The Wisdom of Crowds or Group Irrationality? Non-Bayesian Social Learning with Misinformation*

Jaden Yang Chen†

July 4, 2024

Abstract

This paper investigates a non-Bayesian social learning problem where individuals may encounter misinformation and incorrectly specify their signal structures. I introduce a new concept—*group irrationality*—which describes situations where the outcome of social learning contradicts the results of independent learning by each individual. The paper characterizes asymptotic beliefs and demonstrates that, contrary to the "wisdom of crowds" notion, group irrationality is prevalent across many common learning rules. Notably, in many cases, society may still experience incorrect learning even when all individuals are capable of learning the truth independently.

JEL Classification: D81, D83

Keywords: Social learning, misinformation, model misspecification, group irrationality, the wisdom of crowds

*The paper was previously circulated as "Naive Social Learning with Heterogeneous Model Perceptions". It is based on the third chapter of my thesis at Cornell University. I thank my thesis committee, David Easley (chair), Larry Blume and Tommaso Denti.

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

"Human beings are not perfectly designed decision makers [...] despite all these limitations, when our imperfect judgements are aggregated in the right way, our collective intelligence is often excellent."

—James Surowiecki, *The Wisdom of Crowds*

"In crowds it is stupidity and not mother-wit that is accumulated."

—Gustave Le Bon, *The Crowd: A Study of the Popular Minds*

1 Introduction

Information exchange is ubiquitous in the modern age. Through channels like social media, professional networks, and online forums, individuals continually share information and learn from one another. This phenomenon, known as social learning, raises critical questions: Does social learning facilitate correct decisions, or can it lead to worse outcomes?

In the literature, two opposing views exist regarding the effect of social learning. The first perspective warns that social learning can lead to suboptimal outcomes. For instance, in sequential social learning, society may take an incorrect action and form an informational cascade even if all individuals are rational Bayesian learners (Banerjee, 1992; Bikhchandani et al., 1992). The second perspective, however, posits that social learning can effectively aggregate information and achieve *the wisdom of crowds*. Studies have shown that Bayesian social learning can achieve efficient information aggregation under moderate restrictions on learning environments such as signal and network structures.[1] Furthermore, even if individuals are not Bayesian learners, information can still be efficiently aggregated under heuristic learning rules, such as when beliefs are aggregated using a weighted average rule (Golub and Jackson, 2010; Jadbabaie et al., 2012).

The results on non-Bayesian social learning demonstrate the possibility of efficient information aggregation with boundedly rational individuals, which aligns closely with the original idea of the wisdom of crowds.[2] However, the literature mostly assumes that individuals accurately interpret their information. Albeit being a natural benchmark, it overlooks

[1]For example, Smith and Sørensen (2000) show that in sequential social learning, information cascades are non-generic properties; they show that correct learning can occur with unbounded signals. Acemoglu et al. (2011) extend this insight in a broader framework, showing that correct learning can be obtained if signals are unbounded and the network structure satisfies certain expanding property.

[2]The wisdom of crowds emphasizes that society can successfully aggregate information despite each of its member being boundedly rational. Surowiecki (2005) provided many examples where the aggregated opinions of laymen can be accurate, sometimes outperforming expert opinion. A well-known example is that at a country fair, individuals were asked to guess the weight of an ox; surprisingly, their average guess was very close to the ox's true weight, despite most of them having no expert knowledge of cattle.

an important feature of social learning—the presence of *misinformation*, which is prevalent in real life and can significantly distort the dynamics of social learning. When individuals incorrectly interpret their information or are exposed to false information, the success of social learning depends not only on whether dispersed information can be effectively aggregated, but also on the extent to which incorrect information distorts the learning process. For instance, while online reviews can guide consumers toward better purchasing decisions by providing more information, they can also perpetuate misinformation and reinforce biases. To gain a better understanding of social learning, it is imperative to consider scenarios where incorrect information is present. The analysis, however, is not straightforward. In addition to technical challenges, it is unclear which standard to use to evaluate social learning.[3]

In the paper, I present a model of non-Bayesian social learning that allows misinformation. In the model, a finite group of individuals try to learn the true state of the world. The learning process has two stages. The first stage involves social learning, where individuals share their current beliefs with each other. At the end of this stage, the society forms a social belief through an aggregation function F. The paper focuses on aggregation rules that can be approximated by the generalized mean of individuals' beliefs, that is:

$$F\left(\mu_1, ..., \mu_n\right)(\theta) \propto \left(\sum w_i \times \mu_i^p(\theta)\right)^{1/p},$$

where w_i describes the weight of individual i, and p determines the shape of the aggregation rule, called the degree of the aggregator. The second stage is private learning, where individuals receive private signals and update the social belief using Bayes' rule based on their private signals. An important feature of the model is that individuals may be misspecified about their data-generating processes, allowing misinformation to spread through the learning process. The process repeats itself and generates a sequence of beliefs.

To discuss social learning, the paper introduces a new concept—*group irrationality*, which refers to the event that social learning outcome is inconsistent with all individuals' learning outcomes if they were to learn independently. Specifically, group irrationality occurs whenever social learning leads to an asymptotic posterior that assigns positive probability to a state that would be assigned probability 0 in everyone's independent learning. Consequently, society will form a posterior that no individual could have formed on their own. This paper also introduces a related concept, *strong group irrationality*, defined as the event that the so-

[3] With misinformation, correct learning may not be an interesting criterion, because we can easily achieve incorrect learning under some model perceptions, e.g., when all individuals perceive a highly biased signal structure. To make the discussion meaningful, we either restrict the set of model perceptions, or we consider an alternative learning criterion that is more suitable to the situations where individuals may have arbitrary model perceptions.

ciety doesn't learn the true state even if all individuals were able to learn it independently. In real life, numerous examples exist where group irrationality is likely to occur. For instance, a company board may approve decisions that each board member would likely oppose, such as the AOL-Time Warner merger in 2000, which many individual board members later admitted to having reservations about.[4] Other examples include wartime scenarios where soldiers may commit atrocities they would never consider in civilian life and in jury trials, where the collective verdict diverges from the opinions of individual jurors.[5] It is worth noting that group irrationality is different from incorrect learning. The latter emphasizes whether the learning outcome is correct, whereas the former emphasizes the inconsistency between social and individual learning.

To explore group irrationality, this paper presents a series of characterizations of asymptotic belief dynamics. Section 5 provides a benchmark characterization. Proposition 1 shows that in many situations, beliefs tend to concentrate on a state θ that minimizes society's *weighted relative entropy*, i.e., the weighted average of the Kullback-Leibler divergence between each individual's true signal distribution and their perceived distribution in state θ. Essentially, society tends to settle on the state that minimizes the average distance between individuals' perceived and true signal distributions. This characterization partly explains the emergence of group irrationality. To see this, recall a famous result from Berk (1966) that shows in independent learning, asymptotic posteriors will concentrate on the states that minimize the relative entropy of perceived signal distributions. However, the state that minimizes the society's average relative entropy may not minimize the relative entropy for any individual, potentially leading the society to adopt a compromised belief that would not be formed in independent learning, thereby producing group irrationality.

Section 6 provides a more complete characterization of asymptotic beliefs. I define a new notion of divergence—*the weighted p-entropy*. Formally, a state θ has a lower weighted p-entropy than that of state θ' if the expected log of the weighted p-mean of perceived likelihood ratios between θ' and θ is less than 0, where p represents the degree of the belief aggregator. Theorem 1 demonstrates that asymptotic beliefs can be tightly characterized by the weighted p-entropy. Specifically, whenever beliefs converge, they will—and in a certain

[4]Various sources indicate that many high-level executives had reservations regarding the merging decision, which becomes one of the most infamous corporate failures in history. As Bob Pittman, the AOL's former COO, wrote in an email to Fortune, "I think everyone involved in the deal certainly had some doubts, but given that we went forward with the deal, we thought the positives outweighed the negatives." (source: https://www.aol.com/finance/jerry-levin-known-ceo-pushed-154124725.html)

[5]Group irrationality is also consistent what Gustave Le Bon said in his famous book *The Crowd* regarding crowd psychology: "(collective mind) makes them feel, think, and act in a manner quite different from that in which each individual of them would feel, think, and act were he in a state of isolation. There are certain ideas and feelings which do not come into being, or do not transform themselves into acts except in the case of individuals forming a crowd."(Le Bon, 1895)

sense, will only—concentrate on the minimizers of the weighted p-entropy, i.e., states with lower weighted p-entropy compared to other states. Furthermore, Theorem 1 establishes the connection between asymptotic belief dynamics and the shape of the belief aggregator. It reveals that belief dynamics can exhibit distinct patterns depending on the sign of the degree p. For $p > 0$, beliefs may oscillate infinitely often; whereas for $p < 0$, beliefs may converge to multiple limits, each with positive probability. These patterns emerge because the binary relation induced by the weighted p-entropy may fail to be an order on the state space—when $p > 0$, the relation can be incomplete, and no state can dominate all others; when $p < 0$, the relation can be intransitive, allowing multiple states to strictly dominate other states. As p approaches 0, the weighted p-entropy approaches the weighted relative entropy, returning to the benchmark characterization in Proposition 1.

Using previous characterization, Section 7 examines the robustness of correct learning. In contrast to the wisdom of crowds, I show that group irrationality is prevalent in social learning with misinformation. Formally, an aggregation rule is defined as *susceptible to (strong) group irrationality* if, under this aggregation rule, social learning produces (strong) group irrationality with positive probability under some model perceptions. Proposition 2 shows that under certain regularity conditions, *all* aggregation rules are susceptible to group irrationality. Therefore, with specific types of misinformation, social learning can produce beliefs inconsistent with independent learning. Proposition 2 further demonstrates that all aggregation rules with non-zero degree are susceptible to strong group irrationality. In these situations, we can find "innocuous" misinformation such that every individual can learn correctly, but society as a whole cannot. In Section 7, I also address which types of misinformation—i.e., perceived signal structures—allow society to achieve correct learning under various aggregation rules. With these model perceptions, social learning is robust, meaning its success is not heavily dependent on the specific aggregation rules individuals employ. Proposition 3 shows that these perceived models can be characterized using the limit weighted p-entropy as $p \to +\infty$. An implication is that it becomes increasingly challenging to achieve robust social learning as society grows larger.

The remainder of the paper is organized as follows. Section 2 presents examples of group irrationality. Section 3 presents the main model. Section 4 discusses the main assumptions. Sections 5 and 6 present characterizations of asymptotic beliefs. Section 7 discusses implications on group irrationality. Section 8 presents an extension of generalized Bayes' rule in private learning. Sections 9 and 10 are literature review and conclusion.

2 Examples

This section shows examples of how group irrationality can arise. To get the basic idea, consider a society 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 beliefs from the previous period and apply the DeGroot's rule to aggregate information. The social 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 communication, a signal is then realized, but individuals may misspecify the data-generating process. They apply Bayes rule to the social 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, studies have shown that social learning can lead to correct learning in the limit, consistent with the wisdom of crowds. However, if either individual imprecisely specifies the true data-generating process, social learning may produce worse outcomes than that of independent learning. Such group irrationality can even occur with seemingly harmless model perceptions. Below are two examples.

Example 1. (Group irrationality-1) There are three states $\Theta = \{\alpha, \beta, \gamma\}$ and two signals $S = \{H, L\}$. 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

.

Suppose that signals are i.i.d. and the true signal structure is $\hat{l}_1$. At time t, myopic individuals take an action $a_{i,t} \in \{h, l\}$ to maximize the one-period expected payoff of $u(a, \theta)$, where

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

so the optimal action in states α and β is h, and the optimal action in state γ is l. 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 incorrectly specified, he only rearranges the distributions between α and β. A natural conjecture is that this rearrangement should not affect optimal learning because both individuals can independently learn optimally, and that social belief is the average of their beliefs. However, this conjecture is incorrect. For instance, suppose that $\theta^* = \beta$, and hence the optimal action is h. If individuals were to learn independently, the asymptotic posteriors of individuals 1 and 2 would settle on state β and α respectively, so both of them would take the optimal action, action h, in the limit. However, in social learning, the asymptotic posteriors of both individuals will settle on state γ, so optimal learning cannot be achieved.

Example 2. (Group irrationality-2) Suppose that the state space $\Theta = \{G, B\}$, and the signal space is $S = \{S_G, S_B\}$. The perceived data-generating processes are

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

, 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 equal p^*, both individuals correctly understand the signal direction—meaning that they know signal S_θ better indicates state θ. This type of misspecification seems innocuous as individuals can still deduce the true state by comparing the frequency of each signal—If there are more S_G than S_B over time, the true state is G; otherwise, the true state is B. However, in social learning, individuals may *not* learn the truth, and beliefs can oscillate forever as shown later.

3 Model Setup

The state space Θ is finite, and the true state $\theta^* \in \Theta$ is fixed and unknown. A society consists of a finite set of individuals, $N = \{1, 2, 3, ..., n\}$. Every individual $i \in N$ holds a full-support prior $\mu_{i,0} \in \Delta_{++}(\Theta)$ and is trying to learn the true state. Time is discrete $t \in \mathbb{T} = \{1, 2, ...\}$. Each period t consists of two stages of learning—social learning and private learning.

3.1 Social learning stage

The first stage is social learning, in which individuals communicate their beliefs with others. Ultimately, society forms a ***social belief*** v_t which satisfies

$$v_t = F\left(\mu_{1,t-1}, ..., \mu_{n,t-1}\right), \text{where } F : \Delta^n(\Theta) \to \Delta(\Theta). \tag{1}$$

F is referred as the ***belief aggregator***, which provides a reduced-form description of how society aggregates information. One assumption is that society reaches a consensus after communication. One can think of it as the steady state of some underlying social learning process, so it approximates situations where the society experiences adequate exchanges of opinions. It is worth noting that the social belief needn't be consistent with Bayesian updating; instead, it may come from some ***non-Bayesian rules***. One example is that individuals repeatedly apply DeGroot's rule, in which case the social belief is equal to the weighted average of everyone's belief (DeGroot, 1974). This paper primarily focuses on aggregation rules generalizing DeGroot's rule, i.e., society takes average of individuals' beliefs. The following assumption is made:

Assumption 1. (Power-mean tail) *When* $\max_i \mu_i(\theta) \to 0$*, we have*

$$F\left(\mu_1, ..., \mu_n\right)(\theta) \sim \left(\sum w_i \times \mu_i^p(\theta)\right)^{1/p} \tag{2}$$

for some $p \in \mathbb{R}$ *and* $w \in \Delta_{++}(N)$.[6]

Assumption 1 encompasses a broad range of aggregation rules. First, society can aggregate beliefs using any power mean with DeGroot's rule being a special case when $p = 1$. Second, it only requires the social belief to be locally approximated by the power-mean when everyone's belief is close to 0. Throughout this paper, I refer to the power p as the ***degree*** of the aggregation rule and w_i as the ***weight*** of individual i. Below is an example.

Example 3. (Power-mean rule) Suppose that F satisfies

$$F\left(\mu_1, ..., \mu_n\right)(\theta) = \frac{\left(\sum w_i \times \mu_i^p(\theta)\right)^{1/p}}{\sum_{\theta' \in \Theta}\left(\sum w_i \times \mu_i^p(\theta')\right)^{1/p}}, \tag{3}$$

i.e., the social belief is the power-mean of individuals' beliefs with a normalization term.

3.2 Private learning stage

The second stage of learning is private learning. In this stage, a signal profile $s_t = (s_{1,t}, ..., s_{n,t}) \in S$ is generated according to some data-generating process (DGP) $l(s|\theta)$. The signal space

[6] Here, $f(x) \sim g(x)$ as $x \to x_0$ means that $\lim_{x \to x_0} \frac{f(x)}{g(x)} = 1$. When $p = 0$, we define $F(\mu_1, ..., \mu_n)(\theta) \sim \exp\left(\sum w_i \times \log \mu_i(\theta)\right)$.

$S = \times_{i=1}^{n} S_i$ is finite. Signals are not perfectly revealing, thus, $l(s|\theta) \in (0,1)$ for all $s \in S$. The paper assumes the absence of identification problems, meaning that no state θ' induces an identical distribution as the true distribution. Let $l_i(s|\theta)$ denotes the i-th marginal distribution of $l(s|\theta)$. Let $\mathbb{P}$ and $\mathbb{E}$ denote the probability measure and expectation induced by the true distribution, $l(s|\theta^*)$, on the probability space of all signal paths, $(S^\infty, \sigma(S^\infty))$. Signals are independent across time but may be correlated across individuals. Individual i can only observe s_i but not others. Individual i may be ***misspecified*** about their DGPs. Let $\hat{l}_i(s|\theta)$ denote the perceived DGP of individual i, which may differ from the true DGP, $l_i(s|\theta)$. I also assume the absence of identification problems by requiring $\hat{l}_i(\cdot|\theta) \neq \hat{l}_i(\cdot|\theta')$ for all $\theta \neq \theta'$ and all i. The posterior belief $\mu_{i,t}$ satisfies

$$\forall \theta \in \Theta: \qquad \mu_{i,t}(\theta) = BU_i(v_t, s_{i,t})(\theta) = \frac{v_t(\theta) \times \hat{l}_i(s_{i,t}|\theta)}{\sum_{\theta' \in \Theta} v_t(\theta') \times \hat{l}_i(s_{i,t}|\theta')}. \tag{4}$$

That is, each individual updates the social belief v_t based on their private signal $s_{i,t}$ and perceived DGP $\hat{l}_i$ using Bayes' rule. Subsequently, individuals communicate with each other again, and the society forms a new social belief. Based on this new belief, individuals update their posteriors, and the learning process repeats itself infinitely, generating a sequence of beliefs $\boldsymbol{\mu} = (\mu_1, \mu_2, ...)$,where $\mu_t = (\mu_{1,t}, ..., \mu_{n,t})$ denotes the profile of beliefs at time t. This paper aims to characterize asymptotic beliefs. To facilitate discussion, I focus on situations where the following regularity assumption is satisfied.

Assumption 2. (Irreducibility) *For all $t \in \mathbb{T}$, $v_t \in \Delta_{++}(\Theta)$, $\theta \in \Theta$ and $\varepsilon > 0$, there exists $T < \infty$ and signal profile $(s_t, ..., s_{t+T}) \in S$ such that $v_{t+T}(\theta) > 1 - \varepsilon$.*

Alternatively, it is possible for social beliefs to assign high probability on any given state. If this assumption is violated, society will never learn the truth for some choice of the true state, so suboptimal learning outcomes emerge trivially by definition. Technically, Assumption 2 allows extreme beliefs to be reached with positive probability, which enables us to employ the local properties of belief aggregators imposed by Assumption 1.

3.3 Group Irrationality

I then introduce the concept of group irrationality. Intuitively, group irrationality means that society achieves a learning outcome inconsistent with the outcomes that could have been achieved independently by each individual. To formalize this idea, I define the following concepts

$$\mathcal{R}_i(\theta) \equiv \mathbb{E} \log \left(\frac{l_i(s|\theta^*)}{\hat{l}_i(s|\theta)} \right), \quad \Theta_i \equiv \arg\min_{\theta \in \Theta} \mathcal{R}_i(\theta),$$

where $\mathcal{R}_i(\theta)$ is the ***relative entropy*** of state θ under $\hat{l}_i$; Θ_i denotes the set of states that minimize the relative entropy. The relative entropy measures the distance between individual i's perceived and the true signal distributions. The set Θ_i represents the states that best fit the true distribution under individual i's perceived model. We further define that individual i's model perception $\hat{l}_i$ is ***innocuous*** if $\Theta_i = \{\theta^*\}$, that is, the true state is the unique minimizer of the relative entropy. I define group irrationality as follows.

Definition 1. We say that: (i) ***group irrationality*** emerges if $\mu_t(\cup_i \Theta_i)$ doesn't converge to $\mathbf{1}$ as $t \to +\infty$, and (ii) ***strong group irrationality*** emerges if all individuals' model perceptions are innocuous, and $\mu_t(\theta^*)$ does not converge to $\mathbf{1}$ as $t \to +\infty$.[7]

To understand Definition 1, it is helpful to recall a well-known result from Berk (1966) that says in misspecified Bayesian learning, each individual i's posteriors will settle on the best-fitting states, Θ_i.[8] Group irrationality occurs when individuals' asymptotic posteriors assign positive probability to states that are not best-fitting under *any* individual's perceived signal structure, so society ends up forming a posterior that contradicts everyone's signal interpretation. If group irrationality occurs and if all individuals have innocuous model perceptions, then we say that strong group irrationality occurs. In this case, society runs into a highly suboptimal outcome that incorrect social learning emerges although all individuals can independently identify the true state.

Example 4. (Group irrationality) In Example 1, the distributions in states β and α perfectly match the true distribution according to 1 and 2's perceived signal structures respectively, so $\Theta_1 = \{\beta\}$ and $\Theta_2 = \{\alpha\}$. However, in social learning, both individuals' posteriors assign probability 1 to γ in the limit, thus we have group irrationality.

Example 5. (Strong group irrationality) In Example 2, both individuals correctly specify the direction of each signal, so their model perceptions are innocuous with $\Theta_1 = \Theta_2 = \{\theta^*\}$. However, society doesn't learn the true state, so strong group irrationality occurs.

4 Discussion of Assumptions

Social learning rule. This paper employs a reduced-form approach in which social learning is described by a belief aggregator F. This approach abstracts away from details of the learning process and enables us to characterize beliefs in a simple manner. One implicit

[7] Here, $\mu_t = (\mu_{1,t}, ..., \mu_{n,t})$ is belief profile in period t. It converges to $\mathbf{1}$ means each $\mu_{i,t}$ converges to 1.

[8] More formally, suppose individual i were to learn independently by Bayes' rule, then we have $\mu_{i,t}(\Theta_i) \to 1$ as $t \to +\infty$ almost surely.

assumption is that society will reach a consensus after every round of social learning. The emergence of a consensus is a prevalent feature in various social learning models, e.g., Banerjee (1992), Bikhchandani et al. (1992), Gale and Kariv (2003), Rosenberg et al. (2009). We can interpret the social belief as the steady state of some underlying belief-exchange process, so this learning environment approximates situations where there are adequate communications, and individuals hold similar beliefs in the end. Besides, one focus of the paper is to examine the divergence between social and individual learning, so imposing a social belief is a convenient assumption.[9] The setup can also be microfounded by a more structural approach where individuals are connected through a social network and interact with their neighbors in a repeated manner (see an earlier version in Chen (2022)). For example, when individuals adopt the DeGroot's rule, beliefs will converge to a

$$F(\mu_1, ..., \mu_n)(\theta) = \sum_i w_i \times \mu_i(\theta), \tag{5}$$

where the weighting vector w denotes the eigenvector centrality of the network (suppose the network is aperiodic and irreducible). It is worth noting that the paper's main interest is to examine how the consensus evolves overtime instead of how it emerges in communications, which is another feature different from the literature on belief exchange such as Golub and Jackson (2010) and Cerreia-Vioglio et al. (2024). One can also think of other microfoundations, e.g., individuals play coordination games repeatedly, and the social belief corresponds to the equilibrium in each stage game, and this paper investigates the evolution of the equilibrium.

Private learning rule. Different from social learning, the paper assumes that individuals apply Bayes' rule in private learning. One justification is that in social learning, it is difficult for individuals to apply Bayes' rule, which involves complicated inferences regarding other individuals' signals. In contrast, in private learning, individuals are facing a much easier learning problem, making Bayes' rule more plausible. The disparity in updating rules has appeared in quasi-Bayesian social learning papers, e.g., Molavi et al. (2018). Differently, this paper also allows individuals to misspecify their signal-generating processes, which captures the presence of misinformation. Thus, this paper's setup features two types of learning mistakes:

- *Naivety in social learning*: social belief is aggregated through some heuristic rule,

[9] A natural extension is to allow individuals to have heterogeneous beliefs after each round of social learning. I conjecture that the paper's results can generalize to this case. In particular, all results are expected to hold as long as the society achieves consensus asymptotically. The situations not captured by this paper's approach are those where consensus fails to emerge in the limit. However, in such cases, society cannot achieve correct learning, which is still consistent with the paper's implications.

- *Misspecification in private learning*: individuals can incorrectly interpret their private information.

Notably, these two sources play different roles in determining belief dynamics, and neither can substitute for the other. One can find examples where correct learning is obtained when only one source is present but not both. For example, without misspecification, some heuristic learning rules can guarantee correct learning, e.g., the DeGroot's rule. However, as will be shown later, with model misspecification—even if the misspecification appears in an innocuous manner—correct learning will not be achieved for almost all heurstic rules considered in the paper. This suggests many results supporting the wisdom of the crowds hinge on the assumption that all individuals correctly specify their signal structures.

Group irrationality. In general, group irrationality refers to the phenomenon where the group's behavior differs—often in a negative way—from how its members would have independently behaved. This paper focuses on the belief aspect, where group irrationality means that the social learning outcome *contradicts* the individual learning outcomes. Motivated by the fact that posteriors will settle on entropy-minimizing states in individual learning, this paper defines group irrationality as the event where posteriors in social learning attach positive weights to states that do not minimize relative entropy for any individual. Consequently, social learning will produce an asymptotic belief absent in any individual learning.

It is also worth emphasizing the difference between group irrationality and the wisdom of crowds. In models featuring the wisdom of crowds, society may also achieve an outcome different from what individuals could have independently achieved. For instance, each individual might face an identification problem and be unable to learn the truth independently; however, society as a whole does not face identification problems, and the truth can be learned through social learning (Jadbabaie et al., 2012). The wisdom of crowds emphasizes the positive effect of social learning in aggregating dispersed information, whereas group irrationality emphasizes the inconsistency between social belief and private information.

Remark 1. The entropy-based definition relies on the assumption that individuals apply Bayes' rule in the private learning stage, but the idea can be extended to non-Bayesian updating rules. Instead of employing the entropy-minimizing states, we can analogously use the support of asymptotic posteriors and define group irrationality in a similar way. This extension is discussed in Section 8.

5 A Benchmark Characterization

In this section, I first present a benchmark characterization of asymptotic beliefs and then discuss its implications on group irrationality. Subsequently, a more comprehensive characterization will be presented in the next section. I first define the ***weighted relative entropy*** of state θ as

$$\mathcal{I}^w(\theta) \equiv \sum_{i=1}^{n} w_i \times \mathcal{R}_i(\theta), \text{ where } w = (w_1, ..., w_n) \in \Delta(N).$$

The weighted relative entropy has a simple interpretation: It describes society's average distance, in the relative entropy sense, between each individual's true signal distribution and perceived signal distribution in state θ. If θ minimizes the weighted relative entropy, then it means θ achieves the minimum average distance between individuals' perceived signal distributions and true signal distributions. We have the following proposition:

Proposition 1. (Benchmark characterization) *Suppose there is a unique state θ_0 that minimizes $\mathcal{I}^w(\theta)$. As $t \to +\infty$, we have:*

(i) if $p > 0$, then whenever μ_t converges, $\mu_{i,t} \to \delta_{\theta_0}$ for all i except for $\mathbb{P}$-null events;

(ii) if $p < 0$, then it happens with $\mathbb{P}$-positive probability that $\mu_{i,t} \to \delta_{\theta_0}$ for all i,

where p denotes the degree and w denotes the weights of individuals in the aggregation rule.

Proposition 1 shows that posteriors have a tendency to accumulate on the state that minimizes the weighted relative entropy. Specifically, if the belief aggregator has a positive degree ($p > 0$), beliefs can only converge to the point-mass belief on the minimizer of the weighted entropy; if it has a negative degree ($p < 0$), beliefs will converge to the point-mass belief with a positive probability. The reason why positive and negative degrees have different statements will be explained later. Proposition 1 provides a channel through which group irrationality can occur—states that minimize society's weighted relative entropy may not minimize the relative entropy for any individual; consequently, society may settle on a state that would be been assigned probability 0 in independent learning. Using this characterization, we can explain the occurrence of group irrationality in Example 1.

Example 6. In Example 1, it can be verified that

$$\mathcal{I}^w(\gamma) = \frac{1}{2}\log\frac{3}{4} + \frac{1}{2}\log\frac{3}{2} \approx 0.03, \quad \mathcal{I}^w(\alpha) = \mathcal{I}^w(\beta) = \frac{1}{4}\log\frac{5}{9} + \frac{1}{4}\log 5 \approx 0.11.$$

Thus, state γ is the unique state that minimizes the weighted relative entropy. Here, although both α (or β) can perfectly match the true signal distribution, $(1/2, 1/2)$, under individual

2 (or 1)'s perceived signal structure, it induces an extreme distribution, $(9/10, 1/10)$, under the other individual's perception. That is, the average distance from the true distribution in state α or β is very large. Proposition 1 hence implies that asymptotic beliefs can only settle on state γ, the best-fitting state on average, which is consistent with group irrationality.

5.1 Intuition: the log-linear case

I then explain the intuition behind Proposition 1. To grasp the key idea, I focus on discussing the *log-linear rule*, where

$$\log \frac{F(\mu_1, ..., \mu_n)(\theta)}{F(\mu_1, ..., \mu_n)(\theta')} = \sum_i w_i \times \log \frac{\mu_i(\theta)}{\mu_i(\theta')}, \tag{6}$$

which represents the limit case of (3) when $p \to 0$.[10] Under this rule, the log likelihood ratio of the social belief is equal to the simple average of the log likelihood ratio of each individual's private belief. Combining (6) with (1) and (4), the social belief $v = \{v_t\}$ must satisfy the following recursive form

$$\log \frac{v_{t+1}(\theta)}{v_{t+1}(\theta')} = \log \frac{v_t(\theta)}{v_t(\theta')} + \sum_i w_i \times \log \frac{\hat{l}_i(s_{i,t}|\theta)}{\hat{l}_i(s_{i,t}|\theta')}. \tag{7}$$

Taking the time average on both sides of (7), we obtain

$$\frac{1}{t} \log \frac{v_{t+1}(\theta)}{v_{t+1}(\theta')} = \frac{1}{t} \log \frac{v_1(\theta)}{v_1(\theta')} + \sum_i w_i \times \frac{1}{t} \sum_{\tau=1}^{t} \log \frac{\hat{l}_i(s_{i,\tau}|\theta)}{\hat{l}_i(s_{i,\tau}|\theta')}.$$

The strong law of large numbers implies that as $t \to +\infty$, we $\mathbb{P}$-almost surely have

$$\frac{1}{t} \log \frac{v_{t+1}(\theta)}{v_{t+1}(\theta')} \to \sum_i w_i \times \mathbb{E} \log \frac{\hat{l}_i(s_{i,t}|\theta)}{\hat{l}_i(s_{i,t}|\theta')} = \mathcal{I}^w(\theta') - \mathcal{I}^w(\theta).$$

Therefore, if $\mathcal{I}^w(\theta') < \mathcal{I}^w(\theta)$, we must have $v_t(\theta)$ converges to 0. Suppose that θ_0 uniquely minimizes $\mathcal{I}^w(\theta)$, then v_t must assign all weights on θ_0, which implies that $\mu_{i,t}$ will settle on state θ_0 in the limit.

[10] The rule can be microfounded by the log-linear rule in Molavi et al. (2018), where w stands for the eigenvector centrality of the network.

5.2 Situations not captured by Proposition 1

It is worth noting that the characterization in Proposition 1 is very incomplete. Although it helps to explain the emergence of group irrationality, it cannot explain why strong group irrationality can emerge (e.g., in Example 2). This is because if everyone's model perception is innocuous, then the true state must minimize society's weighted relative entropy. Therefore, correct learning aligns with Proposition 1. In fact, the weighted relative entropy is often inadequate to characterize asymptotic beliefs. Below is an example in which learning dynamics are different despite the minimizers of the weighted relative entropy being the same.

Example 7. Consider the same setup as in Example 1. Suppose 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 $x \in (0,1)$. It can be verified that when x is sufficiently small ($x \to 0$) or sufficiently large ($x \to 1$), we have $\mathcal{I}^w(\gamma) < \mathcal{I}^w(\alpha) = \mathcal{I}^w(\beta)$. In both cases, γ uniquely minimizes the weighted relative entropy. However, when x is sufficiently large, v_t converges to δ_γ almost surely; when x is small, v_t almost surely doesn't converge.

6 Characterization of Asymptotic Beliefs

In this section, I present a more comprehensive characterization of asymptotic beliefs that allows us to analyze situations not covered by Proposition 1. I first introduce a new notion of divergence that generalizes the idea of the weighted relative entropy.

Definition 2. For all $\theta, \theta' \in \Theta$, $w \in \Delta_{++}(N)$ and $p \neq 0$, we define $\theta \succeq_p^w \theta'$ whenever

$$\mathcal{D}_p^w(\theta', \theta) = \frac{1}{p} \times \mathbb{E} \log \left(\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)} \right)^p \right) \leq 0, \tag{8}$$

and call that the ***weighted*** p***-entropy*** of θ is lower than that of θ'.

Definition 2 generalizes the idea of the weighted relative entropy: If state θ has a lower weighted p-entropy than that of θ', then the weighted average of society's perceived likelihood ratios in states θ' and θ decreases on expectation. The only difference between these two concepts is how to average perceived likelihoods: the weighted relative entropy employs the average log-likelihood ratio, whereas the weighted p-entropy employs the log of the p-average of the likelihood ratio. I denote by Θ_p^w the states that dominate all other states under $\succeq_p^w$ and refer to it as the set of ***minimizers of the weighted p-entropy***. I also define a strict relation $\succ_p^w$ by imposing a strict inequality in (8) and let $\hat{\Theta}_p^w$ denote the states that strictly dominate others. It is worth noting that $\succeq_p^w$ may not be an order, thus it is possible to have no dominating state or multiple strict-dominating states. Below is the main characterization theorem.

Theorem 1. (Characterization of Asymptotic Beliefs) *Let p and w denote the degree and weights of the belief aggregator respectively. As $t \to +\infty$, we have:*

(i) Whenever μ_t converges, $\mu_{i,t} \to \delta_\theta$ for all i for some $\theta \in \Theta_p^w$ except for $\mathbb{P}$-null events; conversely, if $\theta \in \hat{\Theta}_p^w$, then $\mu_{i,t} \to \delta_\theta$ for all i with a $\mathbb{P}$-strictly positive probability.

(ii) If μ_t oscillates $\mathbb{P}$-almost surely, then $p \geq 0$ and $\hat{\Theta}_p^w = \emptyset$; conversely, if $p > 0$ and $\Theta_p^w = \emptyset$, then μ_t oscillates $\mathbb{P}$-almost surely.

(iii) If μ_t can converge to multiple limits, then $p \leq 0$ and $|\Theta_p^w| > 1$; conversely, if $p < 0$ and $|\hat{\Theta}_p^w| > 1$, then μ_t can converge to multiple limits.[11]

Theorem 1 (i) shows that beliefs will—and in some sense will only—concentrate on the minimizers of the weighted p-entropy. First, whenever beliefs converge, they must converge to the point-mass belief on one of the minimizers; second, if there is a strict minimizer, then beliefs must converge to the point-mass belief on this state with a positive probability. Theorem 1 (ii) and (iii) further show that dynamics display different patterns when p is positive and negative. Specifically, beliefs will oscillate infinitely often only when p is negative and will settle on multiple limits only when p is positive. Below I first present some examples and then discuss the mechanism behind Theorem 1.

Example 8. (Emergence of group irrationality) Theorem 1 can explain the emergence of strong group irrationality in Example 2. In that example, beliefs are aggregated in a linear way so $p = 1$. Suppose the true state is G, then we have

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

[11] Formally, convergence to multiple limits means that there are at least two limit belief profiles μ_∞ and μ'_∞ such that both $\mu_t \to \mu_\infty$ and $\mu_t \to \mu'_\infty$ happen with positive probability.

Fixing p_2, when $p_1 \to 1$, we have $\mathcal{D}_1^w(B,G) \to +\infty$, so $G \not\succeq_1^w B$. Similarly, we can show that $B \not\succeq_1^w G$. Therefore, if one individual has a sufficiently large p_i, no state can dominate the other state, i.e., $\Theta_1^w = \emptyset$. Theorem 1 (ii) implies that beliefs almost surely don't converge.

Example 9. (The weighted relative entropy and p-entropy) In Example 7, γ is the unique state that minimizes the weighted relative entropy both when x is sufficiently small and sufficiently large, but beliefs only converge in the second case. This can also be explained using Theorem 1. If we calculate the weighted p-entropy, we obtain

$$\begin{aligned}\mathcal{D}_1^w(\alpha,\gamma) =& \mathcal{D}_1^w(\beta,\gamma)\\ =& \frac{1}{4} \times \log\left[\left(\frac{3}{4}x+\frac{3}{8}\right)\left(\frac{3}{4}x+\frac{3}{4}\right)\left(\frac{15}{8}-\frac{3}{2}x\right)\left(\frac{9}{4}-\frac{3}{2}x\right)\right].\end{aligned}$$

As $x \to 0$, we have

$$\mathcal{D}_1^w(\alpha,\gamma), \mathcal{D}_1^w(\beta,\gamma) \to 0.02 > 0,$$

so we have $\gamma \not\succeq_1^w \alpha$ and $\gamma \not\succeq_1^w \beta$. Moreover, we can also show that α and β cannot dominate γ, thus we have $\Theta_1^w = \emptyset$. Therefore, beliefs oscillate almost surely by Theorem 1 (ii). However, as $x \to 1$, we have

$$\mathcal{D}_1^w(\alpha,\gamma), \mathcal{D}_1^w(\beta,\gamma) \to -0.08 < 0.$$

Thus, we have $\gamma \succ_1^w \alpha$ and $\gamma \succ_1^w \beta$, which implies that $\hat{\Theta}_1^w = \{\gamma\}$, so we have v_t converges to δ_γ with a strictly positive probability by Theorem 1 (i). In the Appendix A.5, I further show that the convergence happens with probability 1. In this example, $x \to 0$ and $x \to 1$ have identical effects on the weighed relative entropy, but their effects on the weighted p-entropy are asymmetric and hence induce different belief dynamics.

6.1 Proof sketch

Subsequently, I discuss the mechanics behind Theorem 1. The discussion focuses on two aspects: First, why asymptotic dynamics can be characterized by the weighted p-entropy; second, why belief dynamics exhibit different patterns for positive and negative ps.

The role of the weighted p-entropy

First, I demonstrate that dynamics of social belief $\{v_t\}$ can be locally characterized using the the weighted p-entropy. Suppose that society is very confident in some state θ, i.e.,

$v_t(\theta) \approx 1$, then for any other state θ' and any i, we have:

$$\mu_{i,t+1}(\theta') = \frac{v_t(\theta')\, \hat{l}_i(s_{i,t}|\theta')}{\sum_{\hat{\theta}} v_t(\theta')\, \hat{l}_i(s_{i,t}|\theta')} \approx v_t(\theta') \times \frac{\hat{l}_i(s_{i,t}|\theta')}{\hat{l}_i(s_{i,t}|\theta)}. \tag{9}$$

Furthermore, Assumption 1 says that the aggregated social belief can be locally approximated by the p-average of all individuals' beliefs. Combining this assumption with the approximation in (9), we can approximate the dynamics of $\{v_t(\theta')\}$ recursively and obtain

$$v_{t+1}(\theta') = F(\mu_{1,t+1}, ..., \mu_{n,t+1})(\theta') \approx \left[\sum w_i \times \left(\frac{\hat{l}_i(s_{i,t}|\theta')}{\hat{l}_i(s_{i,t}|\theta)}\right)^p\right]^{1/p} \times v_t(\theta').$$

After taking log expectation to this approximation, we obtain

$$\mathbb{E}[\log v_{t+1}(\theta')\,|\mathcal{F}_t] \approx \log v_t(\theta') + \mathcal{D}_p^w(\theta', \theta), \tag{10}$$

where $\mathcal{F}_t \equiv \sigma(s_1, ..., s_t)$ denotes the information up to time t. Therefore, if state θ' has a lower (higher) weighted p-entropy than state θ, then $\{\log v_t(\theta')\}$ is a submartingale (supermartingale) and tends to increase (decrease) locally. We can further show that if there exists some state θ that has a lower weighted p-entropy than any other state, and if the society is confident in θ, beliefs will eventually converge to the point-mass belief on state θ. Assumption 2 ensures that society can be confident in any state, which imply that beliefs converge to the point-mass on any minimizer of the weighted p-entropy with positive probability. Similarly, we can also show that if a state does not minimize the weighted p-entropy, then beliefs will not settle on that state. Thus, we can prove Theorem 1(i).

The differences in dynamics between $p > 0$ and $p < 0$

Another feature of Theorem 1 is that dynamics exhibit different patterns for positive and negative p. In particular, beliefs oscillate when $p > 0$ and can converge to multiple limits when $p < 0$. To see this, we notice that when $p \to 0$, we have

$$\mathcal{D}_p^w(\theta', \theta) \to \mathcal{D}_0^w(\theta', \theta) \equiv \sum_{i=1}^n w_i \times \mathbb{E}\log\left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)}\right) = \mathcal{I}^w(\theta) - \mathcal{I}^w(\theta'),$$

hence the weighted p-entropy approaches the weighted relative entropy, and we use $\succeq_0^w$ to denote the dominance under the weighted relative entropy. Note that $\succeq_0^w$ constitutes an order on Θ. However, in general, $\succeq_p^w$ may not be an order. Moreover, we have the following result:

$$\forall p \geq q: \quad \theta \succeq_p^w \theta' \Longrightarrow \theta \succeq_q^w \theta'.$$

Alternatively, the binary relation $\succeq_p^w$ is stronger with a larger p in the sense that it is more difficult for a state to dominate another as p becomes larger. Recall that $\succeq_0^w$ is an order, therefore $\succeq_p^w$ can be **incomplete** when $p > 0$, in which case the minimizer of the weighted p-entropy may not exist; similarly, $\succeq_p^w$ can be **intransitive** when $p < 0$, in which case there can be multiple strict minimizers of the weighted p-entropy (see Corollary 3 in the Appendix). From the previous discussion, the first case corresponds to non-convergence, whereas the second case corresponds to the multiplicity of limit points.

6.2 Discussion: relation to other divergences

The main characterization relies on the weighted p-entropy. A similar notion of divergence is the p**-divergence** in Bhattacharya et al. (2019) and Frick et al. (2023), where

$$\mathcal{I}_p(\theta', \theta) = \mathbb{E}\left(\frac{\hat{l}_i(s|\theta')}{\hat{l}_i(s|\theta)}\right)^p.$$

This divergence differs from the weighted p-entropy, and they induce distinct orderings over the state space. Furthermore, the role of p also varies. Specifically, Bhattacharya et al. (2019) employs the p-divergence to analyze the concentration rate of power Bayes' rule, where p is equal to the Bayesian power; Frick et al. (2023) employs the divergence to construct a bounded local martingale, and their p is not related to the shape of the learning rule. Moreover, both papers require $p > 0$ whereas this paper allows both positive and negative p. Another related divergence is the **Rényi entropy** in which

$$\mathcal{I}_{\text{Rényi}}(\theta) = \frac{1}{p}\log \mathbb{E}\left(\frac{l_i(s|\theta^*)}{\hat{l}_i(s|\theta)}\right)^p,$$

where relative entropy corresponds to the limit where $p \to 0$. Different from the weighted p-entropy and the p-divergence, the Rényi entropy constitutes an ordering over the state space.

7 Group Irrationality and Learning Robustness

Based on previous characterization, this section examines whether group irrationality is a universal phenomenon in this paper's learning framework. Furthermore, this section discusses whether society can robustly learn the true state, i.e., correct learning does not rely crucially on individuals' model perceptions or society's aggregation rule.

7.1 The prevalence of group irrationality

To facilitate discussion, we need to modify Assumption 2 because it implicitly restricts both belief aggregators and model perceptions, making it inconvenient to fix one and change the other. I introduce the following assumptions.

Assumption 3. (Intermediate Likelihood Ratio) *F satisfies*

$$\frac{F(\mu_1, ..., \mu_n)(\theta)}{F(\mu_1, ..., \mu_n)(\theta')} \in \left[\min_i \frac{\mu_i(\theta)}{\mu_i(\theta')}, \max_i \frac{\mu_i(\theta)}{\mu_i(\theta')}\right]$$

for all $\theta, \theta' \in \Theta$ and $\mu_1, ..., \mu_n \in \Delta_{++}(\Theta)$.

Assumption 4. (State-specific good news) For all i and all $\theta \in \Theta$, there exists some s_θ^i such that $\hat{l}(s_\theta^i|\theta) > \hat{l}(s_\theta^i|\theta')$ for all $\theta' \neq \theta$.

Assumption 3 says that the likelihood ratio of social belief cannot exceed the maximum and minimum likelihood ratios of individuals' beliefs. An example is (3) in which the social belief is equal to the p-th power average of individuals' beliefs. Assumption 4 says the perceived signal structure satisfies that for each state, there exists some "good-news" signal that is most likely to occur in that state. It is easy to verify that Assumptions 3 and 4 imply Assumption 2, thus all previous results hold. With some abuse of language, I say that a belief aggregator F is ***regular*** if it satisfies Assumptions 1 and 3, and a model perception is regular if it satisfies Assumption 4. Throughout the whole section, I focus on regular belief aggregators and model perceptions.[12]

Definition 3. F is ***susceptible to (strong) group irrationality*** if there are regular model perceptions $\hat{l}_1, ..., \hat{l}_n$ under which (strong) group irrationality emerges with $\mathbb{P}$-positive probability.

Susceptibility to group irrationality does not mean that correct learning can never be obtained. It only indicates that group irrationality can occur in the presence of particular forms of misinformation. In other words, we can find model perceptions under which states

[12] Assumptions 3 and 4 are not necessary for the main result. They are merely convenient assumptions that can ensure Assumption 2 and encompass many interesting cases. All results extend to other situations where Assumption 2 holds. When Assumption 2 does not hold, the suboptimality of social learning trivially occurs, as mentioned earlier. Therefore, the qualitative results based on these two assumptions are, in some sense, without loss of generality.

that are not best-fitting for any individual will be assigned a positive probability in the limit. Furthermore, if F is susceptible to strong group irrationality, then there exist model perceptions under which every individual can independently identify the true state, but social learning can still lead to incorrect learning. We have the following result.

Proposition 2. (Prevalence of group irrationality) *Every regular F is susceptible to group irrationality; every regular F with non-zero degree is susceptible to strong group irrationality.*

Proposition 2 shows that group irrationality is prevalent for all regular aggregation rules. This contrasts with the "wisdom of crowds" results in the non-Bayesian learning literature, which demonstrate that correct learning can be achieved for some common aggregation rules, such as the power-average rule (3) with $p \in [-1, 1]$ (Molavi et al., 2018). However, the "wisdom of crowds" results largely depend on the assumption that all individuals have correct model perceptions. In contrast, Proposition 2 indicates that if some individuals have incorrect model perceptions, even if their perceptions are innocuous, correct learning no longer holds under various aggregation rules. An example of this is the linear learning rule as illustrated in Example 2. Below is another example.

Example 10. (Harmonic-mean rule) Consider the same setup as in Example 2. The only difference is that individuals adopt the following belief aggregation rule:

$$F(\mu_1, \mu_2)(\theta) \propto \frac{1}{2}\frac{1}{\mu_{1,t-1}(\theta)} + \frac{1}{2}\frac{1}{\mu_{2,t-1}(\theta)}.$$

That is, individuals take harmonic mean (with normalization) of each other's belief. It can be verified that if some p_i is sufficiently large, we have $\hat{\Theta}_p^w = \{G, B\}$, so beliefs can settle on the incorrect state with a strictly positive probability.

Proposition 2 also suggests that aggregation rules with degree 0 may prevent strong group irrationality. One example is the log-linear rule (6). We say that an aggregation rule is ***immune to strong group irrationality*** if, for all innocuous model perceptions, strong group irrationality occurs with zero probability. We have the following result.

Corollary 1. (Log-linear rule and strong group irrationality) *The log-linear rule is immune to strong group irrationality. Conversely, if a regular aggregation rule is immune to strong group irrationality, then it must approximate the log-linear rule near extreme beliefs.*[13]

The second part comes directly from the second half of Proposition 2. To see the first part, notice that under the log-linear rule, all individuals' posteriors will converge to the point-mass belief on the state that minimizes the society's weighted relative entropy, assuming that

[13]Formally, $F(\mu_1, ..., \mu_n) \sim \exp\left(\sum w_i \times \log \mu_i(\theta)\right)$ as $\max_i \mu_i(\theta) \to 0$.

such state is unique (see the discussion in Section 5). With innocuous model perceptions, the true state θ^* uniquely minimizes the relative entropy for all i, so the minimizer of the weighted relative entropy is the true state as well. Consequently, correct learning almost surely occurs with innocuous model perceptions.

Remark 2. Corollary 1 provides another rationale for the log-linear rule proposed by Molavi et al. (2018). It implies that if we allow for model misspecifications, then the log-linear rule is, in some sense, the *only* aggregation rule can prevent strong group irrationality. For any other regular aggregation rule, we can always find situations in which the society fails to learn the truth, even if each individual is able to learn it independently.

7.2 Which model perceptions can ensure correct learning?

In the last subsection, we discussed whether it is possible for group irrationality to occur under *some* model perceptions; however, we may still find a large class of model perceptions under which correct learning can be achieved. This section characterizes model perceptions that can ensure correct learning against a variety of aggregation rules. This discussion can help us to gain a better understanding of the robustness of correct learning.

Definition 4. Model perceptions $\hat{l} = \left(\hat{l}_1, ..., \hat{l}_n\right)$ are ***robustly innocuous*** if with those perceptions, society can almost surely learn the true state for all regular belief aggregators.

With robustly innocuous model perceptions, correct learning is robust in the sense that it does not depend significantly on specific aggregation rules. Robust innocuousness is a stronger concept than innocuousness: $\hat{l}$ being robustly innocuous requires everyone's model perception to be innocuous (except for some tie situations). To fully characterize robustly innocuous model perceptions, I introduce the following concept:

Definition 5. For all $\theta, \theta' \in \Theta$, we define

$$\theta \succeq_\infty \theta' \Longleftrightarrow \mathcal{D}_\infty\left(\theta', \theta\right) = \mathbb{E}\left[\max_{i \in N} \log\left(\frac{\hat{l}_i\left(s_i|\theta'\right)}{\hat{l}_i\left(s_i|\theta\right)}\right)\right] \leq 0, \tag{11}$$

and say that θ ***robustly dominates*** θ'. We further define Θ_∞ (and $\hat{\Theta}_\infty$) as the set of states that (strictly) dominate other states under $\succeq_\infty$.

The robust domination represents the limit case of the domination induced by the weighted p-entropy as $p \to +\infty$. In words, if θ robustly dominates θ', society's maximum perceived likelihood ratio between these θ' and θ decreases on expectations. Using Definition 5, we can characterize robustly innocuous perceptions.

Proposition 3. *$\hat{l}$ is robustly innocuous if it satisfies $\theta^* \in \hat{\Theta}_\infty$ and only if it satisfies $\theta^* \in \Theta_\infty$.*

Proposition 3 provides an almost necessary and sufficient condition for model perceptions to be robustly innocuous. Note that the robust domination (11) is defined using the maximum over all individuals. A natural implication is that with independent signals, it becomes more challenging for the true state to robustly dominate other states as the size of the society increases. This leads to an interesting result—it is harder to achieve correct learning in a larger society. Below is an example.

Example 11. (Harder to learn in a larger society) Consider the binary case in Example 2, where signals are i.i.d. with true probability p^*, where $p^* > 1/2$. Suppose there are n individuals, each with an identical model perception $\hat{p} > 1/2$ (thus, I use $\hat{p}$ to represent the model perception profile $\hat{l}$). Assume the true state is G. From Proposition 3, $\hat{p}$ is robustly innocuous if

$$\mathcal{D}_\infty (B, G) = (1 - 2p^{*n}) \times \log \left(\frac{\hat{p}}{1 - \hat{p}} \right) < 0. \tag{12}$$

We have the following observations. First, no model perception is robustly innocuous in a large society. When the size of the society $n > \log_{p^*}^{1/2}$, (12) fails to hold for all feasible $\hat{p}$. Furthermore, if we allow p^* to change, robust innocuousness requires individuals to have more precise private signals in a larger society. To see this, from (12), robust innocuousness requires $p^* > \sqrt[n]{\frac{1}{2}} \equiv p(n)$. When $n = 1$, $p(n) = 1/2$, meaning signals only need to be informative. As $n \to +\infty$, $p(n) \to 1$, meaning signals need to be nearly perfectly informative.

Example 11 does not imply that it is harder to achieve correct learning under a specific aggregation rule as the society grows larger. Instead, it suggests that it easier to find *some* aggregation rules under which correct learning collapses as the society expands. Moreover, all model perceptions are not robustly innocuous when the society is sufficiently large. Below is the statement for more general cases.

Corollary 2. (No robust learning in large society) *Suppose that individuals have i.i.d. signal distribution l and homogenous model perception $\hat{l}$. Then, for all feasible l and $\hat{l}$, there exists some $n_0 < +\infty$ such that $\hat{l}$ is not robustly innocuous when the size of the society $n \geq n_0$.*[14]

Corollary 2 further contrasts the idea of the wisdom of crowds by showing that a large society is vulnerable to incorrect learning. This vulnerability arises because errors can be amplified in a large society when individuals learn in a non-Bayesian manner.

[14]The assumption of homogeneous model perception can be relaxed. For example, we can allow individuals' model perceptions to come from any finite set.

8 Extension: Other Updating Rules in Private Learning

In the model setup, I assume that individuals apply Bayes' rule when updating beliefs during private learning. It is reasonable to consider situations where individuals deviate from the Bayesian paradigm. This section discusses how the paper's characterization can be extended to more general situations.

8.1 Learning with generalized Bayes' rule

Let's consider an alternative setup where, in the social learning stage, individuals apply the same updating rule; in the private learning stage, individuals follow the ***generalized Bayes' rule***. That is, their posteriors are given by

$$\forall \theta \in \Theta : \quad \mu_{i,t}(\theta) = GBU_i(v_t, s_{i,t})(\theta) = \frac{v(\theta) \times \psi_i(s|\theta)}{\sum_{\theta' \in \Theta} v(\theta') \times \psi_i(s|\theta')},$$

where $\psi_i : S_i \times \Theta \to \mathbb{R}_{++}$ is called the ***pseudo-likelihood function***. When ψ_i is a signal distribution, it corresponds to Bayes' rule as in the benchmark model. The generalized Bayes' rule also allows for other interesting cases as below.

Example 12. Suppose that the learning rule satisfies

$$\psi_i(s|\theta) = \hat{l}_i^{\alpha}(s|\theta), \text{ where } \alpha > 0, \tag{13}$$

then it becomes the power Bayes' rule, and the posterior is called *power posterior* (Walker and Hjort, 2001). As α decreases, individuals attach larger weights to their prior beliefs. When $\alpha = 0$, individuals stick to their priors; as $\alpha \to +\infty$, posteriors put almost all weights on the maximum likelihood states.

Example 13. Suppose that the pseudo-likelihood function is

$$\psi_i(s|\theta) = \exp(-\rho_i(s, \theta)), \tag{14}$$

where $\rho_i(s, \theta)$ stands for a loss function. The posterior is referred to as *Gibbs posterior* (Jiang and Tanner, 2008; Bissiri et al., 2016). The motivation is that individuals select a posterior to minimize the expected loss that depends on both state and signal.

8.2 Characterization of asymptotic beliefs

For all previous cases, we can redefine the relative entropy and the weighted p-entropy by substituting the likelihood function $\hat{l}_i$ with the pseudo-likelihood function ψ_i. For example, we can define the ***generalized relative entropy*** as follows:

$$\mathcal{R}(\theta;\psi_i) \equiv \mathbb{E}\log\left(\frac{l_i(s|\theta^*)}{\psi_i(s|\theta)}\right),$$

then define group irrationality by employing the minimizers of the generalized relative entropy. Due to the identical structure, all previous characterization results can be extended analogously. An interesting application of this is the power posterior rule discussed in Example 12. First, I impose the following assumption.

Assumption 5. *(No indifference) For all $\theta \neq \theta'$, we have $\mathcal{I}^w(\theta) \neq \mathcal{I}^w(\theta')$.*

In other words, no two states have identical weighted relative entropy. We have the following proposition.

Proposition 4. (Characterization with power Bayes' rule) *Suppose that Assumptions 1 and 3 to 5 hold. If individuals follow the power Bayes' rule (13), there exists some $\alpha_0 \in \mathbb{R}_+$ such that as $t \to +\infty$, we have*

$$\forall i: \quad \mu_{i,t} \to \delta_{\theta_0}\ \mathbb{P}\text{-}a.s. \quad \text{where } \theta_0 = \arg\min_{\theta\in\Theta}\mathcal{I}^w(\theta)$$

whenever $\alpha \in (0,\alpha_0)$, where α stands for the posterior power.

When α is small, it indicates that individuals are cautious about incorporating their private signals into their posteriors. Proposition 4 states that when all individuals are sufficiently cautious, their posteriors will almost surely converge to the state that minimizes the weighted relative entropy, aligning with the benchmark characterization in Proposition 1. This provides another interpretation of the benchmark characterization—it characterizes asymptotic beliefs with adequately cautious agents. An implication from Proposition 4 is that society can successfully avoid strong group irrationality by being sufficiently cautious with their private signals. In Example 2, if individuals apply the power Bayes' rule with small exponent instead of standard Bayes' rule, beliefs will not oscillate infinitely often, and both individuals will almost surely learn the true state.

9 Related Literature

The paper is related to two threads of literature. The first is the literature on non-Bayesian social learning. The two most similar papers are Jadbabaie et al. (2012) and Molavi et al. (2018), which study a social learning problem where individuals receive infinitely many signals and exchange beliefs with their neighbors. Jadbabaie et al. (2012) show that when individuals incorporate others' beliefs in a simple-average manner, correct learning will be obtained. Molavi et al. (2018) axiomatize a log-linear learning rule and provides conditions under which learning is correct. The main difference is that this paper introduces misinformation by allowing individuals to have an incorrect interpretation of their signals. It turns out that misinformation can produce group irrationality even under rules that can achieve correct learning without misspecification. This paper is also related to a large body of literature on belief exchange in social networks. For example, DeMarzo et al. (2003) and Golub and Jackson (2010) examine a social learning problem where individuals communicate beliefs back and forth in the DeGroot (1974)'s manner; Cerreia-Vioglio et al. (2024) study a general class of belief aggregation rules; Banerjee and Compte (2023) investigate a problem where individuals endogeneously choose the updating rules. Differently, this paper takes a more reduced-form approach and starts with the outcome of communications, instead of studying how beliefs evolve during the communication process. Other papers involve non-Bayesian learning include Li and Tan (2020) studies a learning problem where individuals apply Bayes rule in the local network; Eyster and Rabin (2010), Guarino and Jehiel (2013) and Dasaratha and He (2020) looked at sequential social learning with non-Bayesian agents.

The second is the literature on social learning with model misspecification. The most technically similar paper is Frick et al. (2023), which employs the p-divergence to construct local martingales, similar to the role of the weighted p-entropy in this paper; their differences were discussed in Section 6. Bohren (2016) and Bohren and Hauser (2021) study a sequential social learning model in which individuals have misspecified beliefs about the learning environments. Arieli et al. (2023) study a sequential social learning problem where individuals can either overestimate or underestimate predecessors' signal informativeness. Other related papers include studies on social learning with model uncertainty, such as Hare et al. (2020), Chen (2022), and Huang (2023), which explore social learning problems where individuals consider multiple, possibly incorrect, models.

10 Conclusion

This paper examines a non-Bayesian social learning problem where individuals may be misinformed about their data-generating processes. It extends the standard non-Bayesian learning literature by formalizing the impact of misinformation on social learning. In this work, I introduce a new concept—group irrationality—which captures scenarios where the social learning outcome diverges from what would be achieved if every individual learned independently. Contrary to the "wisdom of crowds" results, this paper demonstrates that group irrationality is common across a wide range of typical aggregation rules. This prevalence arises because social learning involves not just the aggregation of dispersed information but also the spread of misinformation. Given that individuals may face different types of misinformation, their interactions can lead to social beliefs that contradict each individual's private information, resulting in group irrationality. Notably, as the size of society increases, the detrimental effects of misinformation become more pronounced, making it more difficult to guarantee correct learning. In terms of future research, it would be interesting to explore group irrationality in various learning contexts. For instance, the paper mostly focuses on aggregation rules that resemble a generalized average, thus it is natural to examine other class of aggregation rules; this paper looks at a passive learning problem, but we could also look at active learning problems where data are endogenous etc.

Bibliography

Acemoglu, Daron, Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar (2011) "Bayesian learning in social networks," *The Review of Economic Studies*, 78 (4), 1201–1236. 2

Arieli, Itai, Yakov Babichenko, Stephan Müller, Farzad Pourbabaee, and Omer Tamuz (2023) "The Hazards and Benefits of Condescension in Social Learning," *arXiv preprint arXiv:2301.11237*. 26

Banerjee, Abhijit and Olivier Compte (2023) "Consensus and Disagreement: Information Aggregation under (not so) Naive Learning," *arXiv e-prints*, arXiv–2311. 26

Banerjee, Abhijit V (1992) "A simple model of herd behavior," *The Quarterly Journal of Economics*, 107 (3), 797–817. 2, 11

Berk, Robert H (1966) "Limiting behavior of posterior distributions when the model is incorrect," *The Annals of Mathematical Statistics*, 37 (1), 51–58. 4, 10

Bhattacharya, Anirban, Debdeep Pati, and Yun Yang (2019) "Bayesian fractional posteriors," *The Annals of Statistics*, 47 (1), 39 – 66. 19

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992) "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, 100 (5), 992–1026. 2, 11

Bissiri, Pier Giovanni, Chris C Holmes, and Stephen G Walker (2016) "A general framework for updating belief distributions," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 78 (5), 1103–1130. 24

Bohren, J Aislinn (2016) "Informational herding with model misspecification," *Journal of Economic Theory*, 163, 222–247. 26

Bohren, J Aislinn and Daniel N Hauser (2021) "Learning with heterogeneous misspecified models: Characterization and robustness," *Econometrica*, 89 (6), 3025–3077. 26

Cerreia-Vioglio, Simone, Roberto Corrao, and Giacomo Lanzani (2024) "Dynamic opinion aggregation: long-run stability and disagreement," *Review of Economic Studies*, 91 (3), 1406–1447. 11, 26

Chen, Jaden Yang (2022) "Sequential Learning under Informational Ambiguity," *Available at SSRN 3480231*. 26

Chen, Yang (2022) *Essays on Learning under Model Uncertainty* Ph.D. dissertation, Cornell University. 11

Dasaratha, Krishna and Kevin He (2020) "Network structure and naive sequential learning," *Theoretical Economics*, 15 (2), 415–444. 26

DeGroot, Morris H (1974) "Reaching a consensus," *Journal of the American Statistical Association*, 69 (345), 118–121. 8, 26

DeMarzo, Peter M, Dimitri Vayanos, and Jeffrey Zwiebel (2003) "Persuasion bias, social influence, and unidimensional opinions," *The Quarterly journal of Economics*, 118 (3), 909–968. 26

Eyster, Erik and Matthew Rabin (2010) "Naive herding in rich-information settings," *American Economic Journal: Microeconomics*, 2 (4), 221–43. 26

Frick, Mira, Ryota Iijima, and Yuhta Ishii (2023) "Belief convergence under misspecified learning: A martingale approach," *The Review of Economic Studies*, 90 (2), 781–814. 19, 26

Gale, Douglas and Shachar Kariv (2003) "Bayesian learning in social networks," *Games and Economic Behavior*, 45 (2), 329–346. 11

Golub, Benjamin and Matthew O Jackson (2010) "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, 2 (1), 112–49. 2, 11, 26

Guarino, Antonio and Philippe Jehiel (2013) "Social learning with coarse inference," *American Economic Journal: Microeconomics*, 5 (1), 147–174. 26

Hare, James Z, Cesar A Uribe, Lance Kaplan, and Ali Jadbabaie (2020) "Non-Bayesian social learning with uncertain models," *IEEE Transactions on Signal Processing*, 68, 4178–4193. 26

Huang, Wanying (2023) "Learning about informativeness." 26

Jadbabaie, Ali, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi (2012) "Non-Bayesian social learning," *Games and Economic Behavior*, 76 (1), 210–225. 2, 12, 26

Jiang, Wenxin and Martin A. Tanner (2008) "Gibbs posterior for variable selection in high-dimensional classification and data mining," *The Annals of Statistics*, 36 (5), 2207 – 2231. 24

Le Bon, Gustave (1895) *The crowd: A study of the popular minds*. 4

Li, Wei and Xu Tan (2020) "Locally Bayesian learning in networks," *Theoretical Economics*, 15 (1), 239–278. 26

Molavi, Pooya, Alireza Tahbaz-Salehi, and Ali Jadbabaie (2018) "A theory of non-Bayesian social learning," *Econometrica*, 86 (2), 445–490. 11, 14, 21, 22, 26

Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille (2009) "Informational externalities and emergence of consensus," *Games and Economic Behavior*, 66 (2), 979–994. 11

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

Surowiecki, James (2005) *The wisdom of crowds*: Anchor. 2

Walker, Stephen and Nils Lid Hjort (2001) "On bayesian consistency," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 63 (4), 811–821. 24

A Proofs

A.1 Auxiliary Results

I first present a lemma about the properties of the weighted p-entropy.

Lemma 1. (Properties of the weighted p-entropy) *For all $\theta \neq \theta'$, we have:*

(i) $\mathcal{D}_p^w(\theta', \theta) \to \mathcal{D}_0^w(\theta', \theta) \equiv \mathcal{I}^w(\theta) - \mathcal{I}^w(\theta')$ as $p \to 0$,

(ii) $\mathcal{D}_p^w(\theta', \theta) > \mathcal{D}_q^w(\theta', \theta)$ whenever $p > q$.

Proof. (i) For simplicity in notation, I denote by $x_i = \hat{l}_i(s_i|\theta') / \hat{l}_i(s_i|\theta)$. From L' Hopital's rule, we have

$$\begin{aligned}\lim_{p\to 0} \mathcal{D}_p^w(\theta', \theta) &= \lim_{p\to 0} \frac{1}{p} \times \mathbb{E} \log \left(\sum_{i=1}^n w_i \times x_i^p \right) = \lim_{p\to 0} \mathbb{E} \left(\frac{\sum_{i=1}^n w_i \times x_i^p \log x_i}{\sum_{i=1}^n w_i \times x_i^p} \right) \\ &= \sum_{i=1}^n w_i \times \log x_i = \mathcal{D}_0^w(\theta', \theta).\end{aligned}$$

(ii) Suppose that $p > q$, then Jensen's inequality implies that

$$\begin{aligned}\mathcal{D}_p^w(\theta', \theta) &= \frac{1}{p} \mathbb{E} \log \left(\sum_{i=1}^n w_i \times x_i^p \right) = \frac{1}{p} \mathbb{E} \log \left(\sum_{i=1}^n w_i \times x_i^{q\times p/q} \right) \\ &> \frac{1}{p} \mathbb{E} \log \left[\sum_{i=1}^n w_i \times x_i^q \right]^{p/q} = \frac{1}{q} \mathbb{E} \log \left[\sum_{i=1}^n w_i \times x_i^q \right]^{1/q} = \mathcal{D}_q^w(\theta', \theta),\end{aligned}$$

where the inequality holds strictly because there is no identification problem, i.e., $\hat{l}_i(\cdot|\theta) \neq \hat{l}_i(\cdot|\theta')$ for all $\theta \neq \theta'$ and all i. □

Lemma 1 (i) says that the weighted relative entropy approximates the weighted p-entropy when p is near 0; Lemma 1 (ii) says that the weighted p-entropy is monotonic in p. We have the following corollary.

Corollary 3. (Relation to the weighted relative entropy) *For all $\theta \neq \theta'$ and $p > 0 > q$, we have:*

$$\theta \succeq_p^w \theta' \Longrightarrow \theta \succ_0^w \theta' \Longrightarrow \theta \succ_q^w \theta',$$

where $\succ_0^w$ is the strict order induced by the weighted relative entropy.

From Corollary 3, $\succeq_p^w$ may be incomplete when $p > 0$ since it becomes more difficult for a state to dominate another state, and hence some states may not be comparable; similarly,

$\succeq_p^w$ may be intransitive when $p < 0$ since it is easier for the dominance relation to hold, so it is possible that two states strictly dominate each other.

A.2 Proof of Theorem 1 (i)

Proof of the "only if " direction.

To prove this direction, I first prove the following lemmas.

Lemma 2. *Whenever μ_t converges, the social belief v_t and every individual's posterior $\mu_{i,t}$ all converge to the same Dirac belief δ_θ except for $\mathbb{P}$-null events.*

Proof. Suppose that μ_t converges to some limit $\mu_\infty = (\mu_{1,\infty}, ..., \mu_{n,\infty})$. From the continuity of F, we have $v_t = F(\mu_t) \to F(\mu_\infty) \equiv v_\infty$, so v_t also converges. Recall that $\mu_{i,t} = BU_i(v_t, s_{i,t})$, and there exist signals that alter the likelihood ratio between any two states, so we must have $v_\infty = \delta_\theta$ for some θ. Therefore, $\mu_{i,\infty} = v_\infty = \delta_\theta$ for all i. □

Lemma 3. *Whenever v_t converges to some Dirac belief δ_θ, then we have $\theta \in \Theta_p^w$ except for $\mathbb{P}$-null events.*

Proof. Suppose instead that $\theta \notin \Theta_p^w$, i.e., there exists another state $\theta' \neq \theta$ such that $\theta \not\succeq_p^w \theta'$, that is,

$$\frac{1}{p}\mathbb{E}\log\left(\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)}\right)^p\right) > 0. \tag{15}$$

From Assumption 1, we know that

$$v_{t+1}(\theta') = F(BU_1(v_t, s_{1,t}), ..., BU_n(v_t, s_{n,t}))(\theta') \sim \left[\sum w_i \times \left(\frac{v_t(\theta')\hat{l}_i(s_{i,t}|\theta')}{\sum_{\hat{\theta}} v_t(\hat{\theta})\hat{l}_i(s_{i,t}|\hat{\theta})}\right)^p\right]^{1/p}$$

$$\sim v_t(\theta')\left[\sum w_i \times \left(\frac{\hat{l}_i(s_{i,t}|\theta')}{\hat{l}_i(s_{i,t}|\theta)}\right)^p\right]^{1/p} \quad \text{as } v_t(\theta) \to 1. \tag{16}$$

Therefore, for all $\varepsilon > 0$, there exists some $\delta > 0$ such that if $v_t(\theta) \geq 1 - \delta$, then

$$\frac{v_{t+1}^p(\theta')}{v_t^p(\theta') \times \left[\sum w_i \times \left(\frac{\hat{l}_i(s_{i,t}|\theta')}{\hat{l}_i(s_{i,t}|\theta)}\right)^p\right]} \in [\exp(-|p| \times \varepsilon), \exp(|p| \times \varepsilon)], \tag{17}$$

which implies that

$$\log\left[v_{t+1}\left(\theta'\right)\right] \geq \log\left[v_t\left(\theta'\right)\right] + \frac{1}{p}\log\left[\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i\left(s_{i,t}|\theta'\right)}{\hat{l}_i\left(s_{i,t}|\theta\right)}\right)^p\right] - \varepsilon. \tag{18}$$

I then show that v_t can't converge to δ_θ with a strictly positive probability. Denote by $X_t = \log\left[v_t\left(\theta'\right)\right]$. When $v_t \to \delta_\theta$ we must have $X_t \to -\infty$. Therefore, $\mathbb{P}\left(v_t \to \delta_\theta\right) > 0$ only if $\mathbb{P}\left(X_t \to -\infty\right) > 0$, which implies that there exists some $t_0 < \infty$ such that $\mathbb{P}\left(X_t \leq -C \text{ for all } t > t_0\right) > 0$, where C is a positive constant. However, on event $E \equiv \{X_t \leq -C \text{ for all } t > t_0\}$, we almost surely have

$$\begin{aligned}\frac{1}{t-t_0}X_t &\geq \frac{1}{t-t_0}X_{t_0} + \frac{1}{t-t_0} \times \sum_{t'=t_0+1}^{t} \frac{1}{p}\log\left[\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i\left(s_{i,t}|\theta'\right)}{\hat{l}_i\left(s_{i,t}|\theta\right)}\right)^p\right] - \varepsilon \\ &\to \frac{1}{p}\mathbb{E}\log\left[\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i\left(s_{i,t}|\theta'\right)}{\hat{l}_i\left(s_{i,t}|\theta\right)}\right)^p\right] - \varepsilon \quad \text{as } t \to \infty.\end{aligned}$$

Here, we can choose ε to be sufficiently small such that $\frac{1}{p}\mathbb{E}\log\left[\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i(s_{i,t}|\theta')}{\hat{l}_i(s_{i,t}|\theta)}\right)^p\right] > \varepsilon > 0$. In this case, we almost surely have $X_t \to +\infty$, which implies that E can only be a null event. Therefore, when $\theta \notin \Theta_p^w$, we must have $\mathbb{P}\left(v_t \to \delta_\theta\right) = 0$. As a consequence, whenever $v_t \to \delta_\theta$, we must have $\theta \in \Theta_p^w$ except for null events. □

The only if direction of Theorem 1 (i) follows directly from Lemmas 2 and 3.

Proof of the "if" direction.

Lemma 4. *Let X be a bounded random variable. Suppose that $\mathbb{E}\log X < (\text{or} >)\ 0$, then there exists some $\rho > 0$ such that $\mathbb{E}X^\rho < (\text{or} >)\ 1$.*

Proof. Since X is bounded, we can apply the dominated convergence theorem and get

$$\lim_{\rho\to 0+}\mathbb{E}\left(\frac{X^\rho - 1}{\rho}\right) = \mathbb{E}\left(\lim_{\rho\to 0+}\frac{X^\rho - 1}{\rho}\right) = \mathbb{E}\left(\lim_{\rho\to 0+} X^\rho \times \log X\right) = \mathbb{E}\log X < (\text{or} >)\ 0.$$

So, there exists $\rho > 0$ such that $\mathbb{E}\left(\frac{X^\rho - 1}{\rho}\right) < (\text{or} >)\ 0$, or $\mathbb{E}X^\rho < (\text{or} >)\ 1$. □

Now we start proving the "if" direction. Suppose that $\theta \in \hat{\Theta}_p^w$, then we have

$$\frac{1}{p}\mathbb{E}\log\left(\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i\left(s_{i,t}|\theta'\right)}{\hat{l}_i\left(s_{i,t}|\theta\right)}\right)^p\right) < 0 \quad \text{for all } \theta' \neq \theta. \tag{19}$$

We want to show that $\mu_{i,t}$ will converge to δ_θ with positive probability. We first consider the case where p is positive.

Case 1: Suppose that $p > 0$**.** Then Lemma 4 and (19) imply that there exists some $\rho > 0$ such that

$$\mathbb{E}\left(\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)}\right)^p\right)^\rho < 1 \text{ for all } \theta' \neq \theta.$$

Then (17) implies that for all $q > 0$, there exists some $\varepsilon > 0$ such that when $v_t \in B_\varepsilon(\delta_\theta) \equiv \{v : v(\theta) \geq 1 - \varepsilon\}$, we have

$$v_{t+1}^{p\rho}(\theta') \leq v_t^{p\rho}(\theta') \times (1+q)\left[\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i(s_{i,t}|\theta')}{\hat{l}_i(s_{i,t}|\theta)}\right)^p\right]^\rho.$$

Lemma 5. *There exists* $\varepsilon' > 0$ *such that when* $v_1 \in B_{\varepsilon'}(\delta_\theta)$*, we have* $\mathbb{P}(v_t \to \delta_\theta) > 0$*.*

Proof. For some $\varepsilon > 0$, we define a stopping time $T = \inf\{t : v_t \notin B_\varepsilon(\delta_\theta)\}$, which is the first time that v_t escapes from $B_\varepsilon(\delta_\theta)$, and define $Y_t(\theta') = v_{t\wedge T}^{p\rho}(\theta')$. From the fact that $\mathbb{E}\left(\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)}\right)^p\right)^\rho < 1$, there exists some $q > 0$ such that $\mathbb{E}\left(\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)}\right)^p\right)^\rho < \frac{1}{1+q}$. Therefore, when ε is sufficiently small, we have

$$\mathbb{E}\left(Y_{t+1}(\theta')\,|\mathcal{F}_t\right) \leq Y_t(\theta'),$$

so $\{Y_t(\theta')\}$ is a bounded and nonnegative supermartingale and thus converges to some limit random variable $Y_\infty(\theta')$. We also have

$$\mathbb{P}\left(Y_T(\theta') \geq \varepsilon^{p\rho}\right) \leq \frac{\mathbb{E}\left(Y_T(\theta')\right)}{\varepsilon^{p\rho}} \leq \frac{Y_1(\theta')}{\varepsilon^{p\rho}} \leq \left(\frac{\varepsilon'}{\varepsilon}\right)^{p\rho}, \tag{20}$$

where the equality comes from Markov inequality and optional stopping theorem.

$$\mathbb{P}(T < \infty) \leq \mathbb{P}\left(\cup_{\theta'\neq\theta}\{Y_T(\theta') \geq \varepsilon^{p\rho}\}\right) \leq \sum_{\theta'\neq\theta}\mathbb{P}\left(Y_T(\theta') \geq \varepsilon^{p\rho}\right) \leq |\Theta| \times \left(\frac{\varepsilon'}{\varepsilon}\right)^{p\rho} \tag{21}$$

Suppose ε' is sufficiently small relative to ε and ε itself is also sufficiently small, then we have $\mathbb{P}(T < \infty) < 1$, and hence $\mathbb{P}(T = \infty) > 0$, which implies $\mathbb{P}(v_t \in B_\varepsilon(\delta_\theta) \text{ for all } t \geq 0) > 0$. On $\{T = \infty\}$, v_t is trapped in $B_{\varepsilon'}(\delta_\theta)$ forever and $v_t(\theta')$ almost surely converges. From Lemma 2, the only possible limit for v_t is δ_θ, so we must have $\mathbb{P}(v_t \to \delta_\theta) > 0$. □

By Assumption 2, for all $v_1 \in \Delta_{++}(\Theta)$, v_t enters $B_{\varepsilon'}(\delta_\theta)$ with a strictly positive probability, so for all $v_1 \in \Delta_{++}(\Theta)$, we have $\mathbb{P}(v_t \to \delta_\theta) > 0$.

Case 2: Suppose that $p < 0$. The case is similar to the first case. The only change is that the local supermartingale takes a different form. When $p < 0$, we have

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

Lemma 4 implies that there exists some $\rho > 0$ such that

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

From (16), we know that as $v_t(\theta) \to 1$, we have

$$v_{t+1}\left(\theta'\right)^{|p|\rho} \sim v_t\left(\theta'\right)^{|p|\rho} \times \left[\frac{1}{\sum_{i=1}^{n} w_i \times \left(\frac{\hat{l}_i(s_i|\theta)}{\hat{l}_i(s_i|\theta')}\right)^{|p|}}\right]^{\rho}. \tag{22}$$

We now define $Z_t\left(\theta'\right) = v_{t\wedge T}\left(\theta'\right)^{|p|\rho}$ and can show that $\{Z_t\left(\theta'\right)\}$ is a supermartingale. The rest of the proof is almost identical to the proof of Case 1.

A.3 Proof of Theorem 1 (ii)

Proof. (i) The "if" direction: Suppose that beliefs almost surely don't converge, then from Theorem 1 (i), we have $\hat{\Theta}_p^w = \emptyset$. Note that $\Theta_0^w \neq \emptyset$, where Θ_0^w denotes the minimizer of the weighted relative entropy, so we must have $p \geq 0$ from Corollary 3. (ii) The "only if" direction: Suppose $\Theta_p^w = \emptyset$ and $p > 0$, then beliefs can't converge from Theorem 1 (i). □

A.4 Proof of Theorem 1 (iii)

Proof. (i) The "if" direction: Suppose that μ_t can converge to multiple limit points. Theorem 1 (i) implies that $|\Theta_p^w| > 1$. Suppose that $p > 0$. Corollary 3 implies that for all $\theta \in \Theta_p^w$, we have $\theta \succ_0^w \theta'$ for all $\theta' \neq \theta$. Notice that there is at most one $\succ_0^w$-maximizer, so there is at most one $\succeq_p^w$-maximizer for all $p > 0$, that is, $|\Theta_p^w| \leq 1$ for $p > 0$. As a consequence, we can only have $p < 0$. (ii) The "only if" direction: suppose that $|\hat{\Theta}_p^w| > 1$, which is only possible when $p < 0$, then Theorem 1 (i) implies that for every $\theta \in \hat{\Theta}_p^w$, we have $\mu_{i,t} \to \delta_\theta$ for all i with a positive probability. □

A.5 Proof of almost sure convergence in Example 9

Claim. In Example 9, when x is large enough, we have $\mu_{i,t} \to \delta_\gamma$ for all i almost surely.

Proof. Denoting by $\mathcal{F}_t \coloneqq \sigma(s_0, ..., s_{t-1})$, we notice that

$$\begin{aligned}
&\mathbb{E}\left[\log v_{t+1}(\gamma) \,|\mathcal{F}_t\right] \\
&= \log v_t(\gamma) + \mathbb{E}\log\left(\frac{1}{2}\frac{\hat{l}_1(s_{1t}|\gamma)}{\sum_{\hat{\theta}\in\Theta} v_t\left(\hat{\theta}\right)\hat{l}_1\left(s_{1t}|\hat{\theta}\right)} + \frac{1}{2}\frac{\hat{l}_2(s_{2t}|\gamma)}{\sum_{\hat{\theta}\in\Theta} v_t\left(\hat{\theta}\right)\hat{l}_2\left(s_{2t}|\hat{\theta}\right)}\right) \qquad (23)\\
&\geq \log v_t(\gamma) + \frac{1}{2}\mathbb{E}\log\left(\frac{\hat{l}_1(s|\gamma)}{\sum_{\hat{\theta}\in\Theta} v_t\left(\hat{\theta}\right)\hat{l}_1\left(s|\hat{\theta}\right)}\right) + \frac{1}{2}\mathbb{E}\log\left(\frac{\hat{l}_2(s|\gamma)}{\sum_{\hat{\theta}\in\Theta} v_t\left(\hat{\theta}\right)\hat{l}_2\left(s|\hat{\theta}\right)}\right)\\
&\geq \log v_t(\gamma) + \min_{v\in\Delta(\Theta)}\left[\frac{1}{2}\mathbb{E}\log\left(\frac{\hat{l}_1(s|\gamma)}{\sum_{\hat{\theta}\in\Theta} v\left(\hat{\theta}\right)\hat{l}_1\left(s|\hat{\theta}\right)}\right) + \frac{1}{2}\mathbb{E}\log\left(\frac{\hat{l}_2(s|\gamma)}{\sum_{\hat{\theta}\in\Theta} v\left(\hat{\theta}\right)\hat{l}_2\left(s|\hat{\theta}\right)}\right)\right]. \qquad (24)
\end{aligned}$$

Denote by $g(v) \coloneqq \frac{1}{2}\mathbb{E}\log\left(\frac{\hat{l}_1(s|\gamma)}{\sum_{\hat{\theta}\in\Theta} v(\hat{\theta})\hat{l}_1(s|\hat{\theta})}\right) + \frac{1}{2}\mathbb{E}\log\left(\frac{\hat{l}_2(s|\gamma)}{\sum_{\hat{\theta}\in\Theta} v(\hat{\theta})\hat{l}_2(s|\hat{\theta})}\right)$. 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 $v(\alpha)$ 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 can be shown that when x is sufficiently close to 1 (e.g., $x = 9/10$), the minimizer is $v(\gamma) = 1$. In this case, we have $\min_{v\in\Delta(\Theta)} g(v) = 0$. From (24), when x is sufficiently large, we have

$$\mathbb{E}\left[\log v_{t+1}(\gamma)\,|\mathcal{F}_t\right] \geq \log v_t(\gamma) + \min_{v\in\Delta(\Theta)} g(v) = \log v_t(\gamma),$$

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 $v_t \to \delta_\gamma$ almost surely, which implies that $\mu_{i,t} \to \delta_\gamma$ almost surely. □

A.6 Proof of Proposition 2

The first part is straightforward. In particular, if a belief aggregator has zero degree, it is easy to see that group irrationality occurs.[15] The proof thus focuses on the second part of Proposition 2 and shows that all regular aggregators with non-zero degree are susceptible to strong group irrationality.

To prove this claim, it suffices to show that there exist innocuous model perceptions $\hat{l} = \left\{\hat{l}_1, ..., \hat{l}_n\right\}$ and some state $\theta \neq \theta^*$ such that: (i) when $p > 0$, we have $\theta^* \not\succsim_p^w \theta$, and (ii) when $p < 0$, we have $\theta \succ_p^w \theta^*$.[16] The proof is by construction. Suppose $\Theta = \{\theta_1, ..., \theta_{K+1}\}$. Without loss of generality, suppose $\theta^* = \theta_1$. I construct a signal structure such that the number of signals is equal to the number of states, so $S_i = \left\{s^1, ..., s^{K+1}\right\}$ and

$$\hat{l}_i\left(s_j|\theta_k\right) = \begin{cases} \varepsilon_i^k & j \neq k \\ 1 - K\varepsilon_i^k & j = k \end{cases}, \tag{25}$$

where ε_k is a sufficiently small number. We have the following lemma.

Lemma 6. *For all $\varepsilon_i^1 \in (0,1)$, there exists some $\varepsilon_i^2, ..., \varepsilon_i^{K+1} \in (0, \varepsilon_i)$ such that $\Theta_i = \{\theta^*\}$.*

Proof. Fixing ε_1, we have

$$\begin{aligned}\mathcal{R}_i\left(\theta^*\right) - \mathcal{R}_i\left(\theta_k\right) =& l_i\left(s_i^1|\theta^*\right) \log \frac{1 - K\varepsilon_i^k}{\varepsilon_i^1} + l_i\left(s_i^k|\theta^*\right) \log \frac{\varepsilon_i^k}{1 - K\varepsilon_i^1} \\ &+ \sum_{j \neq 1,k} l_i\left(s_i^j|\theta^*\right) \log \frac{1 - K\varepsilon_i^k}{1 - K\varepsilon_i^1} \to -\infty \quad \text{as } \varepsilon_i^k \to 0.\end{aligned}$$

Therefore, when $\varepsilon_i^2, ..., \varepsilon_i^{K+1}$ are sufficiently small relative to ε_i^1, θ^* is the unique state that minimizes the relative entropy for individual i. □

Lemma 6 ensures that we can construct an innocuous model perceptions with form (25). For convenience, we consider a particular selection of feasible $\varepsilon_i^2, ..., \varepsilon_i^{K+1}$ in Lemma 6 and let $\varepsilon_i^k\left(\varepsilon_i^1\right)$ denote a continuous function such that $\Theta_i = \{\theta^*\}$.[17] Suppose that all individuals

[15] As discussed in the main text, asymptotic beliefs are characterized by Θ_0^w—the set of states that minimize the weighted relative entropy—which needn't minimize the relative entropy for any individual.

[16] This is because by Theorem 1, we know that (i) implies $\theta^* \notin \Theta_p^w$, so beliefs will not converge to the point-mass on θ^*, and (ii) implies that there is another state $\theta \in \hat{\Theta}_p^w$, so beliefs will settle on θ with a strictly positive probability.

[17] The existence is ensured by Michael selection theorem.

hold identical innocuous perception in the form of (25). When $p > 0$, we have

$$\mathcal{D}_p^w(\theta_k, \theta^*) = \frac{1}{p}\mathbb{E}\log\left(\sum_{i=1}^n w_i \times \left(\frac{\hat{l}_i(s_i|\theta_k)}{\hat{l}_i(s_i|\theta^*)}\right)^p\right)$$
$$= \frac{1}{p}\sum_{s\in S} l(s|\theta^*) \times \log\left(\sum_{i:s_i=s^1} w_i \times \left(\frac{1-K\varepsilon_i^k(\varepsilon_i^1)}{\varepsilon_i^1}\right)^p + \sum_{i:s_i=s^k} w_i \times \left(\frac{\varepsilon_i^k}{1-K\varepsilon_i^1}\right)^p\right.$$
$$\left.+\sum_{i:s_i\neq s^1,s^k} w_i \times \left(\frac{1-K\varepsilon_i^k(\varepsilon_i^1)}{1-K\varepsilon_i^1}\right)^p\right) \to +\infty \qquad \text{as } \varepsilon_1^1, ..., \varepsilon_n^1 \to 0.$$

This implies that when $p > 0$, we have $\theta^* \not\succeq_p^w \theta_k$ for all k. Therefore, $\theta^* \notin \Theta_p^w$, implying that beliefs will not converge to the point-mass on θ^* by Theorem 1.

Case 2: Symmetrically, when $p < 0$, we also have

$$\mathcal{D}_p^w(\theta^*, \theta_k) \to -\infty \qquad \text{as } \varepsilon_1^1, ..., \varepsilon_n^1 \to 0$$

This implies that when $p < 0$, we have $\theta_k \succ_p^w \theta^*$ for all k. Also note that $\theta^* \succ_0^w \theta_{k'}$ for all k', so we have $\theta_k \succ_p^w \theta_{k'}$ for all k' by Corollary 3. Therefore, $\theta_k \in \hat{\Theta}_p^w$, and hence beliefs will converge to the point-mass on θ_k with a positive probability. Combing previous two cases, we see that any belief aggregator, regardless of positive or negative p, is susceptible to strong group irrationality.

A.7 Proof of Proposition 3

Proof. (i) The "if" direction: By definition, for all $\theta \in \Theta$, we have

$$\log\left(\frac{\mu_{i,t+1}(\theta)}{\mu_{i,t+1}(\theta^*)}\right) = \log\left(\frac{F(\mu_{1,t}, ..., \mu_{n,t})(\theta)}{F(\mu_{1,t}, ..., \mu_{n,t})(\theta^*)}\right) + \log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta)}{\hat{l}_i(s_{i,t+1}|\theta^*)}\right).$$

Assumption 3 implies that

$$\max_i \log\left(\frac{\mu_{i,t+1}(\theta)}{\mu_{i,t+1}(\theta^*)}\right) \leq \max_i \log\left(\frac{\mu_{i,t}(\theta)}{\mu_{i,t}(\theta^*)}\right) + \max_i \log\left(\frac{\hat{l}_i(s_{i,t+1}|\theta)}{\hat{l}_i(s_{i,t+1}|\theta^*)}\right).$$

Suppose $\theta^* \succeq_\infty \theta$ for all $\theta \neq \theta^*$, then the strong law of large numbers implies that $\max_i \log\left(\frac{\mu_{i,t}(\theta)}{\mu_{i,t}(\theta^*)}\right) \to -\infty$ almost surely for all $\theta \neq \theta^*$, which further implies that $\mu_{i,t}(\theta^*) \to 1$ almost surely for all i. (ii) The "only if" direction: Suppose instead that there exists some $\theta \neq \theta^*$ such that $\theta^* \not\succeq_\infty \theta$, then we have $\mathcal{D}_\infty(\theta, \theta^*) > 0$. Note that $\mathcal{D}_p^w(\theta, \theta') \to \mathcal{D}_\infty(\theta, \theta')$ as $p \to +\infty$, so we also have $\mathcal{D}_p^w(\theta, \theta^*) > 0$ for some $p < +\infty$ and $w \in \Delta_{++}(N)$. Therefore, $\theta^* \notin \Theta_p^w$, and hence beliefs can't converge almost surely to δ_{θ^*} by Theorem 1. $\square$

A.8 Proof of Corollary 2

Proof. Let θ be any state that is not the true state. Let $X_i \equiv \log\left(\frac{\hat{l}(s_i|\theta)}{\hat{l}(s_i|\theta^*)}\right)$. By assumption, X_i's are i.i.d. Denote by $\underline{X}$ and $\overline{X}$ the minimum and maximum possible values of X_i, so we have $\underline{X} < 0 < \overline{X}$. Let ε be a number between 0 and $\overline{X}$, and we have

$$\begin{aligned}\mathcal{D}_\infty(\theta,\theta^*) = \mathbb{E}\left(\max_i X_i\right) &\geq \mathbb{P}\left(\max_i X_i \leq \varepsilon\right)\underline{X} + \mathbb{P}\left(\max_i X_i > \varepsilon\right)\varepsilon \\ &= \mathbb{P}^n\left(X_i \leq \varepsilon\right)\underline{X} + \left[1-\mathbb{P}^n\left(X_i \leq \varepsilon\right)\right]\varepsilon \to \varepsilon > 0 \quad \text{as } n \to +\infty.\end{aligned}$$

Therefore, we have $\theta^* \not\succeq_\infty \theta$ when n is sufficiently large. By Proposition 3, $\hat{l}$ is not robustly innocuous. $\square$

A.9 Proof of Proposition 4

I define the *generalized weighted p-entropy* as follows

$$\mathcal{D}_w^p(\theta',\theta;\psi) = \frac{1}{p}\times\mathbb{E}\log\left(\sum_{i=1}^n w_i \times \left(\frac{\psi_i(s_i|\theta')}{\psi_i(s_i|\theta)}\right)^p\right).$$

It is easy to verify that Theorem 1 still holds. Suppose that $\psi_i(s|\theta) = \hat{l}_i^{\,\alpha}(s|\theta)$. Then for all $\theta',\theta \in \Theta$, we have

$$\begin{aligned}\frac{1}{\alpha}\mathcal{D}_w^p(\theta',\theta;\alpha) &= \frac{1}{\alpha p}\times\mathbb{E}\log\left(\sum_{i=1}^n w_i\times\left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)}\right)^{\alpha p}\right) \\ &\xrightarrow{\alpha\to 0}\sum_{i=1}^n w_i\times\mathbb{E}\log\left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)}\right) = \mathcal{I}^w(\theta) - \mathcal{I}^w(\theta'),\end{aligned} \tag{26}$$

where $\mathcal{D}_w^p(\theta',\theta;\alpha)$ represents the generalized weighted p-entropy with α-Bayes' rule. Let $\succeq_{p,\alpha}^w$ denote binary relation induced by it. Therefore, (26) implies that when α is sufficiently small, $\succeq_{p,\alpha}^w$ and $\succeq_0^w$ induces the same ordering over Θ—that is, $\theta' \succeq_{p,\alpha}^w \theta$ if and only if $\theta' \succeq_0^w \theta$ for all $\theta',\theta\in\Theta$. From the assumption that $\mathcal{I}^w(\theta)$ is distinct for all states, we have $\Theta_0^{**} = \Theta_0^* = \{\theta_0\}$. Therefore, when α is small, $\Theta_{p,\alpha}^{**} = \Theta_{p,\alpha}^* = \{\theta_0\}$, where $\Theta_{p,\alpha}^*$ (and $\Theta_{p,\alpha}^{**}$) represents the (strict) minimizer of the generalized weighted p-entropy with power α. From Theorem 1, we know that beliefs converge to δ_{θ_0} with a strictly positive probability. It remains to show that the convergence happens with probability 1.

Lemma 7. *When α is sufficiently small, there exists some $\eta > 0$ and $\varepsilon > \varepsilon' > 0$ such that*

$$\mathbb{P}\left(v_t \in B_\varepsilon\left(\delta_{\theta_0}\right) \text{ for all } t \geq 1 | v_1 \in B_{\varepsilon'}\left(\delta_{\theta_0}\right)\right) \geq \eta, \tag{27}$$

and for all $\theta \neq \theta_0$

$$\mathbb{P}\left(\max\left\{v_t\left(\hat{\theta}\right) : \hat{\theta} \succ_0^w \theta\right\} \geq \varepsilon \text{ for some } t < \infty | v_1 \in B_{\varepsilon'}\left(\delta_\theta\right)\right) \geq \eta, \tag{28}$$

where $\succ_0^w$ denotes the strict ranking induced by the weighted relative entropy, and θ_0 is the minimizer of the weighted relative entropy.

Proof. The proof for (27) comes directly from Theorem 1 which implies that $v_t \to \delta_{\theta_0}$ with a strictly positive probability. To prove (28), we follow a similar approach as in the proof of Theorem 1. Choose any state $\theta \neq \theta_0$ and let θ' be another state such that $\theta \succ_0^w \theta'$. As $v_t(\theta) \to 1$, we have

$$v_{t+1}^p\left(\theta'\right) \sim v_t^p\left(\theta'\right) \times \left[\sum_{i=1}^n w_i \left(\frac{\hat{l}_i\left(s_{i,t}|\theta'\right)}{\hat{l}_i\left(s_{i,t}|\theta\right)}\right)^{\alpha p}\right].$$

Belwo, I focus on the case where $p > 0$ as the case where $p < 0$ follows exactly symmetrically as in Theorem 1's proof. When $p > 0$, we have $\mathbb{E}\left[\sum_{i=1}^n w_i \left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)}\right)^{\alpha p}\right] < 1$ when α is sufficiently small.[18] Define $T = \inf\left\{t : v_t \notin B_\varepsilon\left(\delta_\theta\right)\right\}$, so $Y_t\left(\theta'\right) = v_{t\wedge T}^p\left(\theta'\right)$ is a super-martingale. Following the approach as in the proof of Theorem 1, we have

$$\mathbb{P}\left(\cup_{\theta \succ_0 \theta'}\left\{Y_T\left(\theta'\right) \geq \varepsilon^p\right\}\right) \leq |\Theta| \times \left(\frac{\varepsilon'}{\varepsilon}\right)^p.$$

By definition, θ doesn't minimizes the weighted relative entropy—and hence doesn't minimize the weighted p entropy with small α—Theorem 1 implies that beliefs cannot converge to the point-mass on θ, thus $\mathbb{P}(T < \infty) = 1$. Therefore, we have $\mathbb{P}\left(\cup_{\theta' \neq \theta}\left\{Y_T\left(\theta'\right) \geq \varepsilon^p\right\}\right) = 1$, which implies that

$$\mathbb{P}\left(\cup_{\theta' \succ_0 \theta}\left\{Y_T\left(\theta'\right) \geq \varepsilon^p\right\}\right) \geq 1 - \mathbb{P}\left(\cup_{\theta \succ_0 \theta'}\left\{Y_T\left(\theta'\right) \geq \varepsilon^p\right\}\right) \geq 1 - |\Theta| \times \left(\frac{\varepsilon'}{\varepsilon}\right)^p. \tag{29}$$

Letting $\varepsilon' = \frac{\varepsilon}{(|\Theta|+1)^{1/p}}$, the R.H.S. of (29) becomes $\frac{1}{|\Theta|+1}$. Let $\eta = \frac{1}{|\Theta|+2}$, so when ε is

[18] Let $f(m) = \mathbb{E}\left[\sum_{i=1}^n w_i \left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)}\right)^m\right]$. It is easy to verify that $f(0) = 1$ and $f'(0) = \mathcal{D}_0^w\left(\theta', \theta\right) < 0$, so $f(m) < 1$ when m is small.

sufficiently small,

$$\mathbb{P}\left(\max\left\{v_t\left(\hat{\theta}\right):\hat{\theta}\succ_0^w\theta\right\}\geq\varepsilon \text{ for some } t<\infty|v_1\in B_{\varepsilon'}\left(\delta_\theta\right)\right)\geq\eta>0,$$

so the claim is proved. □

For any history $\mathcal{F}_t=\sigma\left(s_1,...,s_t\right)$, there exists some $\theta_1\in\Theta$ such that $v_t\left(\theta_1\right)\geq\frac{1}{|\Theta|}>\varepsilon$, when ε is sufficiently small. Assumptions 3 and 4 jointly imply that

$$\mathbb{P}\left(v_{t_1}\in B_{\varepsilon'}\left(\delta_{\theta_1}\right) \text{ for some } t_1\geq t|\mathcal{F}_t\right)\geq\delta.$$

Lemma 7 implies that there exists some $\theta_2\succ_0^w\theta_1$ such that

$$\mathbb{P}\left(v_{T_1}\left(\theta_2\right)\geq\varepsilon \text{ for some } T_1\geq t_1|v_{t_1}\in B_{\varepsilon'}\left(\delta_{\theta_1}\right)\right)\geq\eta.$$

From Assumptions 3 and 4 again, we get

$$\mathbb{P}\left(v_{t_2}\in B_{\varepsilon'}\left(\delta_{\theta_2}\right) \text{ for some } t_2\geq T_1|v_{T_1}\left(\theta_2\right)\geq\varepsilon\right)\geq\delta.$$

In summary, beliefs transit from $B_{\varepsilon'}\left(\delta_{\theta_1}\right)$ to $B_{\varepsilon'}\left(\delta_{\theta_2}\right)$ with probability $\eta\delta>0$. After repeating the process k times, beliefs transit sequentially from the neighborhoods of Dirac beliefs on $\theta_1,\theta_2,...,\theta_k$, where $\theta_k\succ_0^w...\succ_0^w\theta_1$ and each transit occurs with probability at least $\eta\delta>0$. Because $|\Theta|$ is finite, beliefs will enter $B_{\varepsilon'}\left(\delta_{\theta_0}\right)$ after at most $|\Theta|-1$ rounds, and it will remain inside with probability $\geq\eta$ from (27). So, for all $\mathcal{F}_t$, we have

$$\mathbb{P}\left(v_{t'}\in B_{\varepsilon'}\left(\delta_{\theta_0}\right) \text{ for some } t'\geq t|\mathcal{F}_t\right)\geq\left(\eta\delta\right)^{|\Theta|-1}\Rightarrow\mathbb{P}\left(v_t\to\delta_{\theta_0}|\mathcal{F}_t\right)\geq\left(\eta\delta\right)^{|\Theta|}>0.$$

Levy's 0-1 Law then implies that $\mathbb{P}\left(v_t\to\delta_{\theta_0}\right)=1$.