\frac{\sum_{\theta' \in \Theta} v_t(\theta') \hat{l}_i\left(s_{i,t}|\theta'\right)}{\hat{l}_i\left(s_{i,t}|\theta_0\right)} \right)^{|p|}}{\sum w_i \times \left(\frac{\sum_{\theta' \in \Theta} v_t(\theta') \hat{l}_i\left(s_{i,t}|\theta'\right)}{\hat{l}_i\left(s_{i,t}|\theta\right)} \right)^{|p|}} \right]^{\rho} \to \mathbb{E} \left(\frac{1}{\sum w_i \times \left(\frac{\hat{l}_i(s_i|\theta_0)}{\hat{l}_i(s_i|\theta)} \right)^{|p|}} \right)^{\rho} < 1.$$

Therefore, $\{m_t(\theta)\}$ is a bounded supermartingale when ε is sufficiently small. Following the proof in Theorem 3.3, we know that beliefs will remain in $B_\varepsilon\left(\delta_{\theta_0}\right)$ with a strictly positive probability, hence $\mu_t \to \delta_{\theta_0}$ with a strictly positive probability.

BIBLIOGRAPHY

Acemoglu, D., Chernozhukov, V., and Yildiz, M. (2016). Fragility of asymptotic agreement under bayesian learning. *Theoretical Economics*, 11(1):187–225. 8

Acemoglu, D., Dahleh, M. A., Lobel, I., and Ozdaglar, A. (2011). Bayesian learning in social networks. *The Review of Economic Studies*, 78(4):1201–1236. 2, 90

Ackert, L. F. and Athanassakos, G. (1997). Prior uncertainty, analyst bias, and subsequent abnormal returns. *Journal of Financial Research*, 20(2):263–273. 54

Alicke, M. D. and Govorun, O. (2005). The better-than-average effect. *The Self in Social Jsudgment*, 1:85–106. 42

Anderson, L. R. and Holt, C. A. (1997). Information cascades in the laboratory. *The American economic review*, pages 847–862. 19

Athanassakos, G. and Kalimipalli, M. (2003). Analyst forecast dispersion and future stock return volatility. *Quarterly Journal of Business and Economics*, pages 57–78. 54

Babcock, L. and Loewenstein, G. (1997). Explaining bargaining impasse: The role of self-serving biases. *Journal of Economic perspectives*, 11(1):109–126. 59

Babcock, L., Loewenstein, G., Issacharoff, S., and Camerer, C. (1995). Biased judgments of fairness in bargaining. *The American Economic Review*, 85(5):1337–1343. 45, 59

Babcock, L., Wang, X., and Loewenstein, G. (1996). Choosing the wrong pond: Social comparisons in negotiations that reflect a self-serving bias. *The Quarterly Journal of Economics*, 111(1):1–19. 45

Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of economics*, 107(3):797–817. 2, 13, 90

Barron, K. (2021). Belief updating: does the 'good-news, bad-news' asymmetry extend to purely financial domains? *Experimental Economics*, 24(1):31–58. 43

Battigalli, P., Catonini, E., Lanzani, G., and Marinacci, M. (2019a). Ambiguity attitudes and self-confirming equilibrium in sequential games. *Games and Economic Behavior*, 115:1–29. 46

Battigalli, P., Cerreia-Vioglio, S., Maccheroni, F., and Marinacci, M. (2015). Self-confirming equilibrium and model uncertainty. *American Economic Review*, 105(2):646–77. 8, 46, 56

Battigalli, P., Francetich, A., Lanzani, G., and Marinacci, M. (2019b). Learning and self-confirming long-run biases. *Journal of Economic Theory*, 183:740–785. 8, 46

Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics*, 117(3):871–915. 58

Berk, R. H. (1966). Limiting behavior of posterior distributions when the model is incorrect. *The Annals of Mathematical Statistics*, pages 51–58. 44, 63, 68, 99, 103, 146, 149, 165

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, 100(5):992–1026. 2, 13, 90

Blume, L., Easley, D., et al. (1998). Rational expectations and rational learning. *Organizations with incomplete information: Essays in economic analysis: A tribute to Roy Radner*, pages 61–109. 42

Blume, L. E. and Easley, D. (1982). Learning to be rational. *Journal of Economic Theory*, 26(2):340–351. 47

Bohren, J. A. (2016). Informational herding with model misspecification. *Journal of Economic Theory*, 163:222–247. 2, 9, 13, 47, 93

Bohren, J. A. and Hauser, D. N. (2021). Learning with heterogeneous misspecified models: Characterization and robustness. *Econometrica*, 89(6):3025–3077. 2, 9, 13, 47, 93

Bose, S. and Renou, L. (2014). Mechanism design with ambiguous communication devices. *Econometrica*, 82(5):1853–1872. 14

Bowen, R., Dmitriev, D., and Galperti, S. (2021). Learning from shared news: when abundant information leads to belief polarization. Technical report, National Bureau of Economic Research. 93

Brous, P. A. and Kini, O. (1993). A reexamination of analysts' earnings forecasts for takeover targets. *Journal of Financial Economics*, 33(2):201–225. 42, 55

Brunnermeier, M. K. and Parker, J. A. (2005). Optimal expectations. *American Economic Review*, 95(4):1092–1118. 58

Bunke, O. and Milhaud, X. (1998). Asymptotic behavior of bayes estimates under possibly incorrect models. *The Annals of Statistics*, 26(2):617–644. 57, 63

Butler, K. C. and Lang, L. H. (1991). The forecast accuracy of individual analysts: Evidence of systematic optimism and pessimism. *Journal of Accounting Research*, pages 150–156. 42

Camerer, C. and Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach. *American economic review*, 89(1):306–318. 51

Caplin, A. and Leahy, J. (2001). Psychological expected utility theory and anticipatory feelings. *The Quarterly Journal of Economics*, 116(1):55–79. 58

Cassar, L. and Klein, A. H. (2019). A matter of perspective: How failure shapes distributive preferences. *Management Science*, 65(11):5050–5064. 52

Çelen, B. and Kariv, S. (2004). Distinguishing informational cascades from herd behavior in the laboratory. *American Economic Review*, 94(3):484–498. 19

Chang, J. W. and Choi, H. M. (2017). Analyst optimism and incentives under market uncertainty. *Financial Review*, 52(3):307–345. 54

Chang, Y.-C. and Cheng, H.-W. (2015). Information environment and investor behavior. *Journal of Banking & Finance*, 59:250–264. 42

Chen, J. Y. (2019). Sequential learning under informational ambiguity. Available at SSRN: `https://ssrn.com/abstract=3480231`. 46, 93

Chen, J. Y. (2020). Biased learning under ambiguous information. *Available at SSRN 3738440*. 8

Coutts, A. (2019). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics*, 22(2):369–395. 43

Das, S., Levine, C. B., and Sivaramakrishnan, K. (1998). Earnings predictability and bias in analysts' earnings forecasts. *Accounting Review*, pages 277–294. 54

Dasaratha, K. and He, K. (2020). Network structure and naive sequential learning. *Theoretical Economics*, 15(2):415–444. 10, 93

Deffains, B., Espinosa, R., and Thöni, C. (2016). Political self-serving bias and redistribution. *Journal of Public Economics*, 134:67–74. 45, 52

DeMarzo, P. M., Vayanos, D., and Zwiebel, J. (2003). Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics*, 118(3):909–968. 10, 92

Easterwood, J. C. and Nutt, S. R. (1999). Inefficiency in analysts' earnings forecasts: Systematic misreaction or systematic optimism? *The Journal of Finance*, 54(5):1777–1797. 42, 55

Eichberger, J., Grant, S., and Kelsey, D. (2007). Updating choquet beliefs. *Journal of Mathematical Economics*, 43(7-8):888–899. 14

Eichberger, J., Grant, S., and Kelsey, D. (2016). Randomization and dynamic consistency. *Economic Theory*, 62(3):547–566. 16

Eil, D. and Rao, J. M. (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2):114–38. 43

Epstein, L. G. and Schneider, M. (2007). Learning under ambiguity. *The Review of Economic Studies*, 74(4):1275–1303. 8, 88, 143

Ertac, S. (2011). Does self-relevance affect information processing? experimental evidence on the response to performance and non-performance feedback. *Journal of Economic Behavior & Organization*, 80(3):532–545. 43

Esponda, I. and Pouzo, D. (2016). Berk–nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica*, 84(3):1093–1130. 9, 57

Eyster, E. and Rabin, M. (2010). Naive herding in rich-information settings. *American economic journal: microeconomics*, 2(4):221–43. 10, 93

Ford, J., Kelsey, D., and Pang, W. (2013). Information and ambiguity: herd and contrarian behaviour in financial markets. *Theory and Decision*, 75(1):1–15. 8

Frick, M., Iijima, R., and Ishii, Y. (2020a). Misinterpreting others and the fragility of social learning. *Econometrica*, 88(6):2281–2328. 9, 47, 93

Frick, M., Iijima, R., and Ishii, Y. (2020b). Stability and robustness in misspecified learning models. 2, 9, 47, 93, 104, 174, 187

Fryer Jr, R. G., Harms, P., and Jackson, M. O. (2019). Updating beliefs when evidence is open to interpretation: Implications for bias and polarization. *Journal of the European Economic Association*, 17(5):1470–1501. 8, 46

Fudenberg, D., Lanzani, G., and Strack, P. (2021). Limit points of endogenous misspecified learning. *Econometrica*, 89(3):1065–1098. 47, 164

Fudenberg, D., Romanyuk, G., and Strack, P. (2017). Active learning with a misspecified prior. *Theoretical Economics*, 12(3):1155–1189. 47

Gale, D. and Kariv, S. (2003). Bayesian learning in social networks. *Games and economic behavior*, 45(2):329–346. 90

Galperti, S. (2019). Persuasion: The art of changing worldviews. *American Economic Review*, 109(3):996–1031. 60

Ghirardato, P., Maccheroni, F., and Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*, 118(2):133–173. 33

Gilboa, I. and Marinacci, M. (2016). Ambiguity and the bayesian paradigm. In *Readings in formal epistemology*, pages 385–439. Springer. 46

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18:pp–141. 15

Golub, B. and Jackson, M. O. (2010). Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–49. 10, 92

Graham, J. R., Harvey, C. R., and Huang, H. (2009). Investor competence, trading frequency, and home bias. *Management Science*, 55(7):1094–1106. 42

Grinblatt, M. and Keloharju, M. (2009). Sensation seeking, overconfidence, and trading activity. *The Journal of Finance*, 64(2):549–578. 42

Guarino, A. and Jehiel, P. (2013). Social learning with coarse inference. *American Economic Journal: Microeconomics*, 5(1):147–74. 10

Haisley, E. C. and Weber, R. A. (2010). Self-serving interpretations of ambiguity in other-regarding behavior. *Games and Economic Behavior*, 68(2):614–625. 45

Hansen, L. and Sargent, T. J. (2001). Robust control and model uncertainty. *American Economic Review*, 91(2):60–66. 26

Heath, C. and Tversky, A. (1991). Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty*, 4(1):5–28. 51

Heidhues, P., Kőszegi, B., and Strack, P. (2018). Unrealistic expectations and misguided learning. *Econometrica*, 86(4):1159–1214. 47

Herrera, H. and Hörner, J. (2012). A necessary and sufficient condition for information cascades. 2

Hurwicz, L. (1951). Some specification problems and applications to econometric models. *Econometrica*, 19(3):343–344. 33

Jadbabaie, A., Molavi, P., Sandroni, A., and Tahbaz-Salehi, A. (2012). Non-bayesian social learning. *Games and Economic Behavior*, 76(1):210–225. 92, 94, 100

Kalai, E. and Lehrer, E. (1993). Rational learning leads to nash equilibrium. *Econometrica*, pages 1019–1045. 38, 137

Karelaia, N. and Hogarth, R. M. (2010). The attraction of uncertainty: Interactions between skill and levels of uncertainty in market-entry games. *Journal of Risk and Uncertainty*, 41(2):141–166. 51

Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892. 16, 35

Köszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association*, 4(4):673–707. 58

Li, W. and Tan, X. (2020). Locally bayesian learning in networks. *Theoretical Economics*, 15(1):239–278. 10, 93

Loewenstein, G. and Adler, D. (1995). A bias in the prediction of tastes. *The Economic Journal*, 105(431):929–937. 59

Machina, M. J. and Siniscalchi, M. (2014). Ambiguity and ambiguity aversion. In *Handbook of the Economics of Risk and Uncertainty*, volume 1, pages 729–807. Elsevier. 14, 46

Marinacci, M. (2002). Learning from ambiguous urns. *Statistical Papers*, 43(1):143. 8, 46

Marinacci, M. (2015). Model uncertainty. *Journal of the European Economic Association*, 13(6):1022–1100. 8, 46

Marinacci, M. and Massari, F. (2019). Learning from ambiguous and misspecified models. *Journal of Mathematical Economics*, 84:144–149. 8, 46

Mobius, M. M., Niederle, M., Niehaus, P., and Rosenblat, T. S. (2011). Managing self-confidence: Theory and experimental evidence. Available at SSRN: https://www.nber.org/papers/w17014. 43

Molavi, P., Tahbaz-Salehi, A., and Jadbabaie, A. (2018). A theory of non-bayesian social learning. *Econometrica*, 86(2):445–490. 10, 92, 94, 100, 109, 110

Nyarko, Y. (1991). Learning in mis-specified models and the possibility of cycles. *Journal of Economic Theory*, 55(2):416–427. 47

Ortoleva, P. (2012). Modeling the change of paradigm: Non-bayesian reactions to unexpected news. *American Economic Review*, 102(6):2410–36. 60

Pires, C. (2002). A rule for updating ambiguous beliefs. *Theory and Decision*, 53(2):137–152. 14

Pronin, E., Lin, D. Y., and Ross, L. (2002). The bias blind spot: Perceptions of bias in self versus others. *Personality and Social Psychology Bulletin*, 28(3):369–381. 59

Rabin, M. and Schrag, J. L. (1999). First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics*, 114(1):37–82. 46, 87

Rosenberg, D., Solan, E., and Vieille, N. (2009). Informational externalities and emergence of consensus. *Games and Economic Behavior*, 66(2):979–994. 90

Rosenberg, D. and Vieille, N. (2019). On the efficiency of social learning. *Econometrica*, 87(6):2141–2168. 10, 29, 126

Saito, K. (2012). Subjective timing of randomization and ambiguity. Technical report, mimeo. 16

Siniscalchi, M. (2011). Dynamic choice under ambiguity. *Theoretical Economics*, 6(3):379–421. 86

Smith, L. and Sørensen, P. (2000). Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398. 2, 6, 13, 19, 27, 29, 116

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The review of economic studies*, 23(3):165–180. 86

Taber, C. S. and Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science*, 50(3):755–769. 42

Taylor, S. E. and Brown, J. D. (1988). Illusion and well-being: a social psychological perspective on mental health. *Psychological bulletin*, 103(2):193. 42