The Wisdom of Crowds or Group Irrationality? Non-Bayesian Social Learning With Misspecification*

Jaden Yang Chen†

December 21, 2025

Abstract

This paper examines a non-Bayesian social learning problem in which individuals may have misspecified models of their signal distributions. I introduce the concept of *group irrationality*, in which social learning fails even though each individual would learn the truth in isolation. I characterize asymptotic beliefs using a new divergence measure, the weighted p-entropy. The analysis reveals that group irrationality is prevalent, but correct learning can be restored when the social learning rule is sufficiently close to log-linear or when individuals sufficiently underreact to private information.

JEL Classification: D81, D83
Keywords: Social learning, 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. I also thank the editor, Navin Kartik, and four anonymous referees for thoughtful comments.

†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 Mind*

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 a central question: does information exchange facilitate accurate learning, or can it distort beliefs and lead to systematically worse outcomes?

The literature offers two contrasting perspectives. One line of work warns that social learning may generate suboptimal outcomes: because individuals respond to others' actions, interactions can impede information aggregation.[1] Another view is more optimistic, emphasizing that social learning can aggregate dispersed information and deliver the "wisdom of crowds." In Bayesian models, efficient aggregation is achievable under suitable conditions on signals and network structure.[2] Even with *non-Bayesian* learners, certain heuristic updating rules, such as weighted-average belief updating, can still produce correct learning.[3] These non-Bayesian results demonstrate that efficient information aggregation is sometimes attainable even with boundedly rational individuals, echoing the original intuition behind the "wisdom of crowds" phenomenon.[4]

However, most existing models assume that individuals correctly interpret their information. While natural as a benchmark, this overlooks an important feature of real-world learning: *model misspecification*. Individuals may misperceive their own signals, and once

[1] In sequential social learning, information cascades can emerge (Banerjee, 1992; Bikhchandani et al., 1992). In models with repeated interactions, learning can converge slowly (Harel et al., 2021; Huang et al., 2024).

[2] See Smith and Sørensen (2000) on sequential learning with unbounded signals; Kartik et al. (2024) on "excludability" condition in multi-state environments; Acemoglu et al. (2011) on the expanding condition on networks; and Mossel et al. (2014) on asymptotic learning in networks.

[3] See Golub and Jackson (2010) and Jadbabaie et al. (2012).

[4] Surowiecki (2004) documents many examples in which aggregated opinions of non-experts closely match the truth, sometimes outperforming experts; for instance, the country-fair experiment where the average guess of an ox's weight was remarkably accurate despite most participants lacking expertise.

misspecification is present, the performance of social learning depends not only on how dispersed information is aggregated but also on how distorted signals propagate. For instance, online reviews may help consumers make better choices but can also transmit misinformation and reinforce systematic biases.

This paper develops a model of non-Bayesian social learning with misspecification. A finite set of individuals attempt to learn the true state through a repeated two-stage process. In the first stage, individuals communicate their beliefs, and society forms a social belief through an aggregation function F. The analysis focuses on aggregation rules that can be locally approximated by the generalized mean of individual beliefs:

$$F(\mu_1, \ldots, \mu_n)(\theta) \;\sim\; \left(\sum_{i=1} w_i \, \mu_i(\theta)^p \right)^{1/p},$$

where w_i describes the weight placed on individual i, and p determines the shape (or order) of the aggregator. We can interpret F as a reduced-form description of how diverse opinions are synthesized; for example, through committee deliberation, expert weighting, or algorithmic aggregation of information. In the second stage, individuals receive private signals and update the social belief using Bayes' rule, possibly under a misspecified model of the data generating process (DGP). The process then repeats indefinitely.

This paper introduces a new concept, *group irrationality*, which occurs when society fails to learn the true state even though each individual would do so under independent learning. A weaker notion, *weak group irrationality*, arises when, in the limit, society assigns positive probability to a state that would be assigned probability zero in everyone's independent learning. In short, (weak) group irrationality captures situations where the collective outcome is inconsistent with individual outcomes. In real life, numerous examples suggest the presence of group irrationality. For instance, a company board may approve decisions that its members would likely oppose individually, such as the AOL-Time Warner merger in 2000, which many board members later admitted they had reservations about.[5] Other examples include wartime environments in which soldiers commit actions they would never contemplate individually, or jury decisions in which the collective verdict sharply diverges from the private assessments of individual jurors.[6] Below, I provide an example to illustrate

[5] Various sources indicate that many high-level executives had reservations regarding the merging decision, which became one of the most infamous corporate failures in history. As Bob Pittman, 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)

[6] This idea echoes a classic insight from crowd psychology. In *The Crowd*, Gustave Le Bon writes that the collective mind leads individuals to "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" (Le Bon, 1895).

group irrationality.

Example 1 (Group Irrationality)**.** Consider a society with two individuals. The state space is $\Theta = \{G, B\}$, and the true state $\theta^* \in \Theta$ is unknown to both. Each individual i holds a full-support prior $\mu_{i,0}$. In each period $t \in \{1, 2, ...\}$, they first communicate their most recent beliefs and aggregate them using DeGroot's rule

$$v_t = \frac{1}{2}\mu_{1,t-1} + \frac{1}{2}\mu_{2,t-1},$$

where $\mu_{i,t-1}$ denotes individual i's beliefs at time $t-1$, and v_t denotes the social belief. After communication, each individual i receives a signal $s_{i,t} \in S = \{S_G, S_B\}$. They then update the social belief using Bayes' rule under their perceived DGP

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

The true DGP has the same form with parameter $p^* > 1/2$.

- **Correct Specification:** If $p_i = p^*$ for all i, then individuals learn the true state. This echoes the classic "wisdom of crowds" result under DeGroot-style learning.

- **Misspecification:** Suppose $p_i \neq p^*$ for some i. Existing work does not address whether correct learning still obtains.[7] At first glance, misspecification might seem harmless: individuals still interpret signals in the right direction, so each would learn the truth independently.[8] However, this paper shows that under DeGroot-style aggregation, the true state may fail to be learned, and beliefs can even oscillate indefinitely; in other words, group irrationality can arise.

This paper develops a set of characterizations for the long-run behavior of beliefs under misspecified social learning. I introduce a divergence measure, *the weighted p-entropy*, that captures how individuals' misspecified models interact with the social aggregation rule. Theorem 1 shows that this divergence plays a role analogous to relative entropy in classical Bayesian learning: it determines which states can receive weight in any limiting belief and, in particular, implies that long-run beliefs must concentrate on states that minimize this

[7]Related work examines social learning in other settings; for instance, Bohren and Hauser (2021) study sequential learning with misspecification, whereas Molavi et al. (2018) analyze repeated interactions under correct specification.

[8]If, in the long run, there are more S_G than S_B, individuals know that the true state is G; otherwise, the true state is B.

divergence (Berk, 1966). A key difference from the classical setting is that the weighted p-entropy does not induce a complete ranking over states, so beliefs may fail to converge or converge to multiple limits. I show that the full set of asymptotic outcomes can be characterized in the binary-state case (Proposition 1) and extend the analysis to general state spaces (Propositions 2 and 3).

Building on these characterizations, I examine when and why group irrationality arises in social learning. Proposition 4 asks whether some aggregation rules are immune to group irrationality. The result is negative: many aggregation rules are vulnerable. In particular, weak group irrationality can occur under essentially any aggregation rule, and group irrationality arises whenever the rule has nonzero order. Proposition 5 then considers the converse question of whether some model perceptions guarantee correct learning under all aggregation rules. Such perceptions exist, but only under highly restrictive conditions.[9] One implication is that when signals are independent across individuals and society is sufficiently large, no perception can ensure correct learning.

The preceding results show that group irrationality is widespread when neither aggregation rules nor perceptions are restricted. I therefore identify two sufficient conditions under which correct learning can nonetheless be recovered. First, when individuals hold consistent model perceptions—so each would learn the truth in isolation—correct learning arises whenever the aggregation rule has *sufficiently small order* (Proposition 6). This is because low curvature makes the social belief less sensitive to extreme posteriors, preventing such beliefs from dominating the aggregation. Second, for any given aggregation rule, correct learning also obtains when individuals' perceived models are consistent and *sufficiently uninformative* (Proposition 7). Underreaction keeps posterior movements modest, making the distribution of beliefs less extreme and thereby limiting the distortions introduced during aggregation.[10] These two conditions are also, in a sense, necessary for correct learning. Learning failures become more likely as the aggregation rule has larger order (Proposition 8) or as perceived models are more informative (Corollary 6).[11] Together, these results point to a common mechanism: group irrationality emerges when aggregation is overly sensitive to extreme beliefs or when individual beliefs are excessively dispersed, and correct learning can be restored by dampening either force.

[9] It requires that the true state dominates all other states even in $p \to +\infty$, a requirement that becomes extremely restrictive in large societies.

[10] In some examples, correct learning fails under correct specification but succeeds with underreaction, suggesting that misspecification can even improve social learning.

[11] Specifically, Proposition 8 shows that we can rank aggregation rules by their order, where learning is easier to occur under rules with smaller order. Corollary 6 shows that under certain conditions, learning fails whenever perceived signals are sufficiently informative.

1.1 Related Literature

Bayesian Social Learning. Classical social learning models typically assume that agents are rational and update beliefs in a Bayesian manner. In a *sequential* environment, Banerjee (1992) and Bikhchandani et al. (1992) show that information cascades and herding can emerge, while Smith and Sørensen (2000) establish that learning is asymptotically correct when signals are unbounded. Subsequent work extends this analysis to richer environments, including networked interactions (e.g., Acemoglu et al. 2011; Arieli and Mueller-Frank 2019) and general state spaces and preferences (e.g., Kartik et al. 2024; Arieli and Mueller-Frank 2021). Notably, the framework in this paper also accommodates arbitrary finite state spaces, moving beyond the binary settings common in much of the literature. There is also a substantial literature on Bayesian social learning with *repeated interactions* in networks (e.g., Gale and Kariv 2003; Rosenberg et al. 2009; Mossel et al. 2015; Harel et al. 2021; Huang et al. 2024). These models examine settings in which agents repeatedly observe neighbors or exchange information and update beliefs over time. While the present paper also features repeated interactions, the key distinction is that agents here are not Bayesian learners.

Non-Bayesian Social Learning. This paper also relates to the literature on non-Bayesian social learning. The closest antecedents are Jadbabaie et al. (2012) and Molavi et al. (2018), who study environments in which agents receive infinitely many signals and repeatedly exchange beliefs with their neighbors. Jadbabaie et al. (2012) show that simple averaging of neighbors' beliefs leads to correct learning, while Molavi et al. (2018) provide an axiomatic foundation for log-linear learning rules and characterize conditions for asymptotic learning; I discuss them in more detail later in the paper. Other strands emphasize different departures from full rationality: Bala and Goyal (1998) and Li and Tan (2020) analyze quasi-Bayesian learning in networks with repeated interactions, and Eyster and Rabin (2010), Guarino and Jehiel (2013), and Dasaratha and He (2020) study sequential learning with naive or otherwise boundedly rational agents. This paper is also related to the broader literature on belief exchange and opinion aggregation. Classic contributions such as DeGroot (1974), DeMarzo et al. (2003), and Golub and Jackson (2010) study how averaging rules shape belief dynamics, while more recent work examines richer aggregation mechanisms; for example, Cerreia-Vioglio et al. (2024) on nonlinear aggregators, and Banerjee and Compte (2024) on endogenous aggregation rules. Relative to these papers, the key novelty of this paper is to introduce model misspecification into a repeated non-Bayesian learning environment. I show that when individuals systematically misperceive their own signals, group irrationality can emerge even under aggregation rules that would guarantee correct learning in correctly specified settings.

Misspecified Social Learning. This paper is also related to the literature on social learning with model misspecification. The most technically related paper is Frick et al. (2023), who employ the p-divergence to construct local martingales. This paper also uses a local approach but relies on a different divergence, the weighted p-entropy; see Section 3 for more discussion. Bohren (2016) and Bohren and Hauser (2021) study a sequential social learning problem in which agents hold misspecified beliefs about the information environment, while Arieli et al. (2023) analyze settings in which individuals misperceive the informativeness of predecessors' signals. Gagnon-Bartsch and Rosato (2024) study a social learning problem with "taste projection", where agents misinterpret others' preferences. Frick et al. (2020) study a social learning problem with continuous state space and show that correct learning is not robust with respect to small misspecification. A complementary line of work studies model uncertainty more broadly: Hare et al. (2020), Chen (2022), and Huang (2023) analyze environments in which agents entertain multiple, potentially incorrect, models and update beliefs accordingly.

The remainder of the paper proceeds as follows. Section 2 presents the main model. Section 3 introduces the notion of the weighted p-entropy. Section 4 presents characterizations of asymptotic beliefs. Section 5 discusses implications for group irrationality. Section 6 presents an extension of generalized Bayes' rule in private learning. Section 7 is the conclusion.

2 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 seeks to learn the true state. Time is discrete, $t \in \mathbb{T} = \{1, 2, ...\}$, and each period t consists of two stages: a social-learning stage followed by a private-learning stage.

2.1 Social Learning Stage

In the first stage, individuals communicate their beliefs. Society then forms a social belief

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

Here F is referred to as the *belief aggregator*, which is a reduced-form description of how society pools individual beliefs into a single collective belief. We can view F as a belief-pooling procedure, such as deliberation, weighted scoring, or algorithmic aggregation, through which diverse opinions are synthesized into a common position. In practice, for instance, a board,

jury, or expert panel may use a fixed rule to convert members' beliefs into a unified verdict, forecast, or recommendation.

The aggregator need not coincide with the Bayesian posterior; it may arise from *non-Bayesian* heuristics. A canonical example is DeGroot's rule (DeGroot, 1974), under which the social belief is a weighted arithmetic mean of individual beliefs. This paper studies a broader class of averaging rules that generalizes DeGroot's formulation. The following assumption describes the aggregator's local behavior in the tail.

Assumption 1 (Power-Mean Tail)**.** *If* $\min_i \mu_i(\hat{\theta}) \to 1$ *for some state* $\hat{\theta}$*, then for all* $\theta \neq \hat{\theta}$*,*

$$F(\mu_1, \ldots, \mu_n)(\theta) \;\sim\; \left(\sum_{i=1}^{n} w_i \, \mu_i(\theta)^p \right)^{1/p},$$

for some fixed $p \in \mathbb{R}$ *and* $w \in \Delta_{++}(N)$*.*[12]

Assumption 1 states that, as the group becomes nearly certain about some state, the aggregated belief can be approximated by a power mean of individual beliefs. I refer to p as the *order* of the aggregation rule and to w_i as individual i's *influence weight*. I also write $p\,(F)$ and $w_i\,(F)$ when emphasizing the dependence on the aggregator.

The weight w_i captures the importance of individual i's belief; for example, a committee may give more influence to members with greater expertise or seniority. The order p governs the sensitivity of the aggregator to extreme opinions.[13] When $|p|$ is large, the aggregator becomes highly sensitive to extreme components—high beliefs when $p > 0$ and low beliefs when $p < 0$—as when a vocal minority disproportionately shapes the group's decision, whereas when $|p|$ is small, it is less affected by extremes and behaves more like a balanced average. This class encompasses a broad range of aggregation rules including arithmetic ($p = 1$), geometric ($p \to 0$), and max/min ($p \to \pm\infty$) aggregation. Below is an example:

Example 2 (Power-Mean Rule)**.** Suppose that F satisfies:

$$F\left(\mu_1, ..., \mu_n\right)(\theta) = \frac{\left(\sum_i w_i \times \mu_i\left(\theta\right)^p\right)^{1/p}}{\sum_{\theta' \in \Theta}\left(\sum_i w_i \times \mu_i\left(\theta'\right)^p\right)^{1/p}}, \tag{2}$$

i.e., the social belief is the normalized power-mean of individual beliefs. This is the "power-mean rule" in Molavi et al. (2018); their log-linear rule corresponds to the limit $p \to 0$.

[12] Here $f(x) \sim g(x)$ means $\lim f/g = 1$ along the relevant limit. When $p = 0$, interpret the right-hand side as $\exp(\sum_i w_i \log \mu_i(\theta))$.

[13] In generalized-mean aggregators, the order p controls curvature: larger p corresponds to a more convex rule that amplifies large belief components, while smaller p makes the rule more concave and sensitive to small components (see Hardy et al. (1952)). This curvature interpretation is also related to the role of the elasticity of substitution in CES aggregators.

2.2 Private Learning Stage

In the second stage, individuals receive private signals. Signals are independent across time but may be correlated across individuals. In each period t, a signal profile $s_t = (s_{1,t}, ..., s_{n,t}) \in S$ is generated according to a DGP $l(s|\theta)$. The signal space $S = \times_{i=1}^{n} S_i$ is finite, and every signal occurs with strictly positive probability in every state: $l(s|\theta) \in (0,1)$ for all $s \in S$ and $\theta \in \Theta$. I assume the absence of identification problems: no state $\theta' \neq \theta^*$ induces the same distribution as the true one $l(s|\theta^*)$. Let $\mathbb{P}$ and $\mathbb{E}$ denote, respectively, the probability measure and expectation induced by the true distribution $l(\cdot|\theta^*)$ on $(S^\infty, \sigma(S^\infty))$.

Each individual i observes only their own signal $s_{i,t}$, and may be *misspecified* about its distribution. Let $l_i(s|\theta)$ be individual i's true DGP and $\hat{l}_i(s|\theta)$ the perceived one, which may differ from $l_i(s|\theta)$. I also assume $\hat{l}_i(\cdot|\theta) \neq \hat{l}_i(\cdot|\theta')$ for all $\theta \neq \theta'$ and all i, so individuals do not confound distinct states under their perceived models.

Given the social belief v_t, individual i's posterior is obtained through Bayesian updating:

$$\forall \theta \in \Theta: \qquad \mu_{i,t}(\theta) = BU_i(v_t, s_{i,t})(\theta) \equiv \frac{v_t(\theta) \times \hat{l}_i(s_{i,t}|\theta)}{\sum_{\theta' \in \Theta} v_t(\theta') \times \hat{l}_i(s_{i,t}|\theta')}, \tag{3}$$

where BU denotes the Bayesian update operator. That is, individuals update the social belief using their private signals and their perceived DGPs. After updating, individuals communicate again, a new social belief is formed, and the process repeats indefinitely. This generates a stochastic sequence of belief profiles $\boldsymbol{\mu} = (\mu_1, \mu_2, ...)$, where $\mu_t = (\mu_{1,t}, ..., \mu_{n,t})$.

To simplify the analysis, I impose the following regularity condition.

Assumption 2 (Irreducibility)**.** *Fix any $\delta, \varepsilon \in (0,1)$. There exists a finite $T < \infty$ such that for all $t \in \mathbb{T}$, all $\theta \in \Theta$, and all v_t satisfying $v_t(\theta) \geq \delta$, there exists a signal profile $(s_t, \ldots, s_{t+T})$ for which $v_{t+T}(\theta) > 1 - \varepsilon$.*

This assumption states that social beliefs can be driven arbitrarily close to any state in a finite number of steps. If it were violated, society could be unable to place sufficiently high weight on some states, making correct learning impossible for those states and generating trivial learning failures. Technically, Assumption 2 ensures that extreme beliefs can be reached with positive probability, allowing us to apply the local approximation in Assumption 1. The uniform bound on the hitting time allows us to establish almost-sure claims in the limit. Importantly, this is not an assumption on the aggregator or on individuals' perceptions. It can be guaranteed by primitive conditions (e.g., Assumptions 4 and 5 in Section 5).[14]

[14]The assumption can fail, for example, when aggregation rules impose hard truncation, when some states are never favored by any signal, when interior belief sets are absorbing, or in non-i.i.d. environments where the probability of transitions vanishes over time. Outside such cases, the condition is satisfied by a wide class of aggregation rules and signal structures.

2.3 Group Irrationality

I now introduce the concept of group irrationality. Intuitively, group irrationality means that social learning becomes incorrect even though each individual would learn the truth if learning independently. To formalize this idea, define

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

where $\mathcal{R}_i(\theta)$ is the *relative entropy* (or Kullback-Leibler divergence) of individual i's perceived signal distribution in state θ from the true signal distribution $l_i(s|\theta^*)$. Relative entropy measures how well a candidate distribution approximates the true one, so Θ_i is the set of states that best fit the true distribution under individual i's perceived model.

A classical result of Berk (1966) implies that if individual i were to learn independently using Bayes' rule, then their posterior would converge to the set of best-fitting states: $\mu_{i,t}(\Theta_i) \to 1$ almost surely as $t \to +\infty$. This motivates the following definition.

Definition 1 (Consistent Model Perception)**.** Individual i's model perception $\hat{l}_i$ is called *consistent* if $\Theta_i = \{\theta^*\}$ for every possible true state $\theta^* \in \Theta$.

Under consistent model perceptions, the true state is the unique minimizer of relative entropy, so each individual would correctly learn the truth in isolation. I now formalize group irrationality.

Definition 2 (Group Irrationality)**.** Fix any true state $\theta^* \in \Theta$.[15] We say that:

(i) *weak group irrationality* emerges if $\mu_t(\cup_i \Theta_i)$ does not converge to $\mathbf{1}$ as $t \to \infty$.[16]

(ii) *group irrationality* emerges if all individuals' model perceptions are consistent and $\mu_t(\theta^*)$ does not converge to $\mathbf{1}$ as $t \to \infty$.

To interpret these definitions, recall that if individuals learned independently, each would converge to Θ_i. Weak group irrationality arises when social learning leads the group to place positive weight on states that lie outside all Θ_i's—states that no individual regards as best fitting their own signals. If weak group irrationality occurs and all individuals have consistent perceptions, then group irrationality occurs. In that case, society reaches a highly suboptimal outcome: social learning fails to identify the true state even though every individual would learn it in isolation.

[15] Throughout the paper, θ^* denotes the (arbitrary but fixed) true state. All results hold for any choice of θ^*.

[16] Recall that $\mu_t = (\mu_{1,t}, \ldots, \mu_{n,t})$ is the belief profile in period t. Convergence to $\mathbf{1}$ means that for each individual i, $\mu_{i,t}(\cup_i \Theta_i) \to 1$ as $t \to \infty$.

2.4 Discussion of Assumptions

Social Learning Rule. The social-learning stage is modeled through a belief aggregator F, which maps individual beliefs into a collective belief. The use of a reduced-form aggregator follows a large literature on opinion pooling (e.g., Genest and Zidek, 1986) and non-Bayesian social learning, in which individuals communicate their beliefs but do not fully internalize the higher-order inference required for Bayesian updating. In this spirit, F is treated as a primitive that describes how a society, committee, or information platform synthesizes heterogeneous beliefs into one position.[17]

This formulation also aligns with several real-world settings. Rating agencies, for instance, combine evaluations from multiple analysts according to a predetermined pooling rule.[18] Similarly, many committees, such as central banks or forecasting panels, form a consensus view by weighting or averaging the assessments of their members rather than reconstructing a full Bayesian posterior (e.g., Clemen, 1989; Timmermann, 2006). In algorithmic environments, a structurally related mechanism appears in federated learning (McMahan et al., 2017), where a central server aggregates locally trained model updates using the FedAvg algorithm, which implements a simple weighted averaging rule.[19]

Private Learning Rule. By contrast, individuals are assumed to use Bayes' rule in the private-learning stage, as in many non-Bayesian social learning models (e.g., Molavi et al., 2018). This asymmetry has two motivations. First, applying Bayes' rule to others' signals in the social-learning stage requires complex higher-order inference, whereas conditioning on one's own signal is comparatively straightforward. Second, many non-Bayesian updating rules can be represented as Bayesian updating under a misspecified model (Bohren and Hauser, 2024), so the private-learning stage also accommodates non-Bayesian updating through an appropriate perceived model.

Under this setup, two distinct sources of learning error arise: (i) *naivety in social learning*—the social belief is formed via a heuristic aggregator rather than through Bayesian integration of all information; and (ii) *misspecification in private learning*—individuals hold

[17] Several microfoundations can generate reduced-form pooling. One possibility is belief aggregation through the minimization of a loss function; for example, linear pooling arises from minimizing mean squared error (MSE) loss; see Timmermann (2006) for a survey on forecast combination. Another possibility is repeated local averaging on a network, where under standard regularity conditions society converges to a consensus that can be expressed as an average of individual beliefs (e.g., Golub and Jackson, 2010).

[18] As noted in Moody's (2010):" A Credit Rating is MIS's current opinion regarding the relative future creditworthiness of a credit commitment, a debt or debt-like security or contract, or an issuer of such obligation, as determined by a rating committee and expressed using its established Aaa to C alpha-numeric rating scale, or other Credit Rating scales as identified from time to time by MIS. "

[19] Although the purposes differ—federated learning aggregates parameter updates to optimize a loss function, whereas the present framework aggregates subjective beliefs—the underlying idea of merging decentralized information using a simple averaging rule is closely parallel.

an incorrect model of their signals. These two channels are not substitutes. Without misspecification, some heuristic rules (e.g., the DeGroot rule) deliver correct learning; with misspecification, even if individuals have consistent model perceptions, correct learning can fail for a broad class of heuristic rules.

Group Irrationality. In this paper, group irrationality means that society fails to learn the true state even though each individual would learn it in isolation. Likewise, weak group irrationality captures situations in which the social-learning outcome contradicts private-learning outcomes. In particular, the social posterior assigns positive probability to states that are not entropy minimizers for any individual, so society converges to a belief that no individual would reach on their own.[20] [21]

It is useful to contrast this with the classical "wisdom of crowds" phenomenon. In models such as Jadbabaie et al. (2012), each individual faces an identification problem and therefore cannot learn the state alone; yet aggregation resolves the identification problem at the group level, enabling society to correctly learn the truth. In that sense, the divergence between private and social learning is beneficial: the group reaches a belief no individual could reach independently, but the belief is fully consistent with all agents' information. By contrast, (weak) group irrationality captures the opposite pattern—social learning is *inconsistent* with individuals' own information. Here, social learning differs from individual learning, not because the group resolves an identification problem, but because the pooling rule distorts how private information is combined.

3 The Weighted p-Entropy

This section introduces a divergence measure that plays a central role in characterizing asymptotic belief dynamics.

Definition 3 (Weighted p-Entropy)**.** For all $\theta, \theta' \in \Theta$, $w \in \Delta_{++}(N)$ and $p \neq 0$, define[22]

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

[20] This usage is complementary to the perspective in Sobel (2014), which shows that under perfect pooling, group actions may rationally be more extreme than individual recommendations; extremeness of the group action is therefore not, by itself, evidence of irrationality.

[21] The entropy-based definition relies on individuals applying Bayes' rule in the private-learning stage. The idea extends to non-Bayesian updating by replacing "entropy minimizers" with the support of asymptotic posteriors under the alternative rule; see Section 6.

[22] Recall that $\mathbb{E}$ denotes the expectation under the true probability measure, i.e., the one induced by $l(\cdot|\theta^*)$.

We say that θ *weakly dominates* θ' *under the weighted* p*-entropy*, written as $\theta \succeq_p^w \theta'$, whenever $\mathcal{D}_p^w(\theta', \theta) \leq 0$. In this case, the weighted p-entropy of θ is said to be less than that of θ'.

Let Θ_p^w denote the set of states that dominate all others under $\succeq_p^w$, which I refer to as the *minimizers of the weighted* p*-entropy*. Define $\theta \succ_p^w \theta'$ if $\mathcal{D}_p^w(\theta', \theta) < 0$, and let $\hat{\Theta}_p^w$ be the set of states that strictly dominate all other states.

To build intuition, it is useful to compare weighted p-entropy to its limit case as $p \to 0$, the weighted relative entropy.

Definition 4 (Weighted Relative Entropy)**.** Define the *weighted relative entropy* of state θ as

$$\mathcal{I}^w(\theta) \equiv \sum_{i=1}^{n} w_i \times \mathcal{R}_i(\theta).$$

Let $\mathcal{D}_0^w(\theta', \theta) \equiv \mathcal{I}^w(\theta) - \mathcal{I}^w(\theta')$. Define $\theta \succeq_0^w \theta'$ (resp. $\theta \succ_0^w \theta'$) if $\mathcal{D}_0^w(\theta', \theta) \leq 0$ (resp. < 0).

Weighted relative entropy has a natural interpretation: it measures society's weighted average Kullback-Leibler distance between the true signal distribution and the perceived distribution in state θ. The weighted p-entropy generalizes this by altering how likelihood ratios are aggregated across individuals. It satisfies the following properties.

Lemma 1 (Continuity and Monotonicity)**.** *For all* $\theta \neq \theta'$ *and all* $w \in \Delta_{++}(N)$*:*

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

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

Part (i) shows that weighted p-entropy converges to weighted relative entropy as $p \to 0$. Part (ii) implies a monotonicity property: for any two states, it becomes increasingly difficult for $\theta \succeq_p^w \theta'$ to hold when p is larger. This yields the following comparison.

Corollary 1 (Relation to the Weighted Relative Entropy)**.** *For all* $\theta \neq \theta'$ *and* $p > 0 > q$*,*

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

Notice that $\succeq_0^w$ induced by weighted relative entropy is a complete and transitive order. By contrast, $\succeq_p^w$ need not be an order for $p \neq 0$. In particular,

- when $p > 0$, the relation may be *incomplete*, so the minimizer set Θ_p^w may be empty.
- when $p < 0$, the relation may be *intransitive*, so multiple strict minimizers may exist.

This yields the following corollary.

Corollary 2. *(i)* $\Theta_p^w = \emptyset$ *only if* $p > 0$*, and (ii)* $\left|\hat{\Theta}_p^w\right| > 1$ *only if* $p < 0$*.*

Corollary 2 shows that the existence and uniqueness of minimizers depend critically on the sign of p. As shown in the next section, the cardinality of Θ_p^w and $\hat{\Theta}_p^w$ directly influences the long-run belief dynamics. Thus aggregators with positive and negative p lead to qualitatively different asymptotic learning patterns.

Remark 1 (Relation to Other Divergences). The weighted p-entropy is distinct from several divergence measures used in the learning literature. A few comparisons are useful:

- *p-Divergence.* Used in Frick et al. (2023) and, in a different context, Bhattacharya et al. (2019), the p-divergence is

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

  This is a single-agent divergence: it does not depend on how beliefs are aggregated across individuals and does not incorporate any weights w_i. In Frick et al. (2023), the parameter $p > 0$ is tied to the construction of a likelihood-ratio martingale, while in Bhattacharya et al. (2019) it governs the concentration rate under power-Bayes updating. By contrast, in this paper p arises from the shape of the social aggregation rule, and the divergence explicitly reflects cross-agent pooling.[23]

- *Rényi Entropy.* Rényi entropy of order p is defined as

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

  Relative entropy corresponds to the limit as $p \to 0$. Unlike the weighted p-entropy (which may be incomplete or intransitive depending on p), Rényi entropy induces a complete ordering over states and is not linked to belief aggregation across individuals.

4 Characterization of Asymptotic Beliefs

This section characterizes asymptotic beliefs. Henceforth, all results are stated under Assumptions 1 and 2, which are maintained throughout without further mention.

[23] Also, the aggregation across agents is non-linear; it takes the form of the expected log of a weighted average of individual likelihood ratios.

4.1 Benchmark Characterizations

I begin with a theorem that characterizes convergence. It answers two basic questions: (i) if beliefs converge, what are the possible limit beliefs? and (ii) under what conditions do beliefs converge to those limits with positive probability?

Theorem 1 (Characterization of Convergence)**.** *Let p and w denote the order and weights of the aggregation rule F, respectively.*[24] *Then: (i) Conditional on the event that μ_t converges, we have $\mu_{i,t} \to \delta_\theta$ for all i for some $\theta \in \Theta_p^w$, except on $\mathbb{P}$-null events. (ii) Conversely, for each $\theta \in \hat{\Theta}_p^w$, we have $\mathbb{P}(\mu_{i,t} \to \delta_\theta \text{ for all } i) > 0$.*

Theorem 1 shows that beliefs will—and, in some sense, can only—concentrate on minimizers of the weighted p-entropy. First, conditional on convergence, beliefs must converge to a point mass on some minimizer. Second, whenever a state strictly minimizes weighted p-entropy, convergence to that state occurs with positive probability.

Building on this result, I now characterize when oscillations or multiple belief limits arise. To simplify exposition, I say that μ_t *can converge to multiple limits* if there exist distinct beliefs, $\mu_\infty^1, ..., \mu_\infty^K$ with $K > 1$ such that, for each i, $\mu_{i,t}$ converges to each μ_∞^k with strictly positive probability.[25]

Corollary 3 (Oscillations and Multiple Limits)**.** *(i) If μ_t oscillates $\mathbb{P}$-almost surely, then $\hat{\Theta}_p^w = \emptyset$; conversely, if $\Theta_p^w = \emptyset$, then μ_t oscillates $\mathbb{P}$-almost surely. (ii) If μ_t can converge to multiple limits, then $|\Theta_p^w| > 1$; conversely, if $|\hat{\Theta}_p^w| > 1$, then μ_t can converge to multiple limits. (iii) Almost-sure oscillations can occur only if $p \geq 0$, while convergence to multiple limits can occur only if $p \leq 0$.*

Corollary 3 yields two main insights. First, parts (i) and (ii) show that oscillations and multiplicity of limits are governed by the *cardinality* of Θ_p^w and $\hat{\Theta}_p^w$. Second, part (iii) shows that belief dynamics also depend on the *sign* of the aggregation rule: almost-sure oscillations arise only for $p \geq 0$, while multiple limits arise only for $p \leq 0$. This distinction follows directly from Corollary 2, which shows that $\Theta_p^w = \emptyset$ only if $p > 0$, and $|\hat{\Theta}_p^w| > 1$ only if $p < 0$.

Example 3 (Group Irrationality Revisited)**.** Theorem 1 explains the emergence of group irrationality in Example 1. In that example, beliefs are aggregated linearly, so $p = 1$. For

[24] Henceforth, p and w refer to the order and weights of the aggregation rule, unless stated otherwise.

[25] Formally, this means that $\mathbb{P}(\mu_{i,t} \to \mu_\infty^k) > 0$ for all k and all i. By Theorem 1, beliefs can only converge to Dirac beliefs, so this is equivalent to the existence of distinct states $\theta_1, ..., \theta_K$ such that $\mathbb{P}(\mu_{i,t} \to \delta_{\theta_k} \text{ for all } i) > 0$ for all k.

any $\theta \neq \theta'$,

$$\mathcal{D}_1^w(\theta, \theta') = \mathbb{E} \log \left(w_1 \times \frac{\hat{l}_1(s_1|\theta)}{\hat{l}_1(s_1|\theta')} + w_2 \times \frac{\hat{l}_2(s_2|\theta)}{\hat{l}_2(s_2|\theta')} \right).$$

Fixing p_2 and letting $p_1 \to 1$,

$$\frac{\hat{l}_1(s_1|\theta)}{\hat{l}_1(s_1|\theta')} = \begin{cases} \frac{p_1}{1-p_1} \to +\infty & s_1 = S_\theta \\ \frac{1-p_1}{p_1} \to 0 & s_1 = S_{\theta'} \end{cases},$$

whereas $\frac{\hat{l}_2(s_2|\theta)}{\hat{l}_2(s_2|\theta')}$ remains unchanged and is strictly positive. Thus, $\mathcal{D}_1^w(\theta, \theta') \to +\infty$ as $p_1 \to 1$. Therefore, when p_1 is close to 1, both $\mathcal{D}_1^w(B, G) > 0$ and $\mathcal{D}_1^w(G, B) > 0$ hold, implying $\Theta_1^w = \emptyset$. By Corollary 3, beliefs oscillate almost surely.

4.1.1 Proof Sketch of Theorem 1

The key idea is to show that the local evolution of the social belief v_t can be approximated using the weighted p-entropy. Fix a state θ, and suppose that at time t, society is highly confident in θ, in the sense that $v_t(\theta) \approx 1$.

Local Approximation. For any state $\theta' \neq \theta$ and any i, Bayesian updating implies

$$\mu_{i,t}(\theta') = \frac{v_t(\theta')\, \hat{l}_i(s_{i,t} \mid \theta')}{\sum_{\hat{\theta}} v_t(\hat{\theta})\, \hat{l}_i(s_{i,t} \mid \hat{\theta})} \approx v_t(\theta') \times \frac{\hat{l}_i(s_{i,t} \mid \theta')}{\hat{l}_i(s_{i,t} \mid \theta)}, \tag{4}$$

where the approximation uses that the denominator is dominated by the term corresponding to θ. Assumption 1 further implies that the social belief can be locally approximated by a power mean. Combining this assumption with the approximation (4) yields the recursive approximation

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

After taking log expectation to this approximation, we obtain

$$\mathbb{E}\left[\log v_{t+1}(\theta') \,|\mathcal{F}_t\right] \approx \log v_t(\theta') + \mathcal{D}_p^w(\theta', \theta), \tag{5}$$

where $\mathcal{F}_t \equiv \sigma(s_1, ..., s_{t-1})$ denotes the information by time t.

Establish Convergence. Therefore, if $\mathcal{D}_p^w(\theta', \theta) < 0$, the process $\{\log v_t(\theta')\}$ behaves locally like a supermartingale and tends to decrease. If this inequality holds for all $\theta' \neq \theta$, then the point mass δ_θ is locally stable: once the social belief v_t enters an ε-neighborhood of δ_θ, it converges to δ_θ with positive probability.

The irreducible assumption ensures that such neighborhoods are reached with positive probability from any initial belief, which establishes part (ii) of Theorem 1. Part (i) follows from a symmetric argument together with the fact that the only stationary points of the belief dynamics are point-mass beliefs.

4.2 More Complete Characterizations

While Theorem 1 sharply characterizes the *possible* limit beliefs, it does not determine whether convergence occurs almost surely, or whether convergence and non-convergence can coexist with positive probability. This subsection provides a more complete characterization of asymptotic belief dynamics.

4.2.1 Binary-State Cases

When the state space is binary, we can even obtain a *complete* characterization of asymptotic beliefs. I first impose a regularity condition.

Assumption 3 (No Indifference)**.** $\mathcal{D}_p^w(\theta, \theta') \neq 0$ *for all* $\theta \neq \theta'$.

Note that p and w denote the order and the weights of the aggregation rule, rather than arbitrary parameters. Assumption 3 rules out indifference between any two distinct states: the divergence $\mathcal{D}_p^w$ is always either strictly positive or strictly negative. Under this assumption, weak and strict dominance coincide, and hence $\Theta_p^w = \hat{\Theta}_p^w$.

The following proposition characterizes the resulting asymptotic belief dynamics.

Proposition 1 (Complete Characterization for the Binary-State Case)**.** *Suppose* $|\Theta| = 2$*, and Assumption 3 holds. Then: (i) if* $\Theta_p^w = \emptyset$*, then* μ_t *oscillates* $\mathbb{P}$*-almost surely with* $\liminf \mu_{i,t}(\theta) = 0$ *and* $\limsup \mu_{i,t}(\theta) = 1$ *for all* i *and* θ*. (ii) If* $\Theta_p^w \neq \emptyset$*, then* μ_t *converges* $\mathbb{P}$*-almost surely. Moreover,*

$$\mathbb{P}\left(\cup_{\theta \in \Theta_p^w} \{\mu_{i,t} \to \delta_\theta \text{ for all } i\}\right) = 1,$$

and for each $\theta \in \Theta_p^w$*,* $\mathbb{P}(\mu_{i,t} \to \delta_\theta \text{ for all } i) > 0$.

Part (i), as well as the fact that convergence to each δ_θ with $\theta \in \Theta_p^w$ occurs with positive probability, follows directly from Theorem 1. The new contribution of Proposition 1 is to

Minimizer	Dominance Relation	Asymptotic Dynamics
$\Theta_p^w = \emptyset$	$\mathcal{D}_p^w(\theta_1, \theta_2) > 0$ and $\mathcal{D}_p^w(\theta_2, \theta_1) > 0$	a.s. oscillations
$\Theta_p^w = \{\theta_1\}$	$\mathcal{D}_p^w(\theta_1, \theta_2) > 0$ and $\mathcal{D}_p^w(\theta_2, \theta_1) < 0$	a.s. convergence to $\mu_\infty = \delta_{\theta_1}$
$\Theta_p^w = \{\theta_2\}$	$\mathcal{D}_p^w(\theta_1, \theta_2) < 0$ and $\mathcal{D}_p^w(\theta_2, \theta_1) > 0$	a.s. convergence to $\mu_\infty = \delta_{\theta_2}$
$\Theta_p^w = \{\theta_1, \theta_2\}$	$\mathcal{D}_p^w(\theta_1, \theta_2) < 0$ and $\mathcal{D}_p^w(\theta_2, \theta_1) < 0$	a.s. convergence to $\mu_\infty \in \{\delta_{\theta_1}, \delta_{\theta_2}\}$

Table 1: Complete Characterization in $\Theta = \{\theta_1, \theta_2\}$

establish that convergence, when it occurs, does so *almost surely*: convergence and non-convergence cannot both occur with positive probability. When $\Theta = \{\theta_1, \theta_2\}$, there are exactly four possible asymptotic behaviors, as summarized in Table 1.

The proof strategy parallels that of Theorem 1. In a neighborhood of an extreme belief δ_θ, belief dynamics can be locally approximated using the weighted p-entropy, allowing us to determine whether δ_θ is locally stable or repelling based on the sign of $\mathcal{D}_p^w(\cdot, \cdot)$. Assumption 2 ensures that neighborhoods of extreme beliefs are reachable in a uniformly bounded number of steps, with probability bounded away from zero. This allows local dominance properties to be extended to global, almost-sure conclusions.

Figure 1 illustrates the case $\Theta_p^w = \theta_2$. Since $\mathcal{D}_p^w(\theta_2, \theta_1) > 0$, the belief δ_{θ_1} is repelling: the process escapes any neighborhood of δ_{θ_1}. Once outside this neighborhood, irreducibility implies that the process eventually enters a neighborhood of δ_{θ_2}. Because $\mathcal{D}_p^w(\theta_1, \theta_2) < 0$, this neighborhood is locally stable, and the belief is captured with positive probability. Since all transition probabilities are uniformly bounded away from zero, standard Borel-Cantelli arguments imply almost sure convergence to δ_{θ_2}.

Example 4 (Harmonic-Mean Rule)**.** Consider the same environment as in Example 1, but suppose beliefs are aggregated according to the *harmonic mean*:

$$F(\mu_1, \mu_2)(\theta) \propto 1 / \left[\frac{w_1}{\mu_1(\theta)} + \frac{w_2}{\mu_2(\theta)} \right].$$

This corresponds to $p = -1$. As in Example 3, one can show that when p_1 is sufficiently large, $\Theta_{-1}^w = \Theta = \{G, B\}$. Theorem 1 implies that beliefs can converge to both δ_G and δ_B, each with positive probability. Proposition 1 strengthens this result by showing that beliefs can *only* converge to these two limits.

4.2.2 Multiple-State Cases

We now turn to environments with more than two states. In multi-state settings, we can again establish almost sure convergence under suitable dominance conditions. The key object is the following.

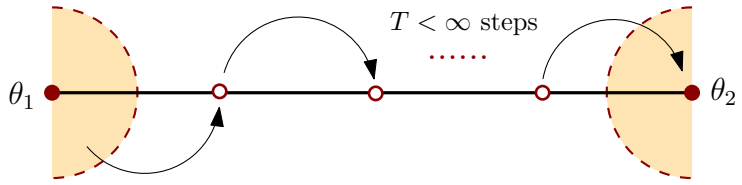

Figure 1: Illustration of Dynamics When $\Theta_p^w = \{\theta_2\}$

Definition 5 (Undominated Sets)**.** For each state $\theta \in \Theta$, define its *undominated set* as

$$\mathcal{U}_\theta \equiv \left\{\theta' \neq \theta : \theta \not\succeq_p^w \theta'\right\},$$

that is, the set of states that are not dominated by θ under the weighted p-entropy.

Intuitively, $\mathcal{U}_\theta$ collects the states that can locally challenge θ in belief dynamics.

Complete Rankings of Undominated Sets. We first consider the case in which undominated sets admit a complete ranking.

Proposition 2 (Multiple States: Complete Ranking of $\mathcal{U}_\theta$)**.** *Suppose that Assumption 3 holds. If there exists an ordering of states $\{\theta_1, ..., \theta_k\}$ such that*

$$\emptyset = \mathcal{U}_{\theta_1} \subset \mathcal{U}_{\theta_2} \subset ... \subset \mathcal{U}_{\theta_k}.$$

then beliefs converge $\mathbb{P}$-almost surely. Moreover,

$$\mathbb{P}\left(\cup_{\theta \in \Theta_p^w} \{\mu_{i,t} \to \delta_\theta \text{ for all } i\}\right) = 1,$$

and for each $\theta \in \Theta_p^w$, $\mathbb{P}\left(\mu_{i,t} \to \delta_\theta \text{ for all } i\right) > 0$.

This proposition implies that when undominated sets can be completely ranked, beliefs converge almost surely to point-mass beliefs on the minimizers of weighted p-entropy; i.e., states with empty undominated sets. The logic parallels that of Frick et al. (2023): a complete ranking rules out dominance cycles within subsimplices of the belief simplex, allowing local dominance relations to extend to global convergence. When Θ_p^w is a singleton, Proposition 2 yields global stability, analogous to their Proposition 1. More generally, the result also covers cases with multiple minimizers, in which convergence remains almost sure but is path-dependent.

Figure 2 illustrates several dominance configurations.[26] In panel (a), the undominated

[26] An arrow from θ to θ' indicates that $\theta' \in \mathcal{U}_\theta$; equivalently, $\mathcal{D}_p^w(\theta', \theta) > 0$, so that δ_θ is locally repelling in the direction of θ'.

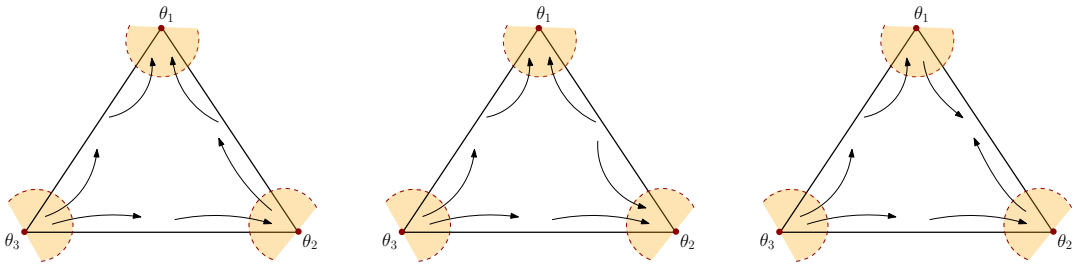

(a) $\Theta_p^w = \{\theta_1\}$ (b) $\Theta_p^w = \{\theta_1, \theta_2\}$ (c) $\Theta_p^w = \emptyset$

Figure 2: Belief Dynamics With Multiple States

sets satisfy

$$\mathcal{U}_{\theta_1} = \emptyset, \quad \mathcal{U}_{\theta_2} = \{\theta_1\}, \quad \mathcal{U}_{\theta_3} = \{\theta_1, \theta_2\},$$

forming a complete ranking. Since $\Theta_p^w = \{\theta_1\}$, beliefs converge almost surely to δ_{θ_1}. In panel (b), a complete ranking also holds, but $\Theta_p^w = \{\theta_1, \theta_2\}$. Beliefs still converge almost surely, but the limit may be either δ_{θ_1} or δ_{θ_2}, each with positive probability.

Example 5 (Weak Group Irrationality)**.** There are three states $\Theta = \{\alpha, \beta, \gamma\}$ and two signals $S = \{H, L\}$. Signals are i.i.d. with distribution $(1/2, 1/2)$. There are two individuals whose model perceptions $\left(\hat{l}_1, \hat{l}_2\right)$ satisfy

$\hat{l}_1(s\|\theta)$	H	L
α	$1-\varepsilon_1$	ε_1
β	$1/2$	$1/2$
γ	$2/3$	$1/3$

$\hat{l}_2(s\|\theta)$	H	L
α	$1/2$	$1/2$
β	$1-\varepsilon_2$	ε_2
γ	$2/3$	$1/3$

,

where $\varepsilon_2 > \varepsilon_1 > 0$ is a small number. Beliefs are aggregated via a power-mean rule

$$F(\mu_1, \mu_2)(\theta) \propto \left[\frac{1}{2}\mu_1(\theta)^p + \frac{1}{2}\mu_2(\theta)^p\right]^{1/p}.$$

For sufficiently small $|p|$, weighted p-entropy approximates weighted relative entropy, yielding

$$\mathcal{U}_\alpha = \{\beta, \gamma\}, \quad \mathcal{U}_\beta = \{\gamma\}, \quad \mathcal{U}_\gamma = \emptyset.$$

The ranking condition holds, so beliefs converge almost surely to δ_γ. Yet if individuals learned independently, their asymptotic beliefs would differ: agent 1 would converge to β, agent 2 to α. Social learning instead leads both agents to γ, a state that best fits neither agent's perceived signal structure, so weak group irrationality occurs.

Incomplete Ranking of Undominated Sets. Proposition 1 characterizes almost-sure convergence when undominated sets can be completely ranked. In many applications, especially when the state space is large, such a complete ranking may fail to exist. For example, in Figure 2(c),

$$\mathcal{U}_{\theta_1} = \{\theta_2\}, \quad \mathcal{U}_{\theta_2} = \{\theta_1\}, \quad \mathcal{U}_{\theta_3} = \{\theta_1, \theta_2\}.$$

In this case, the undominated sets cannot be completely ranked by set inclusion. Theorem 1 then implies only that beliefs oscillate almost surely (since $\Theta_p^w = \emptyset$), but it does not pin down which states receive positive probability in the long run. The next result provides a sharper characterization in this setting.

Definition 6 (Asymptotic Support)**.** For individual i, the set $\Theta_\infty^i \subset \Theta$ is called the *asymptotic support* of $\{\mu_{i,t}\}$ if it is the minimal set such that $\mu_{i,t}\left(\Theta \backslash \Theta_\infty^i\right) \to 0$.

That is, the asymptotic support is the smallest set of states such that the belief mass outside it vanishes asymptotically. We have the following result.[27]

Proposition 3 (Multiple States: Incomplete Ranking of $\mathcal{U}_\theta$)**.** *Suppose that Assumption 3 holds. Let $\hat{\Theta} \subset \Theta$ satisfy:*

(i) The collection $\left\{\mathcal{U}_\theta \cap \hat{\Theta}\right\}_{\theta \in \hat{\Theta}}$ can be completely ranked by set inclusion;
(ii) For every $\theta \in \hat{\Theta}$ with $\mathcal{U}_\theta \cap \hat{\Theta} = \emptyset$, we have $\mathcal{U}_\theta \neq \emptyset$.

Then, for every individual i, the set $\Theta \backslash \hat{\Theta}$ contains at least one (possibly path-dependent) element in Θ_∞^i $\mathbb{P}$-almost surely.

Proposition 3 provides a partial characterization of long-run beliefs when undominated sets do not admit a complete ranking. Its content can be summarized by a simple "leakage" logic. Suppose that beliefs were to asymptotically concentrate on a subset of states $\hat{\Theta}$. Condition (i) ensures that, within this restricted state space, the undominated sets admit a complete ranking, so local dominance relations would force convergence to a state that is undominated *within* $\hat{\Theta}$. Condition (ii), however, rules this out by requiring that every such internally undominated state is still dominated by at least one state *outside* $\hat{\Theta}$. As a result, belief mass cannot remain confined to $\hat{\Theta}$ in the long run and must eventually "leak" to states outside the subset, implying that at least one state in $\Theta \backslash \hat{\Theta}$ belongs to the asymptotic support almost surely.

[27] A related characterization appears in Beker and Chen (2023). That paper studies consumption dynamics in a general equilibrium model rather than learning, and the dominance relations used there differ substantially from those considered here.

Figure 2 (c) provides a concrete illustration. Take $\hat{\Theta} = \{\theta_2, \theta_3\}$. Then,

$$\mathcal{U}_{\theta_2} \cap \hat{\Theta} = \emptyset, \quad \mathcal{U}_{\theta_3} \cap \hat{\Theta} = \{\theta_2\},$$

so the undominated sets restricted to $\hat{\Theta}$ admit a complete ranking, satisfying condition (i). At the same time, $\mathcal{U}_{\theta_2} \neq \emptyset$, so θ_2 remains dominated by a state outside $\hat{\Theta}$, satisfying condition (ii). Proposition 3 therefore implies $\Theta \backslash \hat{\Theta} = \{\theta_1\}$ must belong to the asymptotic support almost surely. Repeating the argument with $\hat{\Theta} = \{\theta_1, \theta_3\}$ shows that θ_2 must also belong to the asymptotic support.

5 Group Irrationality and Learning Robustness

This section studies when social learning leads to correct aggregation and when it fails. I first show that group irrationality is a generic possibility in two complementary senses. Fixing almost any regular aggregation rule, there exist consistent model perceptions under which society fails to learn the truth (Section 5.1). Conversely, fixing any consistent model perceptions, there are, under mild conditions, aggregation rules for which correct learning also fails (Section 5.2). I then identify conditions under which correct learning is restored (Section 5.3). In particular, correct learning obtains when the aggregation rule is close to log-linear or when individuals sufficiently underreact to social information.

5.1 Does Learning Occur Under All Consistent Models?

This subsection shows that, fixing the aggregation rule, correct learning need not obtain even when all individuals' model perceptions are consistent. In particular, there often exist consistent model perceptions under which group irrationality emerges.

To separate restrictions on aggregation rules from those on model perceptions, I replace Assumption 2 with two more transparent conditions.

Assumption 4 (Intermediate Likelihood Ratio)**.** *The aggregation rule 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 5 (State-Specific Good News)**.** *For every individual i and every state $\theta \in \Theta$, there exists a signal $s^i_\theta \in S_i$ such that $\hat{l}(s^i_\theta|\theta) > \hat{l}(s^i_\theta|\theta')$ for all $\theta' \neq \theta$.*

Assumption 4 requires that the social likelihood ratio lies between the most optimistic and most pessimistic individual likelihood ratios. A canonical example is the power-mean rule in Example 2. Assumption 5 requires that, under each perceived model, every state admits a signal that is most indicative of that state. These two assumptions jointly imply irreducibility, so all previous results continue to apply. For convenience, I focus on aggregation rules and perceptions that satisfy the following regularity conditions.

Definition 7 (Regular Aggregation Rules & Perceptions)**.** We say that: (i) F is a regular aggregation rule if it satisfies Assumptions 1 and 4. (ii) $\hat{l} = (\hat{l}_1, ..., \hat{l}_n)$ is a profile of regular model perceptions if it satisfies Assumption 5.

I now formalize the notion that an aggregation rule may admit group irrationality under some consistent misspecifications.

Definition 8 (Group Irrationality Susceptibility)**.** An aggregation rule F is *susceptible to (weak) group irrationality* if there exist a true model l and a regular and consistent model perception $\hat{l}$ such that (weak) group irrationality emerges with $\mathbb{P}$-positive probability.

Susceptibility does not mean that correct learning never occurs. Rather, it means that for a given aggregation rule, one can find *some* consistent model perceptions and true models under which (weak) group irrationality occurs with positive probability.

Proposition 4 (Prevalence of Group Irrationality)**.** *Every regular aggregation rule is susceptible to weak group irrationality. Moreover, every regular aggregation rule with nonzero order is susceptible to group irrationality.*

Proposition 4 shows that group irrationality is generic across regular aggregation rules. This contrasts with the "wisdom of crowds" results in the non-Bayesian learning literature, which establish correct learning for certain aggregation rules, such as the power-average rule with $p \in [-1, 1]$, under correct specification (Molavi et al., 2018). It shows that once individuals' model perceptions may differ from the truth, even if they remain directionally correct, correct learning can break down.

Proposition 4 also suggests that aggregation rules of zero order may be special. I say that an aggregation rule is *immune to group irrationality* if it is not susceptible to group irrationality. That is, for all true models and all consistent and regular model perceptions, group irrationality occurs with probability zero.

Corollary 4 (Log-Linear Rule and Group Irrationality)**.** *The log-linear rule is immune to group irrationality. Conversely, if a regular aggregation rule is immune to group irrationality, then it must approximate the log-linear rule near extreme beliefs.*[28]

[28] That is, for any $\hat{\theta}$ and any $\theta \neq \hat{\theta}$, $F(\mu_1, ..., \mu_n) \sim \exp\left(\sum w_i \times \log \mu_i(\theta)\right)$ as $\min_i \mu_i(\hat{\theta}) \to 0$.

The converse direction follows directly from the second part of Proposition 4. To see the first claim, note that under the log-linear rule, social beliefs converge to the point mass on the state that minimizes the society's weighted relative entropy, provided the minimizer is unique. With consistent model perceptions, the true state θ^* uniquely minimizes each individual's relative entropy and therefore also minimizes the weighted sum. As a result, correct learning occurs almost surely under the log-linear rule whenever model perceptions are consistent.

Remark 2 (Relation to Molavi et al. (2018)). The key difference from Molavi et al. (2018) is that the present paper allows for model misspecification.

- Under *correct specification*, Molavi et al. (2018) show that correct learning obtains for $p \in [-1, 1]$. Intuitively, when p is small in magnitude, the aggregation rule introduces only mild distortions, which do not alter the dominance position of the true state and therefore preserve correct learning (see Proposition 10 in Appendix B).

- Under *misspecification*—even if those perceptions are consistent—these distortions can be amplified when $p \neq 0$, altering weighted p-entropy comparisons and thereby changing long-run learning outcomes.[29] As a result, misspecification combined with nonzero p can generate group irrationality.

There is, however, continuity between the two frameworks: as misspecification becomes small, the asymptotic dominance relations and belief dynamics in the present model converge to those in Molavi et al. (2018) (see Proposition 11 in Appendix B).

5.2 Does Learning Occur Under All Aggregation Rules?

The previous subsection shows that, fixing almost any regular aggregation rule, correct learning can fail even when all individuals hold consistent model perceptions. This subsection asks the converse question: fixing model perceptions, can correct learning be guaranteed across a wide class of aggregation rules? The goal is to understand whether there exist model perceptions under which correct learning is robust to the choice of aggregation rule.

Definition 9 (Robust Consistency)**.** Fix the true model l. We say that a profile of model perceptions $\hat{l} = \left(\hat{l}_1, ..., \hat{l}_n\right)$ is *robustly consistent* if, under these perceptions, society learns the true state almost surely for all regular aggregation rules.

[29] When $p = 0$, consistent model perceptions do not affect dominance relations: the weighted relative entropy is additive across individuals, so if the true state minimizes each individual's divergence, it also minimizes the social divergence. For $p \neq 0$, this additivity breaks down, allowing distortions to emerge and potentially reverse the ranking of states.

Robust consistency is a strong requirement. It demands not only that each individual's perception is consistent, but also that correct learning is insensitive to how beliefs are aggregated. To characterize when this holds, I introduce a notion of dominance that captures the most extreme aggregation distortions.

Definition 10 (Robust Dominance)**.** For all $\theta, \theta' \in \Theta$, define

$$\mathcal{D}_\infty(\theta', \theta) \equiv \mathbb{E}\left[\max_{i \in N} \log\left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)}\right)\right]. \tag{6}$$

We say that θ *robustly dominates* θ', written as $\theta \succeq_\infty \theta'$, whenever $\mathcal{D}_\infty(\theta', \theta) \leq 0$. Let Θ_∞ (resp. $\hat{\Theta}_\infty$) denote the set of states that (strictly) dominate all others under $\succeq_\infty$.

Robust dominance corresponds to the limiting dominance relation induced by weighted p-entropy as $p \to +\infty$. Intuitively, θ robustly dominates θ' if, even under the most distortionary aggregation that emphasizes the most optimistic individual, the perceived likelihood ratio still favors θ on average.

Using this concept, robust consistency admits a simple characterization.

Proposition 5. *Model perception $\hat{l}$ is robustly consistent if $\theta^* \in \hat{\Theta}_\infty$ and only if $\theta^* \in \Theta_\infty$.*

Proposition 5 provides an almost necessary and sufficient condition for correct learning to be robust across aggregation rules. A key implication is that robust dominance depends on the *maximum* perceived likelihood ratio across individuals. As a result, robust consistency becomes harder to satisfy as the number of individuals increases. This observation leads to an interesting implication: robust learning easily fails in large societies.

Example 6 (Larger Society, Harder to Learn)**.** Consider the binary-state environment of Example 1. Signals are i.i.d. with true probability $p^* > 1/2$, and there are n individuals with identical model perception $\hat{p} > 1/2$. The true state is G. By Proposition 5, $\hat{p}$ is robustly consistent if and only if

$$\mathcal{D}_\infty(B, G) = (1 - 2p^{*n}) \times \log\left(\frac{\hat{p}}{1-\hat{p}}\right) < 0. \tag{7}$$

Two implications follow. First, no model perception is robustly consistent once the society is sufficiently large: if $n > \log_{p^*}(1/2)$, the inequality (7) fails for all feasible $\hat{p}$. Second, for a fixed society size, robust consistency requires increasingly informative signals. Indeed, the condition reduces to $p^* > \sqrt[n]{1/2} \equiv p(n)$, which converges to one as $n \to \infty$.

Example 6 does not claim that correct learning necessarily fails under a fixed aggregation rule as the society grows. Rather, it shows that as the society expands, it becomes easier to find *some* aggregation rule under which learning breaks down. In large societies, robustly consistent perceptions may not exist at all. Below is the general statement.

Corollary 5 (No Robust Learning in Large Societies)**.** *Suppose individuals receive i.i.d. signals with DGP l and share a homogeneous model perception $\hat{l}$. Then, for any feasible l and $\hat{l}$, there exists $n_0 < \infty$ such that $\hat{l}$ is not robustly consistent for all $n \geq n_0$.*[30]

Corollary 5 further contrasts the classical "wisdom of crowds" intuition by showing that a larger society can make correct learning more difficult. The intuition is that robust dominance is governed by the most extreme individual. As the society grows, extreme likelihood ratios become more likely to arise, making it increasingly difficult for the true state to dominate all alternatives uniformly across aggregation rules.

5.3 When Does Correct Learning Occur?

The previous discussion provides negative results showing that learning can collapse under some aggregation rules or under some consistent model perceptions. However, these results rely on allowing a broad class of aggregation rules or perceptions. A natural question is whether correct learning can be restored once attention is restricted along either dimension.

This subsection focuses on the binary-state case.[31] The following proposition provides sufficient conditions under which correct learning is restored.

Proposition 6 (Conditions for Correct Learning)**.** *Suppose $|\Theta| = 2$. Then,*

(i) Fixing any consistent perception $\hat{l}$ and weights w, there exists $\varepsilon > 0$ such that whenever $|p| < \varepsilon$, correct learning occurs $\mathbb{P}$-almost surely.

(ii) Fixing any regular aggregation rule F, there exists a consistent perception $\hat{l}$ such that correct learning occurs $\mathbb{P}$-almost surely.

Proposition 6 shows that correct learning can be achieved once either aggregation rules or model perceptions are appropriately restricted.

First, fixing any consistent model perception, correct learning obtains when the aggregation order p is sufficiently close to zero. This reflects the fact that the log-linear rule ($p = 0$) is immune to group irrationality, together with the continuity of aggregation rules in p.

[30] Here, "feasible" means that l and $\hat{l}$ satisfy the regularity conditions in Section 2. The assumption of homogeneous model perceptions can be relaxed; e.g., we can allow individuals' model perceptions to come from any finite set.

[31] Focusing on binary states allows for a complete characterization of asymptotic beliefs; see Section 4.2. Similar insights are conjectured to extend to more general state spaces.

Second, fixing any aggregation rule F, there exist consistent model perceptions under which correct learning occurs almost surely. The proposition, however, is purely existential and does not identify which perceptions restore learning or why. To make this mechanism transparent, I now focus on a structured class of consistent model perceptions.

Learning under α-Consistent Model Perceptions

Throughout, I focus on a simple binary environment where $\Theta = \{\theta_1, \theta_2\}$ and $S_i = \{s_1^i, s_2^i\}$.[32] Suppose that signal s_k^i is more likely to occur in state θ_k under the true DGP l_i. I introduce the following class of perceptions.

Definition 11 (α-Consistent Model Perceptions)**.** For every individual i, the model perception $\hat{l}_i$ is *α-consistent* if

$$\hat{l}_i\left(s|\theta_1\right) = \begin{cases} c\pi^\alpha & s = s_1^i \\ c\left(1-\pi\right)^\alpha & s = s_2^i \end{cases} \text{ and } \hat{l}_i\left(s|\theta_2\right) = \begin{cases} c\left(1-\pi\right)^\alpha & s = s_1^i \\ c\pi^\alpha & s = s_2^i \end{cases}, \tag{8}$$

where $\pi \in (1/2, 1)$, $\alpha > 0$ and c is the normalization term.

Here, α governs perceived signal informativeness. When $\alpha < 1$, individuals underreact to private information; as $\alpha \to 0$, perceived likelihood ratios approach one and signals become nearly uninformative. For simplicity, I assume that individuals share a common α.

Proposition 7 (Learning Under α-Consistent Perceptions)**.** *Suppose $|\Theta| = |S_i| = 2$ for all i. Fixing any $\overline{p} \in (0, +\infty)$ and for all regular F satisfying $|p\left(F\right)| \leq \overline{p}$, there exists some $\overline{\alpha} > 0$ such that correct learning occurs $\mathbb{P}$-almost surely if individuals have α-consistent model perceptions with $\alpha \in (0, \overline{\alpha})$.*

Proposition 7 shows that correct learning can be restored by sufficiently dampening individuals' perceived signal informativeness. Two implications are worth highlighting. First, rather than requiring restrictions on a single aggregation rule, the result holds uniformly over any compact set of aggregation orders, regardless of how large that set is. Second, achieving correct learning requires α to be sufficiently small—that is, underreaction can help learning. Notably, the result imposes no restriction on the weight vector w: correct learning obtains for all possible weights once perceptions are sufficiently uninformative.

Figure 3 illustrates the learning dynamics under α-consistent perceptions. The middle hatched region (Region I) corresponds to pairs (α, p) for which correct learning occurs almost surely. The upper and lower regions (Regions II and III) correspond, respectively, to belief

[32] The insights can be extended to more general cases; see Section 6.

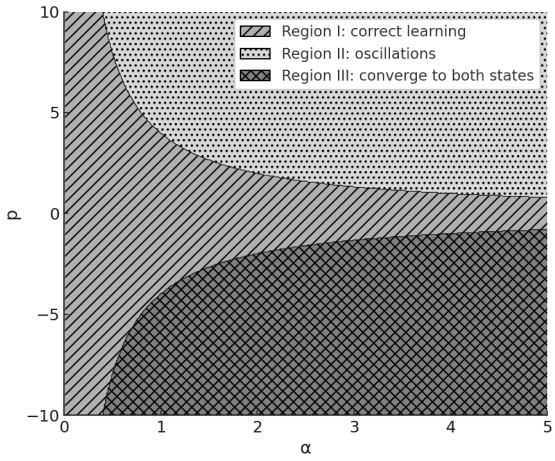

Figure 3: Dynamics Under α-Consistent Perceptions

Note: State space and signal space are both binary, $\Theta = \{\theta_1, \theta_2\}$ and $S = \{s_1, s_2\}$. There are two individuals with equal weights, $w_1 = w_2 = 1/2$. Signals are i.i.d. with $l^*(\theta_k|s_k) = 2/3$. Individuals' model perception is $\hat{l}(\theta_k|s_k) \propto (2/3)^\alpha$, where $\alpha > 0$.

oscillations and convergence to either state. As $\alpha \to 0$, the region supporting correct learning expands to include nearly all values of p. In particular, for any fixed interval $[\underline{p}, \overline{p}]$, there exists $\overline{\alpha} > 0$ such that correct learning occurs for all p in this interval whenever $\alpha < \overline{\alpha}$.

Remark 3 (The Benefits of Misspecification). Figure 3 also illustrates that misspecification can outperform correct specification. The vertical line $\alpha = 1$ corresponds to correct specification. Along this line, oscillations occur almost surely for large p, while multiple belief limits arise for sufficiently negative p. By contrast, when individuals underestimate signal informativeness ($\alpha < 1$), correct learning can occur even in parameter regions where it collapses under correct specification. In this sense, underreaction to private information can improve social learning outcomes.

From Proposition 7, we know that sufficiently small α can restore correct learning. The converse is also true: when α is large, correct learning may fail. To formalize this idea, I impose the following condition.

Assumption 6 (Binary i.i.d. Signals)**.** *For every individual i, $S_i = \{s_1, s_2\}$. Signals are i.i.d. across individuals with $l_i(s_k \mid \theta_k) = \pi^* > \frac{1}{2}$ for all i and k.*

This assumption ensures that signals are independent across individuals. If signals were perfectly correlated, social learning would reduce to individual learning, and correct learning would trivially obtain. Under independent signals, however, aggregation distortions can accumulate. The following result shows that sufficiently large α undermines learning.

Corollary 6 (Incorrect Learning Under Large α)**.** *Suppose that Assumption 6 holds, and individuals have identical α-consistent model perceptions. If the size of society satisfies $n > \log_{\pi^*}^{1/2}$, then for every $p \neq 0$, there exists $\alpha^* < +\infty$ such that correct learning does not $\mathbb{P}$-almost surely occur if $\alpha > \alpha^*$.*

Corollary 6 provides a partial converse to Proposition 7. When society has sufficiently many independent information sources, learning collapses if individuals perceive signals as overly informative. Taken together, these results show that incorrect social learning arises from overreaction to private information: *underreaction* (small α) can restore correct learning, whereas *overreaction* (large α) undermines it.

Ranking of Aggregation Rules

Figure 3 also suggests that aggregation rules can be partially ordered by how conducive they are to correct learning. Intuitively, learning is more likely to succeed under aggregation rules whose order is closer to zero. If $p_1 > p_2 > 0$ and correct learning occurs under (p_1, α), then it also occurs under (p_2, α). This intuition can be extended more generally.

Proposition 8 (Ranking of Aggregation Rules)**.** *Suppose $|\Theta| = 2$. Fix an aggregation rule F and consistent model perceptions $\hat{l} = \left(\hat{l}_1, ..., \hat{l}_n\right)$ that jointly satisfy Assumption 3. If correct learning occurs almost surely under F, then it also occurs almost surely under any F' that satisfies: (i) $w(F') = w(F)$; and (ii) $p(F')$ and $p(F)$ have the same sign, and $|p(F')| < |p(F)|$.*

This proposition shows that, holding weights fixed, aggregation rules whose order is closer to zero are weakly more favorable for correct learning, provided the sign of the order is unchanged. The restriction on weights is essential, since different weights directly affect $\mathcal{D}_p^w$ and would confound comparisons across aggregation rules. The requirement that the orders share the same sign can be relaxed. In particular, if correct learning is achieved under an aggregation rule F with $p(F) > 0$, then it also occurs under any rule F' that satisfies $p(F') \in (\hat{p}(F), p(F))$ for some cutoff $\hat{p}(F) < 0$. However, the cutoff $\hat{p}(F)$ generally depends on both the aggregation rule and the model perception, and there is no guarantee that $\hat{p}(F) + p(F) = 0$, since weighted p-entropy is not symmetric around zero.

5.4 Discussion and Intuition

This subsection provides intuition for *two distinct channels* through which group irrationality can arise in the model. The first channel operates through the aggregation rule, when the

order p places excessive weight on extreme beliefs. The second channel operates through individuals' model perceptions, when perceived signals are overly informative. Although these channels are analytically separable, they interact in shaping long-run learning outcomes.

(i) Group Irrationality From Extreme Aggregation (Large $|p|$)

When the aggregation order $|p|$ is large, the social belief becomes highly sensitive to extreme individual beliefs. Although such extreme beliefs occur with low probability, they receive disproportionate weight in the aggregation process, which can distort learning dynamics.

To see this, consider the case $p \to +\infty$. When the social belief v_t lies in a small neighborhood of a point mass δ_θ, the updating dynamics can be locally approximated by

$$v_{t+1}(\theta') \approx \left[\sum_i w_i \times \left(\frac{\hat{l}_i(s_{i,t}|\theta')}{\hat{l}_i(s_{i,t}|\theta)}\right)^p\right]^{1/p} \times v_t(\theta'). \tag{9}$$

As p grows, the power mean converges to the maximum across individuals, so the evolution of beliefs is effectively driven by the most optimistic individual at each date.

With independent signals and a large society, it becomes likely that *some* individual draws a signal that favors θ' over θ, making the maximum ratio exceed one and increasing $v_t(\theta')$ in the neighborhood of δ_θ. This mechanism can render each δ_θ locally repelling and generate persistent oscillations. Symmetrically, as $p \to -\infty$ the p-mean approaches the minimum, which can instead make multiple δ_θ locally attracting and lead to path-dependent convergence to different point masses.

(i) Group Irrationality From Overreaction to Signals (Large α)

A second route to learning failure arises not from extreme aggregation, but from overly informative perceived models. For expositional clarity, I illustrate this mechanism using α-consistent perceptions, under which the local approximation (9) becomes

$$v_{t+1}(\theta') \approx \left[\sum_i w_i \times \left(\frac{l_i(s_{i,t}|\theta')}{l_i(s_{i,t}|\theta)}\right)^{\alpha p}\right]^{1/p} \times v_t(\theta'),$$

so the effective distortion is governed by the product αp. When α is large, perceived likelihood ratios become more extreme, and even moderate nonzero p can overweight those extremes in aggregation. This amplifies local "pushes" away from the true state in dominance comparisons, making it easier for incorrect states to become locally attractive or for cycles to emerge. In contrast, when α is small (underreaction), likelihood ratios are com-

pressed toward one, which dampens these extreme influences and can restore the dominance relations that favor correct learning.

Summary

The key lesson of this section is that group irrationality can arise through two parallel amplification mechanisms. One operates at the social level: aggregation rules with large p overweight extreme individual beliefs. The other operates at the individual level: highly informative perceived models generate extreme posteriors. Either mechanism can destabilize learning on its own, and together they can reinforce each other. Correct learning is restored when these amplification forces are sufficiently muted—either because aggregation places limited weight on extremes, or because individuals underreact to private information.

6 Extension: Other Updating Rules in Private Learning

In the baseline model, individuals apply Bayes' rule when updating beliefs during the private learning stage. It is natural, however, to consider situations in which individuals deviate from the Bayesian paradigm. This section discusses how the paper's characterization can be extended to more general private learning rules.

6.1 Learning With Generalized Bayes' Rule

Consider an alternative setup in which the social learning stage remains unchanged, but individuals follow a *generalized Bayes' rule* in the private learning stage. In particular, individual 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 coincides with a signal distribution, this rule reduces to standard Bayesian updating. More generally, the generalized Bayes' rule accommodates a variety of non-Bayesian learning behaviors, including the following examples.

Example 7 (Power Posterior)**.** Suppose that the learning rule satisfies

$$\psi_i(s|\theta) = \hat{l}_i(s|\theta)^{\alpha}, \text{ where } \alpha > 0, \tag{10}$$

in which case individuals follow the power Bayes' rule, and the resulting posterior is referred

to as a *power posterior* (Walker and Hjort, 2001). As α decreases, individuals place greater weight on their prior beliefs. When $\alpha = 0$, individuals never update; as $\alpha \to +\infty$, posteriors concentrate almost entirely on the maximum-likelihood states.

Example 8 (Gibbs Posterior)**.** Suppose instead that the pseudo-likelihood function takes the form

$$\psi_i(s|\theta) = \exp(-\rho_i(s,\theta)), \tag{11}$$

where $\rho_i(s,\theta)$ is a loss function. The resulting posterior is known as the *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.

6.2 Characterization of Asymptotic Beliefs

All previous results extend naturally to this generalized setting. In particular, one can redefine relative entropy and weighted p-entropy by replacing the perceived likelihood $\hat{l}_i$ with the pseudo-likelihood function ψ_i. For example, define the *generalized relative entropy* as follows

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

and let the *weighted generalized relative entropy* be

$$\mathcal{I}^w(\theta;\psi) \equiv \sum_{i=1}^{n} w_i \times \mathcal{R}(\theta;\psi_i).$$

Group irrationality can then be defined using the minimizers of the weighted generalized relative entropy. Owing to the identical structure of the learning dynamics, all previous characterization results extend analogously.

Asymptotic Beliefs Under the Power Bayes' Rule

An important example is the power Bayes' rule, $\psi_i(s|\theta) = \hat{l}_i(s|\theta)^{\alpha}$, where $\hat{l}_i$ stands for individual i's model perception. I introduce the following assumption.

Assumption 7 (No Indifference at Zero)**.** $\mathcal{I}^w(\theta) \neq \mathcal{I}^w(\theta')$ *for all* $\theta \neq \theta'$.

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

Proposition 9 (Characterization Under the Power Bayes' Rule)**.** *Suppose further that Assumption 7 holds, and individuals follow the power Bayes' rule with parameter* α*. Then,*

there exists $\alpha_0 > 0$ *such that for every* i,

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

whenever the Bayesian power satisfies $\alpha \in (0, \alpha_0)$. [33]

When α is small, individuals are cautious in incorporating private signals into their posteriors. Proposition 9 shows that when all individuals are sufficiently cautious, beliefs converge almost surely to the state minimizing the weighted relative entropy. A special case arises when individuals hold α-consistent model perceptions: in this case, $\theta_0 = \theta^*$, and correct learning is obtained for sufficiently small α, as shown in Proposition 7.[34] More generally, the proposition allows for settings in which the entropy-minimizing state need not coincide with the true state. This proposition only focuses on the small-α case. When α is large, and when signals are independent across individuals, learning collapses by the same logic underlying Corollary 6.

7 Conclusion

This paper studies social learning when individuals update beliefs under model misspecification and society aggregates beliefs using a non-Bayesian rule. I introduce the concept of group irrationality to capture situations in which social learning outcomes diverge from those that would arise if individuals learned independently. Contrary to the classical "wisdom of crowds" view, group irrationality is prevalent across a wide range of standard aggregation rules, because social learning transmits not only information but also belief distortions. The analysis shows that long-run learning outcomes are governed by how extreme perceived likelihood ratios are amplified through aggregation, as captured by the weighted p-entropy. When such amplification is strong, social beliefs may oscillate or converge to incorrect states; when it is sufficiently muted—either through less extreme aggregation or more cautious private updating—correct learning is restored. An important direction for future research is to study group irrationality in broader learning environments, including alternative aggregation mechanisms and active learning settings in which data acquisition is endogenous.

[33] By Assumption 3, there is a unique minimizer of the weighted relative entropy, so θ_0 is unique.

[34] Here, employing the power Bayes' rule with power α is equivalent to having an α-consistent model perception; see Bohren and Hauser (2024) for the equivalence between non-Bayesian rules and model misspecification.

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, 6

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*. 7

Arieli, Itai and Manuel Mueller-Frank (2019) "Multidimensional social learning," *The Review of Economic Studies*, 86 (3), 913–940. 6

——— (2021) "A general analysis of sequential social learning," *Mathematics of Operations Research*, 46 (4), 1235–1249. 6

Bala, Venkatesh and Sanjeev Goyal (1998) "Learning from neighbours," *The Review of Economic Studies*, 65 (3), 595–621. 6

Banerjee, Abhijit and Olivier Compte (2024) "Consensus and disagreement: Information aggregation under (not so) naive learning," *Journal of Political Economy*, 132 (8), 2790–2829. 6

Banerjee, Abhijit V (1992) "A simple model of herd behavior," *The Quarterly Journal of Economics*, 107 (3), 797–817. 2, 6

Beker, Pablo F and Jaden Yang Chen (2023) "If You're NOT So Smart, Why Are You Rich? Robust Market Selection with General Recursive Preferences," *Available at SSRN 4616770*. 21

Berk, Robert H (1966) "Limiting behavior of posterior distributions when the model is incorrect," *The Annals of Mathematical Statistics*, 37 (1), 51–58. 5, 10

Bhattacharya, Anirban, Debdeep Pati, and Yun Yang (2019) "Bayesian fractional posteriors," *The Annals of Statistics*, 47 (1), 39 – 66. 14

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, 6

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. 32

Bohren, J Aislinn (2016) "Informational herding with model misspecification," *Journal of Economic Theory*, 163, 222–247. 7

Bohren, J Aislinn and Daniel N Hauser (2021) "Learning with heterogeneous misspecified models: Characterization and robustness," *Econometrica*, 89 (6), 3025–3077. 4, 7

——— (2024) "Behavioral Foundations of Model Misspecification." 11, 33

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. 6

Chen, Jaden Yang (2022) "Sequential Learning under Informational Ambiguity," *Available at SSRN 3480231*. 7

Clemen, Robert T (1989) "Combining forecasts: A review and annotated bibliography," *International journal of forecasting*, 5 (4), 559–583. 11

Dasaratha, Krishna and Kevin He (2020) "Network structure and naive sequential learning," *Theoretical Economics*, 15 (2), 415–444. 6

DeGroot, Morris H (1974) "Reaching a consensus," *Journal of the American Statistical Association*, 69 (345), 118–121. 6, 8

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. 6

Eyster, Erik and Matthew Rabin (2010) "Naive herding in rich-information settings," *American Economic Journal: Microeconomics*, 2 (4), 221–43. 6

Frick, Mira, Ryota Iijima, and Yuhta Ishii (2020) "Misinterpreting others and the fragility of social learning," *Econometrica*, 88 (6), 2281–2328. 7

——— (2023) "Belief convergence under misspecified learning: A martingale approach," *The Review of Economic Studies*, 90 (2), 781–814. 7, 14, 19

Gagnon-Bartsch, Tristan and Antonio Rosato (2024) "Quality is in the eye of the beholder: taste projection in markets with observational learning," *American Economic Review*, 114 (11), 3746–3787. 7

Gale, Douglas and Shachar Kariv (2003) "Bayesian learning in social networks," *Games and Economic Behavior*, 45 (2), 329–346. 6

Genest, Christian and James V Zidek (1986) "Combining probability distributions: A critique and an annotated bibliography," *Statistical Science*, 1 (1), 114–135. 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, 6, 11

Guarino, Antonio and Philippe Jehiel (2013) "Social learning with coarse inference," *American Economic Journal: Microeconomics*, 5 (1), 147–174. 6

Hardy, G. H., J. E. Littlewood, and G. Pólya (1952) *Inequalities*, Cambridge: Cambridge University Press. 8, 38

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. 7

Harel, Matan, Elchanan Mossel, Philipp Strack, and Omer Tamuz (2021) "Rational groupthink," *The Quarterly Journal of Economics*, 136 (1), 621–668. 2, 6

Huang, Wanying (2023) "Learning about informativeness." 7

Huang, Wanying, Philipp Strack, and Omer Tamuz (2024) "Learning in repeated interactions on networks," *Econometrica*, 92 (1), 1–27. 2, 6

Jadbabaie, Ali, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi (2012) "Non-Bayesian social learning," *Games and Economic Behavior*, 76 (1), 210–225. 2, 6, 12

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. 32

Kartik, Navin, SangMok Lee, Tianhao Liu, and Daniel Rappoport (2024) "Beyond unbounded beliefs: How preferences and information interplay in social learning," *Econometrica*, 92 (4), 1033–1062. 2, 6

Le Bon, Gustave (1895) *The crowd: A study of the popular mind*. 3

Li, Wei and Xu Tan (2020) "Locally Bayesian learning in networks," *Theoretical Economics*, 15 (1), 239–278. 6

McMahan, Brendan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas (2017) "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*, 1273–1282, PMLR. 11

Molavi, Pooya, Alireza Tahbaz-Salehi, and Ali Jadbabaie (2018) "A theory of non-Bayesian social learning," *Econometrica*, 86 (2), 445–490. 4, 6, 8, 11, 23, 24, 53, 54

Moody's (2010) "Best Practices Guidance for the Credit Rating Process,"Technical report, Moody's Investors Service. 11

Mossel, Elchanan, Allan Sly, and Omer Tamuz (2014) "Asymptotic learning on Bayesian social networks," *Probability Theory and Related Fields*, 158 (1), 127–157. 2

——— (2015) "Strategic learning and the topology of social networks," *Econometrica*, 83 (5), 1755–1794. 6

Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille (2009) "Informational externalities and emergence of consensus," *Games and Economic Behavior*, 66 (2), 979–994. 6

Smith, Lones and Peter Sørensen (2000) "Pathological outcomes of observational learning," *Econometrica*, 68 (2), 371–398. 2, 6

Sobel, Joel (2014) "On the relationship between individual and group decisions," *Theoretical Economics*, 9 (1), 163–185. 12

Surowiecki, James (2004) *The wisdom of crowds*: Anchor. 2

Timmermann, Allan (2006) "Forecast combinations," *Handbook of economic forecasting*, 1, 135–196. 11

Walker, Stephen and Nils Lid Hjort (2001) "On Bayesian consistency," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 63 (4), 811–821. 32

A Proofs

A.1 Proof of Lemma 1

Proof. (i) Denote by $x_i = \hat{l}_i(s_i|\theta')/\hat{l}_i(s_i|\theta)$. By L'Hôpital's rule,

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

(ii) Let $x_i = \hat{l}_i(s_i\mid\theta')/\hat{l}_i(s_i\mid\theta)$. Suppose $p > q$. Then

$$\mathcal{D}_p^w(\theta',\theta) = \frac{1}{p}\mathbb{E}\log\left(\sum_{i=1}^n w_i x_i^p\right) = \mathbb{E}\log M_p(x),$$

where $M_p(x)\equiv\left(\sum_{i=1}^n w_i x_i^p\right)^{1/p}$ is the weighted power mean, with the convention $M_0(x) = \prod_i x_i^{w_i}$. The power mean $M_p(x)$ is monotone increasing in p for all real p (see, e.g., Hardy et al. (1952)). Hence, for any $p > q$, $M_p(x)\geq M_q(x)$, with strict inequality unless the x_i's are identical. Taking logarithms and expectations yields

$$\mathcal{D}_p^w(\theta',\theta) = \mathbb{E}\log M_p(x) > \mathbb{E}\log M_q(x) = \mathcal{D}_q^w(\theta',\theta),$$

where the strict inequality follows from the assumption that $\hat{l}_i(\cdot\mid\theta)\neq\hat{l}_i(\cdot\mid\theta')$ for $\theta\neq\theta'$. □

A.2 Proof of Theorem 1

A.2.1 Proof of the "Only If" Direction

I now prove the "only if" direction. That is, whenever beliefs converge, they must converge to a point mass on some minimizer of the weighted p-entropy.

Lemma 2 (Limit Belief)**.** *Conditional on that μ_t converges, both the social belief v_t and each individual's posterior $\mu_{i,t}$ converge to the same Dirac belief δ_θ except on $\mathbb{P}$-null events.*

Proof. Suppose that $\mu_t\to\mu_\infty = (\mu_{1,\infty},...,\mu_{n,\infty})$. By continuity of F, $v_t = F(\mu_{t-1})\to F(\mu_\infty)\equiv v_\infty$. Since individual posteriors satisfy $\mu_{i,t} = BU_i(v_t, s_{i,t})$, and there exist signals that alter the likelihood ratio between any two states, we must have $v_\infty = \delta_\theta$ for some θ. Therefore, $\mu_{i,\infty} = v_\infty = \delta_\theta$ for all i. □

Next I show that beliefs can only occur at states in Θ_p^w. The key step is a local approximation of belief dynamics near a point mass.

Lemma 3 (Local Approximation)**.** *Fix $\theta \in \Theta$. For every $\delta > 0$, there exists $\varepsilon > 0$ such that for all $\theta' \neq \theta$,*

$$\frac{v_{t+1}(\theta')}{v_t(\theta') \times \left[\sum w_i \times \left(\frac{\hat{l}_i(s_{i,t}|\theta')}{\hat{l}_i(s_{i,t}|\theta)}\right)^p\right]^{1/p}} \in [1-\delta, 1+\delta], \tag{12}$$

whenever $v_t(\theta) \in B_\varepsilon(\delta_\theta) \equiv \{v : v(\theta) > 1-\varepsilon\}$.

Proof. By Assumption 1, as $v_t(\theta) \to 1$ we have

$$v_{t+1}(\theta') \sim \left[\sum_{i=1}^n w_i \left(\frac{v_t(\theta')\hat{l}_i(s_{i,t} \mid \theta')}{\sum_{\hat{\theta}\in\Theta} v_t(\hat{\theta})\hat{l}_i(s_{i,t} \mid \hat{\theta})}\right)^p\right]^{1/p} \sim v_t(\theta')\left[\sum_{i=1}^n w_i \left(\frac{\hat{l}_i(s_{i,t} \mid \theta')}{\hat{l}_i(s_{i,t} \mid \theta)}\right)^p\right]^{1/p}, \tag{13}$$

where $a_t \sim b_t$ means $a_t/b_t \to 1$. The claim follows immediately. □

Then I prove the main result.

Lemma 4 (Converge to the Minimizer)**.** *Conditional on the event that v_t converges to some Dirac belief δ_θ, we must have $\theta \in \Theta_p^w$ except on $\mathbb{P}$-null events.*

Proof. Suppose toward a contradiction that $v_t \to \delta_\theta$ but $\theta \notin \Theta_p^w$. Then there exists $\theta' \neq \theta$ such that

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

Since $v_t(\theta') \to 0$, there exists $C > 0$ and t_0 such that $\log v_t(\theta') \leq -C$ for all $t \geq t_0$. By Lemma 3, for any $\varepsilon > 0$ and all sufficiently large t,

$$\log v_{t+1}(\theta') \geq \log v_t(\theta') + \frac{1}{p}\log\left(\sum_{i=1}^n w_i \left(\frac{\hat{l}_i(s_{i,t} \mid \theta')}{\hat{l}_i(s_{i,t} \mid \theta)}\right)^p\right) - \varepsilon. \tag{14}$$

Iterating (14) from t_0 to $t-1$ and dividing by $t - t_0$ yields

$$\frac{1}{t-t_0}\log v_t(\theta') \geq \frac{1}{t-t_0}\sum_{t'=t_0}^{t-1}\frac{1}{p}\log\left(\sum_{i=1}^n w_i \left(\frac{\hat{l}_i(s_{i,t'} \mid \theta')}{\hat{l}_i(s_{i,t'} \mid \theta)}\right)^p\right) - \varepsilon + o(1).$$

By the strong law of large numbers, the time average converges almost surely to $\mathcal{D}_p^w(\theta', \theta)$. Choosing $\varepsilon > 0$ sufficiently small, we obtain

$$\liminf_{t\to\infty} \frac{1}{t}\log v_t(\theta') > 0 \quad \text{a.s.},$$

which implies $\log v_t(\theta') \to +\infty$ almost surely. This contradicts $v_t(\theta') \to 0$. Therefore, convergence to δ_θ is impossible unless $\theta \in \Theta_p^w$. □

The "only if" direction of Theorem 1 then follows directly previous lemmas.

A.2.2 Proof of the "If" Direction

I now prove the converse direction. That is, if a state is a strict minimizer of the weighted p-entropy, beliefs will converge to the point mass on that state with positive probability.

Lemma 5. *Let X be a bounded random variable. Suppose $\mathbb{E}\log X < 0$, then there exists some $\rho > 0$ such that $\mathbb{E}X^\rho < 1$.*

Proof. Since X is bounded, the dominated convergence theorem implies

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

Hence, for sufficiently small $\rho > 0$, we have $\mathbb{E}X^\rho < 1$. □

The next lemma shows that each entropy minimizer is locally attracting.

Lemma 6 (Minimizer is Local Attracting)**.** *For any $\theta \in \hat{\Theta}_p^w$ and any $\delta > 0$, there exists $\varepsilon > 0$ such that if $v_1 \in B_\varepsilon(\delta_\theta)$, then $\mathbb{P}(v_t \to \delta_\theta) > 1 - \delta$.*

Proof. Fix $\theta \in \hat{\Theta}_p^w$. Then for all $\theta' \neq \theta$,

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

We want to show that $\mu_{i,t}$ will converge to δ_θ with positive probability. For expositional clarity, the argument is presented for $p > 0$; the cases $p \leq 0$ follow by symmetric reasoning.

Step 1: Construction of a supermartingale. Lemma 5 and (15) imply that there exist constants $\rho, q > 0$ such that

$$\mathbb{E}\left[\sum_{i=1}^n w_i \left(\frac{\hat{l}_i(s_i \mid \theta')}{\hat{l}_i(s_i \mid \theta)}\right)^p\right]^\rho < \frac{1}{1+q} \quad \text{for all } \theta' \neq \theta. \tag{16}$$

By Lemma 3, there exists $\varepsilon_1 > 0$ such that whenever $v_t \in B_{\varepsilon_1 \times |\Theta|}(\delta_\theta)$,

$$v_{t+1}(\theta')^{p\rho} \leq (1+q) \times v_t(\theta')^{p\rho}\left[\sum_{i=1}^n w_i \left(\frac{\hat{l}_i(s_{i,t} \mid \theta')}{\hat{l}_i(s_{i,t} \mid \theta)}\right)^p\right]^\rho. \tag{17}$$

Let $r_1 \equiv \varepsilon_1 \times |\Theta|$ and define

$$T \equiv \inf\{t : v_t \notin B_{r_1}(\delta_\theta)\}, \qquad Y_t(\theta') \equiv v_{t\wedge T}(\theta')^{p\rho}.$$

Then $\{Y_t(\theta')\}$ is a nonnegative supermartingale and hence converges almost surely with $\mathbb{E}[Y_T(\theta')] \leq Y_1(\theta')$.

Step 2: Staying in the neighborhood. Suppose $v_1 \in B_{\varepsilon_2}(\delta_\theta)$ for some $\varepsilon_2 > 0$. Then $Y_1(\theta') \leq \varepsilon_2^{p\rho}$. By Markov's inequality,

$$\mathbb{P}(Y_T(\theta') \geq \varepsilon_1^{p\rho}) \leq \frac{\mathbb{E}[Y_T(\theta')]}{\varepsilon_1^{p\rho}} \leq \left(\frac{\varepsilon_2}{\varepsilon_1}\right)^{p\rho}. \tag{18}$$

If $T < \infty$, then necessarily there exists some $\theta' \neq \theta$ such that $v_T(\theta') \geq \varepsilon_1$, and hence $Y_T(\theta') \geq \varepsilon_1^{p\rho}$. Therefore,

$$\begin{aligned}\mathbb{P}(T < \infty) &\leq \mathbb{P}\left(\bigcup_{\theta'\neq\theta}\{Y_T(\theta') \geq \varepsilon_1^{p\rho}\}\right) \\ &\leq \sum_{\theta'\neq\theta}\mathbb{P}(Y_T(\theta') \geq \varepsilon_1^{p\rho}) \leq |\Theta|\left(\frac{\varepsilon_2}{\varepsilon_1}\right)^{p\rho}.\end{aligned} \tag{19}$$

Choosing ε_2 sufficiently small ensures $\mathbb{P}(T = \infty) \geq 1 - \delta$. On thes event $\{T = \infty\}$, v_t remains in $B_{r_1}(\delta_\theta)$ forever. Since $Y_t(\theta')$ converges almost surely for each $\theta' \neq \theta$, the social belief v_t converges almost surely as well. By Assumption 2, the only possible limit is δ_θ. Hence, $v_t \to \delta_\theta$ with probability at least $1 - \delta$. □

Finally, by Assumption 2, for any initial belief $v_1 \in \Delta_{++}(\Theta)$ and any $\varepsilon > 0$, v_t enters $B_\varepsilon(\delta_\theta)$ with a strictly positive probability. By the lemma above and the Markov property, this implies that beliefs converge to δ_θ with positive probability, completing the proof of the "if" direction.

A.3 Proof of Proposition 1

A.3.1 Auxiliary Results

Below I first state some auxiliary results. Notice that these results hold for general state space Θ, so they will also be invoked when proving results for multiple states.

Lemma 7 (Attracting States in Θ_p^w)**.** *Suppose Assumption 3 holds. Then, for any $\theta \in \Theta_p^w$ and any $\delta > 0$, there exists $\varepsilon > 0$ such that $\mathbb{P}(v_t \to \delta_\theta) > 1 - \delta$ whenever $v_0 \in B_\varepsilon(\delta_\theta)$.*

Proof. It comes directly from Lemma 6 and the fact that $\Theta_p^w = \hat{\Theta}_p^w$ under Assumption 3. □

Lemma 8 (Repelling States Outside Θ_p^w)**.** *For any $\varepsilon > 0$ and any $t \in \mathbb{T}$, we have:*

$$\forall \theta \notin \Theta_p^w : \quad \mathbb{P}\left(v_\tau(\theta) \notin B_\varepsilon(\delta_\theta) \text{ for some } \tau \geq t | v_t \in B_\varepsilon(\delta_\theta)\right) = 1.$$

In other words, for any $\theta \notin \Theta_p^w$, v_t almost surely escapes from any neighborhood of δ_θ.

Proof. Fix $\theta \notin \Theta_p^w$ and $\varepsilon > 0$. Suppose, toward a contradiction, that with positive probability the belief process $\{v_t\}$ remains in $B_\varepsilon(\delta_\theta)$ for all sufficiently large t. On this event, there are two possibilities.

Case (i): v_t converges. Then $v_t \to \delta_\theta$. This is impossible: by Theorem 1, convergence to δ_θ when $\theta \notin \Theta_p^w$ can occur only on a $\mathbb{P}$-null event.

Case (ii): v_t does not converge. Then v_t oscillates within $B_\varepsilon(\delta_\theta)$. In particular, there exists some $\theta' \neq \theta$ such that

$$\limsup_{t \to \infty} v_t(\theta') > 0,$$

so there exists $\delta' > 0$ for which $v_t(\theta') \geq \delta'$ infinitely often.

By Assumption 2, whenever $v_t(\theta') \geq \delta'$, there exist $T < \infty$ and $\eta > 0$ such that for any $\varepsilon' > 0$,

$$\mathbb{P}(v_{t+T}(\theta') \geq 1 - \varepsilon' \,|\, \mathcal{F}_t) \geq \eta.$$

Choose $\varepsilon' > 0$ sufficiently small so that $v_{t+T}(\theta') \geq 1 - \varepsilon'$ implies $v_{t+T} \notin B_\varepsilon(\delta_\theta)$. Hence, each time $v_t(\theta') \geq \delta'$, the conditional probability of exiting $B_\varepsilon(\delta_\theta)$ within T steps is at least η. Since $v_t(\theta') \geq \delta'$ occurs infinitely often, Lévy's extension of the Borel–Cantelli lemma implies that v_t exits $B_\varepsilon(\delta_\theta)$ almost surely, contradicting the assumption.

Therefore, with probability one, the process $\{v_t\}$ leaves $B_\varepsilon(\delta_\theta)$ in finite time. □

Lemma 9 (Probability of Visiting Neighborhoods)**.** *For any $\varepsilon > 0$, there exists $\delta > 0$ such that*

$$\forall t \in \mathbb{T} : \quad \mathbb{P}\left(v_\tau \in \bigcup_{\theta \in \Theta} B_\varepsilon(\delta_\theta) \text{ for some } \tau \geq t \,\middle|\, \mathcal{F}_t\right) > \delta.$$

That is, the probability that v_t will enter some neighborhood $B_\varepsilon(\delta_\theta)$ is always bounded away from zero.

Proof. Fix $\varepsilon > 0$. For any $t \in \mathbb{T}$, there exists $\theta \in \Theta$ such that $v_t(\theta) \geq 1/|\Theta|$. By Assumption 2, there exist an integer $T < \infty$ and a constant $\delta > 0$ such that whenever $v_t(\theta) \geq 1/|\Theta|$,

$$\mathbb{P}\big(v_{t+T}(\theta) \geq 1 - \varepsilon \,\big|\, \mathcal{F}_t\big) \geq \delta.$$

In particular, on this event we have $v_{t+T} \in B_\varepsilon(\delta_\theta)$. Therefore,

$$\mathbb{P}\Big(v_\tau \in \bigcup_{\hat{\theta}\in\Theta} B_\varepsilon(\delta_{\hat{\theta}}) \text{ for some } \tau \geq t \,\Big|\, \mathcal{F}_t \Big) \;\geq\; \mathbb{P}\big(v_{t+T} \in B_\varepsilon(\delta_\theta) \,\big|\, \mathcal{F}_t\big) \;\geq\; \delta,$$

which proves the claim. □

A.3.2 Proof of Proposition 1

Proof. We consider the following three cases.

Case (i): $\Theta_p^w = \emptyset$. By Corollary 3, beliefs oscillate almost surely. Lemma 8 further implies for each state $\theta \in \Theta$, $\liminf v_t(\theta) = 0$ and $\limsup v_t(\theta) = 1$, so Proposition 1(i) is proved.

Case (ii): $\Theta_p^w = \Theta$. Fix $\delta \in (0,1)$. Since $\Theta_p^w = \Theta$, Lemma 7 implies that for each $\theta \in \Theta$ there exists $\varepsilon > 0$ such that whenever $v_t \in B_\varepsilon(\delta_\theta)$,

$$\mathbb{P}(v_{t+k} \to \delta_\theta \text{ as } k \to \infty \,|\, \mathcal{F}_t) \geq 1 - \delta.$$

On the other hand, by Lemma 9, for this ε there exists $\delta_0 > 0$ such that for all $t \in \mathbb{T}$,

$$\mathbb{P}\left(\exists\, \tau \geq t :\; v_\tau \in \bigcup_{\theta\in\Theta} B_\varepsilon(\delta_\theta) \,\middle|\, \mathcal{F}_t \right) \geq \delta_0. \tag{20}$$

Combining the two statements and using the Markov property yields

$$\mathbb{P}\left(\bigcup_{\theta\in\Theta} \{v_t \to \delta_\theta\} \,\middle|\, \mathcal{F}_t \right) \geq \delta_0(1-\delta) \equiv \delta_1 > 0 \qquad \text{for all } t.$$

By Lévy's 0–1 law, we have

$$\mathbb{P}\left(\bigcup_{\theta\in\Theta} \{v_t \to \delta_\theta\} \right) = 1.$$

Case (iii): Θ_p^w is a singleton set. Suppose $\Theta = \{\theta_1, \theta_2\}$ with $\theta_1 \in \Theta_p^w$ and $\theta_2 \notin \Theta_p^w$. Fix $\delta \in (0,1)$. By Lemma 7, there exists $\varepsilon > 0$ such that whenever $v_t \in B_\varepsilon(\delta_{\theta_1})$,

$$\mathbb{P}(v_{t+k} \to \delta_{\theta_1} \text{ as } k \to \infty \,|\, \mathcal{F}_t) \geq 1 - \delta.$$

By Lemma 8, for $\theta_2 \notin \Theta_p^w$ the process almost surely exits every neighborhood of δ_{θ_2}, and hence it cannot remain in $B_\varepsilon(\delta_{\theta_2})$ forever. Combining this with Assumption 2 (and finiteness/i.i.d.

of signals), there exists $\delta_0 > 0$ such that for all t,

$$\mathbb{P}(\exists\, \tau \geq t :\ v_\tau \in B_\varepsilon(\delta_{\theta_1}) \,|\, \mathcal{F}_t) \geq \delta_0.$$

Therefore, we obtain the uniform bound

$$\mathbb{P}(v_t \to \delta_{\theta_1} \mid \mathcal{F}_t) \ \geq\ \delta_0(1-\delta) \ \equiv\ \delta_1 \ >\ 0 \qquad \text{for all } t,$$

which implies that $\mathbb{P}(v_t \to \delta_{\theta_1}) = 1$ by standard 0–1 arguments.

Combining Cases (ii) & (iii), we conclude:

$$\mathbb{P}\left(\cup_{\theta \in \Theta_p^w} \{v_\tau \to \delta_\theta\}\right) = 1,$$

which proves Proposition 1(ii). □

A.4 Proof of Proposition 2

Proposition 2 follows from the next two lemmas.

Lemma 10 (Escaping Probability to Undominated Sets)**.** *Suppose Assumption 3 holds. For any $\delta > 0$, there exist $\varepsilon_1 > \varepsilon_2 > 0$ such that for all $t \in \mathbb{T}$,*

$$\forall \theta \notin \Theta_p^w : \quad \mathbb{P}\left(\max_{\theta' \in \mathcal{U}_\theta} v_\tau(\theta') \geq \varepsilon_1 \text{ for some } \tau \geq t \,\middle|\, v_t \in B_{\varepsilon_2}(\delta_\theta)\right) \geq 1-\delta. \tag{21}$$

Remark 4. Lemma 8 shows that when $\theta \notin \Theta_p^w$, the belief process almost surely escapes any neighborhood of δ_θ. Lemma 10 further shows that, with arbitrarily high probability, the escaping belief assigns nontrivial mass to some state θ' that is in the undominated set of θ, provided the process starts in a sufficiently small neighborhood of δ_θ.

Proof. Fix $\theta \notin \Theta_p^w$. Partition the remaining states into

$$\mathcal{U}_\theta \equiv \{\theta' \neq \theta : \theta \not\succeq_p^w \theta'\}, \qquad \mathcal{L}_\theta \equiv \{\theta' \neq \theta : \theta \succeq_p^w \theta'\}.$$

By construction, $\mathcal{U}_\theta \cup \mathcal{L}_\theta = \Theta \setminus \{\theta\}$. Below, I focus on the case where $p > 0$ as the case where $p \leq 0$ follows exactly symmetrically.

Step 1: Behavior of dominated states $\mathcal{L}_\theta$. Fix $\theta' \in \mathcal{L}_\theta$. As $v_t(\theta) \to 1$, the local approximation implies

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

Since $\theta \succ_p^w \theta'$, Lemma 5 implies that there exists $\rho > 0$ such that

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

Let $\varepsilon_1 > 0$ be a sufficiently small number and choose $\varepsilon_2 < \varepsilon_1$. Define the exit time

$$T \equiv \inf\{t : v_t \notin B_{r_1}(\delta_\theta)\}, \quad \text{where} \quad r_1 \equiv \varepsilon_1 \times |\Theta|,$$

and for $\theta' \in \mathcal{L}_\theta$ define

$$Y_t(\theta') \equiv v_{t \wedge T}(\theta')^{p\rho}.$$

Then $\{Y_t(\theta')\}$ is a nonnegative supermartingale. By Markov's inequality,

$$\mathbb{P}(Y_T(\theta') \geq \varepsilon_1^{p\rho}) \leq \frac{\mathbb{E}[Y_T(\theta')]}{\varepsilon_1^{p\rho}} \leq \left(\frac{\varepsilon_2}{\varepsilon_1} \right)^{p\rho}.$$

Taking a union bound over $\theta' \in \mathcal{L}_\theta$,

$$\mathbb{P}\left(\bigcup_{\theta' \in \mathcal{L}_\theta} \{Y_T(\theta') \geq \varepsilon_1^{p\rho}\} \right) \leq |\Theta| \left(\frac{\varepsilon_2}{\varepsilon_1} \right)^{p\rho}. \tag{22}$$

Step 2: Exit must place mass on undominated states $\mathcal{U}_\theta$. Since $\theta \notin \Theta_p^w$, Theorem 1 implies that v_t cannot converge to δ_θ. Hence the process cannot remain forever in $B_{r_1}(\delta_\theta)$, and therefore $\mathbb{P}(T < \infty) = 1$.

At time T, there must exist some $\theta' \neq \theta$ such that $v_T(\theta') \geq \varepsilon_1$, i.e.,

$$\mathbb{P}\left(\bigcup_{\theta' \in \mathcal{U}_\theta \cup \mathcal{L}_\theta} \{Y_T(\theta') \geq \varepsilon_1^{p\rho}\} \right) = 1.$$

Combining this with (22),

$$\mathbb{P}\left(\bigcup_{\theta' \in \mathcal{U}_\theta} \{Y_T(\theta') \geq \varepsilon_1^{p\rho}\} \right) \geq 1 - |\Theta| \left(\frac{\varepsilon_2}{\varepsilon_1} \right)^{p\rho}.$$

Choosing ε_2 sufficiently small (relative to ε_1) ensures the right-hand side is at least $1-\delta$. Recall that $Y_t(\theta') = v_{t\wedge T}(\theta')^{p\rho}$ with $p\rho > 0$. Hence,

$$Y_T(\theta') \geq \varepsilon_1^{p\rho} \quad \Longleftrightarrow \quad v_T(\theta') \geq \varepsilon_1.$$

Therefore, probability bounds established for Y_T translate directly into corresponding bounds for the original belief process v_T. Since the process is Markov, the argument applies at any time t with $v_t \in B_{\varepsilon_2}(\delta_\theta)$, completing the proof. □

Lemma 11 (Probability of Reaching Entropy Minimizer)**.** *For any $\theta \in \Theta$, there exist $\delta > 0$ and $\varepsilon > \varepsilon' > 0$ such that*

$$\forall t \in \mathbb{T}: \quad \mathbb{P}\left(v_\tau \in \cup_{\theta' \in \Theta_p^w} B_\varepsilon\left(\delta_{\theta'}\right) \text{ for some } \tau \geq t | v_t \in B_{\varepsilon'}\left(\delta_\theta\right)\right) > \delta,$$

i.e., conditional on that the social belief is sufficiently close to any point-mass belief, it visits the neighborhood of the point mass on some $\theta' \in \Theta_p^w$ with positive probability.

Proof. Since $\mathcal{U}_{\theta_1} \subset \mathcal{U}_{\theta_2} \subset \cdots \subset \mathcal{U}_{\theta_k}$, we can partition the state space Θ into finitely many level sets $\{\Theta_1, \ldots, \Theta_M\}$ such that:

- for any $m < n$, $\theta \in \Theta_m$ and $\theta' \in \Theta_n$ imply $\mathcal{U}_\theta \subset \mathcal{U}_{\theta'}$;
- for any m and any $\theta, \theta' \in \Theta_m$, we have $\mathcal{U}_\theta = \mathcal{U}_{\theta'}$.

By construction, $\Theta_1 = \Theta_p^w$. Lemma 10 and Assumption 2 imply that there exists a sequence $\{\varepsilon_m\}_{m=1}^M$ with

$$\varepsilon_1 > \varepsilon_2 > \cdots > \varepsilon_M > 0$$

and $\delta_0 > 0$ such that for every $m > 1$ and every $\theta \in \Theta_m$,[35]

$$\mathbb{P}\left(\bigcup_{m'<m} \bigcup_{\theta' \in \Theta_{m'}} \left\{v_\tau \in B_{\varepsilon_{m'}}(\delta_{\theta'}) \text{ for some } \tau \geq t\right\} \,\middle|\, v_t \in B_{\varepsilon_m}(\delta_\theta)\right) \geq 1 - \delta_0.$$

Since the number of levels is finite, iterating this argument at most M times yields

$$\mathbb{P}\left(\bigcup_{\theta \in \Theta_1} \left\{v_\tau \in B_{\varepsilon_1}(\delta_\theta) \text{ for some } \tau \geq t\right\} \,\middle|\, v_t \in B_{\varepsilon_m}(\delta_\theta)\right) \geq (1-\delta_0)^M \equiv \delta > 0.$$

[35] Lemma 10 implies that, starting sufficiently close to δ_θ with $\theta \in \Theta_m$, the belief process assigns positive mass to some state in $\mathcal{U}_\theta \subset \bigcup_{m'<m} \Theta_{m'}$ with probability arbitrarily close to one. Assumption 2 then ensures that any belief placing positive mass on such a state reaches a neighborhood of the corresponding Dirac belief with uniformly positive probability. The lower bound $1-\delta_0$ can be taken as the product of the (lowest) escaping probability and transition probability.

This establishes the claim. □

Lemma 7 implies that there exists $\varepsilon > 0$ such that for all $\theta \in \Theta_p^w$, $\mathbb{P}\left(v_t \to \delta_\theta | v_t \in B_\varepsilon\left(\delta_\theta\right)\right)$ is bounded away from zero. Combining Lemmas 9 to 11, we have $\mathbb{P}\left(\cup_{\theta\in\Theta_p^w}\left\{v_t \to \delta_\theta\right\}|\mathcal{F}_t\right)$ is bounded away from zero for all $t \in \mathbb{T}$ and all possible realizations of signals. This implies that $\mathbb{P}\left(\cup_{\theta\in\Theta_p^w}\left\{v_t \to \delta_\theta\right\}\right) = 1$ by Lévy's 0-1 law.

A.5 Proof of Proposition 3

Proof. Let $E \equiv \left\{v_t\left(\Theta\backslash\hat{\Theta}\right) \to 0\right\}$. On event E, the social belief assigns probability one to $\hat{\Theta}$ in the limit. Conditional on E, there are two possibilities: $E_1 \equiv \{v_t \text{ converges}\}$ and $E_2 \equiv \{v_t \text{ does not converge}\}$. We now show that E is a $\mathbb{P}$-null event by showing that each possibility happens with zero probability.

Step 1: Convergence on E is impossible. We first show that $\mathbb{P}\left(E \cap E_1\right) = 0$. By assumption, $\mathcal{U}_\theta \neq \emptyset$ for all $\theta \in \hat{\Theta}$, which implies that no state in $\hat{\Theta}$ belongs to Θ_p^w. By Theorem 1, whenever v_t converges, it must converge to a Dirac belief on some $\theta \in \Theta_p^w$. This is incompatible with event E. Hence, convergence on E occurs only on a $\mathbb{P}$-null event.

Step 2: Non-convergence on E is also impossible. I now show that $\mathbb{P}\left(E \cap E_2\right) = 0$. On $E \cap E_2$, there almost surely exists a state $\theta_0 \in \hat{\Theta}$ such that $\limsup v_t\left(\theta_0\right) = 1$ and $\liminf v_t\left(\theta_0\right) = 0$.[36] Fix such a θ_0. Hence, for any $\varepsilon_0 > 0$,

$$v_t \in B_{\varepsilon_0}(\delta_{\theta_0}) \quad \text{i.o.}$$

By Lemma 10, there exist small constants $\varepsilon > \varepsilon_0 > 0$ and $\delta \in (0,1)$ such that for any t,

$$\mathbb{P}\left(\max_{\theta\in\mathcal{U}_{\theta_0}} v_\tau\left(\theta\right) \geq \varepsilon \text{ for some } \tau \geq t \;\middle|\; v_t \in B_{\varepsilon_0}\left(\delta_{\theta_0}\right)\right) \geq 1 - \delta.$$

Hence, on $E \cap E_2$, we have almost surely

$$\sum_{t=1}^{\infty} \mathbb{P}\left(\max_{\theta\in\mathcal{U}_{\theta_0}} v_\tau\left(\theta\right) \geq \varepsilon \text{ for some } \tau \geq t | \mathcal{F}_t\right) = \infty.$$

This fact implies that on event $E \cap E_2$, there exists almost surely a state $\theta_1 \in \mathcal{U}_{\theta_0} \cap \hat{\Theta}$ such that $v_t(\theta_1) > \varepsilon$ infinitely often. Together with the irreducibility assumption, this further

[36] The informativeness of signals, together with irreducibility, implies that boundary beliefs are unstable: suitable signal realizations can move likelihood ratios in either direction, ruling out interior trapping.

implies that for any $\varepsilon_1 > 0$,

$$v_t \in B_{\varepsilon_1}(\delta_{\theta_1}) \quad \text{i.o.,} \quad \text{where} \quad \theta_1 \in \mathcal{U}_{\theta_0} \cap \hat{\Theta}.$$

By induction, on $E \cap E_2$, there exists almost surely a state $\theta^\dagger$ such that for all $\varepsilon^\dagger > 0$,

$$v_t \in B_{\varepsilon^\dagger}(\delta_{\theta^\dagger}) \quad \text{i.o.,} \quad \text{where} \quad \mathcal{U}_{\theta^\dagger} \cap \hat{\Theta} = \emptyset.$$

Since $\mathcal{U}_{\theta^\dagger}$ is nonempty and $\mathcal{U}_{\theta^\dagger} \subset \Theta \backslash \hat{\Theta}$, Lemma 10 implies that on $E \cap E_2$, there exists almost surely some $\varepsilon > 0$ such that

$$v_t\left(\Theta \backslash \hat{\Theta}\right) \geq \varepsilon \quad \text{i.o.,}$$

contradicting the definition of E. Therefore, $\mathbb{P}(E \cap E_2) = 0$.

Combining the two cases, we conclude that

$$\mathbb{P}\left(\limsup_{t \to \infty} v_t\left(\Theta \backslash \hat{\Theta}\right) > 0\right) = 1,$$

which proves the proposition. □

A.6 Proof of Proposition 4

Proof. The proof proceeds by construction. (i) For any regular F with zero order, weak group irrationality can arise because the minimizer of the weighted relative entropy need not coincide with the minimizer of any individual relative entropy. (e.g., Example 5 when $p \to 0$). (ii) For any regular F with nonzero degree, group irrationality can occur. In particular, one can construct a true DGP and consistent model perceptions such that group irrationality arises with positive probability. Corollary 6 provides an explicit construction: when individuals sufficiently overreact to private signals and the true DGP is not overly informative, society fails to learn the truth with positive probability.[37] □

A.7 Proof of Corollary 4

Proof. The second part comes from Proposition 4. I now show the first part; that is, the log-linear rule is immune to group irrationality. Under the log-linear rule, social belief sequence

[37] In the binary environment, Corollary 6 implies that if the true signal distribution satisfies $\pi^* \in (1/2, \sqrt[n]{1/2})$ and α is sufficiently large, then group irrationality occurs with positive probability. The argument also extends to multi-state environments because we can effectively reduce the problem to a binary one, e.g., by choosing signal distributions in the remaining states to be arbitrarily close to those of one of the binary states.

$\{v_t\}$ satisfies the following recursion.

$$\log\left(\frac{v_{t+1}(\theta)}{v_{t+1}(\theta')}\right) = \log\left(\frac{v_t(\theta)}{v_t(\theta')}\right) + \sum_i w_i \times \log\left(\frac{\hat{l}_i(s_{i,t} \mid \theta)}{\hat{l}_i(s_{i,t} \mid \theta')}\right). \tag{23}$$

Taking the time average on both sides of (23), we obtain

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

The strong law of large numbers implies that

$$\frac{1}{t}\log\left(\frac{v_{t+1}(\theta)}{v_{t+1}(\theta')}\right) \to \sum_i w_i \times \mathbb{E}\log\left(\frac{\hat{l}_i(s_{i,t} \mid \theta)}{\hat{l}_i(s_{i,t} \mid \theta')}\right) = \mathcal{I}^w(\theta') - \mathcal{I}^w(\theta), \quad \mathbb{P}-\text{a.s.}$$

Therefore, if $\mathcal{I}^w(\theta') < \mathcal{I}^w(\theta)$, then $v_t(\theta) \to 0$. Under all consistent model perceptions, θ^* uniquely minimizes $\mathcal{I}^w(\theta)$, then $v_t(\theta^*) \to 1$ almost surely, which further implies that $\mu_{i,t}(\theta^*) \to 1$ almost surely. Therefore, group irrationality cannot occur. □

A.8 Proof of Proposition 5

Proof. The proof proceeds in the following two directions.

(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 4 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

$$\forall \theta \neq \theta^*: \quad \max_i \log\left(\frac{\mu_{i,t}(\theta)}{\mu_{i,t}(\theta^*)}\right) \to -\infty \quad \mathbb{P}-\text{a.s.},$$

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. □

A.9 Proof of Corollary 5

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

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

Therefore, we have $\theta^* \not\succeq_\infty \theta$ when n is sufficiently large. By Proposition 5, $\hat{l}$ is not robustly consistent. □

A.10 Proof of Proposition 6

Proof. Suppose $|\Theta| = \{\theta_1, \theta_2\}$ with θ_1 being the true state.

Part (i): Since $\hat{l} = (\hat{l}_1, ..., \hat{l}_n)$ is consistent,

$$\mathcal{R}_i(\theta_1) < \mathcal{R}_i(\theta_2) \quad \text{for all } i,$$

which implies that $\mathcal{I}^w(\theta_1) < \mathcal{I}^w(\theta_2)$, or equivalently,

$$\mathcal{D}_0^w(\theta_1, \theta_2) > 0 > \mathcal{D}_0^w(\theta_2, \theta_1).$$

By Lemma 1, $\mathcal{D}_p^w \to \mathcal{D}_0^w$ as $p \to 0$. Therefore, there exists $\varepsilon > 0$ such that

$$\mathcal{D}_p^w(\theta_1, \theta_2) > 0 > \mathcal{D}_p^w(\theta_2, \theta_1) \quad \text{for all } p \text{ with } |p| < \varepsilon.$$

That is, both $\theta_1 \succ_p^w \theta_2$ and $\theta_2 \not\succ_p^w \theta_1$ hold, hence correct learning obtains by Proposition 1.

Part (ii): It follows from the proof of Proposition 7 or from Proposition 9.[38] □

[38] Proposition 7 assumes a binary signal space, but the argument extends to a general signal space. Its main result is essentially a special case of Proposition 9.

A.11 Proof of Proposition 7

Proof. Suppose $\Theta = \{\theta_1, \theta_2\}$ with θ_1 being the true state. Every individual i has an α-consistent model perception

$$\hat{l}_i(s|\theta) \propto q_i(s|\theta)^\alpha,$$

where q_i itself is a consistent model perception. First observe that

$$\frac{1}{\alpha}\mathcal{D}_p^w(\theta_2, \theta_1) = \frac{1}{\alpha p} \times \mathbb{E}\log\left(\sum w_i \times \left(\frac{q_i(s_i|\theta_2)}{q_i(s_i|\theta_1)}\right)^{\alpha p}\right) \equiv \phi(p, w, \alpha).$$

Define

$$\begin{aligned}\phi_0(p, w) \equiv \lim_{\alpha\to 0}\phi(p, w, \alpha) &= \sum_i w_i \times \mathbb{E}\log\left(\frac{q_i(s_i|\theta_2)}{q_i(s_i|\theta_1)}\right) \\ &< \max_i \mathbb{E}\log\left(\frac{q_i(s_i|\theta_2)}{q_i(s_i|\theta_1)}\right) \equiv c < 0,\end{aligned}$$

where the inequality comes from $\mathbb{E}\log\left(\frac{q_i(s_i|\theta_2)}{q_i(s_i|\theta_1)}\right) < 0$ for each i (since q_i is also consistent). Note that $p \in [-\overline{p}, \overline{p}] \equiv \mathcal{P}$ and $w \in \Delta$, which are both compact. Let

$$\Gamma(\alpha) \equiv \max_{(p,w)\in\mathcal{P}\times\Delta} |\phi(p, w, \alpha) - \phi_0(p, w)|.$$

The maximum theorem implies that $\Gamma(\alpha)$ is a continuous function. Since $\phi(p, w, \alpha) \to \phi_0(p, w)$ pointwise and continuously, we have $\Gamma(\alpha) \to 0$ as $\alpha \to 0$. Therefore, there exists $\overline{\alpha} > 0$ such that

$$\forall \alpha \in (0, \overline{\alpha}): \quad \max_{(p,w)\in\mathcal{P}\times\Delta} |\phi(p, w, \alpha) - \phi_0(p, w)| < -\frac{c}{2}.$$

Since $\phi_0(p, w) < c$, this implies for all $(p, w) \in \mathcal{P} \times \Delta$ and $\alpha \in (0, \overline{\alpha})$, we have $\phi(p, w, \alpha) < 0$, or equivalently, $\theta_1 \succ_p^w \theta_2$. Symmetrically, we can also prove that under similar condition, $\theta_2 \not\succ_p^w \theta_1$. Therefore, correct learning occurs almost surely by Proposition 1. □

A.12 Proof of Corollary 6

Proof. Suppose that $\Theta = \{\theta_1, \theta_2\}$ with θ_1 being the true state. The true distribution is $(\pi^*, 1-\pi^*)$ where $\pi^* > 1/2$ denotes the true probability of signal s_1. For each i, the

perceived model satisfies $\hat{l}_i(s_k|\theta_k) = c\pi^\alpha$. As $\alpha \to +\infty$,

$$\frac{1}{\alpha}\mathcal{D}_p^w(\theta,\theta') \to \begin{cases} \mathbb{E}\log\left(\max_i \frac{q(s_i|\theta)}{q(s_i|\theta')}\right) & p>0 \\ \mathbb{E}\log\left(\min_i \frac{q(s_i|\theta)}{q(s_i|\theta')}\right) & p<0 \end{cases}.$$

where $q(s|\theta)$ is the signal structure satisfying $q(s_k|\theta_k) = \pi$. Let's first consider the case $p > 0$. Notice that

$$\begin{aligned}\mathbb{E}\log\left(\max_i \frac{q(s_i|\theta_2)}{q(s_i|\theta_1)}\right) &= \pi^{*n} \times \log\frac{1-\pi}{\pi} + (1-\pi^{*n}) \times \log\frac{\pi}{1-\pi} \\ &= (1-2\pi^{*n})\log\frac{\pi}{1-\pi}.\end{aligned}$$

Similarly,

$$\mathbb{E}\log\left(\max_i \frac{q(s_i|\theta_1)}{q(s_i|\theta_2)}\right) = (1-2(1-\pi^*)^n)\log\frac{\pi}{1-\pi}.$$

If $1-2\max\{\pi^*, 1-\pi^*\} > 0$ (i.e., $n > \log_{\pi^*}^{1/2}$), then $\mathcal{D}_p^w(\theta,\theta') > 0$ for all $\theta \neq \theta'$, when α is adequately large. Thus, $\hat{\Theta}_p^w = \emptyset$, and oscillations occur almost surely. Symmetrically, when $p < 0$ and α is large, we have $\hat{\Theta}_p^w = \Theta$, so beliefs converge to both δ_{θ_1} and δ_{θ_2}, each with positive probability. □

A.13 Proof of Proposition 8

Proof. Suppose $\Theta = \{\theta_1, \theta_2\}$ with θ_1 being the true state. Under Assumption 3, if correct learning occurs almost surely under F, then $\theta_1 \succ_p^w \theta_2$ and $\theta_2 \not\succ_p^w \theta_1$ by Proposition 1.

Suppose that $p > p' > 0$, where p and p' denote the power of F and F', respectively. Since $\theta_1 \succ_p^w \theta_2$, we have $\mathcal{D}_p^w(\theta_2,\theta_1) < 0$. By monotonicity, we have

$$\mathcal{D}_{p'}^w(\theta_2,\theta_1) < \mathcal{D}_p^w(\theta_2,\theta_1) < 0,$$

implying that $\theta_1 \succ_{p'}^w \theta_2$. Since model perceptions $\hat{l}_i$'s are consistent, θ_1 uniquely minimizes the weighted relative entropy; that is, $\mathcal{D}_0^w(\theta_1,\theta_2) > 0$. Applying the monotonicity property again, we obtain

$$\mathcal{D}_{p'}^w(\theta_1,\theta_2) > \mathcal{D}_0^w(\theta_1,\theta_2) > 0,$$

implying that $\theta_2 \not\succ_{p'}^w \theta_1$. In summary, both $\theta_1 \succ_{p'}^w \theta_2$ and $\theta_2 \not\succ_{p'}^w \theta_1$ hold, so correct learning occurs almost surely under F' as well. The analysis for the case $0 > p' > p$ is symmetric. □

A.14 Proof of Proposition 9

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

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

All previous characterizations still hold if we redefine the dominance relation based on $\mathcal{D}_w^p(\theta', \theta; \psi)$. Suppose that $\psi_i(s|\theta) = \hat{l}_i(s|\theta)^\alpha$. Then for all $\theta', \theta \in \Theta$, we have

$$\begin{aligned} \frac{1}{\alpha}\mathcal{D}_w^p(\theta', \theta; \alpha) &= \frac{1}{\alpha p} \times \mathbb{E} \log \left(\sum_{i=1}^n w_i \times \left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)} \right)^{\alpha p} \right) \\ &\overset{\alpha \to 0}{\longrightarrow} \sum_{i=1}^n w_i \times \mathbb{E} \log \left(\frac{\hat{l}_i(s_i|\theta')}{\hat{l}_i(s_i|\theta)} \right) = \mathcal{I}^w(\theta) - \mathcal{I}^w(\theta'), \end{aligned} \tag{24}$$

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, (24) implies that when α is sufficiently small, $\succeq_{p,\alpha}^w$ and $\succeq_0^w$ induces the same ordering over Θ. Assumption 7 implies that $\mathcal{I}^w(\theta)$ is distinct for all state, so there exists a strict ranking

$$\mathcal{I}^w(\theta_0) < \mathcal{I}^w(\theta_1) < ... < \mathcal{I}^w(\theta_k),$$

where θ_0 is the minimizer of the weighted relative entropy, and k is the number of states. Thus, when α is sufficiently small, the undominated sets defined under $\succeq_{p,\alpha}^w$ can be completely ranked with

$$\emptyset = \mathcal{U}_{\theta_0} \subset \mathcal{U}_{\theta_1} \subset ... \subset \mathcal{U}_{\theta_k}.$$

Proposition 2 then implies that $\mu_{i,t} \to \delta_{\theta_0}$ almost surely for all i. □

B Relation to Molavi et al. (2018)

This appendix illustrates more explicitly how the results in this paper relate to Theorem 7 of Molavi et al. (2018). I first show that under correct specification, the weighted p-entropy reproduces their correct-learning threshold $p \in [-1, 1]$. I then show that when misspecification is small, the belief dynamics converge to those in their model.

Correct Specification

Proposition 10 (Correct Specification). *Suppose that $\hat{l}_i = l_i$ for all i, and Assumption 7 holds. Then, $\Theta_p^w = \{\theta^*\}$ for all $p \in [-1, 1]$.*

Proof. When $p \in [0, 1]$, we have

$$\begin{aligned}\mathcal{D}_p^w(\theta, \theta^*) &< \frac{1}{p} \times \log\left(\sum_{i=1}^n w_i \times \mathbb{E}\left(\frac{l_i(s_i|\theta)}{l_i(s_i|\theta^*)}\right)^p\right) \\ &< \frac{1}{p} \times \log\left(\sum_{i=1}^n w_i \times \left(\mathbb{E}\left(\frac{l_i(s_i|\theta)}{l_i(s_i|\theta^*)}\right)\right)^p\right) = 0,\end{aligned}$$

where the inequalities come from the strict concavity of the log function and x^p when $p \in [0, 1]$. It follows that $\theta^* \succ_p^w \theta$ for all $\theta \neq \theta^*$, so $\theta^* \in \Theta_p^w$. Under Assumption 7, $\Theta_0^w = \{\theta^*\}$. By Corollary 2, Θ_p^w can only be a singleton, so $\Theta_p^w = \{\theta^*\}$.

Similarly, when $p \in [-1, 0)$,

$$\begin{aligned}\mathcal{D}_p^w(\theta^*, \theta) &= \frac{1}{p} \times \mathbb{E}\log\left(\sum_{i=1}^n w_i \times \left(\frac{l_i(s_i|\theta)}{l_i(s_i|\theta^*)}\right)^{|p|}\right) > \frac{1}{p} \times \log\left(\sum_{i=1}^n w_i \times \mathbb{E}\left(\frac{l_i(s_i|\theta)}{l_i(s_i|\theta^*)}\right)^{|p|}\right) \\ &> \frac{1}{p} \times \log\left(\sum_{i=1}^n w_i \times \left(\mathbb{E}\left(\frac{l_i(s_i|\theta)}{l_i(s_i|\theta^*)}\right)\right)^{|p|}\right) = 0.\end{aligned}$$

Thus, there is no $\theta \neq \theta^*$ satisfying $\theta \succeq_p^w \theta^*$, so every $\theta \neq \theta^*$ does not belong to Θ_p^w. When p is negative, $\Theta_p^w \neq \emptyset$ by Corollary 2, so we also have $\Theta_p^w = \{\theta^*\}$. □

Proposition 10 establishes that for $p \in [-1, 1]$, the true state θ^* is the *unique* minimizer of the weighted p-entropy under the correct specification. This fact eliminates almost-sure oscillations when $p > 0$ and multiple limit points when $p < 0$, thereby ruling out both channels of group irrationality. Under additional conditions,[39] the characterizations in the paper further imply the almost-sure convergence to the Dirac belief on the true state, echoing the correct learning result in Molavi et al. (2018).

Small Misspecification

The next proposition shows that when individuals' perceived models approach the truth, the learning dynamics converge to those described in Molavi et al. (2018).

[39] In the binary case, Proposition 1 directly implies correct learning. In the multi-state case, one may require additional structure; e.g., complete ranking of undominated sets as in Proposition 2.

Proposition 11 (Small Misspecification)**.** *Suppose* $|\Theta| = 2$ *and Assumption 7 holds at the true DGP* l*. Let* $\{\hat{l}_n\}$ *be a sequence of regular model perceptions converging to* l*.*[40] *Then, for any regular aggregation rule with* $p \in [-1, 1]$*, there exists* $N < \infty$ *such that correct learning occurs* $\mathbb{P}$*-almost surely for all* $n \geq N$*.*

Proof. Suppose $\Theta = \{\theta_1, \theta_2\}$ with θ_1 being the true state. Note that the divergence is continuous in model perceptions. That is, for all $\theta, \theta' \in \Theta$,

$$\mathcal{D}^w_{p,n}(\theta, \theta') \to \mathcal{D}^w_{p,true}(\theta, \theta') \quad \text{as } n \to +\infty,$$

where $\mathcal{D}^w_{p,n}$ and $\mathcal{D}^w_{p,true}$ denote the divergence under $\hat{l}_n$ and l, respectively.

By Proposition 10, $\Theta^w_{p,true} = \{\theta_1\}$. Hence, for sufficiently large n, we have $\Theta^w_{p,n} = \{\theta_1\}$; that is, $\theta_1 \succ^w_{p,n} \theta_2$ and $\theta_2 \not\succ^w_{p,n} \theta_1$. Since the aggregation rule and model perceptions are regular, the irreducibility condition holds. Following the logic of Proposition 1, correct learning occurs almost surely for n sufficiently large. □

The proposition focuses on the binary-state case, but it is conceivable that the paper's characterization extends to more general settings under additional structure; for example, when the undominated set can be completely ranked, as in Proposition 2.

[40] Since Θ and S are finite, the convergence is simply coordinate convergence, i.e., $\hat{l}_{i,n}(s|\theta) \to l_i(s|\theta)$ as $n \to +\infty$ for all i, s, θ.