Speed of Common Learning with Biased Memory*

Ruichao Pan
Department of Economics, Sciences Po

May 17, 2024

Abstract

When multiple agents have access to many draws of private signals, results from Frick, Iijima, and Ishii (2023) show that we can characterize the speed of individual and common learning away from the limit. However, in a data-rich setting like this, it is challenging for agents to remember and utilize precisely all of the signals they have received in the past. In this paper, we consider a setting where agents experience biased memory when making inferences, meaning certain signals are more likely to be remembered than others. Our first main result characterizes the speed of individual and common learning under this novel setting. We then show that, contrary to classical wisdom, under certain information structures, common learning occurs slower than individual learning; that is, higher-order uncertainty vanishes slower than first-order uncertainty. We also provide an intuitive necessary condition for this result: the state that agents individually find difficult to distinguish from the true state has to be different from the state that agents expect others to find difficult to distinguish.

*I would like to extend my deepest gratitude to my supervisor, Professor Franz Ostrizek, whose unwavering support, insightful guidance, and generous allocation of time have been invaluable throughout the completion of this master's thesis. I am also indebted to Xinyi Cao, Yakun Zheng, and Juyang Wang for their companionship and friendship over the course of my two-year master's journey. Additionally, I thank Mathilde Blanchon, Ramy Shaban, Huiyann Foo, Romain Angotti, and Nicolas Dussaux for academic inspiration and being the best problem set buddies. Last but not least, I thank my family for their unconditional supoort.

1 Introduction

In a learning setting where multiple agents have access to many draws of private information, economists have been interested in understanding how agents commonly learn the underlying state of the world despite facing uncertainty to varying degrees. A classical result from Cripps et al. (2008)shows that agents can always achieve approximate common knowledge (in the sense of Monderer and Samet (1989)) of the true state once they observe an infinitely large number of signal draws. Recent advancments by Frick, Iijima and Ishii (2023) provide a way of characterizing the speed of individual and common learning away from the limit, when agents observe a large but finite number of signal draws. Despite common learning appearing more demanding, as agents' first-order uncertainty about the state and their higher-order uncertainty about other players' beliefs about the state both have to vanish for common learning to be successful, a key lesson in this literature has been that common learning occurs at the same rate as individual learning. The intuition behind this result is that, for any information structure, higher-order uncertainty always vanishes faster than first-order uncertainty.

However, in a "data-rich" setting like this, it is challenging for agents to remember and utilize precisely all of the signals they have received in the past. In this paper, we consider a setting where agents cannot remember all the signals they observed in the past; particularly, when making inferences, they experience biased memory, where certain signals are more likely to be remembered or recalled than others. Our first main result characterizes the speed of individual and common learning under this novel setting. We then show that, under certain information structures, common learning occurs slower than individual learning; that is, higher-order uncertainty is vanishing slower than first-order uncertainty. We also provide an intuitive necessary condition for when common learning occurs slower than individual learning: the state that each agent individually finds difficult to distinguish from the true state has to be different from the state that each agent expects others to find difficult to distinguish.

Section 2 provides a brief literature review. Our results closely relate to the learning literature, contributing to our understanding of both the speed of common learning and the importance of higher-order uncertainty. We also discuss how the assumption of biased memory is related to the behavioral economics literature and provide real-world examples.

Section 3 introduces the setup of the learning environment. For technical simplicity, we consider an information structure $\mathcal{I}$ with binary signals, where higher frequencies of high signals are more indicative of a higher state (state set is finite). Agents receive independent draws of private signals from $\mathcal{I}$, which specifies a (full support) joint distribution over agents' private signals in each state but may feature arbitrary correlation across different players' signals.

Section 4 characterizes the speed of common learning with biased memory. For each information structure $\mathcal{I}$, some biased memory functions could cause common learning to fail. We are primarily interested in the rest of the memory function that does not lead to a breakdown in common learning, which we term "mild memory functions." Under mild memory functions, we consider the probability that agents have a common p-belief of the true state after$t$signal draws from$\mathcal{I}$, and analyze how fast p converges to one as t grows large. Utilizing Sanov's theorem from large deviation theory, we derive a "multi-agent memory-adjusted learning efficiency index" that characterizes the rate at which agents individually and commonly learn the state. The proof of Propostion 1 and Theorem 1 demonstrate that individual learning and common learning occur at different rates. Proposition 2 provides a necessary condition for when common learning can occur slower than individual learning. We also include a subsection 4.3 with numerical exmaples to better facilitate understanding of the intuition behind our main results.

Section 5concludes the article with a brief discussion on future research agendas and questions.

2 Related Literature

First and foremost, our paper closely relates to the learning literature. We adopt a similar learning setting as inCripps et al. (2008) and Frick, Iijima and Ishii (2023). As mentioned above, we contribute to their work by incorporating biased memory into the canonical setting, offering novel insights into the speed of individual and common learning. Contrary to classical wisdom, we demonstrate that with biased memory, under certain information structures, common learning can occur slower than individual learning: even when first-order uncertainty has vanished for the slowest learning agent, higher-order uncertainty persists, delaying common learning. Other studies have explored different settings where there is an extreme breakdown of common learning. Steiner and Stewart (2011) and Cripps et al. (2013) examine settings with correlated signals between draws, while Acemoglu, Chernozhukov and Yildiz (2016) study settings featuring identification problems due to uncertainty about the information structure. Our approach is less extreme and offers deeper insights; even with i.i.d. signals between draws and no uncertainty about the information structure, with certain memory functions, common learning or individual learning breaks down, whereas with other memory functions, common learning is preserved but occurs more slowly than individual learning.

We follow the approaches in Moscarini and Smith (2002) and Frick, Iijima and Ishii (2023) in deriving the learning efficiency index and extend it to a multi-agent learning environment with biased memory.Many others, such as Vives (1993), Hann-Caruthers, Martynov and Tamuz (2018), Rosenberg and Vieille (2019), and Dasaratha and He (2023), have considered individual learning efficiency indices in a social learning setup but have not focused on the role of higher-order uncertainty. Harel et al. (2021) discuss how higher-order beliefs influence agents' inference and provide an upper bound on the speed of learning. However, their setting involves long-lived agents repeatedly observing both private signals and actions of other players. Our result does not require agents to observe actions; repeated private signals are sufficient.

Our paper also contributes to a large literature on higher-order beliefs (for exmaple, Rubinstein (1989), Carlsson and van Damme (1993), and Weinstein and Yildiz (2007)). Our result reinstates the importance of higher-order uncertainty on learning outcomes. Even when agents have access to finitely many private signal draws and are rational in making inferences, biased memory could cause additional confusion regarding agents' estimates about others' beliefs, slowing down the speed of common learning. To improve efficiency in common learning, merely improving first-order uncertainty is not sufficient.

he biased memory that we consider is documented in several papers. For example, it can resemble the ego-boosting memory bias, which is documented and generalized in Fudenberg, Lanzani and Strack (0), where signals reveal information on ego-relevant characteristics, such as a successful IQ test. Zimmermann (2020)'s experiment finds that after taking an IQ test, subjects who received negative feedback were less likely to recall the feedback compared to subjects who received positive feedback. Walters and Fernbach (2021) finds that investors are 10% less likely to recall an investment that led to loss compared to ones that led to gain. Our paper provides a learning outcome for this type of ego-boosting memory bias when we consider high signals to be "positive feedback" and agents are endowed with a biased memory function that remembers higher signals more frequently than lower signals.

3 Setup

The learning environment addressed in this paper is characterized by a fixed, finite set of agents, denoted by N, and a fixed, finite set of states, denoted by Θ. Agents possess a common full-support prior belief over the states, $p_0 \in \Delta(\Theta)$.

Agents receive independent and identically distributed (i.i.d.) signal draws from an **Information Structure** $\mathcal{I}$ consisting of binary private signals $s_i \in \{L, H\}$. The space of private signals for each agent $i \in N$ is S_i, with $S := \prod_{i \in N} S_i$ denoting the set of all signal profiles. Let $\mu^\theta \in \Delta(S)$ represent the joint signal distribution conditional on each state $\theta \in \Theta$, where arbitrary correlations between agents' signals are allowed, and $\mu_i^\theta \in \Delta(S_i)$ denotes the marginal distribution over agent i's signals. For technical simplicity, assume that μ_i^θ has full support, and for each $\theta > \theta'$, $\mu_i^\theta(s_i = H) > \mu_i^{\theta'}(s_i = H)$, indicating that higher frequencies of high signals are more indicative of a higher state. Agents observe repeated i.i.d. signal draws from information structure $\mathcal{I}$. Let $t \in \mathbb{N}$ denote the number of signal draws. Then, for each information structure $\mathcal{I}$ and $t \in \mathbb{N}$, $\mathbb{P}_t^{\mathcal{I}} \in \Delta(\Theta \times S^t)$ denotes the probability over states and sequences of signal observations when state θ is drawn according to prior p_0, and signal profiles $s^t = (s_\tau)_{\tau=1,\ldots,t}$ are generated from μ^θ, conditional on state θ. Agent i's observed signals up to t are denoted by $x_i^t = (x_{i\tau})_{\tau=1,\ldots,t}$.

Let $\nu_{it} \in \Delta(S_i)$ denote agent i's proportion of high signal realizations up to time t:

$$\nu_{it} := \frac{1}{t} \sum_{\tau=1}^{t} \mathbf{1}_{\{s_{i\tau} = H\}}$$

Since signals are binary, ν_{it} serves as a sufficient statistic for agent i's empirical signal distribution.

In addition to this rational benchmark, an agent's memory of past signals is distorted by a collection of signal-dependent memory functions, denoted as $f : S \to [0, 1]$, where $f(s_i)$ specifies the probability with which the agent remembers a past realization of a signal. The agent is unaware of their memory bias and naively updates their beliefs as if the signals they remembered are the only ones they received. The posterior belief of the signal distribution is governed by a memory function m:

$$m(\nu_{it}) = \frac{f(s_i = H)\nu_{it}}{f(s_i = H)\nu_{it} + f(s_i = L)(1 - \nu_{it})}$$

For any $t \in \mathbb{N}$, $p \in (0, 1)$, memory function m, and event $E \subseteq \Theta \times S^t$, let $B_t^p(E)$ denote the event that E is p-believed at t, i.e., that all agents assign probability at least p to E after t draws from $\mathcal{I}$. That is,

$$B_t^p(E) := \bigcap_{i \in N} B_{it}^p(E), \quad \text{where} \quad B_{it}^p(E) := \Theta \times \left\{ s_i^t \in S_i^t : \mathbb{P}_t^{\mathcal{I}} \left(E \mid m(s_i^t) \right) \geq p \right\} \times \prod_{j \neq i} S_j^t$$

We say that all agents individually learn the true state when, for all $p \in (0,1)$, $\theta \in \Theta$, and memory function m, we have:

$$\lim_{t\to\infty} \mathbb{P}_t^{\mathcal{I}} \left(B_t^p(\theta) \mid \theta\right) = 1$$

Common learning additionally considers agents' higher-order beliefs. Let

$$C_t^p(E) := \bigcap_{k\in\mathbb{N}} \left(B_t^p\right)^k (E)$$

denote the event that E is commonly p-believed at$t$, where $\left(B_t^p\right)^1 (E) := B_t^p(E)$ and $\left(B_t^p\right)^k (E) := B_t^p \left(\left(B_t^p\right)^{k-1} (E)\right)$ for all $k \geq 2$.

We say that all agents commonly learn the true state, when for all $p \in (0,1)$, $\theta \in \Theta$, and memory function m, we have

$$\lim_{t\to\infty} \mathbb{P}_t^{\mathcal{I}} \left(C_t^p(\theta) \mid \theta\right) = 1$$

4 Multi-Agent Learning Efficiency with Biased Memory

4.1 Preliminaries

We first recall a few standard statistical measures from information theory that are crucial for identifying the speed of learning in this setting.

Definition 1. Fix any agent i and state θ, the **Kullback-Leibler divergence** of distribution ν_i relative to μ_i^θ is defined as:

$$\mathrm{KL}\left(\nu_i, \mu_i^\theta\right) := \sum_{s_i\in S_i} \nu_i\left(s_i\right) \log \frac{\nu_i\left(s_i\right)}{\mu_i^\theta\left(s_i\right)}$$

A classical result due to Berk (1966), and subsequently Esponda and Pouzo (2016), demonstrates that a misspecified agent's "long-run" belief (in our case as $t \to \infty$) assigns probability 1 to the state that minimizes the KL-distance between the agent's perceived signal distribution and the theoretical signal distribution. For our naive agents with biased memory, this result applies since, contrary to the rational benchmark, their perceived signal distribution may not match the theoretical signal distribution under the true state. As t grows large, an agent's belief will concentrate on the state that best explains the signal distribution observed.

Definition 2. Fix any agent i and state θ, for any $\theta^{'} \neq \theta$, the **Chernoff distance** between agent i's marginal signal distributions in states $\theta$and $\theta^{'}$ is defined as:

$$d\left(\mu_i^\theta, \mu_i^{\theta'}\right) := \min_{\nu_i\in\Delta(X_i)} \max\left\{\mathrm{KL}\left(\nu_i, \mu_i^\theta\right), \mathrm{KL}\left(\nu_i, \mu_i^{\theta'}\right)\right\} \tag{1}$$

The Chernoff distance is very useful in characterizing how difficult it is for an agent to statistically distinguish θ' fromθ. Note that any minimizer$\nu_i$of $\mathrm{KL}\left(\nu_i, \mu_i^\theta\right) = \mathrm{KL}\left(\nu_i, \mu_i^{\theta'}\right)$ must satisfy this equation. Thus, $d\left(\mu_i^\theta, \mu_i^{\theta'}\right)$ essentially measures the distance fromμ_i^θ and

$\mu_i^{\theta'}$ to their KL-midpoint, and smaller values of $d\left(\mu_i^\theta, \mu_i^{\theta'}\right)$ capture that agent i's private signal distributions in states θ and θ' are closer to each other.

Next, we recall two lemmas from the common learning literature that will be helpful in deriving our main results.

Lemma 1. *(Frick, Iijima, and Ishiii, 2023) Fix any $\theta \in \Theta$ and distinct $i, j \in N$. For each t and realized empirical signal distribution $\nu_{it} \in \Delta(S_i)$, we have*

$$\mathrm{KL}\left(\mathbb{E}\left[\nu_{jt} \mid \theta, \nu_{it}\right], \mu_j^\theta\right) \leq \mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right)$$

or equivalently,

$$\mathrm{KL}\left(\nu_{it} G_{ij}^\theta, \mu_j^\theta\right) \leq \mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right)$$

Moreover, the inequality is strict whenever μ^θ has full support and $\nu_i \neq \mu_i^\theta$.

This lemma states that when agent$i$ is forming an estimate of agent$j$'s signal observations conditional on i's own observation and state θ, this estimate is less atypical than i's own signal observations. This is the key driving force behind Frick, Iijima and Ishii (2023)'s result: because agent$i$'s expectation of agent j's signal distribution is closer to the theoretical distribution than i's own signal observation, whenever agent i privately learns the state, there will be no higher-order uncertainty about j's learning outcome.

Lemma 2. *(Cripps, Ely, Mailath, and Samuelson, 2008) For any information structure $\mathcal{I}$and agents $i$and j, consider the matrix $G_{ij}^\theta \in \mathbb{R}^{S_i \times S_j}$ with (s_i, s_j)-th entry*

$$G_{ij}^\theta(s_i, s_j) = \mu^\theta(s_j \mid s_i)$$

then, $\mathbb{E}\left[\nu_{jt} \mid \theta, \nu_{it}\right] = \nu_{it} G_{ij}^\theta$ and $\mu_i^\theta G_{ij}^\theta = \mu_j^\theta$.

Moreover, let $\|\cdot\|$ denote the sup norm for finite-dimensional real vectors. For any $\varepsilon > 0$ and $q < 1$, there is Tsuch that for all $t \geq T$,$\theta \in \Theta$, and s_i^t,

$$\mathbb{P}_t^{\mathcal{I}}\left(\left\{\left\|\nu_{it} G_{ij}^\theta - \nu_{jt}\right\| < \varepsilon, \forall j \neq i\right\} \mid s_i^t, \theta\right) > q$$

This lemma states that if agent$i$'s empirical signal distribution at time$t$ is ν_{it}, conditional on state θ,i's expectation of j's empirical distribution is given by $\mathbb{E}\left[\nu_{jt} \mid \theta, \nu_{it}\right] = \nu_{it} G_{ij}^\theta$ (treating $\nu_{it} \in \Delta(S_i) \subseteq \mathbb{R}^{1 \times S_i}$ as a vector). Additionally, each agent believes that their expectation of the frequencies of the signals observed by their opponent is likely to be nearly correct.

Recognize that these two lemmas are applicable to agents exhibiting any form of memory distortion. Since our agents are unaware of their memory biases, they draw inferences about other agents' signal observations as if what they remembered is all that happened.

4.2 Speed of Common Learning under Biased Memory

Fix any state $\theta \in \Theta$. To characterize the speed of individual and common learning with biased memory under state θ, we first need to identify the state that agents find most difficult to distinguish from state θ.

we construct $\hat{\nu}_i(\theta^{'})$ to capture the KL-midpoint between state $\theta$and $\theta^{'}$:

$$\hat{\nu_i}(\theta^{'}) := \text{argmin}_{\nu_i \in \Delta(s_i)} \max \left\{ \text{KL}\left(\nu_i, \mu_i^{\theta}\right), \text{KL}\left(\nu_i, \mu_i^{\theta'}\right) \right\}$$

Take any memory function m, denote the state that agent $i$with memory-adjusted belief finds most difficult to distinguish from state θas $\tilde{\theta_i}$. Formally, this is:

$$\tilde{\theta_i} := \text{argmin}_{\theta' \in \Theta \setminus \{\theta\}} \text{KL}\left(\hat{\nu_i}(\theta^{'}), m(\mu_i^{\theta})\right)$$

Denote the state that, given agent i's empirical signal distribution, agent i expects agent j to find most difficult to distinguish from state θas θ_j^*. Formally, this is:

$$\tilde{\theta_j} := \text{argmin}_{\theta' \in \Theta \setminus \{\theta\}} \text{KL}\left(\hat{\nu_j}(\theta^{'}), m(\mu_i^{\theta}) G_{ij}^{\theta}\right)$$

We are going to show later that common learning could fail under some memory functions as $t$grows large. We call these memory functions "wild". For the rest of the memory functions that do not induce a failure in common learning, we refer to them as "mild". For now, we introduce a definition that restricts the focus of our results to mild memory functions only.

Definition 3. Fix any state $\theta \in \Theta$ and information structure $\mathcal{I}$, a memory function is **mild** if for all agent $i \in N$,

(1)

$$\text{KL}\left(m(\mu_i^{\theta}), \mu_i^{\theta}\right) < \text{KL}\left(\hat{\nu_i}(\tilde{\theta_i}), \mu_i^{\theta}\right)$$

and,

(2)

$$\text{KL}\left(m(\mu_i^{\theta}), \mu_i^{\theta}\right) < \text{KL}\left(G^{-1}(\hat{\nu_j}(\tilde{\theta_j})), \mu_i^{\theta}\right)$$

Memory functions that violate either (1) or (2) or both are called **wild**.

Our first result shows that the speed of individual learning under a mild memory function m and information structure $\mathcal{I}$is given by a multi-agent memory-adjusted individual learning efficiency index. Formally,

$$\underline{\lambda_B}^{\theta}(\mathcal{I}, m) := \min_{i \in N} \lambda_i^{\theta}(\mathcal{I}, m) \quad \text{where } \lambda_i^{\theta}(\mathcal{I}, m) := \text{KL}\left(m^{-1}(\hat{\nu_i}(\tilde{\theta_i})), \mu_i^{\theta}\right)$$

Intuitively, $\lambda_i^{\theta}(\mathcal{I}, m)$ captures how difficult agent $i$with memory-adjusted belief finds it to distinguish state $\theta$from the state $\theta^{'}$ that generates the most similar signal distributions, and $\underline{\lambda_B}^{\theta}(\mathcal{I}, m)$ simply considers the slowest learning agent.

On the other hand, the speed of common learning is given by a multi-agent memory-adjusted common learning efficiency index:

$$\underline{\lambda_C}^{\theta}(\mathcal{I}, m) := \min_{i \in N, j \in N} \lambda_{ij}^{\theta}(\mathcal{I}, m) \quad \text{where } \lambda_{ij}^{\theta}(\mathcal{I}, m) := \text{KL}\left(m^{-1}\left(G^{-1}(\hat{\nu_j}(\tilde{\theta_j}))\right), \mu_i^{\theta}\right)$$

Intuitively, $\lambda_{ij}^{\theta}(\mathcal{I}, m)$ captures how difficult it is for agent $i$to believe that agent j could distinguish state $\theta$from the state $\theta^{'}$ that generates the most similar signal distributions from agent j's point of view, and $\underline{\lambda_B}^{\theta}(\mathcal{I}, m)$ simply considers the slowest learning pair of agent $i$and j.

One might think that the learning efficiency index could also be expressed in the memory-adjusted signal space but not in the memory-free signal space. However, that is not the case:

Remark 1. $\mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right) \geq \mathrm{KL}\left(m(\nu_{it}), m(\mu_i^\theta)\right)$

$$
\begin{aligned}
\mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right) &= \sum_{s_i \in \Delta(S_i)} \nu_{it} \log \frac{\nu_{it}}{\mu_i} \\
&= \sum_{s_i \in \Delta(S_i)} m(\nu_{it}) \frac{f(s_i = H)\nu_{it} + f(s_i = L)(1-\nu_{it})}{f(s_i = H)} \log \frac{m(\nu_{it}) \frac{f(s_i=H)\nu_{it} + f(s_i=L)(1-\nu_{it})}{f(s_i=H)}}{m(\mu_i) \frac{f(s_i=H)\mu_i + f(s_i=L)(1-\mu_i)}{f(s_i=H)}} \\
&= \mathrm{KL}\left(m(\nu_{it}) \frac{f(s_i = H)\nu_{it} + f(s_i = L)(1-\nu_{it})}{f(s_i = H)}, m(\mu_i^\theta) \frac{f(s_i = H)\mu_i + f(s_i = L)(1-\mu_i)}{f(s_i = H)}\right) \\
&\geq \mathrm{KL}\left(m(\nu_{it}), m(\mu_i^\theta)\right)
\end{aligned}
$$

with equality if $\nu_{it} = \mu_i^\theta$

Since the memory function alters the **marginal distribution** of signals that agent observe, it also alters the KL-distance. In order to access the rate of learning, we need to express the index in terms of the true (before memory distortion) signal distributions. Therefore, we cannot express or approximate it in the memory-adjusted space.

Theorem 1. *Fix any information structure $\mathcal{I}$, state $\theta \in \Theta$, and $p \in (0,1)$, if memory function m is wild, common learning fails.*

If memory function is mild, individual learning occurs at rate $\underline{\lambda_B}^\theta(\mathcal{I}, m)$, that is:

$$
\mathbb{P}_t^{\mathcal{I}}\left(B_t^p(\theta) \mid \theta\right) = 1 - \exp\left[-\underline{\lambda_B}^\theta(\mathcal{I}, m)t + o(t)\right]
$$

and common learning occurs at rate $\min\{\underline{\lambda_B}^\theta(\mathcal{I}, m), \underline{\lambda_C}^\theta(\mathcal{I}, m)\}$:

$$
\mathbb{P}_t^{\mathcal{I}}\left(C_t^p(\theta) \mid \theta\right) = 1 - \exp\left[-\min\{\underline{\lambda_B}^\theta(\mathcal{I}, m), \underline{\lambda_C}^\theta(\mathcal{I}, m)\}t + o(t)\right]
$$

When$\underline{\lambda_B}^\theta(\mathcal{I}, m) > \underline{\lambda_C^\theta}(\mathcal{I}, m)$, common learning occurs at a rate slower than private learning.

The proof of Theorem 1 utilizes the Sanov's theorem from large deviation theory. It states that, for any set $D \subseteq \Delta(S)$ that is the closure of its interior,

$$
\mathbb{P}_t^{\mathcal{I}}\left(\nu_t \in D \mid \theta\right) = \exp\left[-\inf_{\nu \in D} \mathrm{KL}\left(\nu, \mu^\theta\right)t + o(t)\right]
$$

The proof is straight forward when we can identify the events that gives rise to individual and common learning. We show it in the next proposition:

Proposition 1. *Fix any state $\theta \in \Theta$, mild memory function m, $p \in (0,1)$, $\theta^{'} \neq \theta$, and $\lambda \in \left(0, min\{\underline{\lambda_B}^\theta(\mathcal{I}, m), \underline{\lambda_C}^\theta(\mathcal{I}, m)\}\right)$. There exists T such that for all $i \in N$, $j \neq i$ and $t \geq T$,*

$$
\mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right) \leq \lambda \Longrightarrow \mathbb{P}_t^{\mathcal{I}}\left(\{\theta\} \cap \mathrm{KL}\left(\mathbb{E}\left[\nu_{jt} \mid \theta, m(\nu_{it})\right], \mu_j^{\theta^{'}}\right) > d\left(\mu_j^\theta, \mu_j^{\tilde{\theta}_j}\right) \mid s_i^t\right) \geq p
$$

Proposition 1 says that whenever the empirical signal frequencies that agent $i$received are close enough to the theoretical distribution under the true state θ, for a large enough t, agent $i$believes with high probability that the state is θ, and in expectation, with high probability, agent j observed signal frequencies are further away from all states except state θ. Since agent i is confident that j's observed signal distribution is further away from all state except θ, it is equivalent to saying $i$is confident in j learning state θ. Thus, this "close enough" threshold characterizes the speed of common learning.

I want to stress two technical challenges in proving this proposition.

Firstly, memory function does not preserve the order in KL-distance. That is, take two empirical signal frequencies ν_{it}^1 and ν_{it}^2, any $\theta \in \Theta$, and any mild memory function m,

$$\mathrm{KL}\left(\nu_{it}^1, \mu_i^\theta\right) > \mathrm{KL}\left(\nu_{it}^2, \mu_i^\theta\right) \nRightarrow \mathrm{KL}\left(m(\nu_{it}^1), \mu_i^\theta\right) > \mathrm{KL}\left(m(\nu_{it}^2), \mu_i^\theta\right)$$

However, in proving Proposition 1, we want to characterize a tight threshold in the empirical signal frequencies such that if the signal frequencies that an agent observed are closer to the true distribution than this threshold, learning succeeds; whereas, if the observed signal frequencies are further away from the true distribution than this threshold, learning fails.

Take the speed of individual learning as an example. We would like to find a threshold in the empirical frequencies such that whenever signals are outside that threshold, after memory transformation, they are further away from the true distribution than the KL-midpoint between the true state and the most difficult-to-distinguish state, causing learning to fail. By Sanov's theorem from large deviation theory, this threshold characterizes the speed of individual learning. However, it seems impossible to find a tight threshold around the theoretical distribution when the memory function could alter the ranking in KL-distance.

To address this issue, we utilize the convexity of $\mathrm{KL}(\nu_{it}, \mu_i^\theta)$ in the pair$(\nu_{it}, \mu_i^\theta)$and the non-negativity of KL-distance, $\mathrm{KL}(\nu_{it}, \mu_i^\theta) \geq 0$, with equality if and only if$\nu_{it} = \mu_i^\theta$. We can construct a neighborhood around $\mu_i^\theta, \nu_{it}^\theta \in [\underline{\nu_{it}}, \bar{\nu_{it}}]$, satisfies:

$$\mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right) \leq \mathrm{KL}\left(m^{-1}(\hat{\nu}_i(\tilde{\theta}_i)), \mu_i^\theta\right) = \underline{\lambda_B}^\theta(\mathcal{I}, m)$$

Since the memory function is monotone in ν_{it}, after memory transformation,$\hat{\nu_i}(\tilde{\theta_i}) \notin [m(\underline{\nu_{it}}), m(\bar{\nu_{it}})]$. Together with the fact that $\hat{\nu_i}(\tilde{\theta_i})$ is by construction the memory-adjusted most tricky KL-midpoint, $\underline{\lambda_B}^\theta(\mathcal{I}, m)$ gives us the speed of individual learning.

Secondly, the memory function causes an asymmetrical shift in the observed signal frequencies. Imagine there is a neighborhood around $\mu_i^\theta, \nu_{it}^\theta \in [\underline{\nu_{it}}, \bar{\nu_{it}}]$such that $\mathrm{KL}(\nu_{it}, \mu_i^\theta) \leq \underline{\lambda_C}^\theta(\mathcal{I}, m)$ and at the boundary,

$$\mathrm{KL}(\underline{\nu_{it}}, \mu_i^\theta) = \mathrm{KL}(\bar{\nu_{it}}, \mu_i^\theta) = \underline{\lambda_C}^\theta(\mathcal{I}, m)$$

Since the memory function could distort high signals and low signals differently, after memory transformation, $\mathrm{KL}(m(\underline{\nu_{it}}), \mu_i^\theta) \neq \mathrm{KL}(m(\bar{\nu_{it}}), \mu_i^\theta)$. Despite being at either boundary, in expectation,

$$\mathrm{KL}\left(\mathbb{E}\left[\nu_{jt} \mid \theta, m(\underline{\nu_{it}})\right], \mu_j^\theta\right) < d\left(\mu_j^\theta, \mu_j^{\underline{\theta}}\right) \text{ and } \mathrm{KL}\left(\mathbb{E}\left[\nu_{jt} \mid \theta, m(\bar{\nu_{it}})\right], \mu_j^\theta\right) < d\left(\mu_j^\theta, \mu_j^{\bar{\theta}}\right)$$

We cannot make a general argument about how close $\mathbb{E}[\nu_{jt} \mid \theta, m(\underline{\nu_{it}})]$ is to the true distribution μ_j^θ. Instead, we argue that for all states other than the true state θ, in expectation, agent i believes that agent j's signal observation is at least further away than the closest Chernoff distance$d(\mu_j^\theta, \mu_j^{\tilde{\theta_j}})$. By the standard argument that beliefs at large t concentrate on states whose signal distributions minimize KL-divergence relative to the perceived empirical signal distribution, agent $i$is confident that agent j can only learn the true state θ.

One important lesson from Theorem 1 is that, with biased memory, common learning and individual learning can occur at different rates. A necessary and sufficient condition for common learning to happen slower than individual learning is the comparison of learning efficiency index, that is, when $\underline{\lambda_B}^\theta(\mathcal{I}, m) > \underline{\lambda_C^\theta}(\mathcal{I}, m)$. However, this condition can be very abstract. In order to provide clearer intuition of this learning result, Proposition 2 then introduces a less strict (only necessary but not sufficient) condition.

Proposition 2. *Fix any information structure $\mathcal{I}$, mild memory function m, state $\theta \in \Theta$,*

$$\underline{\lambda_B}^\theta(\mathcal{I}, m) > \underline{\lambda_C}^\theta(\mathcal{I}, m) \implies \tilde{\theta}_i \neq \tilde{\theta}_j$$

Proposition 2 states that for common learning to occur slower than individual learning, the most challenging state for individual learning ($\tilde{\theta}_i$) must differ from the most challenging state for common learning ($\tilde{\theta}_j$). When agent i estimates agent j's signal observations based on its own signal observations, according to Lemma 1, this estimate tends to be closer to the theoretical distribution under state θ in KL-distance than$i$'s own signal observations. If $\tilde{\theta}_i = \tilde{\theta}_j = \theta^*$, agent \(i\) is most concerned about mislearning the state θ^*, and in expectation, also most concerned about agent j mislearning the stateθ^*. Just as in the rational benchmark, Lemma 1 guarantees that the expectation effect brings the estimate of$j$'s signal frequencies closer to the theoretical distribution, which immediately implies further away from the distribution under the critical state θ^*. Thus, when$\tilde{\theta}_i = \tilde{\theta}_j$, Theorem 1 can nest Frick, Iijima and Ishii (2023) as a special case.

However, when$\theta_i \neq \theta_j$, agent$i$ is primarily concerned about mislearning the state $\tilde{\theta}_i$. Crucially, due to biased memory,$\tilde{\theta}_i$ could differ from the state that would otherwise be closest in KL-distance to state θ under the rational benchmark. Fortunately, the expectation effect brings the estimate of agent$j$'s signal frequencies closer to the theoretical distribution under the true state θ and away from $\tilde{\theta}_i$. Unfortunately, signals could exhibit negative correlation between the two agents, causing the estimates to also be closer to the theoretical distribution under$\tilde{\theta}_j$, adding more confusion in higher-order uncertainty. Note that this does not constitute a sufficient condition for creating a slower speed of common learning. To obtain the necessary and sufficient condition, one needs to compare the learning efficiency index.

4.3 Illustrative Example

In this section, we will provide a general intuition for the theorems and propositions in this paper using a numerical example. Consider a simple setting where there are three states, each characterized by frequencies of high signals: θ_1 with $\mu_i^1 = 0.3$, θ_2 with $\mu_i^2 = 0.5$, and θ_3 with $\mu_i^3 = 0.8$. Suppose that $\theta_2$is the true state.

4.3.1 Speed of Individual and Common Learning under Rational Benchmark

The canonical result from Frick, Iijima and Ishii (2023), establishes that in this setting, the speed of individual learning coincides with the speed of common learning. This speed is characterized by the Chernoff distance between the true state, θ_2, and the state whose signal distribution is closest to the true state, which is θ_1 in our example. It's worth noting that the Chernoff distance betweenθ_1 andθ_2 is equivalent to$\mathrm{KL}(\hat{\nu}_i(\theta_1), \mu_i^{\theta_2})$, where $\hat{\nu}_i(\theta_1)$ represents the KL-midpoint between $\theta_2$and θ_1, as constructed. Throughout this paper, we will predominantly use the formulation based on the KL-distance to maintain consistency with the notation.

Why does $\mathrm{KL}\left(\hat{\nu}_i(\theta_1), \mu_i^{\theta_2}\right)$characterize the speed of individual learning? Imagine a very unlucky agent i, whose empirical signal distribution after $t$signal draws, ν_{it}, has with very low high-signal realizations and is further away from the true distribution $\mu_i^{\theta_2}$ than the KL $-$midpoint between θ_2 and θ_1, that is $\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta_2}\right) > \mathrm{KL}\left(\hat{\nu}_i(\theta_1), \mu_i^{\theta_2}\right)$. Then, in such case, agent $i$observes signal distribution better matches with signals in θ_1 than with the true state θ_2 and agent $i$is going to update his/her belief about the state in favor of the wrong state θ_1. Thus, with large t, in order for agents to correctly learn the true state, the empirical signal distribution has to be close enough to the true distribution than to any other distributions. Hence, in this example, the event $B_t^p(\theta_2)$ is bounded by the KL $-$midpoints between $\theta_2$and θ_1 ($\hat{\nu}_i(\theta_1)$in Figure 1) and between $\theta_2$and θ_3 ($\hat{\nu}_i(\theta_3)$in Figure 1). For all large enough t, one can show that $B_t^p(\theta_2)$ is approximated by

$$B_t^p(\theta_2) \approx \left\{\nu_{it} \in \Big(0.4, 0.65\Big), \forall i = 1, 2\right\}$$

To go from the event$B_t^p(\theta_2)$ to the speed of individual learning, we apply the Sanov's theorem from large deviation theory. It states that, for any set $D \subseteq \Delta(S)$ that is the closure of its interior,

$$\mathbb{P}_t^{\mathcal{I}}\left(\nu_t \in D \mid \theta\right) = \exp\left[-\inf_{\nu \in D} \mathrm{KL}\left(\nu, \mu^{\theta}\right) t + o(t)\right]$$

That is, as t grows large, the probability of event D vanishes exponentially at rate given by the KL$-$distance between the closest (to theoretical signal distribution) element in D and the theoretical signal distribution μ^{θ}. Since in our setting,

$$\mathrm{KL}\left(\hat{\nu}_i(\theta_1), \mu_i^{\theta_2}\right) = \mathrm{KL}\Big((0.4, 0.6), (0.5, 0.5)\Big) < \mathrm{KL}\left((0.65, 0.35), (0.5, 0.5)\right) = \mathrm{KL}\left(\hat{\nu}_i(\theta_3), \mu_i^{\theta_2}\right)$$

the probability $\mathbb{P}_t^{\mathcal{I}}\left(\left(B_t^p(\theta_2)\right)^c \mid \theta_2\right)$ of first-order belief failures vanishes at rate $\mathrm{KL}\left((0.4, 0.6), (0.5, 0.5)\right)$, demonstrated by the blue curve in Figure 1.

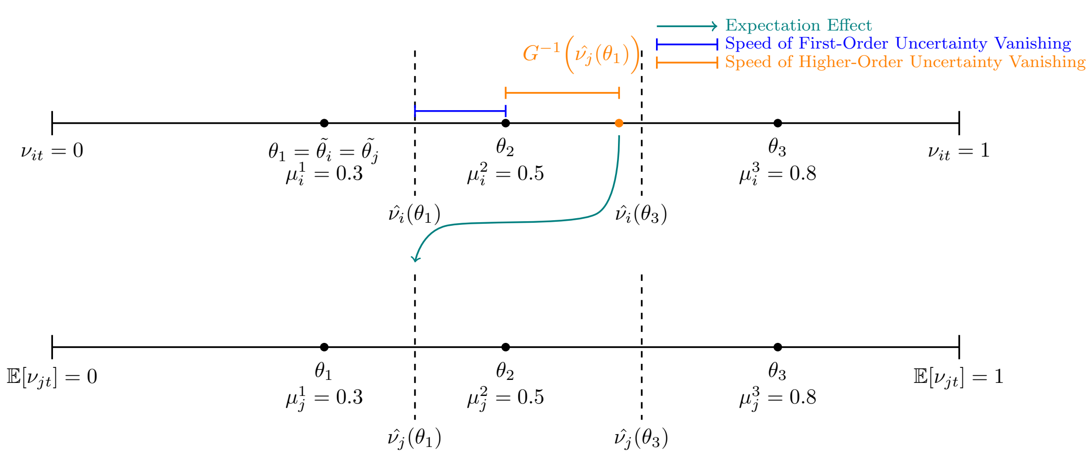

Figure 1: Speed of Individual and Common Learning under Rational Benchmark

In terms of the speed of common learning, Lemma 1 from Frick, Iijima and Ishii (2023) provides a very clear intuition: because agent i's expectation of agent j's signal distribution is closer to the theoretical distribution than agent i's own signal observation, whenever agent $i$privately learn the state, there will be no higher order uncertainty about j's learning outcome. To see it in our numerical example, we introduce a correlation parameter $\rho \in [0,1]$ to capture the extent of correlation across agents' signals. Then, fixing a state θ, the joint probabilities of agents' signals are given by the following table

θ	$s_i = H$	$s_i = L$
$s_j = H$	$\mu_i^\theta \rho$	$\mu_j^\theta (1-\rho)$
$s_j = L$	$\mu_i^\theta (1-\rho)$	$1 - \mu_j^\theta (2-\rho)$

Since the result is quite intuitive when agents' signals are positive correlated, we will demonstrate the case when agents' signals are negatively correlated. Assume, $\rho = 0.1$, under the true state θ_2, the joint probabilities of agents' signals are given by

$\theta = \theta_2$	$s_i = H$	$s_i = L$
$s_j = H$	0.05	0.45
$s_j = L$	0.45	0.05

Then, for any realized signal frequencies ν_{it} and large enough t, Lemma 2 from Cripps et al. (2008) shows that agent $i$becomes confident in θ_2, so i's belief about ν_{ji} concentrates on the expectation $\mathbb{E}[\nu_{jt} \mid \theta, \nu_{it}]$, approximately,

$$\mathbb{E}[\nu_{jt} \mid \theta, \nu_{it}] = \nu_{it}\frac{0.05}{0.5} + (1-\nu_{it})\frac{0.45}{0.5}$$

For agent $i$to be confident that agent j is also correctly learning the true state, we need

$$\mathbb{E}[\nu_{jt} \mid \theta, \nu_{it}] \in (0.4, 0.65) \Longrightarrow \nu_{it} \in (0.3125, 0.625)$$

and we have

$$C_t^p(\theta_2) \approx \Big\{\nu_{it} \in \Big(0.4, 0.65\Big) \cap \Big(0.3125, 0.625\Big), \forall i = 1, 2\Big\}$$

Realize that

$$\mathrm{KL}\left((0.3125, 0.6875), (0.5, 0.5)\right) > \mathrm{KL}\left((0.625, 0.375), (0.5, 0.5)\right)$$

Using Sanov's theorem, the probability $\mathbb{P}_t^{\mathcal{I}}\left(B_t^p(\theta_2) \backslash C_t^p(\theta_2) \mid \theta_2\right)$ of higher-order belief failures vanishes at rate $\mathrm{KL}\left((0.625, 0.375), (0.5, 0.5)\right)$, demonstrated by the orange curve in Figure 1, which is strictly larger than the blue curve. Thus, as $t$goes large, higher-order belief failures become negligible relative to first-order failures, and we have that the individual learning and common learning occur at the same rate.

4.3.2 Failure of Individual and Common Learning under Wild Memory

Once we consider that agents could possess biased memory about past signals that they observed, common learning could fail. We thus include definition 3 to distinguish two types of memory functions. To see intuition, consider the following wild memory function:

Example 1. $f_1(s_i = H) = 0.7$ and $f_2(s_i = H) = 0.3$. This is the case where high signals are more memorable for agent $i$than low signals.

The memory adjusted agent i's belief about signal distribution under state θ_2 is given by the bayesian updating rules:

$$m_1(\mu_i^2) = \frac{f_1(s_i = H)\mu_i^2(s_i = H)}{f_1(s_i = H)\mu_i^2(s_i = H) + f_1(s_i = L)\mu_i^2(s_i = L)} = \frac{0.7 \times 0.5}{0.7 \times 0.5 + 0.3 \times 0.5} = 0.7$$

In this case, the memory function brings the theoretical signal distribution under θ_2 to $m(\mu_i^2) = 0.7$, as demonstrated by the red dot $m_1(\mu_1^2)$ in the Figure 2. Notice that now

$$\mathrm{KL}\left(m(\mu_i^2), \mu_i^2\right) = \mathrm{KL}\left((0.7, 0.3), (0.5, 0.5)\right) > \mathrm{KL}\left((0.65, 0.35), (0.5, 0.5)\right) = \mathrm{KL}\left(\hat{\nu_i}(\theta_3), \mu_i^2\right)$$

This memory function is indeed "wild", as it violates part (1) of definition 3. With such wild memory function, individual learning fails. By a standard law of large number argument, as $t$goes large, agent i's belief about his signal observation under state θ_2 would be concentrated around $m(\mu_i^{\theta_2}) = 0.7$ but not in $\mu_i^2 = 0.5$. Since our agents are unaware of their biased memory and $m(\mu_i^2) = 0.7$ is already further away from μ_i^2 than the KL-mid point between $\theta_2$andθ_3, agent i would believe that he/she has observed signal distribution closer to θ_3 than to the true state θ_2, leading to a learning of the wrong state. The failure of individual learning immediately implies a failure of common learning.

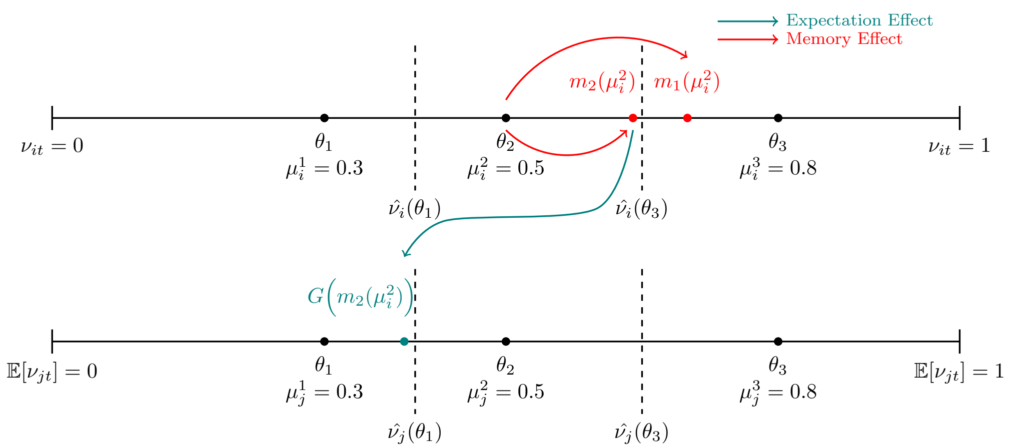

Figure 2: Failure of Individual and Common Learning under Wild Memory

Example 2. $f_2(s_i = H) = 0.8$, $f_2(s_i = H) = 0.45$, and $\rho = 0.1$. This is a case where high signals are more memorable and agents' signals are negatively correlated.

The memory adjusted agent i's belief about the signal distribution under state θ_2 is:

$$m_2(\mu_i^2) = \frac{f_2(s_i = H)\mu_i^2(s_i = H)}{f_2(s_i = H)\mu_i^2(s_i = H) + f_2(s_i = L)\mu_i^2(s_i = L)} = \frac{0.8 \times 0.5}{0.8 \times 0.5 + 0.45 \times 0.5} = 0.64$$

In expectation, agent i's belief about agent j's observed signal frequencies is concentrated on:

$$\mathbb{E}\left[\nu_j \mid \theta_2, m_2(\mu_i^2)\right] = 0.64 \times \frac{0.05}{0.5} + 0.36 \times \frac{0.45}{0.5} = 0.388$$

In this case, the memory function brings the theoretical signal distribution under θ_2 to $m(\mu_i^2) = 0.64$, as demonstrated by the red dot $m_2(\mu_1^2)$ in the Figure 2. This memory effect is not detrimental to the individual learning outcome, as agent $i$would still learn the true state θ_2 as $t$grows large. However, our agents are unaware of their biased memory and they update their belief about other agents' signals as if what they remembered are the only ones that they received. Thus, this expectation effect is directly applied to the memory adjusted theoretical signal distribution of agent i and it brings the expected agent j's signal frequencies to$\mathbb{E}\left[\nu_j \mid \theta_2, m_2(\mu_i^2)\right] = 0.388$.

Notice that now

$$\mathrm{KL}\left(G(m(\mu_i^2)), \mu_j^2\right) = \mathrm{KL}\left((0.388, 0.612), (0.5, 0.5)\right) > \mathrm{KL}\left((0.4, 0.6), (0.5, 0.5)\right) = \mathrm{KL}\left(\hat{\nu_j}(\theta_1), \mu_j^2\right)$$

This memory function is "wild", as it violates part (2) of definition 3. With such wild memory function, despite individual learning succeed, common learning fails . By a standard law of large number argument, as $t$goes large, agent i's memory adjusted expectation of agent j's signal observation under state θ_2 would be concentrated around $\mathbb{E}\left[\nu_j \mid \theta_2, m_2(\mu_i^2)\right] = 0.388$ and is further away from μ_j^2 than the KL-mid point between $\theta_1$andθ_2, agent i would believe that agent j has observed signal distribution closer to θ_1 than to the true state θ_2, leading to a failure of common learning.

4.3.3 Speed of Individual and Common Learning under Mild Memory

When agents have biased memory but the biasness is not extreme, it wouldn't be detrimental to the learning outcomes. For these mild memory bias, we are interested in the implication it has on the speed of learning.

Imagine the following mild memory function: $f(s_i = H) = 0.88$ and $f(s_i = H) = 0.72$. Assume that agents' signals are negatively correlated with $\rho = 0.1$.

With this mild memory function, conditioning on state θ_2, as $t$grows large, agent i's belief about his/her empirical signal distribution is going to concentrate on

$$m(\mu_i^2) = \frac{f(s_i = H)\mu_i^2(s_i = H)}{f(s_i = H)\mu_i^2(s_i = H) + f(s_i = L)\mu_i^2(s_i = L)} = \frac{0.88 \times 0.5}{0.88 \times 0.5 + 0.72 \times 0.5} = 0.55$$

The KL $-$midpoint between $\theta_2$and θ_1 is approximately $\hat{\nu}_i(\theta_1) = 0.4$ and the KL-midpoint between $\theta_2$and θ_3 is approximately $\hat{\nu}_i(\theta_3) = 0.65$. Among these two KL-midpoints, $\hat{\nu}_i(\theta_3)$ is closer to $m(\mu_i^2)$ in KL-distance. Thus, individually, agent $i$is most concerned about the uncertainty from θ_3 and the rate of individual learning is determined by how fast this uncertainty vanishes. In Theorem 1, we show that this rate is given by the inverse of memory function at the (memory adjusted) most trickyKL-midpoint. Numerically,

$$\frac{0.88 \times \nu_{it}}{0.88 \times \nu_{it} + 0.72 \times (1 - \nu_{it})} = 0.65 \implies \nu_{it} \approx 0.603$$

The rate of individual learning is given by:

$$\underline{\lambda_B}^{\theta_2}(\mathcal{I}, m) \approx \mathrm{KL}\Big((0.603, 0.397), (0.5, 0.5)\Big)$$

demonstrated by the blue curve in Figure 3.

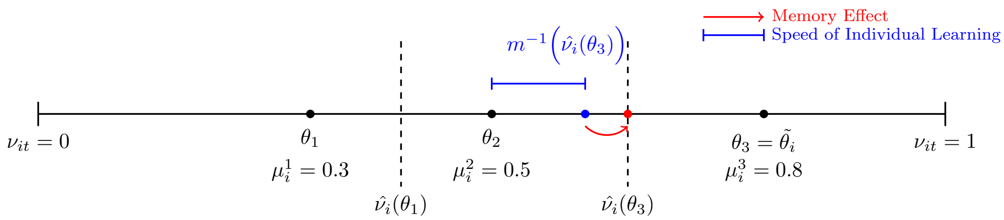

Figure 3: Speed of Individual Learning under Mild Memory

For the speed of common learning, we are concerned about agent i's memory adjusted expectation about agent j's observed signal distribution. As $t$grows large, agent i's belief about agent j's empirical signal distribution is going to concentrate on:

$$\mathbb{E}\left[\nu_j \mid \theta_2, m_2(\mu_i^2)\right] = 0.55 \times \frac{0.05}{0.5} + 0.45 \times \frac{0.45}{0.5} = 0.46$$

Among the two KL-midpoints for agent j, $\hat{\nu_j}(\theta_1)$ is closer to $\mathbb{E}\left[\nu_j \mid \theta_2, m_2(\mu_i^2)\right]$ in KL-distance. Thus, in expectation, agent $i$is most concerned about the uncertainty from agent j about state θ_1 and the rate of common learning is determined by how fast this uncertainty vanishes. In Theorem 1, we show that this rate is given by the inverse of the memory function at the inverse of the expectation effect at the (memory adjusted, in expectation) most tricky KL-midpoint for agent j. Numerically,

$$\Big(\frac{0.88 \times \nu_{it}}{0.88 \times \nu_{it} + 0.72 \times (1-\nu_{it})}\Big) \times \frac{0.05}{0.5} + \Big(1 - \frac{0.88 \times \nu_{it}}{0.88 \times \nu_{it} + 0.72 \times (1-\nu_{it})}\Big) \times \frac{0.45}{0.5} = 0.4 \implies \nu_{it} \approx 0.577$$

The rate of individual learning is given by:

$$\underline{\lambda_C}^{\theta_2}(\mathcal{I}, m) \approx \mathrm{KL}\Big((0.577, 0.423), (0.5, 0.5)\Big)$$

demonstrated by the orange curve in Figure 4.

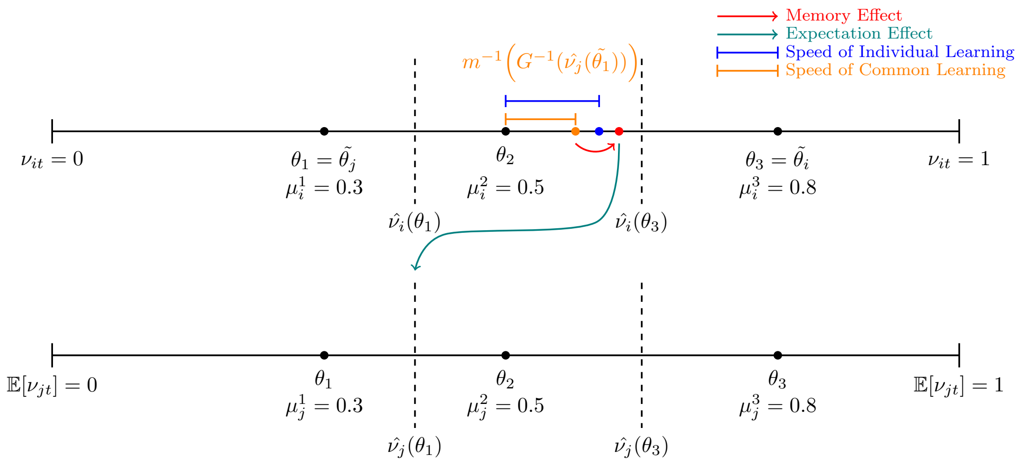

Figure 4: Speed of Individual and Common Learning under Mild Memory

Realize that this is a case of common learning happening slower than individual learning, as higher order uncertainty vanishes faster than first order uncertainty:

$$\underline{\lambda_B}^{\theta_2}(\mathcal{I}, m) \approx \mathrm{KL}\Big((0.603, 0.397), (0.5, 0.5)\Big) > \mathrm{KL}\Big((0.577, 0.423), (0.5, 0.5)\Big) \approx \underline{\lambda_C}^{\theta_2}(\mathcal{I}, m)$$

Proposition 2 applies in this case: under true state θ_2, individually, agent $i$is most concerned about mislearning the state θ_3, whereas in expectation, he/she is most concerned about agent j mislearning θ_1.

$$\tilde{\theta}_i = \theta_3 \neq \theta_1 = \tilde{\theta}_j$$

To see the intuition, let's first consider the case under rational benchmark, without biased memory. Agent i is primarily concerned about mislearning the state that is closest to the true state in terms of KL-distance, which in our case is stateθ_1. Successfully learning state θ_2 requires agents to observe signal frequencies close enough to the theoretical distribution under state$\theta_2$compared to the Chernoff distance between θ_2 and θ_1. At $\hat{\nu_i}(\theta_1)$, agent i is on the boundary where observing slightly lower signals can lead to a failure in individual learning. In this boundary case, by Lemma 1, agent i still expects agent j's observed signal distribution to be closer to θ_2 than i's own observed signals. This expectation effect is bounded by either perfectly positive correlation or perfectly negative correlation, as demonstrated by the teal brackets in Figure 5. Any $\mathbb{E}[\nu_j]$ within the teal brackets is closer to θ_2 and further away from either θ_1 or θ_3 than the boundary case, ensuring that when first-order uncertainty vanishes, there is no higher-order uncertainty.

In our example, due to biased memory, agents are more likely to remember high signals than low signals, which inflates agent i's belief upward. Consequently, agent i is more concerned about mislearning state θ_3 rather than mislearning θ_1 as in the rational benchmark. Despite this bias, in expectation (as demonstrated by the red brackets in Figure 5), agent j's signal distribution is still closer toθ_2 than agent i's own signal observation. However, since the Chernoff distance between θ_2 and θ_3 is larger than that between θ_2 and θ_1, for some boundary conditions near $\hat{\nu_i}(\theta_3)$, a negatively correlated signal structure could bring agent

i's estimates of agent j's signal observation closer to θ_1. This creates additional higher-order uncertainty stemming from θ_1.

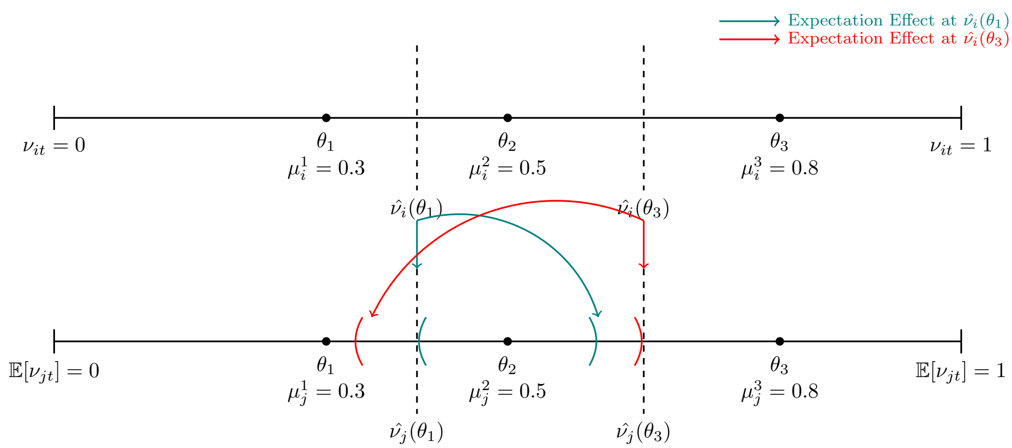

Figure 5: Expectation Effect at the Boundary

5 Conclusion

In conclusion, this paper explores the dynamics of learning in a multi-agent environment with biased memory. By extending existing frameworks and drawing on concepts from information theory, we investigate the speed of individual and common learning under biased memory functions. Our analysis reveals nuanced interactions between memory biases, signal correlations, and learning outcomes. Notably, we find instances where common learning can occur slower than individual learning due to the asymmetrical effects of memory distortion. These results shed light on the complexities of common learning processes and underscore the importance of accounting for cognitive biases in understanding collective decision-making.

To extend our model further, exploring the implications of these results in information design for coordination problems would be interesting. Given that higher-order uncertainty may vanish more slowly than first-order uncertainty, information designers aiming to facilitate learning or coordination among agents face challenges. Questions arise about the necessity of additional information regarding other players' signals, and if so, when and how this information should be provided.

Additionally, investigating social learning environments where players repeatedly choose actions following each signal draw could yield valuable insights. Observing other players' past signals may reveal information about their signals and about players' memory effects. Characterizing the speed of common learning in settings where agents can respond to others' strategic incentives and adjust their behavior based on revealed memory biases would be an intriguing avenue for future research.

Appendix: Proofs

A Proof of Proposition 1

we first prove the following claims

Claim 1. *There exist* $\kappa^{'} \in \left(0, \underline{\lambda_C}^{\theta}(\mathcal{I}) - \lambda\right)$ *and* $T' > 0$ *such that for all* $t \geq T'$, $\theta \in \Theta$ *and* $\theta^{'} \neq \theta$,

$$\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \leq \lambda + \kappa^{'} \Longrightarrow \mathbb{P}_t^{\mathcal{I}}\left(\mathrm{KL}\left(\mathbb{E}\left[\nu_{jt} \mid \theta, m(\nu_{it})\right], \mu_j^{\theta^{'}}\right) > d\left(\mu_j^{\theta}, \mu_j^{\tilde{\theta}_j}\right)\right) \geq \sqrt{p}$$

Proof. Notice first that, $\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right)$ is convex in the pair $(\nu_{it}, \mu_i^{\theta})$. To see that, take two pairs of probability mass functions (ν_{it}^1, μ_i^1) and (ν_{it}^2, μ_i^2), for all $\alpha \in [0,1]$, we can apply the log sum inequality property to obtain:

$$\left(\alpha\nu_{it}^1(s_i)+(1-\alpha)\nu_{it}^2(s_i)\right)\log\frac{\alpha\nu_{it}^1(s_i)+(1-\alpha)\nu_{it}^2(s_i)}{\alpha\mu_i^1(s_i)+(1-\alpha)\mu_i^2} \leq \alpha\nu_{it}^1(s_i)\log\frac{\alpha\nu_{it}^1(s_i)}{\alpha\mu_i^1(s_i)}+(1-\alpha)\nu_{it}^2(s_i)\log\frac{(1-\alpha)\nu_{it}^2(s_i)}{(1-\alpha)\mu_i^2(s_i)}$$

Summing over all s_i, we get

$$\mathrm{KL}\left(\alpha\nu_{it}^1 + (1-\alpha)\nu_{it}^2, \alpha\mu_i^1 + (1-\alpha)u_i^2\right) \leq \alpha\mathrm{KL}\left(\nu_{it}^1, \mu_i^1\right) + (1-\alpha)\mathrm{KL}\left(\nu_{it}^2, \mu_i^2\right)$$

Also notice that KL-distance is nonnegative,$\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \geq 0$, with equality if and only if $\nu_{it} = \mu_i^{\theta}$ and is continuous in ν_{it}. Thus, for each λ, we can find a neighborhood around μ_i^{θ}, $\nu_{it}^{\theta} \in [\underline{\nu_{it}}, \bar{\nu_{it}}]$ satisfies $\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \leq \lambda$. Because we are choosing $\lambda \in \left(0, min\{\underline{\lambda_B}^{\theta}(\mathcal{I}), \underline{\lambda_C}^{\theta}(\mathcal{I})\}\right) \leq \underline{\lambda_C}^{\theta}(\mathcal{I}) \leq \lambda_{ij}^{\theta}(\mathcal{I})$, the empirical distribution ν_{it}^{*}such that $\mathrm{KL}\left(\nu_{it}^{*}, \mu_i^{\theta}\right) = \lambda_{ij}^{\theta}(\mathcal{I}) = \mathrm{KL}\left(m^{-1}\left(G^{-1}(\hat{\nu_j}(\tilde{\theta_j}))\right), \mu_i^{\theta}\right)$ does not reside in the neighborhood of $[\underline{\nu_{it}}, \bar{\nu_{it}}]$.

Next, realize that memory function is monotone in ν_{it}^{θ}. To see that, take any $\nu_{it}^1, \nu_{it}^2 \in \Delta(S_i)$, if $\nu_{it}^1(s_i = H) > \nu_{it}^2(s_i = H)$, then

$$m(\nu_{it}^1(s_i = H))^{-1} = f(s_i = L)(1-\nu_{it}^1(s_i = H)) \leq f(s_i = L)(1-\nu_{it}^2(s_i = H)) = m(\nu_{it}^2(s_i = H))^{-1}$$

Since $m(\cdot) \geq 0$, this gives us that $m(\nu_{it}^1(s_i = H)) \geq m(\nu_{it}^2(s_i = H))$.

Thus, we also have that $G^{-1}(\hat{\nu_j}(\tilde{\theta_j}))$ do not reside in the neighborhood of$[m(\underline{\nu_{it}}), m(\bar{\nu_{it}})]$ and

$$\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \leq \lambda \Longrightarrow \mathrm{KL}\left(m(\nu_{it}), \mu_i^{\theta}\right) \leq \mathrm{KL}\left(G^{-1}(\hat{\nu_j}(\tilde{\theta_j})), \mu_i^{\theta}\right)$$

By the chain rule for KL −divergence, we have

$$\begin{aligned}\mathrm{KL}\left(m(\nu_{it})G_{ij}^{\theta}, \mu_i^{\theta}G_{ij}^{\theta}\right) &= \mathrm{KL}\left(m(\nu_{it}), \mu_i^{\theta}\right) + \sum_{\nu_{it}\in\mathrm{supp}(\nu_i)} m(\nu_{it})\mathrm{KL}\left(G_{ij}^{\theta}, G_{ij}^{\theta}\right) \\ &= \mathrm{KL}\left(m(\nu_{it}), \mu_i^{\theta}\right)\end{aligned}$$

Thus, the matrix transformation of G_{ij}^{θ} preserves the order in KL $-$distance, and we have:

$$\mathrm{KL}\left(m(\nu_{it}), \mu_i^{\theta}\right) \leq \mathrm{KL}\left(G^{-1}(\hat{\nu_j}(\tilde{\theta_j})), \mu_i^{\theta}\right) \Longrightarrow \mathrm{KL}\left(m(\nu_{it})G_{ij}^{\theta}, \mu_i^{\theta}G_{ij}^{\theta}\right) \leq \mathrm{KL}\left(\hat{\nu_j}(\tilde{\theta_j}), \mu_j^{\theta}\right)$$

Recall that $\hat{\nu_j}(\tilde{\theta_j})$ is the KL-midpoint between θ and $\tilde{\theta_j}$, then for $\nu_{it}^{\theta} \in [\underline{\nu_{it}}, \bar{\nu_{it}}]$, it is evident that

$$\mathrm{KL}\left(m(\nu_{it})G_{ij}^{\theta}, \mu_i^{\theta}G_{ij}^{\theta}\right) \leq \mathrm{KL}\left(\hat{\nu_j}(\tilde{\theta_j}), \mu_i^{\theta}\right) = d\left(\mu_i^{\theta}, \mu_i^{\tilde{\theta_j}}\right) \leq \mathrm{KL}\left(m(\nu_{it})G_{ij}^{\theta}, \mu_j^{\tilde{\theta_j}}\right)$$

Notice that by construction, $\tilde{\theta_j}$ is the state agent i expects agent j to find most difficult to distinguish from state θ. Then, for all other states $\theta^{'} \in \Theta\backslash\{\theta, \tilde{\theta_j}\}$,

$$d\left(\mu_j^{\theta^{'}}, \mu_j^{\tilde{\theta_j}}\right) < \mathrm{KL}\left(m(\nu_{it})G_{ij}^{\theta^{'}}, \mu_j^{\tilde{\theta_j}}\right) < \mathrm{KL}\left(m(\nu_{it})G_{ij}^{\theta^{'}}, \mu_j^{\theta^{'}}\right)$$

Since $\mathrm{KL}\left(\cdot, \mu_i\right)$ is continuous for each full-support $\mu_i \in \Delta\left(S_i\right)$ and $\Delta\left(S_i\right)$ is compact, there exists $\eta > 0$ such that for all $j \neq i$, $\nu_i \in \Delta\left(S_i\right)$, $\theta \in \Theta$, *and* $\theta^{'} \neq \theta$,

$$\mathrm{KL}\left(\nu_i, \mu_i^{\theta}\right) \leq \lambda \Longrightarrow \mathrm{KL}\left(m(\nu_i)G_{ij}^{\theta}, \mu_j^{\theta^{'}}\right) > d\left(\mu_j^{\theta}, \mu_j^{\tilde{\theta_j}}\right) + \frac{\eta}{2}$$

Given this, there exists $\kappa^{'} \in \left(0, \underline{\lambda_C}^{\theta}(\mathcal{I}) - \lambda\right)$ such that,

$$\mathrm{KL}\left(\nu_i, \mu_i^{\theta}\right) \leq \lambda + \kappa^{'} \Longrightarrow \mathrm{KL}\left(m(\nu_i)G_{ij}^{\theta}, \mu_j^{\theta^{'}}\right) > d\left(\mu_j^{\theta}, \mu_j^{\tilde{\theta_j}}\right) + \eta$$

Moreover, since agents are unaware of their biased memory and they take expectation only based on their remembered signal realization, there exists $\varepsilon^{'} > 0$ such that,

$$\left[\mathrm{KL}\left(\nu_i, \mu_i^{\theta'}\right) \leq \lambda + \kappa^{'} \text{ and } \left\|\nu_i G_{ij}^{\theta'} - \nu_j\right\| \leq \varepsilon^{'}\right] \Longrightarrow \mathrm{KL}\left(\mathbb{E}\left[\nu_{jt} \mid \theta^{'}, m(\nu_{it})\right], \mu_j^{\theta^{''}}\right) > d\left(\mu_j^{\theta^{'}}, \mu_j^{\tilde{\theta_j}}\right)$$

Together with Lemma 2, this proves Claim 1. $\square$

Claim 2. There exists $\kappa^{''} \in \left(0, \underline{\lambda_B}^{\theta}(\mathcal{I}) - \lambda\right)$ *and* T'' *such that for all* $t \geq T''$ *and* $i \in N$,

$$\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \leq \lambda \quad \Longrightarrow \quad \mathbb{P}_t^{\mathcal{I}}\left(\left\{\theta : \mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \leq \lambda + \kappa^{''}\right\} \mid s_i^t\right) \geq \sqrt{p}$$

Proof. Take any $t \geq 1$ ands_i^t such that $\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \leq \lambda$, then for any $\theta' \neq \theta$, we have $\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta^{'}}\right) > \lambda + \kappa^{''}$ and

$$\begin{aligned} \log \frac{\mathbb{P}_t^{\mathcal{I}}\left(\theta' \mid s_i^t, m\right)}{\mathbb{P}_t^{\mathcal{I}}\left(\theta \mid s_i^t, m\right)} &= \log \frac{p_0\left(\theta'\right)}{p_0(\theta)} + t \sum_{s_i \in S_i} m(\nu_{it}) \log \frac{\mu_i^{\theta'}\left(s_i\right)}{\mu_i^{\theta}\left(s_i\right)} \\ &= \log \frac{p_0\left(\theta'\right)}{p_0(\theta)} + t\left(\mathrm{KL}\left(m(\nu_{it}), \mu_i^{\theta}\right) - \mathrm{KL}\left(m(\nu_{it}), \mu_i^{\theta^{'}}\right)\right) \end{aligned}$$

Denote $\eta := \left(\mathrm{KL}\left(m(\nu_{it}), \mu_i^{\theta}\right) - \mathrm{KL}\left(m(\nu_{it}), \mu_i^{\theta'}\right)\right)$ and take exponential on both side gives us:

$$\frac{\mathbb{P}_t^{\mathcal{I}}\left(\theta' \mid s_i^t, m\right)}{\mathbb{P}_t^{\mathcal{I}}\left(\theta \mid s_i^t, m\right)} \leq \frac{p_0\left(\theta'\right)}{p_0(\theta)} e^{t\eta}$$

Next, we show that $\eta < 0$. we can find the neighborhood $\nu_{it}^{\theta} \in [\underline{\nu_{it}}, \bar{\nu_{it}}]$ satisfies $\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \leq \lambda$. Because we are choosing $\lambda \in \left(0, min\{\underline{\lambda_B}^{\theta}(\mathcal{I}), \underline{\lambda_C}^{\theta}(\mathcal{I})\}\right) \leq \underline{\lambda_B}^{\theta}(\mathcal{I}) \leq \lambda_i^{\theta}(\mathcal{I})$, the empirical distribution such that $\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) = \lambda_i^{\theta}(\mathcal{I}) = \mathrm{KL}\left(m^{-1}(\hat{\nu}_i(\tilde{\theta})), \mu_i^{\theta}\right)$ does not reside in the neighborhood of $[\underline{\nu_{it}}, \bar{\nu_{it}}]$. Using similar arguments as in claim 1, for any empirical distribution within $\nu_{it}^{\theta} \in [\underline{\nu_{it}}, \bar{\nu_{it}}]$ that gives us $\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \leq \lambda$, we have $\mathrm{KL}\left(m(\nu_{it}), \mu_i^{\theta}\right) < \mathrm{KL}\left(m(\nu_{it}), \mu_i^{\tilde{\theta}}\right) < \mathrm{KL}\left(m(\nu_{it}), \mu_i^{\theta'}\right)$.

We can sum over $\theta^{'} \neq \theta$ and rearrange to obtain:

$$\frac{\mathbb{P}_t^{\mathcal{I}}\left(\theta \mid s_i^t, m\right)}{1 - \mathbb{P}_t^{\mathcal{I}}\left(\theta \mid s_i^t, m\right)} \geq \frac{p_0\left(\theta\right)}{1 - p_0\left(\theta\right)} e^{-t\eta}$$

Hence, by choosing $T'' > 0$ large enough, we get Claim 2. □

Combining these two claims, choose $T = \max\{T', T''\}$ and $\kappa = \min\{\kappa', \kappa''\}$ then, whenever $t \geq T$ and $\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \leq \lambda$, we have,

$$\begin{aligned}
&\mathbb{P}_t^{\mathcal{I}}\left(\left(\{\theta\} \cap \mathrm{KL}\left(\mathbb{E}\left[\nu_{jt} \mid \theta, m(\nu_{it})\right], \mu_j^{\theta^{'}}\right) > d\left(\mu_j^{\theta}, \mu_j^{\tilde{\theta}_j}\right)\right) \mid s_i^t\right) \\
&\geq \mathbb{P}_t^{\mathcal{I}}\left(\mathrm{KL}\left(\mathbb{E}\left[\nu_{jt} \mid \theta, m(\nu_{it})\right], \mu_j^{\theta^{'}}\right) > d\left(\mu_j^{\theta}, \mu_j^{\tilde{\theta}_j}\right) \mid s_i^t, \theta\right) \mathbb{P}_t^{\mathcal{I}}\left(\theta \mid s_i^t\right) \\
&\geq \sqrt{p} \times \mathbb{P}_t^{\mathcal{I}}\left(\theta' \mid s_i^t\right) \\
&\geq p
\end{aligned}$$

where the second inequality follows from Claim 1 and the last inequality follows from Claim 2.

B Proof of Theorem 1

Fix any state $\theta \in \Theta$, memory function is wild, if either (1) $\mathrm{KL}\left(m(\mu_i^{\theta}), \mu_i^{\theta}\right) \geq \mathrm{KL}\left(\hat{\nu}_i(\tilde{\theta}_i), \mu_i^{\theta}\right)$ or (2) $\mathrm{KL}\left(m(\mu_i^{\theta}), \mu_i^{\theta}\right) \geq \mathrm{KL}\left(G^{-1}(\hat{\nu}_j(\tilde{\theta}_j)), \mu_i^{\theta}\right)$.

By a law of large numbers argument, agent $i's$ memory-adjusted belief about ν_{it} would concentrate on the memory-adjusted theoretical signal distribution $m(\mu_i^{\theta})$. Under case (1), agent $i's$ belief is further away, in $\mathrm{KL}-$distance, from μ_i^{θ} than $\hat{\nu}_i(\tilde{\theta}_i)$. Recall that $\hat{\nu}_i(\tilde{\theta}_i)$ is the $\mathrm{KL}-$midpoint between the signal distribution under true state $\theta$and the state $\tilde{\theta}_i$ that agent i with biased memory finds the most challenging to distinguish, we have that

$$\mathrm{KL}\left(m(\mu_i^{\theta}), \mu_i^{\theta}\right) \geq \mathrm{KL}\left(\hat{\nu}_i(\tilde{\theta}_i), \mu_i^{\theta}\right) > \mathrm{KL}\left(m(\mu_i^{\theta}), \mu_i^{\tilde{\theta}_i}\right)$$

This implies, at large t, agent i believes that he or she received signal observation closer to state $\tilde{\theta}_i$ than to the true state θ, which led to a failure of learning the true state θ.

If we are not in case (1) but under case (2), then agent $i$is individually confident in state $\theta$and its belief about agent j's signal observation concentrates in $m(\mu_i^\theta)G_{ij}^\theta$. As we showed in claim 1 of Proposition 1, the matrix transformation of G_{ij}^θ preserves the order in KL $-$distance, we thus have that:

$$\mathrm{KL}\left(m(\mu_i^\theta), \mu_i^\theta\right) \geq \mathrm{KL}\left(G^{-1}(\hat{\nu_j}(\tilde{\theta_j})), \mu_i^\theta\right) \Longrightarrow \mathrm{KL}\left(m(\mu_i^\theta)G_{ij}^\theta, \mu_j^\theta\right) \geq \mathrm{KL}\left(\hat{\nu_j}(\tilde{\theta_j}), \mu_j^\theta\right)$$

This implies, at large t, even if agent $i$is individually confident in state θ, his expectation of agent j's signal observation is closer to state $\tilde{\theta_j}$ than to the true state θ, which led to a failure of common learning the true state θ.

To see the speed of common learning, take any $\lambda \in \left(0, \min\{\underline{\lambda_B}^\theta(\mathcal{I}), \underline{\lambda_C}^\theta(\mathcal{I})\}\right)$, apply Proposition 1 to state θ. Then, Claim 2 of Proposition 1 tells us that for each $i \in N$, there exists $T > 0$ such that, for all $t \geq T$, whenever $\mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right) \leq \lambda$, for $p$arbitrarily close to 1,

$$\mathbb{P}_t^{\mathcal{I}}\left(\left\{\theta : \mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right) \leq \underline{\lambda_B}^\theta(\mathcal{I})\right\} \mid s_i^t\right) \geq p$$

State $\theta$is thus individually $p-$believed.

Claim 1 of Proposition 1 tells us that for $j \neq i$ and all $t \geq T$, whenever $\mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right) \leq \lambda$, in agent i's expectation, agent $j's$ signal distribution is further away from all $\theta^{'} \neq \theta$. Together, Proposition 1 implies that, for all $t \geq T$, whenever $\mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right) \leq \lambda$, state $\theta$is $p-$evident for all agents. Therefore,

$$\begin{aligned}\limsup_{t\to\infty} \frac{1}{t}\log\left(1 - \mathbb{P}_t^{\mathcal{I}}\left(C_t^p(\theta) \mid \theta\right)\right) &\leq \limsup_{t\to\infty} \frac{1}{t}\log\left(\sum_i \mathbb{P}_t^{\mathcal{I}}\left(\left\{\mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right) > \lambda\right\} \mid \theta\right)\right) \\ &= \max_i \limsup_{t\to\infty} \frac{1}{t}\log \mathbb{P}_t^{\mathcal{I}}\left(\left\{\mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right) > \lambda\right\} \mid \theta\right) \\ &= -\lambda\end{aligned}$$

where the last equality follows from Sanov's theorem. Since it holds for all $\lambda \in \left(0, \min\{\underline{\lambda_B}^\theta(\mathcal{I}), \underline{\lambda_C}^\theta(\mathcal{I})\}\right)$, this establishes that

$$\limsup_{t\to\infty} \frac{1}{t}\log\left(1 - \mathbb{P}_t^{\mathcal{I}}\left(C_t^p(\theta) \mid \theta, m\right)\right) \leq -\min\{\underline{\lambda_B}^\theta(\mathcal{I}), \underline{\lambda_C}^\theta(\mathcal{I})\}$$

For the speed of individual learning, take $i \in N$ and any $\lambda > \underline{\lambda_B}^\theta(\mathcal{I})$, following a similar argument from Claim 2 of Proposition 1, we can construct a neighborhood around μ_i^θ, $\nu_{it} \in [\underline{\nu_{it}}, \bar{\nu_{it}}]$ satisfies $\mathrm{KL}\left(\nu_{it}, \mu_i^\theta\right) \leq \underline{\lambda_B}^\theta(\mathcal{I})$. For all $\lambda > \underline{\lambda_B}^\theta(\mathcal{I})$, we can find a set of signal distributions $\nu_{it}^{'} \notin [\underline{\nu_{it}}, \bar{\nu_{it}}]$ with $\underline{\lambda_B}^\theta(\mathcal{I}) < \mathrm{KL}\left(\nu_{it}^{'}, \mu_i^\theta\right) < \lambda$, such that

$$\mathrm{KL}\left(\hat{\nu_i}(\tilde{\theta_i}), \mu_i^\theta\right) < \mathrm{KL}\left(m(\nu_{it}^{'}), \mu_i^\theta\right) \Longrightarrow \mathrm{KL}\left(m(\nu_{it}), \mu_i^{\tilde{\theta_i}}\right) < \mathrm{KL}\left(m(\nu_{it}^{'}), \mu_i^\theta\right)$$

Then, for all large enough t,$B_{it}^p(\theta) \cap \left\{\nu_{it}^{'}\right\} = \emptyset$, because by standard arguments,i 's beliefs at large $t$concentrate on states whose signal distributions minimize KL$-$divergence relative to ν_{it}. Thus,

$$\liminf_{t\to\infty} \frac{1}{t} \log\left(1-\mathbb{P}_t^{\mathcal{I}}\left(B_{it}^{p}(\theta) \mid \theta\right)\right) \leq \liminf_{t\to\infty} \frac{1}{t} \log \mathbb{P}_t^{\mathcal{I}}\left(\left\{\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) > \lambda\right\} \mid \theta\right) = -\lambda$$

where the last equality follows from Sanov's theorem. Since it holds for all $\lambda \in \left(0, \underline{\lambda_B}^{\theta}(\mathcal{I})\right)$, this establishes that

$$\limsup_{t\to\infty} \frac{1}{t} \log\left(1-\mathbb{P}_t^{\mathcal{I}}\left(B_{t}^{q}(\theta) \mid \theta, m\right)\right) \leq \underline{\lambda_B}^{\theta}(\mathcal{I})$$

C Proof of Proposition 2

Proof by contrapositive: suppose that $\tilde{\theta}_i = \tilde{\theta}_j$, denote it θ^*.

Thus, for each $\lambda < \underline{\lambda_B}^{\theta}(\mathcal{I}, m)$, we can find a neighborhood around μ_i^{θ}, $\nu_{it}^{\theta} \in [\underline{\nu_{it}}, \bar{\nu_{it}}]$ satisfies $\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \leq \lambda$. Because we are choosing $\lambda < \underline{\lambda_B}^{\theta}(\mathcal{I}, m)$, the empirical distribution ν_{it}^{*} such that $\mathrm{KL}\left(\nu_{it}^{*}, \mu_i^{\theta}\right) = \underline{\lambda_B}^{\theta}(\mathcal{I}, m) = \mathrm{KL}\left(m^{-1}(\hat{\nu_i}(\theta^*)), \mu_i^{\theta}\right)$ does not reside in the neighborhood of $[\underline{\nu_{it}}, \bar{\nu_{it}}]$.

Recall that memory function is monotone in ν_{it}, we thus have that$\hat{\nu_i}(\theta^*)$ do not reside in the neighborhood of$[m(\underline{\nu_{it}}), m(\bar{\nu_{it}})]$, and

$$\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \leq \lambda \Longrightarrow \mathrm{KL}\left(m(\nu_{it}), \mu_i^{\theta}\right) \leq \mathrm{KL}\left(\hat{\nu_i}(\theta^*), \mu_i^{\theta}\right)$$

Applying Lemma 1, we have that

$$\mathrm{KL}\left(\mathbb{E}\left[\nu_{jt} \mid \theta, m(\nu_{it})\right], \mu_j^{\theta}\right) < \mathrm{KL}\left(m(\nu_{it}), \mu_i^{\theta}\right) \leq \mathrm{KL}\left(\hat{\nu_i}(\theta^*), \mu_i^{\theta}\right)$$

Thus, whenever $\mathrm{KL}\left(\nu_{it}, \mu_i^{\theta}\right) \leq \lambda$, higher order uncertainty about $\theta$and θ^* vanishes faster than first order uncertainty.

By definition of $\tilde{\theta}_j$ being the most tricky state in expectation and the assumption that $\tilde{\theta}_i = \tilde{\theta}_j = \theta^*$

$$\begin{aligned}\mathrm{KL}\left(\mathbb{E}\left[\nu_{jt} \mid \theta, m(\nu_{it})\right], \mu_j^{\theta}\right) &< \mathrm{KL}\left(\mathbb{E}\left[\nu_{jt} \mid \theta, m(\nu_{it})\right], \mu_j^{\theta^*}\right) \\ &< \mathrm{KL}\left(\mathbb{E}\left[\nu_{jt} \mid \theta, m(\nu_{it})\right], \mu_j^{\theta'}\right) \quad \forall \theta' \in \Theta \backslash \{\theta, \theta^*\}\end{aligned}$$

Therefore, higher order uncertainty about state other than θ^* could only vanish even faster. This argument holds for all $\lambda \in \left(0, \underline{\lambda_B}^{\theta}(\mathcal{I})\right)$.

It's then straightforward to see that

$$\underline{\lambda_B}^{\theta}(\mathcal{I}, m) \geq \underline{\lambda_C}^{\theta}(\mathcal{I}, m)$$

References

Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz. 2016. "Fragility of asymptotic agreement under Bayesian learning." *Theoretical Economics*, 11(1): 187–225.

Berk, Robert H. 1966. "Limiting Behavior of Posterior Distributions when the Model is Incorrect." *The Annals of Mathematical Statistics*, 37(1): 51 – 58.

Carlsson, Hans, and Eric van Damme. 1993. "Global Games and Equilibrium Selection." *Econometrica*, 61(5): 989–1018.

Cripps, Martin, Jeffrey Ely, George Mailath, and Larry Samuelson. 2013. "Common learning with intertemporal dependence." *International Journal of Game Theory*, 42(1): 55–98.

Cripps, Martin W., Jeffrey C. Ely, George J. Mailath, and Larry Samuelson. 2008. "Common Learning." *Econometrica*, 76(4): 909–933.

Dasaratha, Krishna, and Kevin He. 2023. "Aggregative Efficiency of Bayesian Learning in Networks."

Esponda, Ignacio, and Demian Pouzo. 2016. "BerkâNash Equilibrium: A Framework for Modeling Agents With Misspecified Models." *Econometrica*, 84(3): 1093–1130.

Frick, Mira, Ryota Iijima, and Yuhta Ishii. 2023. "Learning Efficiency of Multiagent Information Structures." *Journal of Political Economy*, 131(12): 3377–3414.

Fudenberg, Drew, Giacomo Lanzani, and Philipp Strack. 0. "Selective Memory Equilibrium." *Journal of Political Economy*, 0(ja): null.

Hann-Caruthers, Wade, Vadim V. Martynov, and Omer Tamuz. 2018. "The speed of sequential asymptotic learning." *Journal of Economic Theory*, 173: 383–409.

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

Monderer, Dov, and Dov Samet. 1989. "Approximating common knowledge with common beliefs." *Games and Economic Behavior*, 1(2): 170–190.

Moscarini, Giuseppe, and Lones Smith. 2002. "The Law of Large Demand for Information." *Econometrica*, 70(6): 2351–2366.

Rosenberg, Dinah, and Nicolas Vieille. 2019. "On the Efficiency of Social Learning." *Econometrica*, 87(6): 2141–2168.

Rubinstein, Ariel. 1989. "The Electronic Mail Game: Strategic Behavior under "Almost Common Knowledge."." *American Economic Review*, 79(3): 385–91.

Steiner, Jakub, and Colin Stewart. 2011. "Communication, timing, and common learning." *Journal of Economic Theory*, 146(1): 230–247.

Vives, Xavier. 1993. "How Fast do Rational Agents Learn?" *The Review of Economic Studies*, 60(2): 329–347.

Walters, Daniel J., and Philip M. Fernbach. 2021. "Investor memory of past performance is positively biased and predicts overconfidence." *Proceedings of the National Academy of Sciences*, 118(36): e2026680118.

Weinstein, Jonathan, and Muhamet Yildiz. 2007. "A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements." *Econometrica*, 75(2): 365–400.

Zimmermann, Florian. 2020. "The Dynamics of Motivated Beliefs." *American Economic Review*, 110(2): 337–61.