Social Learning in General Information Environments*

Wenji Xu†

Job Market Paper

December 18, 2019

Please find the latest version here.

Abstract

I study a social learning model in which agents make decisions sequentially and learn about an unknown payoff-relevant state through two sources – a signal about the state itself (a state-signal) and a signal about the actions taken by previous agents (an action-signal). Our objective is to provide general conditions on the action-signals that lead the agents to eventually behave as if they know the state, i.e., that lead to information aggregation. When the agents' action-signals are what I call *weakly separating*, it is shown that information aggregation occurs when the agents' state-signals are unboundedly informative in the sense of Smith and Sorensen (2000). This result provides a unifying criterion to evaluate when information aggregation occurs. I also provide sufficient conditions for information aggregation when the state-signals are boundedly informative, and necessary conditions for information aggregation. The theory is illustrated with applications to privacy protection on financial platforms, regulation of third-party information provision in social learning environments, and the design of social learning environments when agents suffer "source amnesia".

Keywords: Information aggregation, Learning, Social networks, Disclosure

JEL Code: D47, D82, D83

*I thank especially to my advisor, Philip Reny, and my thesis committee, Benjamin Brooks, Emir Kamenica, and Roger Myerson, for their constant support and guidance. I thank Ozan Candogan, William Cong, Elliot Lipnowski, Andrew Gianou, Andrew Mcclellan, Doron Ravid, Konstantin Sonin, Daniel Rappoport, Kai Hao Yang, and Weijie Zhong as well as participants at conferences for valuable comments and discussion. All errors are my own.

†Department of Economics, The University of Chicago. Email: wenji@uchicago.edu.

1 Introduction

In many economic and social settings, people have limited private information about a payoff-relevant state of the world and rely on their information of the actions of others as a vital source of additional information. For instance, in assessing the quality of a new product, consumers draw inferences from other consumers' purchasing decisions; in judging whether a political candidate is suitable for office, voters obtain information from the voting behavior of others. Likewise, in financial markets, investors learn about the market fundamentals by observing other investors' behavior.

A classic literature studies the possibility of information aggregation in such settings, characterizing conditions under which individuals are able to take the correct state-dependent actions in the long run.[1] A common assumption in the literature is that individuals are connected through a network. In particular, each individual is assumed to know the identities of his neighbors in the network and perfectly observe his neighbors' actions.[2] These network models are suited for settings where learning happens through person-to-person interactions. However, people are often situated in more general environments where they are imperfectly informed about others' identities and actions – for example, consumers consult sales volumes or bestseller lists to learn about other consumers' purchase decisions, while investors examine financial reports to learn about other investors' decisions. A comprehensive understanding of social learning in general information environments is lacking, which greatly limits the applicability of the results that the literature reports.

This paper provides a framework to analyze social learning in general information environments. I consider a large number of agents, who arrive sequentially and obtain information about the state of the world from two sources: a signal about the state itself (a state-signal) and a signal about the actions taken by previous agents (an action-signal). I show that information aggregation occurs if the state-signals can induce extreme posterior beliefs and the action-signals are informative enough to *separate* certain types of action-histories (this property of the action-signals is formally captured in two closely related concepts of "strong separation" and "weak separation"). The characterization result provides a unifying criterion to evaluate whether information aggregation occurs in social learning. In addition to environments studied in the prior literature, the result implies that information aggregation can occur in many under-researched environments – e.g., those in which each agent observes

[1] For surveys, see Bikhchandani et al. (1998); Chamley (2004); Golub and Sadler (2016).

[2] Most of the literature focuses on the special case in which agents observe all past actions, i.e., they are connected in a complete network.

certain coarse summary statistic about past actions.

The theory is illustrated with three important applications. First, I consider how digital platforms can design its information disclosure to ensure information aggregation and, at the same time, protect the privacy of the users. I construct an action-signal information structure that features “dynamically-adjusted benchmarks”. I show that this information structure satisfies my definition of weak separation, and thus give rise to information aggregation. At the same time, the information structure is less informative than the widely-used disclosure method of releasing the volumes of prior actions.[3] Thus, it is better at protecting the users' privacy.

Second, I study the regulation of third party information provision in social learning environments. I show that when action-signals satisfy strong separation, the necessary and sufficient condition for information aggregation to be robust to third party manipulation is that the state-signals are unboundedly informative. This result is then used to analyze whether and how third-party information provision in social learning environments should be regulated.

Finally, I analyze social learning when agents have limited memory. I show that when action-signals satisfy weak separation, information aggregation is robust to certain distortions which *reduce* the informativeness of the action-signals about other agents' identities. I then argue that this type of distortion is precisely the major source of cognitive limitation of agents as documented by laboratory evidence. In light of the result, I discuss how the information environment can be designed so that it is easy for agents with limited memory to aggregate information.

Formally, in my model, a countably infinite number of agents, $n \in \{1, 2, ...\}$, sequentially choose an action a_n, starting with agent 1, to maximize their expected utility. Each agent's utility depends on the state θ and his action a_n. Agent n's information about θ consists of two sources: a *state-signal* s_n about θ and an *action-signal* m_n about histories of other agents' actions. The state-signal s_n is generated according to a distribution $\psi(\cdot|\theta)$ independently conditional on the state θ; the action-signal m_n is generated according to a distribution $\mu_n(\cdot|a_{<n})$ independently conditional on the history of actions $a_{<n}$ before agent n arrives. This information structure generalizes the various information structures employed in the literature. It also captures other important environments, such as those in which agents observe coarse summary statistics of prior actions (e.g., labor markets), and those in which

[3]The informativeness of the action-signals are measured according to the Blackwell order of informativeness.

agents observe the majority choice of actions of cohorts of agents (e.g., product markets where bestseller lists are published).

I provide a systematic analysis of long-run efficiency of learning in the game, as measured by the notion of information aggregation. We say that there is *information aggregation* if, as n goes to infinity, equilibrium actions a_n converge (in probability) to the action that yields the highest utility given the state θ. In the analysis of information aggregation, both the informativeness of the state-signals and that of the action-signals play a key role.

The informativeness of state-signals matters for information aggregation mostly in terms of whether an agent's state-signal can have an unbounded influence on his posterior belief. This concept is first introduced by Smith and Sørensen (2000). I refer to state-signals that satisfy this property as being *unboundedly informative*.[4] If a state-signal is unboundedly informative, then it can move the posterior anywhere no matter where in the interior the prior starts. In contrast, if the state-signal is boundedly informative, then, for any given prior, the posterior will be bounded away from the extremes of 0 or 1. Smith and Sørensen show that, in the special case wherein each agent perfectly observes all past actions, information aggregation occurs if and only if the state-signals are unboundedly informative.

In terms of the action-signals, this paper introduces the key concepts of weak and strong separation, based on whether the action-signals are informative enough so that agents can "separate" certain types of action-histories. The action-signals are *weakly-separating* if each agent n can find some (potentially random) set J_n of prior agents such that every agent only appear in a finite number of these J_n sets, and, in the limit as n goes to infinity, each agent n's action-signal separates action-histories that differ in terms of the optimal decisions implied by the volumes of actions in J_n. If, as n goes to infinity, agent n's action-signal always separates action-histories that differ in terms of the volumes of actions in J_n (instead of the implied optimal decisions), we say that the action-signals are *strongly-separating*. Sections 5.1 and 5.2 illustrate the main idea of the separation concepts in the special case in which the J_n sets are deterministic. Then Section 5.3 presents the general notions of weak and strong separation which allow the J_n sets to be stochastic. Many reasonable action-signals are weakly- or strongly-separating. For example, the canonical environment in this literature, where each agent observes all past actions, are both weakly- and strongly-separating. Environments in which each agent observes some coarse summary statistic of recent actions can be weakly separating but not strongly separating. Two simple examples in which action-signals are neither weakly-separating nor strongly-separating are: one in which

[4] Smith and Sørensen refer to the property as "private beliefs being unbounded."

all future agents only observe the actions of the first k number of agents, and one in which each agent observes his immediate predecessor's action with error probability $\epsilon > 0$.

The main results of this paper are presented in five theorems. Theorems 1, 2, and 3 show that when state-signals are unboundedly informative and action-signals are strongly- or weakly- separating, information aggregation occurs in all equilibria. The result unifies the prior literature and clarifies the underlying forces for information aggregation: agents needs to, at least, figure out the optimal decisions implied by the volumes of recent actions. In addition, the results contributes to the literature by identifying a larger class of environments that give rise to information aggregation, e.g., environments in which agents learn from volumes or certain coarse statistics of actions from recent histories and unbuondedly informative state-signals. On the other hand, Theorem 4 shows that when state-signals are boundedly informative, action-signals that are weakly separating may not aggregate information. Then Theorem 5 provides a sufficient condition for action-signals to aggregate information *regardless* of whether the state-signals are boundedly or unboundedly informative. Finally, Theorem 6 provides a necessary condition for information aggregation to occur, which shows that information aggregation is impossible when the action-signals on which (an infinite subset of) agents build their decisions is simply too noisy about action-histories beyond a finite period of time.

The characterization theorems highlight that a substantially larger set of environments than those studied in the previous literature can induce information aggregation, many of which are widely used in practice. I apply these results to study three important applications.

Section 6.1 studies the design of information disclosure on a digital platform in order to both enable information aggregation and protect users' privacy. Two potential disclosure methods are discussed. First, I analyze the commonly-used disclosure policy of releasing the volumes of actions periodically. This policy is shown to correspond to an action-signal technology that satisfies strong separation. Second, I introduce a disclosure policy that features "dynamically-adjusted benchmarks" and show that it corresponds to an action-signal technology that satisfy weak separation. According to the characterization results, both disclosure policies lead to information aggregation under unboundedly informative state-signals. However, the second disclosure policy is shown to be better at protecting users' privacy because it reveals strictly less information about past actions.

Section 6.2 investigates the robustness of information aggregation towards manipulation by third parties who can provide additional information about past actions to the agents. I show that information aggregation is robust to such manipulation if state-signals are unboundedly informative and the action-signals are strongly separating.[5] As a result, when the state-signals are unboundedly informative, a benevolent regulator that cares about long-run learning outcomes only needs to ensure that action-signals that satisfy strong separation are accurately disclosed and frequently updated. The regulator can adopt a laissez-faire policy regarding third party disclosure of other types of information. However, when state-signals are boundedly informative, third parties can always provide additional information to agents who are not well-informed about others' choices and induce herding behavior which blocks information aggregation. As discussed in Section 6.2, this result is broadly in line with the widely studied fact that advertisement and propaganda programs usually target people who are not well-informed, suggesting a possible (indirect) channel of manipulation through social learning for this phenomenon.

Section 6.3 studies the impact of a *reduction* of information in the action-signals on information aggregation. As a corollary of the main characterization theorem, I show that in a large class of action-signals that are strongly separating, information aggregation is robust to information reduction that features a certain type of "identity-blending." The result can be interpreted as a robustness property of information aggregation against agents' limited memory of identities of prior agents, which in cognitive psychology is documented as "source amnesia."[6] This result can be formalized by assuming that each agent has a simplified memory process: instead of remembering every detail of past information, his memory is limited by a coarse partition of the space of events.

The rest of the paper proceeds in the usual order. Section 2 reviews the literature. Section 3 introduces the model. Section 4 characterizes the agents' optimal strategies. Section 5 presents the main results about information aggregation in the long run. Section 6 studies three applications. Section 7 concludes. Proofs and materials omitted from the main text are presented in the Appendix.

[5] The result is correct regardless of the number of such third parties, and regardless of their objectives and the communication protocols through which they can influence the agents' information.

[6] This phenomenon, which is well-studied in psychology, refers to the fact that people have difficulty remember sources of prior events but find it relatively easy to remember the content of the events (Shimamura and Squire, 1987; Zaragoza and Lane, 1994; Schacter, 1999; Brown and Marsh, 2008). In a social learning context, source amnesia can be especially intensified due to "memory interference"(McGeoch, 1932; Anderson and Neely, 1996): when the same decisions are made repeatedly by different people, it becomes difficult to remember exactly who performed which action. Despite its prevalence, source amnesia has received little formal treatment in the theory of social learning.

2 Related Literature

This paper contributes to the large and growing literature on social learning. Banerjee (1992), Bikhchandani et al. (1992), and Welch (1992) initiate the literature on social learning, wherein agents act sequentially and observe the actions of all previous agents. Smith and Sørensen (2000) provide the most comprehensive and complete analysis of this environment. Their results and the importance of the concepts of bounded and unbounded private beliefs, which they introduced, have already been discussed in this paper's introduction and will play an important role in our analysis. Other important contributions in this area include, among others, Lee (1993), Chamley and Gale (1994), and Vives (1997). An excellent general discussion is provided by Bikhchandani et al. (1998). These papers typically focus on the special case wherein agents perfectly observe the history of actions.

Deviating from the full information paradigm, Celen and Kariv (2004) study learning when each individual observes his immediate predecessor. Extending the model further, Acemoglu et al. (2011) study social learning in arbitrary social networks. They characterize the set of environments that lead to information aggregation. A key concept is that of a social network with expanding observations in which no finite group of agents is excessively influential. They show that there is no information aggregation in networks with non-expanding observations. When state-signals are unboundedly informative and the network topology is expanding, information aggregation will occur. Furthermore, for a sizable class of stochastic social networks, there is information aggregation even when state-signals are boundedly informative. Smith and Sørensen (2014) provides a first (and perhaps the only) attempt to study sequential social learning when prior agents' identities are not perfectly known. They focus on a special case in which agents learn about others' actions through a random sample of prior actions. Smith and Sorensen show that if state-signals are unboundedly informative and the distant past is not over-sampled, information aggregation occurs.

My paper contributes to this literature by considering a general information environment, according to which a countable number of agents arrive sequentially and learn about other agents' actions. In this sense, my paper unifies the literature. In addition, my model encompasses many other important environments that are not studied in earlier papers. The main characterization theorems in my paper provide a simple criterion to check whether information aggregation occurs in a general information environment. The results on the general properties of information aggregation and optimal information design are novel to the literature, and they can only be meaningfully stated in a model that allows all information environments.

Other related work on social learning includes Callander and Horner (2009), who show that when agents are differentially informed, it may be optimal to follow the actions of agents that deviate from past average behavior; Hendricks et al. (2012), Mueller-Frank and Pai (2016), and Ali (2018), who incorporate costly search or costly information into social learning; Bohren (2016, 2017), who studies the robustness of information aggregation to small amount of behavioral biases; Banerjee and Fudenberg (2004) and Frick et al. (2019), who study environments in which a continuum of agents, rather than a single agent, arrive at each point in time; and Ellison and Fudenberg (1993), Ellison and Fudenberg (1995), Bala and Goyal (2001), DeMarzo et al. (2003), Golub and Jackson (2010), among others, who focus on non-Bayesian learning with agents using some reasonable rules of thumb.

3 Model

3.1 Environment

A countably infinite number of agents (individuals), indexed by $n \in \mathbb{N}$, arrive sequentially, and each makes a single decision $a_n \in \{0,1\}$. All agents are expected utility maximizers. The state of the world is $\theta \in \{0,1\}$, and agents share a common prior $\pi_i = \mathbb{P}(\theta = i) = 1/2$ for $i \in \{0,1\}$. Agents prefer to choose the action that matches the state, and they have a common utility function $u(a_n, \theta)$ with

$$u(1,1) = u(0,0) = 1 \text{ and } u(1,0) = u(0,1) = 0.$$ [7]

The state θ is initially unknown. Before making the decision a_n, each agent $n \in \mathbb{N}$ privately observes a *state-signal* s_n taking values in an arbitrary Polish space $\mathcal{S}$. The signals $\{s_n\}_{n\in\mathbb{N}}$ are independent and identically distributed conditional on the state θ, according to a distribution $\psi(\cdot|\theta) \in \triangle(\mathcal{S})$. We assume that $\psi(\cdot|\theta = 1)$ and $\psi(\cdot|\theta = 0)$ are mutually absolutely continuous, so no state-signal realization in $\mathcal{S}$ perfectly reveals the state.

[7] The assumptions on the prior and the functional form of the utility are only for simplicity of exposition. The main results of this paper hold under any prior belief that is not degenerate, i.e., $\pi_i \notin \{0,1\}$, and any general utility function such that

$$u(1,1) > u(0,1) \text{ and } u(0,0) > u(1,0).$$

The model can also be generalized so that agents have heterogenoues utilities as long as the utility functions are common knowledge.

We say that the state-signals are *unboundedly informative*, if for every positive number K, there are measurable subsets C and D of $\mathcal{S}$ such that

$$\frac{\psi(C|\theta=1)}{\psi(C|\theta=0)} > K \text{ and } \frac{\psi(D|\theta=1)}{\psi(D|\theta=0)} < \frac{1}{K}.$$

Otherwise, we say that the state-signals are *boundedly informative*.

For concreteness, consider some examples of signal structures. We could have binary signals, with each s_n taking values in $\mathcal{S} = \{0,1\}$ and where

$$\psi(s_n = 0|\theta = 0) = \psi(s_n = 1|\theta = 1) = q > \frac{1}{2}.$$

The realization $s_n = 0$ provides evidence in favor of $\theta = 0$, while $s_n = 1$ provides evidence in favor of $\theta = 1$. This is the signal structure studied by Banerjee (1992) and Bikhchandani et al. (1992), and the state-signals are boundedly informative in this case. We could also consider real-valued signals with $\psi(\cdot|\theta = 0)$ and $\psi(\cdot|\theta = 1)$ being probability measures on $\mathbb{R}$. For instance, suppose $\psi(\cdot|\theta = 0)$ and $\psi(\cdot|\theta = 1)$ are both supported on $\mathcal{S} = [0,1]$, where $\psi(\cdot|\theta = 0)$ has density $2-2s$ and $\psi(\cdot|\theta = 1)$ has density $2s$. In this case, lower signal realizations provide stronger evidence in favor of $\theta = 0$: the signal $s \in [0,1]$ induces the likelihood ratio $\frac{s}{1-s}$. It is easy to check that these state-signals are unboundedly informative.

In addition to the state-signal s_n, each agent n also privately observes an *action-signal* m_n about histories of past actions, taking values in an arbitrary Polish space $\mathcal{M}$. We use $a = (a_1, a_2, \ldots) \in \mathcal{A} := \{0,1\}^\infty$ to represent a full action-history, which specifies the *ordered* sequence of actions taken by each agent. We use $a_{<n} = (a_1, \ldots, a_{n-1}) \in \mathcal{A}_{<n} := \{0,1\}^{n-1}$ to represent a date-n action-history – i.e., the projection of a full action-history onto dates before n. The set of date-1 action-histories is a singleton set $\mathcal{A}_{<1} = \{\varnothing\}$ with a null initial history. For any subset of agents $I \subseteq \mathbb{N}$, we use $a_I = (a_i)_{i\in I}$ to represent an action-history of agents in I. Conditional on each action-history $a_{<n}$, the action-signal m_n is independently generated according to a distribution $\mu_n(\cdot|a_{<n}) \in \triangle(\mathcal{M})$. The information environment according to which agents learn about others agents' actions is summarized by the signal realization space $\mathcal{M}$ and the sequence of transition probabilities $\{\mu_n\}_{n=1}^\infty$. We therefore refer to $\mu := (\{\mu_n\}_{n=1}^\infty, \mathcal{M})$ as the *action-signal technology*, or simply *technology*.

The representation of the action-signal technology nests earlier models in the literature as special cases. For example, if $\mathcal{M} = \mathcal{A}$ and $\mu_n(a_{<n}|a_{<n}) = 1$ for any n and $a_{<n} \in \mathcal{A}_{<n}$, then agents will fully observe the histories as in the canonical model studied by Banerjee (1992), Bikhchandani et al. (1992), and Smith and Sørensen (2000), among others. We could also

capture social network models. Suppose each agent n is connected to a set $B_n \subseteq \{1, 2, ..., n-1\}$ of agents in the network. Let $\mathcal{M} = \mathcal{A}$ and $\mu_n(a_{B_n}|a_{<n}) = 1$ for any n and $a_{<n} \in \mathcal{A}_{<n}$. In this case, the agents perfectly observe the actions of agents in B_n, and know their identities. We can capture stochastic social networks by allowing B_n to depend on an independent randomization device. These are the environments studied in Acemoglu et al. (2011).

The model also captures many important settings that are not included in the prior literature. For instance, suppose $\mathcal{M} = \{0, 1\}$ and $\mu_n(1|a_{<n}) = 1$ if and only if $\sum_{i=n-k}^{n-1} a_i > k/2$, for some fixed odd integer k. This is an environment in which each agent learns which of the actions, 0 or 1, was taken by the majority of the previous k agents, which is related to many real-world markets: book-buyers learn from bestseller lists; film-goers consult box-office rankings, etc. As another example, we can also consider environments in which agents learn from *summary statistics* of actions. Suppose $\mathcal{M} = \{1, ..., k\}$ for some integer k, and for each $m \in \mathcal{M}$, $\mu_n(m|a_{<n}) = 1$ if $\sum_{i=n-k}^{n-1} a_i = m$. In this case, each agent knows the aggregate number of people choosing action 1, among the most recent k agents. The model can also capture *non-partitional* signals: For example, suppose $\mathcal{M} = \{0, 1\}$ and $\mu_n(m|a_{<n}) = q > \frac{1}{2}$ if $a_{n-1} = m$. In this case, agent n observes a noisy signal about his immediate predecessor's action. When agent n observes $m_n = m$, he knows that it is more likely $a_{n-1} = m$, but he cannot rule out the possibility of $a_{n-1} = 1 - m$.

3.2 Strategies and Equilibrium

A (behavioral) strategy $\sigma_n : \mathcal{S} \times \mathcal{M} \to \triangle(\{0, 1\})$ for agent n is a measurable function that maps each possible realization of her state-signal and action-signal into a randomized strategy over his action $a_n \in \{0, 1\}$. A strategy profile of the agents is a sequence of strategies $\sigma = (\sigma_n)_{n \in \mathbb{N}}$. We use the standard notation $\sigma_{-n} = (\sigma_1, ..., \sigma_{n-1}, \sigma_{n+1}, ...)$ to denote the strategies of all agents other than n. A strategy profile σ induces a joint probability distribution $\mathbb{P}_\sigma$ over the states, actions, action-signals, and state-signals, i.e., (θ, a, s, m) where $a = (a_1, a_2, ...)$, $s = (s_1, s_2, ...)$, and $m = (m_1, m_2, ...)$. We use $\mathbb{E}_\sigma[\cdot]$ to denote the operator that takes the expectation of random variables with respect to $\mathbb{P}_\sigma$. If $\mathcal{S}$ or $\mathcal{M}$ have infinite elements, then conditional expectations given the signals are only defined almost everywhere. In such cases, I use $\mathbb{P}(\cdot|s_n)$ and $\mathbb{P}_\sigma(\cdot|m_n)$ to refer to some version of the various conditional probabilities that may differ on zero measure sets.

In our game, agents' payoffs do not depend on what the other agents do. Thus, off-equilibrium beliefs do not matter. This allows us to define an equilibrium concept by simply requiring an unconditional optimization for each agent, which implies utility maximization conditional on beliefs updated according to Bayes rule almost everywhere on the agent's signal space $\mathcal{S} \times \mathcal{M}$. We say that a strategy profile is an equilibrium if each agent maximizes her expected utility, given the strategies of other agents:

Definition 1. *A strategy profile σ is an equilibrium of the social learning game if for each $n \in \mathbb{N}$, σ_n maximizes the expected payoff of agent n, given the strategies of other agents σ_{-n}, i.e.,*

$$\mathbb{E}_{\sigma}[u(a_n, \theta)] \geq \mathbb{E}_{(\sigma'_n, \sigma'_{-n})}[u(a_n, \theta)],$$

for every behavior strategy σ'_n for agent n.

Proposition 1. *There exists an equilibrium.*

Given the sequence of strategies $\{\sigma_1, ..., \sigma_{n-1}\}$, the maximization problem of agent n has a solution. Since each agent acts once in sequence, an inductive argument establishes the existence of an equilibrium, though in general it is not unique because some agent may be indifferent between the two actions. In this case, we can construct two different equilibria in which this agent adopts different strategies and subsequent agents best reply to these different strategies.

3.3 Long-Run Learning Metrics

Our study of equilibrium behavior centers on asymptotic outcomes. In particular, we consider the metric of information aggregation, based on agents' expected utility as the agent-index n approaches infinity. *Information aggregation* occurs if agents' utility approaches what they would obtain with perfect information:

Definition 2. *For any strategy profile σ, say that* ***information aggregation*** *occurs given σ if*

$$\lim_{n \to \infty} \mathbb{E}_{\sigma}[u(a_n, \theta)] = \pi_1 u(1,1) + \pi_0 u(0,0).[8]$$

This represents the best asymptotic outcome we can hope to achieve: For later agents, it is as though the private information of those that came before them is aggregated into a

[8]Equivalently, information aggregation occurs under strategy profile σ if $\lim \mathbb{P}_{\sigma}(a_n = \theta) = 1$.

single, arbitrarily precise signal.[9]

Another useful long-run learning metric is that of information diffusion, which depends on the distribution (and, in particular, the support) of the private beliefs. In the appendix, I discuss how my results should be interpreted if we use diffusion as the metric for the analysis.

4 Preliminary Analysis

The next result shows that agents' equilibrium strategies can be characterized as a function of two posterior beliefs. These posterior play an important role in our analysis.

Proposition 2. *Let σ be a strategy profile for the agents in any equilibrium. Let $s_n \in \mathcal{S}$ and $m_n \in \mathcal{M}$ be any signal realizations. Then agent n's action $a_n \in supp(\sigma_n(s_n, m_n))$ must satisfy*

$$a_n = \begin{cases} 1, & \textit{if } \mathbb{P}(\theta = 1|s_n) + \mathbb{P}_\sigma(\theta = 1|m_n) > 1 \\ 0, & \textit{if } \mathbb{P}(\theta = 1|s_n) + \mathbb{P}_\sigma(\theta = 1|m_n) < 1 \end{cases} \tag{1}$$

and $a_n \in \{0, 1\}$ otherwise, where $\mathbb{P}(\cdot|s_n)$ and $\mathbb{P}_\sigma(\cdot|m_n)$ are some version of the various conditional probabilities.

This proposition establishes an additive decomposition in the optimal decision rule between information obtained from the state-signal s_n and that from the action-signal m_n. The next definition formally distinguishes between the two components of an agent's information.

Definition 3. *We refer to the probability $p_n = \mathbb{P}(\theta = 1|s_n)$ as agent n's* ***private belief****, and $q_n := \mathbb{P}_\sigma(\theta = 1|m_n)$ as agent n's* ***social belief****.*

Note that the private belief p_n does not depend on the strategy profile of other agents, whereas q_n depends on the strategy profile. Hence, we use $\mathbb{P}$ in defining p_n and $\mathbb{P}_\sigma$ in defining q_n. Since s_n's are identically and independently distributed across n, so are the

[9] Information aggregation requires that the probability of taking the correct action converges to 1. Therefore, information aggregation will fail when, as more agents arrive, the limit inferior of the probability of all agents taking the correct action is strictly less than 1 – i.e., there exists an infinite sub-sequence of agents for whom the probabilities of taking the correct action are uniformly bounded away from 1. This is the strongest limiting result achievable in our framework. Information aggregation implies an almost-sure convergence of individual beliefs, but an almost-sure convergence of actions is impossible because for full information aggregation to occur, agents must continue to act based on their signals.

private beliefs p_n's. We use $\mathbb{G}_\theta$ to denote the distribution function of p_n conditional on the state. We can now state the definition of state-signals that are boundely or unboundedly informative, introduced earlier, in terms of the distribution of private beliefs.

Remark 1. *The state-signals are unboundedly informative if* $\{0,1\} \subseteq supp(p_1)$*, and they are boundedly informative otherwise.*[10]

When the state-signals are boundedly informative, there is a cap as to how informative the state-signal can be, i.e., how large or small the private beliefs induced by the state-signal can be. In this case, if an agent is sufficiently confident about the state of the world after observing his action-signal m_n, the state-signal will not affect his choice of action, according to Proposition 2. On the other hand, when the state-signals are unboundedly informative, the agents may receive arbitrarily informative state-signals in favor of either state. Thus, even if an agent is very confident of the state of the world after observing m_n, he may still change his action if the state-signal is sufficiently indicative of the opposite state of the world. In our subsequent analysis, we will see that whether or not the action-signals are boundedly informative plays a key role in determining whether information aggregation occurs.

In most learning models, social beliefs have natural monotonicity properties. For example, the more information an agent has when making his own decisions, the more informative his action will be for subsequent agents who observe the action. It is straightforward to construct examples in which such monotonicity properties do not hold in a sequential social learning environment. For this reason, to establish our main results (in the next section), we will use a different approach that is based on developing lower bounds on the probability of taking the correct action.

5 Main Results

In this section, I present my main results on information aggregation under action-signals that are strongly-separating and weakly-separating. To illustrate the main ideas without complicating the notations, section 5.1 and 5.2 study special cases of action-signal technologies that satisfy what I call deterministic strong and weak separation. I provide the intuition of the proofs for these cases. Then I introduce the general results about (stochastic) strong

[10]The notation $\text{supp}(p_1)$ represents the support of p_1 whose lower bound and upper bound are $\underline{b}, \bar{b}$ where $\underline{b} = \inf\{x \in [0,1] \,|\, \mathbb{P}(p_1 \leq x) > 0\}$ and $\bar{b} = \sup\{x \in [0,1] \,|\, \mathbb{P}(p_1 \leq x) < 1\}$. Thus, when the state-signals are unboundedly informative, we have $\underline{b} = 0, \bar{b} = 1$.

and weak separation in section 5.3.

5.1 Strong Separation: The Deterministic Case

We start by introducing the first key concept of *deterministic strong separation*.

Let $M_n^s(J,k)$ denote the set of agent n's action-signals that can occur only when exactly k agents in the set of agents $J \subseteq \{1,\dots,n-1\}$ take action 1; i.e.,

$$M_n^s(J,k) = \left\{ m \in \mathcal{M} : \forall a_{<n}, \text{ if } \sum_{j\in J} a_j \neq k, \text{ then } m \notin \text{supp}(\mu_n(\cdot|a_{<n}) \right\}.$$ [11]

Intuitively, if agent n observes an action-signal in $M_n^s(J,k)$, then he knows for sure that exactly k agents in J took action 1, i.e., it is impossible to observe an action-signal in $M_n^s(J,k)$ if exactly k' number of agents in J took action 1 and $k' \neq k$.

Definition 4. *Say that action-signals are* ***deterministically strongly-separating*** *if there exists non-empty sets of agents, $J_2, J_3, \dots$ such that no agent is in infinitely many of these sets, each $J_n \subseteq \{1,\dots,n-1\}$, and*

$$\lim_{n\to\infty} \min_{a_{<n}} \mu_n \left(M_n^s \left(J_n, \sum\nolimits_{j\in J_n} a_j \right) |a_{<n} \right) = 1,$$

Intuitively, if the action-signals are deterministically strongly-separating, then for any fixed integer k, in the limit as n goes to infinity, agent n's action-signal reveals, with probability one, the empirical frequency of action 1 taken by some set of agents after the first k agents.

The following theorem shows that unboundedly informative state-signals and deterministically strongly-separating action-signals are sufficient for information aggregation.

Theorem 1. *If the state-signals are unboundedly informative and the action-signals are deterministically strongly-separating, then information aggregation occurs in all equilibria.*

Theorem 1 implies that information aggregation occurs in many commonly studied action-signal technologies under unboundedly informative state-signals. Examples of action-

[11] Equivalently, the set $M_n^s(J,k)$ can expressed as

$$M_n^s(J,k) = \mathcal{M} - \bigcup_{a_{<n}:\sum_{j\in J} a_j \neq k} \text{supp}(\mu_n(\cdot|a_{<n})).$$

technologies that are deterministically strongly separating include: environments in which each individual observes all past actions (Banerjee, 1992; Bikhchandani et al., 1998; Smith and Sørensen, 2000); those in which each observes just his immediate predecessor (Celen and Kariv, 2004); those in which agents are connected in a fixed social network that features expanding observations (Acemoglu et al., 2011); and those in which agents observe samples of prior actions that do not always involve any individual agent. (Smith and Sørensen, 2014).[12]

As noted in the previous section, there is no monotonicity result that links the informativeness of an agent's action about the state to the agent's own information about past actions. Instead, Theorem 1 can be shown by making use of an utilitarian monotonicity related to the (expected) improvement of welfare across agents. In particular, we first consider some agent n whose action-signal is fully informative about the empirical frequency of action 1 of a set of prior agents $J_n \subset \{1, ..., n-1\}$; i.e.,

$$\mu_n\left(M_n^s(J_n, \sum_{j\in J_n} a_j)\middle| a_{<n}\right) = 1$$

for all $a_{<n} \in \mathcal{A}_{<n}$. We can establish the following utility improvement principle: there exists a strict lower bound on the increase in the ex ante expected utility that agent n gets over the expected average utility of agents in J_n (the average is calculated by assigning each agent in J_n equal weights). Intuitively, when agent n knows the empirical frequency of choices in J_n, it is feasible for him to adopt a mixed strategy to choose action 1 with probability exactly the same as the frequency that agents in J_n choose action 1, i.e., the following strategy is feasible for agent n:

$$\sigma_n^c(m_n, s_n) = \frac{\sum_{i\in J_n} a_i}{|J_n|}$$

because $\sum_{j\in J_n} a_j$ is revealed in m_n. This strategy will guarantee that agent n gets an expected utility exactly the same as average utility of agents in J_n.

Notice that the strategy σ_n^c only depends on the action-signal m_n but not the state-signal s_n. Since s_n is independent of m_n, the strategy σ^c is not utilizing an important source of information about the state. Now, consider how agent n may improve upon the strategy σ_n^c by fine tuning his action based on the realization of s_n. In particular, when the state-signal

[12] Acemoglu et al. (2011) also considers social networks in which agents' "neighborhoods" are stochastic but independently generated. Similarly, Smith and Sørensen (2014) considers cases where agents observe samples of past actions that are independently generated. The general notion of strong separation that captures these special cases are discussed in section 5.3.

is unboundedly informative, he may observe a very strong state-signal that points to either state. In such cases, he will always act on the state-signal, and whenever, he does so, the resulting expected gain of utility is strictly positive. This results in a fine-tuned strategy σ_m^f that generates an expected utility that is strictly higher than the average utility of agents in J_n whenever the state-signals are unboundedly informative and the average expected utility of agents in J_n is smaller than $\pi_1 u(1,1) + \pi_0 u(0,0)$ – the utility when the agents have full information about the states.

We can then prove that with deterministically strongly-separating action-signals, similar improvements of expected utility will continue as more agents arrive because each agent only appear in a finite number of other agents' J_n sets and the state-signals are unboundedly informative. Such strict improvement of expected utility will ultimately lead to information aggregation.

5.2 Weak Separation: The Deterministic Case

The concept of deterministic strong separation captures the property of an action-signal technology that allows the agents to know the empirical frequency of actions in some set of agents who arrived recently. In this section, I introduce the concept of deterministic weak separation to capture a larger class of action-signal technologies, which may not allow the agents to know the empirical frequency of actions in some set of agents, but nonetheless allows them to know the *optimal decisions* implied by the empirical frequencies of actions in a set of agents who arrived recently.

Given a strategy profile σ, for any $n \in \mathbb{N}$ and any $J \subseteq \{1, ..., n-1\}$, we choose a version of the conditional probability $\mathbb{P}_\sigma(\theta = 1 | m_n, \sum_{j \in J} a_j)$ such that whenever the event $m_n = m$ and $\sum_{j \in J} a_j = k$ occurs with zero probability, we have $\mathbb{P}_\sigma(\theta = 1 | m_n = m, \sum_{j \in J} a_j = k) = \frac{1}{2}$. Define

$$M_n^w(J, \sigma) := \{m \in \mathcal{M} : \mathbb{P}_\sigma(\theta = 1 | m_n = m, \sum_{j \in J} a_j = k) - 1/2 \text{ has the same sign} \tag{2}$$

$$\text{for every } k = 0, 1, 2, ..., |J|\}, \tag{3}$$

where zero is deemed as having the same sign as both positive and negative values.

We can now introduce the concept of deterministically weakly-separating action-signals, which captures the class of action-signals that "separates" the optimal decisions implied by the histories that specifies the same $\sum_{j \in J} a_j$ and lead to the same m_n.

Definition 5. *Say that action-signals are* ***deterministically weakly separating*** *if for* *every strategy profile σ that satisfies (1), there are nonempty sets of agents, $J_2, J_3, ...,$ such that no agent is in infinitely many of these sets, each $J_n \subseteq \{1, 2, ..., n-1\}$, and*

$$\lim_{n\to\infty} \min_{a_{<n}} \mu_n \left(M_n^w(J_n, \sigma) | a_{<n}\right) = 1.$$

As the following theorem shows, deterministically weakly separating action-signals and unboundedly informative state-signals are sufficient for information aggregation.

Theorem 2. *If the state-signals are unboundedly informative and the action-signals are deterministically weakly separating, information aggregation occurs in all equilibria.*

Thus, information aggregation occurs if the action-signals reveal, at least, the minimum amount of information that allows agents to figure out the optimal decision implied by the different set of histories that specify the same $\sum_{j\in J_n} a_j$ of some set J_n of recent agents (the term "recent" means that no agent appears in an infinite number of the sets $\{J_n\}_{n=1}^{\infty}$).

For generic information environments, deterministic weak separation is easy to verify given (ψ, μ). In particular, for any ψ that inducce an atom-less distribution (G_0, G_1) of the private beliefs, there is a unique equilibrium strategy profile σ as characterized in Proposition 2 which follows a cut-off rule that can be computed explicitly. Section 6.1 provides a concrete example about how weak separation can be verified and how it can be useful in the design of real-world markets.

The fact that deterministic weak separation needs to be verified using the strategy profile characterized in (1) should not be too unexpected: the exact amount of information that agents know about the state depends on the information structure ψ of the state-signals (not just whether they are unboundedly informative). Thus, in some knife-edge cases where action-signals does not satisfy deterministic strong separation, the exact amount of information each agent gets from the state-signal can become pivotal in determining whether information aggregation occurs.

To see how the proof works, assume unboundedly informative state-signals and assume deterministically weakly separating action-signals, and assume by way of contradiction that σ is an equilibrium profile under which information-aggregation fails. Then there is a subsequence of agents $n_k \to \infty$ such that

$$\lim_{k\to\infty} \mathbb{E}_\sigma(u(a_{n_k}, \theta)) = \varliminf_{n\to\infty} \mathbb{E}_\sigma(u(a_n, \theta)) < 1.$$

The first thing to argue is that any agent n far enough out in the sequence can achieve the above infimum utility, at least approximately, without using his state-signal. He can do this by employing the following naive strategy. Whenever agent n observes, for some nonempty $J \subseteq \{1, 2, ..., n-1\}$, an action-signal $m_n = m \in M_n^w(J, \sigma)$, that indicates the conditional probability $\mathbb{P}_\sigma(\theta = 1 | m_n = m, \sum_{j \in J} a_j = k)$ is non-negative for every $k \in \{0, ..., |J|\}$ such that this conditional probability is well-defined, then he chooses action 1 with probability 1. (If $m \in M_n^w(J, \sigma)$ indicates the conditional probability is non-positive for all k, then choose action 0 with probability 1. If $m \notin M_n^w(J, \sigma)$, then choose an action at random.) For agents far enough out in the sequence, the event $m_n \notin M_n^w(J, \sigma)$ occurs with very small probability. Therefore, by adopting the naive strategy, agent n's expected utility must be weakly higher than the average utility of the agents in the set J: this is because in the event that $m_n = m, \sum_{j \in J} a_j = k$, and if this event implies posterior belief of the state higher than 1/2, taking action 1 guarantees that, in expectation, agent n is outperforming an average agent in J across these histories. Then if the events $m_n = m, \sum_{j \in J} a_j = k$ for every k implies the posterior beliefs are all weakly higher than 1/2, taking action 1 upon observing $m_n = m$ guarantees that agent n outperforms an average agent in J in expectation (across all histories that lead to $m_n = m$). Similarly, the strategy also outperforms an average agent in J when m_n indicates that the posterior belief are all weakly lower than 1/2 for all k.

Since far enough out in the sequence, the agents in J exclude any finite number of agents, all of the agents in J receive ex-ante utility that is close to or above the infimum limit utility. At the same time, since $m_n \in M_n^w(J, \sigma)$ with probability approximately equal to 1 far enough in the sequence, agent n's ex ante expected utility from following this naive strategy must also be close to or above the limit infimum (again for n large). We may conclude that any agent n, by using this naive strategy, can achieve a payoff close to or above the infimum limit utility without using his state-signal.

Consider now any subsequence of agents whose limit utility is equal to the infimum limit utility. By what we have just shown, far enough along this subsequence, each of these agents can achieve (arbitrarily) close to their equilibrium utility without using their state-signal. Therefore, if these agents behave *optimally* given their action signals, but ignore their private signals, then they can, a fortiori, achieve an expected utility that is close to their equilibrium utility. But since this limit utility is strictly less than 1, there must be events in which they are choosing the wrong action and their state-signal is strong enough that, were they to use it, they would switch to the correct action. Hence, using also their state-signal when it is sufficiently strong can strictly increase their utility.

The key of the proof is to show that this increase is bounded away of zero and hence these agents, far enough along the subsequence can obtain a utility that is strictly greater than the infimum limit, which contradicts the assumption that that is their equilibrium utility, and concludes the proof.

As introduced earlier, deterministic weak separation is a weaker condition than deterministic strong separation.

Proposition 3. *If the action-signals are deterministically strongly-separating, then they are deterministically weakly-separating.*

5.3 Stochastic Separation

The definitions of deterministic strong and weak separation can be extended to allow the J_n sets to be random but independent of the action-histories, which lead to the key concepts of (stochastic) strong and weak separation.

For any pair of positive integers $k < n$, let $\mathcal{N}_k^{n-1}$ denote the set of all nonempty subsets of $\{1, 2, ..., n-1\}$ such that $\min_{i \in J} i + |J| > k$.

Definition 6. *The action-signals are* ***(stochastically) strongly-separating*** *if there exist a subset $M_n(J,k) \subseteq M_n^s(J,k)$ for every $n \in \mathbb{N}$, $J \subseteq \{1, ..., n-1\}$, and $k \in \{0, ..., |J|\}$, such that*

(a) $\lim_{n\to\infty} \mu_n(\cup_{J \in \mathcal{N}_k^{n-1}} M_n(J, \sum_{j\in J} a_j)|a_{<n}) = 1$ *for every $a \in \mathcal{A}$,*

(b) $\mu_n(M_n(J,k)|a_{<n})$ *is independent of $a_{<n}$ for every n, J, and k,* [13] *and*

(c) *the sets $\{M_n(J,k)\}_{J \subseteq \{1,...,n-1\}, k \in \{0,...,|J|\}}$ are mutually disjoint.*

To interpret this definition, let γ_n be a distribution over subsets of $\{1, ..., n-1\}$ such that

$$\gamma_n(J) = \mu_n(M_n(J, \sum_{j\in J} a_j)|a_{<n}), \forall a_{<n} \in \mathcal{A}_{<n}.$$

For any strongly-separating action-signal technology, γ_n exists because of condition (b). Then we can interpret γ_n as the (stochastic) analog of J_n in the definition of deterministic strong separation. It governs the distribution of the subsets of agents whose action-counts are revealed to agent n. Condition (a) then requires that no agent is excessively influential

[13] To see why the independence condition is important, see section B4.

in an infinite number of subsets of agents in the support of other agents' γ_n distributions. Condition (b) restricts the action-signal such that the occurrence of the event of agent n receiving an action-signal in $M_n(J, \sum_{j \in J_n} a_j)$ is independent of the action-history. This ensures that the randomization over which set J is revealed in the action-signal to agent n does not depend on what actions the agents in J took. Finally, the requirement in condition (c) that the sets $\{M_n(J,k)\}_{J,k}$ are mutually disjoint ensures that agent n has a feasible strategy that only depends on his action-signal and guarantees a utility that is equal to a weighted average of his recent predecessors' utility, where the average is taking with respect to the γ_n distributions.

To illustrate the environments that are captured by stochastic separation but are left out by the earlier notion of deterministic separation, consider the stochastic networks as in Acemoglu et al. (2011): each agent n observes the identities and the actions of agents in a stochastic neighborhood $B(n) \subseteq \{1, 2, ..., n-1\}$ where $B(n)$ is generated according to some probability distribution $\mathbb{Q}_n$ over all subsets of $\{1, 2, ..., n-1\}$, and draws from each $\mathbb{Q}_n$ are independent from each other for all n and from the realizations of the state-signals.

In terms of the notations introduced in this paper, each agent n's action-signal specifies the identities and the actions of the agent in his stochastic neighborhood, i.e.,

$$m_n = (B(n), a_k, \forall k \in B(n))$$

and is generated according to the information structure $\mu_n(\cdot|a_{<n})$ such that

$$\mu_n((B(n), a_{B(n)})|a_{<n}) = \mathbb{Q}_n(B(n)) \tag{4}$$

Example 1. *Consider the special case where $B(n) = \{n-1\}$ or $\{n-2\}$ with equal probabilities under $\mathbb{Q}_n$. It satisfies stochastic strong separation, but fails to satisfy the condition of deterministic strong or weak separation introduced in the earlier sections. For detailed explanation of this example, see Section B5.*

We can also define a similar stochastic analog to deterministic weak-separation.

Definition 7. *The action-signals are* ***(stochastically) weakly-separating*** *if there exist a subset $M_n(J, \sigma) \subseteq M_n^w(J, \sigma)$ for every $n \in \mathbb{N}$ and $J \subseteq \{1, ..., n-1\}$, such that*

(a) $\lim_{n\to\infty} \mu_n(\cup_{J \in \mathcal{N}_k^{n-1}} M_n(J, \sigma)|a_{<n}) = 1$ *for every $a \in \mathcal{A}$,*

(b) $\mu_n(M_n(J, \sigma)|a_{<n})$ *is independent of $a_{<n}$ for every n and J, and*

(c) *the sets* $\{M_n(J,\sigma)\}_{J\subseteq\{1,\dots,n-1\}}$ *are mutually disjoint.*

Proposition 4 shows that the deterministic versions of strong and weak separation analyzed in the previous subsections are special cases of the more general definitions of stochastic strong and weak separation.

Proposition 4. *If the action-signals are strongly-separating, then they are weakly-separating. If the action-signals are deterministically strongly-separating (or deterministically weakly separating), then they are strongly-separating (or weakly-separating).*

In what follows, I will simply refer to stochastic strong (or weak) separation as strong (or weak) separation, eliminating the qualifier "stochastic". Both strong separation and weak separation ensure agents' expected utilities are strictly higher than the convex combination of some of their recent predecessors. A similar argument as before shows that under strongly or weakly separating action-signals and unboundedly informative state-signals, information aggregation occurs.

Theorem 3. *If the action-signals are strongly-separating or weakly-separating and the state-signals are unboundedly informative, then information aggregation occurs in all equilibria.*

The notion of strong separation unifies the prior literature on sequential social learning. It captures various action-signal technologies imposed by different papers as special cases and clarifies the underlying force that give rise to information aggregation. For example, as discussed in section B6, if agents observe other agents' actions according to a social network that features "expanding observations" as studied in Acemoglu et al. (2011), the corresponding action-signal technology satisfies strong separation.

The results on weak separation contributes to the literature by identifying a much weaker condition that can also give rise to information aggregation, which can be practically useful in designing information disclosure (see section 6.1). The relationship between these two concepts and those in the prior literature is summarized in the Venn Diagram in Figure 1.

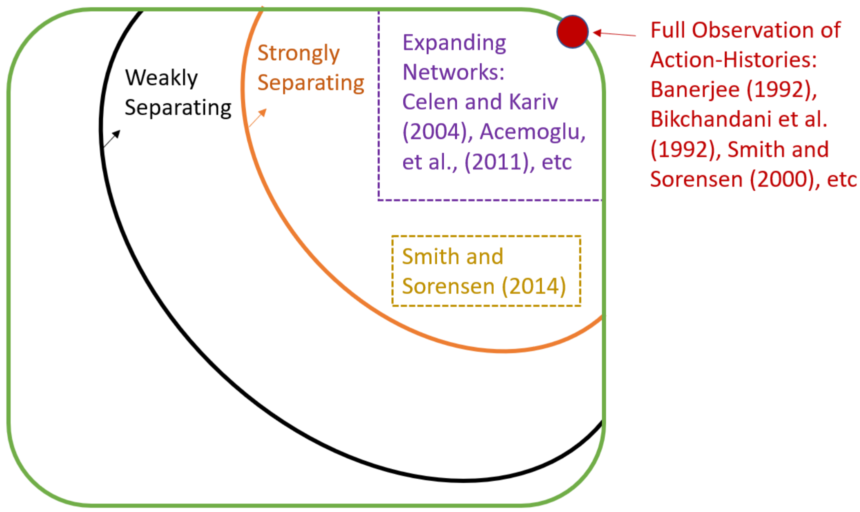

Figure 1: The relationship between different classes of action-signal technologies

5.4 Learning under Boundedly Informative State-Signals

We have obtained sufficient conditions for information aggregation under unboundedly informative state-signals. It is worth asking whether strong or weak separation is sufficient for information aggregation in the absence of unboundedly informative state-signals. At present this question is incompletely answered in both the prior literature that restricted attention to special cases and in my model with a general information environment.

Theorem 4. *Suppose the state-signals are boundedly informative. Then information aggregation fails if the action-signals are perfectly informative about action-histories, i.e.,*

$$\mu_n(a_{<n}|a_{<n}) = 1, \forall n \in \mathbb{N}, \forall a_{<n} \in \mathcal{A}_{<n}.$$

Theorem 4 reiterates the failure of information aggregation identified by classical papers (Banerjee 1992; Bikhchandani et al. 1992; Smith and Sorensen 2000). It suggests that the fully informative action-signals will fail to aggregate information, demonstrating that transparency alone is insufficient to ensure aggregation. This is so because for any state-signal s_n that leads to only non-extreme private beliefs, an extreme social belief induced by the action-signal can overwhelm any state-signal, causing agents to disregard their state-signals when choosing actions.

In some sense the action-signal technology that fully reveals the action-history is a knife-edge case in which the large-sample principle fails because no one is forced to rely on a state-signal. As first discovered in Acemoglu et al. (2011) in a network setting, if we disrupt information for some agents in the population, creating a sub-sequence of "ignorant agents" with uninformative action-signals, this group can provide enough information for the rest of the population to learn the true state. This is generally true in my model with general information environments.

Theorem 5. *Suppose there exists a sequence of agents* $\{n_i\}_{i\in\mathbb{N}}$ *such that*

- *with probability* p_i*, the realization of* m_{n_i} *is completely uninformative about the histories, i.e.,* $\mu_n(m_{n_i}|a_{<n})$ *is the same for all* $a_{<n} \in \mathcal{A}_{<n}$*.*
- $\lim_{i\to\infty} p_i = 0$ *and* $\sum_{i=1}^{\infty} p_i = \infty$
- $\lim_{n\to\infty} \min_{a_{<n}} \mu_n(M_n(\{n_i\}, a_{n_i})|a_{<n}) = 1$ *for each* i*.*

Then information aggregation occurs.

To see why Theorem 5 is correct. We note that the sequence of agents $\{n_i\}_{i\in\mathbb{N}}$ have no information from their action-signals, and thus act on their state-signals. Therefore, the actions of these ignorant agents is an infinite sequence of independent signals about the state. Since all agents learn about actions of the ignorant agents with probability approaching 1, all agents learn the state according to law of large numbers. We should note that the sequence ignorant agents consists of an arbitrarily small portion of the population in the limit, which highlights an important discontinuity in learning outcomes when we compare a transparent environment with an almost-transparent environment.

This theorem generalizes in several directions. Similar to results in a network setting (Acemoglu et al., 2011), for information aggregation to occur, the ignorant agents could have informative action-signals as long as some realizations of the state-signals still dominate the social belief. We can also allow agents to not all perfectly observe the ignorant agents: we only need some infinite sequence of agents to observe action-signals that are informative about the large sample of ignorant agents in a way that is similar to being weakly separating so that they can aggregate information; and then others can learn by observing this infinite sequence of agents, instead of the ignorant agents. The results along these directions are scattered (both in the literature and in my setting). A general characterization of information aggregation when state-signals are boundedly informative is still missing.

5.5 Necessary Condition for Aggregation

In this section, I provide a necessary condition for information aggregation. We say that the action-signal technology is *news-permitting* if the agents can, at least, distinguish some continuation histories of each finite history that occurs with positive probability.

Definition 8. *The action-signal technology is* ***news-permitting*** *(NP) for strategy profile* σ *if* $\forall \epsilon > 0$, $\forall k \in \mathbb{N}$, $\forall a_{<k} \in \mathcal{A}_{<k}$ *with* $\mathbb{P}_\sigma(a_{<k}) > 0$, $\exists n_k > k$ *s.t.* $\forall n > n_k$, $\forall m \in \mathcal{M}$, $\exists b_{<n} \in \mathcal{A}_{<n}$ *s.t.* $b_{<k} = a_{<k}$ *and* $\mu_n(m|b_{<n}) < \epsilon$.

If the action-signal technology is news-permitting, then as more agents arrive, they will eventually, with positive probability, receive some action-signal that allows them to distinguish histories beyond a finite period of time. Hence, as I will show in Theorem 6, when the action-signal technology is not news-permitting, information aggregation fails.

Theorem 6. *If the action-signal technology is not news-permitting for an equilibrium strategy profile* σ*, information aggregation does not occur in* σ*.*

Intuitively, the first finite number of k actions can only contain a bounded amount of information about the state. Thus, if some agent n observes a realization of his action-signal that can be generated with positive probability (uniformly bounded away from zero) given every history consistent with a particular date-k action-history, then this realized action-signal must convey bounded information about the state. Thus, we can bound the probability that this agent makes the correct choice away from 1 and establish the failure of information aggregation with an infinite number of such agents.

As I show in Proposition 5, when the state-signals are unboundedly informative, all action-histories of finite length are reached with positive probability in equilibrium.

Proposition 5. *If the state-signals are unboundedly informative, then for any equilibrium strategy profile* σ, $\forall k \in \mathbb{N}$, $\forall a_{<k} \in \mathcal{A}_{<k}$*, we have* $\mathbb{P}_\sigma(a_{<k}) > 0$*.*

For this reason, when the state-signals are unboundedly informative, if the action-signals are news-permitting for some equilibrium strategy σ, they must be news-permitting for all equilibria. This motivates the following definition.

Definition 9. *The action-signal technology is* ***uniformly news-permitting*** *(UNP) if* $\forall \epsilon >$

0, $\forall k \in \mathbb{N}$, $\forall a_{<k} \in \mathcal{A}_{<k}$, $\exists n_k > k$ *s.t.* $\forall n > n_k$, $\forall m \in \mathcal{M}$, $\exists b_{<n} \in \mathcal{A}_{<n}$ *s.t.* $b_{<k} = a_{<k}$ *and* $\mu_n(m|b_{<n}) < \epsilon$.

Corollary 1. *If the state-signals are unboundedly informative and the action-signal technology is not uniformly news-permitting, information aggregation does not occur in any equilibrium.*

Examples of action-signal technologies that are not news-permitting include those in which all agents observe only the first k actions, those in which each agent observes the action of his immediate predecessor with error probability $\epsilon > 0$, etc. Theorem 6 and Corollary 1 imply that information aggregation fails in such environments.

6 Applications

In this section, I apply the main results in the previous section to three important applications. In Section 6.1, I focus on the design of information disclosure on financial platforms to protect investors' privacy while maintaining long-run efficiency as measured by information aggregation. In Section 6.2, I study whether information aggregation is robust to manipulation by third parties who can provide additional information to the agents. In Section 6.3, I analyze how information loss due to agents' limited memory impacts long-run social learning outcome, and how to design the information environment to make it easy for ordinary people with limited memory to aggregate information.

6.1 Privacy Protection on Digital Platforms

In this section, I present an example in the context of information disclosure on a financial platform. I construct two action-signal technologies: one strongly separating and another one only weakly separating, and then discuss how the example sheds light on the policy discussion of balancing efficiency with privacy protection.

Consider a financial platform for capital investment. Investors $n \in \{1, 2, ...\}$ arrive sequentially, each with a unit endowment. Each investor can either invest his or her endowment in Project A in Chicago or Project B in New York ($a_n = 1$ for investing in A, and $a_n = 0$ for investing in B). Both projects are in the technology sector, and the investors' payoffs are determined by whether Amazon's second headquarter ends up in Chicago or New York ($\theta = 1$ for it being in Chicago, and $\theta = 0$ for it being in New York). Suppose, for simplicity, there are equal probabilities that Amazon chooses either city as the host of its second headquarter, i.e., $\mathbb{P}(\theta = 1) = \mathbb{P}(\theta = 0) = 1/2$, and the investors payoffs are $u(a_n, \theta) = \mathbf{1}\{a_n = \theta\}$, i.e., the projects can only succeed in a city where the headquarter is located. Each investor has some private information $s_n \sim \psi(\cdot|\theta)$ about θ that is unboundedly informative. For technical reasons, the platform can only publish information on its website every week, and any published information in a particular week will be available for all future investors. Assume that t number of investors arrive every week.

In designing the information disclosure policy, the platform has two main concerns: ensuring information aggregation and preserving investor privacy. That is, the platform wants to disclose as little information as possible while ensuring that information aggregation occurs.[14] The concept of strong and weak separation can be used to analyze different disclosure policies. I will discuss two important disclosure policies (i.e., action-signal technologies).

6.1.1 Empirical Frequencies (Strongly Separating)

Suppose the platform adopts an action-signal technology $\mu^s = (\{\mu_n^s\}_{n=1}^{\infty}, \mathcal{M})$ such that for every $c \in \{0, 1, ...\}$ (where c is the index of the week) and every agent $n \in \{ct+1, ..., (c+1)t\}$ arriving in week c, the action-signal is such that:

$$\mu_n^s\left(m_n = \left(\sum_{j=1}^{t} a_j, \sum_{j=t+1}^{2t} a_j, ..., \sum_{j=(c-1)t+1}^{ct} a_j\right) \middle| a_{<n}\right) = 1 \tag{5}$$

Intuitively, for each $k \in \{0, ..., (c-1)\}$, the term $\sum_{j=kt+1}^{(k+1)t} a_j$ corresponds to the number of agents who choose action 1 in the k^{th} week (i.e., the empirical frequency of actions in the k^{th} week). The action-signal, therefore, simply discloses to each agent the empirical frequency of actions in each of the previous weeks. This action-signal technology is "public" in the sense that information disclosed to earlier agents are perfectly observed by later agents. This is a desirable property because, practically, even if the financial platforms delete the information that they published earlier, there may be third parties who record the information and share it with future investors.[15]

This action-signal technology is strongly separating: for each week $c \geq 2$ and each agent n arriving in that week (i.e., $n \in \{ct+1, ..., (c+1)t\}$), take J_n to be the agents arriving in the

[14] Add references for "minimum disclosure".

[15] For example, such information can be found on investment forums.

previous week (i.e., the $(c-1)^{\text{th}}$ week):

$$J_n = \{(c-1)t+1, ..., ct\}.$$[16]

Then it is obvious the action-signal m_n perfectly separates histories that differ in terms of $\sum_{j \in J_n} a_j$, and that no agent appears in an infinite number of the sets $\{J_n\}_{n=1}^{\infty}$. So the action-signal technology satisfies strong separation and will give rise to information aggregation when the state-signals are unboundedly informative.

This disclosure policy is widely used in practice and corresponds to the simple strategy of releasing the "empirical frequency of actions" every week. However, is it possible for the platform to do better at preserving investors' privacy without impacting information aggregation? The answer is: yes, there is a Blackwell less informative action-signal technology μ^w that is weakly separating, and, therefore, also ensures that information aggregation occurs.

6.1.2 Dynamically-Adjusted Benchmarks (Weakly Separating)

In this section, I construct an action-signal technology which is defined using weekly benchmarks $\{\phi_c\}_{c=1}^{\infty}$ that are dynamically adjusted. In week c, each agent's action-signal reveals whether, in *every* previous week, the number of agents who choose action 1 is larger than or smaller than the corresponding benchmark in that week. Since information about all previous weeks are disclosed, the action-signals are public. The public information upto week c is denoted as $x_{<c} = (x_0, x_1 ..., x_{c-1}) \in \{0,1\}^c$, where, for each $k \in \{0, ..., c-1\}$, $x_k = 1$ represents that the aggregate number of agents in week k who choose action 1 is above the week-k benchmark.

The benchmarks are recursively defined. Define the benchmark in week zero as $\phi_0 = t/2$, and for each subsequent week $c \geq 1$, define the benchmark $\phi_c : \{0,1\}^c \to \mathbb{R}$ such that for each $x_{<c} = (x_0, x_1 ..., x_{c-1}) \in \{0,1\}^c$, the benchmark $\phi_c(x_{<c})$ is any real number such that, given the public information disclosed before, if the number of agents in week c who choose action 1 is larger than $\phi_c(x_{<c})$, then it indicates that the state $\theta = 1$ is more likely than $\theta = 0$; otherwise, it indicates that the state $\theta = 0$ is more likely than $\theta = 1$, i.e., $\phi_c(x_{<c})$ satisfies

$$\mathbb{P}_\sigma \left(\theta = 1 \middle| \sum_{j=ct+1}^{(c+1)t} a_j = \lfloor \phi_c(x_{<c}) \rfloor, \mathbf{1}\left\{ \sum_{j=kt+1}^{(k+1)t} a_j > \phi_k(x_{<k}) \right\} = x_k \text{ for all } k = 0, 1, ..., c-1 \right) \leq 1/2,$$

[16] For agents arriving in the fist week, take J_n to be any sets.

and

$$\mathbb{P}_\sigma\left(\theta = 1 \middle| \sum_{j=ct+1}^{(c+1)t} a_j = \lceil \phi_c(x_{<k}) \rceil, \mathbf{1}\left\{\sum_{j=kt+1}^{(k+1)t} a_j > \phi_k(x_{<k})\right\} = x_k \text{ for all } k = 0, 1, \ldots, c-1\right) \geq 1/2.$$

For each agent n in week c, define $\phi_c(a_{<n})$ recursively as

$$\phi_c(a_{<n}) := \phi_c\left(\sum_{j=1}^{t} a_j > \phi_0, \sum_{j=t+1}^{2t} a_j > \phi_1(a_{<t}), \ldots, \sum_{j=(c-1)t+1}^{ct} a_j > \phi_{c-1}(a_{<(c-1)t})\right)$$

The action-signal technology $\mu^w = (\{\mu_n^w\}_{n=1}^\infty, \mathcal{M})$ is such that for every week $c \in \mathbb{N}$ and every agent n arriving in week c, i.e., $n \in \{ct+1, \ldots, (c+1)t\}$, his or her action-signal is such that

$$\mu_n^w\left(m_n = \left(\mathbf{1}\left\{\sum_{j=1}^{t} a_j > \phi_0\right\}, \mathbf{1}\left\{\sum_{j=t+1}^{2t} a_j > \phi_1(a_{<t})\right\}, \ldots, \mathbf{1}\left\{\sum_{j=(c-1)t+1}^{ct} a_j > \phi_{c-1}(a_{<(c-1)t})\right\}\right) \middle| a_{<n}\right) = 1$$

As I show in Proposition 6, this action-signal technology is weakly separating, and thus also gives rise to information aggregation under unboundedly informative state-signals.

Proposition 6. *The action-signal technology μ^w is weakly separating.*

The action-signal technology μ^w is less informative compared to μ^s introduced in the previous subsection, especially when t is large, because empirically frequencies are not fully disclosed in μ^w: the agents only know how the empirical frequencies compare with the benchmarks.

The comparison between μ^s and μ^w shows how the platform can intelligently design its information disclosure so that privacy is preserved without impacting efficiency: the platform can disclose only the relative position of the empirical frequency of actions with respect to some benchmark; and the platform needs to dynamically adjust the benchmark in a clever way to achieve this: given a generic (ψ, μ) the benchmarks can be computed explicitly.[17]

6.2 Third-Party Information Manipulation

Suppose information aggregation occurs in some information environment. Will some third party be able to block it by providing additional information to the agents? Such manipula-

[17]The benchmark is not unique since there may be a range of values of ϕ_c for the two inequalities to hold at the same time. However, for any of these values, the equilibrium outcome will be the same.

tion efforts exist in many settings: from propaganda in political campaigns, to information manipulation in financial markets, to targeted advertisement aimed at increasing sales of a product. From a design perspective, we would want to design information environments in which information aggregation is robust to such third-party manipulation.

Formally, suppose in addition to the state-signal s_n and the action-signal m_n, agent n may observe an additional action-signal $m_n^\dagger \in \mathcal{M}^\dagger$ disclosed by some third party. Furthermore, suppose $m_n^\dagger$ is generated independently according to a distribution $\mu_n^\dagger(\cdot|a_{<n}) \in \triangle(\mathcal{M}^\dagger)$ given any action-history $a_{<n}$. We refer to $\mu^\dagger := (\{\mu_n^\dagger\}_{n=1}^\infty, \mathcal{M}^\dagger)$ as the third parties *manipulation technology*.

One can interpret $\mu^\dagger$ as being implemented by a third party who can commit to any signal structures of $\mu^\dagger$ before the first agent arrives.[18] Alternatively, $\mu^\dagger$ may be the information structure of $m_n^\dagger$'s that appears in some equilibrium of a game where multiple third parties engage in strategic information manipulation of the social learning process. The formulation captures all these possibilities.

Definition 10. *Given (ψ, μ), we say that information aggregation is* ***robust to third-party manipulation*** *if information aggregation occurs in all equilibria under every $\mu^\dagger$.*

A consequence of Theorems 2 and 5 is a fundamental difference between state-signals that are unboundedly informative and those that are boundedly informative. Thus, the analysis will be carried out separately for the two cases.

I first consider cases where the state-signals are unboundedly informative – that is, environments where extreme realizations of the state-signal, that are arbitrarily indicative of each state, can occur and be privately observed by the agents.

Proposition 7. *If the state-signals are unboundedly informative and the action-signals are strongly separating, information aggregation is robust to third-party manipulation.*

Intuitively, with third-party information manipulation, the actual action-signal technology, through which agents observe signals about past action-histories, consists of two components: the original action-signal technology μ and the manipulation technology $\mu^\dagger$. If

[18] In this case, we can think that the signal $m_n^\dagger$'s are directly observed by the agents. Alternatively, as in Gentzkow and Kamenica (2017), a model in which agents directly observe signal realizations is outcome-equivalent to a model in which the third party chooses a manipulation strategy, privately observes the realizations, and sends cheap-talk messages to the agents.

the original action-signal technology is separating, then after giving more information to the agents through $\mu^\dagger$, the resulting action-signal technology remains separating (which can be shown using the definition of weak separateness). Thus, information aggregation occurs under any manipulation strategy.

This implies that when state-signals are unboundedly informative, a regulator needs not worry about third-party manipulation. He can simply focus on making sure that some action-signals that satisfies weak separateness are disclosed accurately and in a timely manner (e.g., publishing and updating data on governmental agencies websites), and he can adopt a laissez faire policy regarding other sources of information disclosure. Such policies are easy to implement, but they should be used very carefully when the regulator is not confident about whether the state-signals are unboundedly informative.

Proposition 8. *When the state-signals are boundedly informative, information aggregation is not robust to third party manipulation.*

This should be obvious given the results in Theorem 4. Intuitively, for information aggregation to occur under boundedly informative state-signals, the existence of a infinite number of "ignorant" agents whose social belief does not dominate the private beliefs is required. However, a third party can always provide full information of past actions to these "ignorant" agents, which will induce herding behavior and block information aggregation.

To illustrate, consider a simple example wherein each agent n observes the entire history with probability $1 - 1/n$ and observes no information about past actions with probability $1/n$ (ignorant). In this case, information aggregation occurs without third party information manipulation because, in expectation, an infinite ($\sum_{n=1}^{\infty} 1/n = \infty$) number of ignorant agents always take actions that are informative about their state-signals, and the entire population ($\lim_{n\to\infty} 1 - 1/n = 1$) will ultimately observe the full history and learn the true state from the actions of the ignorant agents. However, when the third party adopts a manipulation strategy $\mu^\dagger$ such that $m_n^\dagger$ is fully informative about past actions. All agents will be fully informed about the entire history of actions, and herding will occur under boundedly informative state-signals.

Thus, when state-signals are boundedly informative, the regulator needs to carefully examine the third party information provision (e.g., whether the third party is targeting "ignorant" agents), before he takes any regulatory action. Unlike the case with unboundedly informative state-signals, a one-size-fits-all policy choice does not exist.

Furthermore, it should be noted that when state-signals are unboundedly informative, the regulator can always induce information aggregation by ensuring the disclosure of some weakly separating action-signals, whereas when the state-signals are boundedly informative, there exist third-party strategies (e.g., one that reveals the full history) under which information aggregation is doomed to fail (i.e., the regulator cannot induce information aggregation by any disclosure strategy μ).

6.3 Cognitive Limitation

In real world markets, the action-signal technology μ may be subject to distortions caused by the agents themselves: agents may have limited cognitive ability to remember past events, and the information that agents remember is therefore less informative than the information contained in action-signals disclosed in the first place. In this section, I apply the main characterization results to analyze the impact of such distortion on information aggregation.

In a social learning environment, each agent acquires information about past action-histories through the action-signal m_n, which contains two types of information: information about prior choices, and information about the identity of the agent who choose each action. Laboratory evidence demonstrates that information about the sources (location, identity of people, etc) of past events is much more likely to be forgotten than the content of past events. This phenomenon is common among ordinary people, and is referred to as "source-amnesia".[19] In this section, I consider the impact of this type of cognitive limitation on social learning.

We have shown that strongly separating action-signal technologies give rise to information aggregation if the state-signals are unboundedly informative. From the definition of strong separateness, we can see that the identity of the agents within the J_n sets and any information about agents outside J_n are not instrumental for an action-signal technology to be strongly separating. Thus, when agents lose track of information about identities of agents in J_n or information about agents outside J_n, information aggregation should still occur (although potentially at a slower pace). We now establish this simple intuition formally.

Definition 11. *Given an action-signal technology that is strongly separating, a sequence of subsets $\{J_n\}_{n=1}^{\infty}$ that satisfies condition (4) is a **separating sequence** for μ.*

When agents lose track of the identities of other agents, they will not be able to distinguish

[19] See, for example, Shimamura and Squire (1987), Zaragoza and Lane (1994), McGeoch (1932), and Anderson and Neely (1996).

some histories that differ only in terms of the identity of agents. This is formally captured by the concept of identity-blending distortions.

Definition 12. *Given an action-signal technology $\mu = (\{\mu_n\}_{n=1}^{\infty}, \mathcal{M})$ and a sequence of sets $\{J_n\}_{n=1}^{\infty}$ with $J_n \in \{1, ..., n-1\}$ for each n, another action-signal technology $\mu' = (\{\mu'_n\}_{n=1}^{\infty}, \mathcal{M}')$ is an **identity-blending distortion** of μ in the sets $\{J_n\}_{n=1}^{\infty}$ if for any $n \in \mathbb{N}$ there exists a surjection $f_n : \mathcal{M} \to \mathcal{M}'$ with the following properties:*

1. *For any distinct $m_n^1, m_n^2 \in \mathcal{M}$ s.t. $f_n(m_n^1) = f_n(m_n^2)$, we have*

$$\mu_n^{-1}(m_n^1) \bigcup \mu_n^{-1}(m_n^2) \subseteq \left\{ a_{<n} \in \mathcal{A}_{<n} : \sum_{j \in J_n} a_j = k \right\} \text{ for some } k \in \{0, 1, 2, ..., |J_n|\}.$$

2. *For any $a_{<n} \in \mathcal{A}_{<n}$ and any $m_n \in M$,*

$$\mu'_n(m_n | a_{<n}) = \mu_n(f_n^{-1}(m_n) | a_{<n}).$$

Intuitively, an identity-blending information distortion happens when cognitive limited agents at least remember the number of agents in J_n who choose each action but may forget the identity of the agents in J_n or any information that only involve agents outside J_n. The next proposition states that information aggregation under strongly separating action-signals is robust to this type of distortion if the distortion occurs *in a separating sequence* of the action-signals.

Proposition 9. *Suppose μ' is an identity-blending distortion of μ in a separating sequence of μ. Then if information aggregation occurs under μ, it must also occur under μ'.*

This theorem ensures that we do not need to worry too much about people losing track of identities of past agents as long as it is limited to the identities of agents in some separating sequence of the action-signal technology. In addition, we also don't need to worry about people losing track of information that only involves agents outside the separating sequence.

In practice, it may be easier for people to forget the sources of distant events; so action-signal technologies with separating sequence $\{J_n\}_{n=1}^{\infty}$ such that each J_n contains a large number of recent agents (e.g. $J_n = \{n-k, n-k+1, ..., n-1\}$ for large $k \in \mathbb{N}$) would be more robust to identity-blending distortion. This explains why statistics that disclose the empirical frequencies of a large number of recent choices are widely used in practice (e.g., annual employment report).

Finally, the results in this section can also be used to guide the regulation of information disclosure. In particular, as mentioned in the introduction, many industries impose "minimum disclosure" rules to balance efficiency with privacy, requiring that information should not be disclosed if it is not important for efficiency purposes. The results in this section implies that if the regulator wants to hide some information about individual identities for privacy purposes, but wants to maintain long-run efficiency of social learning, he can blend individual identities within the J_n set in arbitrary ways as long as the disclosed information is clear about the empirical frequency of actions in the J_n sets. This provides a simple guideline for designing regulatory rules.

7 Conclusion

In this paper, I study a model that allows all feasible information environments in a sequential social learning setting. This allows us to obtain a general understanding of the conditions under which information aggregation occurs in social learning. First and foremost, I view my findings as providing a simple criterion that can be used in practice to guide the evaluation and design of information disclosure in social learning environments. In addition, the findings further motivate the study, both theoretically and empirically, of non-network models of social learning. Below, I briefly point to three directions whose exploration seems worthwhile.

First, I maintain the assumption that the agents have homogeneous preferences. One might question the descriptive validity of this assumption because in many markets, agents can have very heterogeneous preferences: consumers have different tastes about products; investors differ in terms of their risk-attitudes; voters have diverse ideological preferences, etc. While a full analysis of this question is beyond the scope of this paper, I note that if the heterogeneous preferences are public information, then all my results go through, with similar proofs. In future work, it would be interesting to understand what happens when agents have heterogeneous preferences that are privately observed. As discovered by Smith and Sørensen (2000), in the special case where the action-signals are fully informative about all histories, heterogeneous private types can lead to herding behavior that does not result in information cascade – an outcome that never arises with homogeneous preferences. In ongoing work, I am pursuing a detailed understanding of social learning with private types in any general information environment.

Second, while this paper considers a model of social learning in which agents arrive sequentially in a fixed order and take one-shot actions, it is natural to inquire about the analogs of my characterization theorems for other social learning models when agents engage in *repeated* interactions. For example, it would be interesting to explore what would happen if *long-lived* agents learn about each others' past choices through general action-signals. This would require generalizing the model in Mossel et al. (2015).

Third, while I have assumed that agents interact with others only by observing their actions, my model can be enriched to allow for *payoff interdependence*, wherein agents' utilities depend not only on the state and their own action but on other agents' actions. This is relevant in settings such as voting, investment, etc. Payoff interdependence is also critical in any market that has centralized price-setting because prices act both as a source of information for people (i.e., an action-signal) and as a channel through which earlier agents' actions can affect the utility of subsequent agents. Therefore, a model with general action-signals and payoff interdependence, will lead to a better understanding of the role of prices in social learning environments, such as product markets and financial markets.

References

D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozadaglar. Bayesian learning in social networks. Review of Economic Studies, 2011.

N. Ali. Herding with costly information. Journal of Economic Theory, 2018.

M. Anderson and J. Neely. Interference and inhibition in memory retrieval. Memory, Pages 237-313, 1996.

V. Bala and S. Goyal. Conformism and diversity under social learning. Economic Theory, 2001.

A. Banerjee. A simple model of herd behavior. Quarterly Journal of Economics, 1992.

A. Banerjee and D. Fudenberg. Word-of-mouth learning. Games and Economic Behavior, 2004.

S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. Journal of Political Economy, 1992.

S. Bikhchandani, D. Hirshleifer, and I. Welch. Learning from the behavior of others: Conformity, fads, and informational cascades. Journal of Economic Perspectives, 1998.

A. Bohren. Informational herding with model misspecification. Journal of Economic Theory, 2016.

A. Bohren. Boundedly rational learning: A framework and a robustness result. Working Paper, 2017.

A. Brown and E. Marsh. Evoking false beliefs about autobiographical experience. Psychonomic Bulletin and Review, 2008.

S. Callander and J. Horner. The wisdom of the minority. Journal of Economic Theory, 2009.

B. Celen and S. Kariv. Observational learning under imperfect information. Games and Economic Behavior, 2004.

C. Chamley. Rational herds: Economic models of social learning. Cambridge University Press, 2004.

C. Chamley and D. Gale. Information revelation and strategic delay in a model of investment. Econometrica, 1994.

P. M. DeMarzo, D. Vayanos, and J. Zwiebel. Persuasion bias, social influence, and unidimensional opinions. The Quarterly Journal of Economics, 2003.

G. Ellison and D. Fudenberg. Rules of thumb for social learning. Journal of Political Economy, 1993.

G. Ellison and D. Fudenberg. Word-of-mouth communication and social learning. The Quarterly Journal of Economics, 1995.

M. Frick, R. Iijima, and Y. Ishii. Misinterpreting others and the fragility of social learning. Working Paper, 2019.

M. Gentzkow and E. Kamenica. Bayesian persuasion with multiple senders and rich signal spaces. Games and Economic Behavior, 2017.

B. Golub and M. O. Jackson. Naive learning in social networks and the wisdom of crowds. American Economic Journal: Microeconomics, 2010.

B. Golub and E. Sadler. Learning in social networks. The Oxford Handbook of the Economics of Networks, 2016.

K. Hendricks, A. Sorensen, and T. Wiseman. Observational learning and demand for search goods. American Economic Journal: Microeconomics, 2012.

H. Herrera and J. Horner. A necessary and sufficient condition for information cascades. Working Paper, 2012.

E. Kamenica and M. Gentzkow. Bayesian persuasion. American Economic Review, 2011.

I. H. Lee. On the convergence of informational cascades. Journal of Economic Theory, 1993.

I. Lobel and E. Sadler. Information diffusion in networks through social learning. Working Paper, 2014.

J. McGeoch. Forgeting and the law of disuse. Psychological Review, 1932.

E. Mossel, A. Sly, and O. Tamuz. Strategic learning and the topology of social networks. Econometrica, 2015.

M. Mueller-Frank and M. Pai. Social learning with costly search. American Economic Journal: Microeconomics, 2016.

R. Myerson and P. Reny. Perfect conditional ϵ-equilibria of multi-stage games with infinite sets of signals and actions. Working Paper, 2018.

D. Schacter. The seven sins of memory: insights from psychology and cognitive neuroscience. American Psychologist, 1999.

Shimamura and L. Squire. A neuropsychological study of fact memory and source amnesia. Journal of Experimental Psychology, 1987.

L. Smith and P. Sørensen. Pathological outcomes of observational herding. Econometrica, 2000.

L. Smith and P. N. Sørensen. Rational social learning with random sampling. Working Paper, 2014.

U.S. Department of Health and Human Services. Minimum necessary requirement, 2003. URL https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/minimum-necessary-requirement/index.html.

X. Vives. Learning from others: a welfare analysis. Games and Economic Behavior, 1997.

I. Welch. Sequential sales, learning, and cascades. Journal of Finance, 1992.

M. Zaragoza and S. Lane. Source misattributions and the suggestibility of eyewitness memory. Journal of Experimental Psychology, 1994.

A Proofs and supplementary materials for Section 4

A1 Proof of Proposition 2

Agent n obtains information about the state through two sources: the state-signal s_n and the action-signal m_n. Since conditional on state θ, s_n and m_n are mutually independent, agent n's optimal decision can be stated in terms of a cut-off strategy in terms of his private belief p_n and his social belief q_n.

Proof of Proposition 2. The agent maximizes expected payoffs given her posterior belief will $\mathbb{P}_\sigma(\theta = 1|s_n, m_n)$ of the state. Thus, she chooses action $a_n = 1$ if

$$\begin{aligned}\mathbb{P}_\sigma(\theta = 1|s_n, m_n)u(1,1) \quad &+ \quad \mathbb{P}_\sigma(\theta = 0|s_n, m_n)u(1,0) \\ &> \quad \mathbb{P}_\sigma(\theta = 1|s_n, m_n)u(0,1) + \mathbb{P}_\sigma(\theta = 0|s_n, m_n)u(0,0)\end{aligned}$$

Rewriting it in terms of likelihood ratios, we get

$$\frac{\mathbb{P}_\sigma(\theta = 1|s_n, m_n)}{\mathbb{P}_\sigma(\theta = 0|s_n, m_n)} \quad > \quad \frac{u(0,0) - u(1,0)}{u(1,1) - u(0,1)} = 1 \tag{6}$$

By Bayes' rule, equation (6) is equivalent to

$$\frac{d\mathbb{P}_\sigma(s_n, m_n|\theta = 1)\mathbb{P}(\theta = 1)}{d\mathbb{P}_\sigma(s_n, m_n|\theta = 0)\mathbb{P}(\theta = 0)} > 1, \tag{7}$$

Since $\mathbb{P}(\theta = 1) = \mathbb{P}(\theta = 0) = 1/2$, equation (7) is equivalent to

$$\frac{d\mathbb{P}_\sigma(s_n, m_n|\theta = 1)}{d\mathbb{P}_\sigma(s_n, m_n|\theta = 0)} > 1 \tag{8}$$

By independence of the different sources of the information for agent n, equation (8) is equivalent to

$$\frac{d\mathbb{P}(s_n|\theta = 1)}{d\mathbb{P}(s_n|\theta = 0)} > \frac{d\mathbb{P}_\sigma(m_n|\theta = 0)}{d\mathbb{P}_\sigma(m_n|\theta = 1)} \tag{9}$$

By Bayes' rule again, this can be rewritten as

$$\frac{\mathbb{P}(\theta = 1|s_n)}{\mathbb{P}(\theta = 0|s_n)} > \frac{\mathbb{P}(\theta = 0|m_n)}{\mathbb{P}(\theta = 1|m_n)}.$$

Thus, agent n strictly prefers $a_n = 1$ if

$$\mathbb{P}(\theta = 1|s_n) + \mathbb{P}_\sigma(\theta = 1|m_n) > 1.$$

Similar arguments confirms that agent n strictly prefers $a_n = 0$ if

$$\mathbb{P}(\theta = 1|s_n) + \mathbb{P}_\sigma(\theta = 1|m_n) < 1.$$

□

B Proofs and supplementary materials for Section 5

B1 Proof of Theorem 1

This theorem follows from Theorem 2 and Proposition 3.

B2 Proof of Theorem 2

Lemma 1. *For any private belief distributions (G_0, G_1), we have for all $p \in (0,1)$*

$$\frac{dG_0(p)}{dG_1(p)} = \frac{1-p}{p} \tag{10}$$

Proof. By the definition of private beliefs p_n, we have

$$p_n = \mathbb{P}(\theta = 1|s_n) = \mathbb{P}(\theta = 1|p_n)$$

According to Bayes' rule, we have for any $p_n \in (0,1)$

$$\mathbb{P}(\theta = 1|p_n) = \frac{d\mathbb{P}(p_n|\theta = 1)\mathbb{P}(\theta = 1)}{d\mathbb{P}(p_n|\theta = 1)\mathbb{P}(\theta = 1) + d\mathbb{P}(p_n|\theta = 0)\mathbb{P}(\theta = 0)}$$

Since the two state $\theta = 1$ and $\theta = 0$ are equally likely, and that no signal is completely informative, i.e. $p_n \notin \{0,1\}$, the two equations above imply that

$$p_n = \frac{dG_1(p_n)}{dG_1(p_n) + dG_0(p_n)}$$

and the result follows. $\square$

The following properties of the distribution of the private beliefs (G_0, G_1) will also be useful in the proof.

Lemma 2. *For any* $0 < z < x < 1$,

$$G_0(x) \geq \frac{1-x}{x} G_1(x) + \frac{x-z}{2} G_1(z),$$

and for all $0 < x < w < 1$,

$$1 - G_1(x) \geq \frac{x}{1-x}(1 - G_0(x)) + \frac{w-x}{2}(1 - G_0(w)).$$

Proof. For any $x \in (0,1)$,

$$G_0(x) = \int_{s=0}^{x} dG_0(s) = \int_{s=0}^{x} \frac{1-s}{s} dG_1(s) = \left(\frac{1-x}{x}\right) G_1(x) + \int_{s=0}^{x} \left(\frac{1}{s} - \frac{1}{x}\right) dG_1(s)$$

where the second equality follows from the fact that

$$\frac{dG_0(s)}{dG_1(s)} = \frac{1-s}{s}.$$

For any $0 < z < x < 1$, we have

$$\int_{s=0}^{x} \left(\frac{1}{s} - \frac{1}{x}\right) dG_1(s) \geq \int_{s=0}^{z} \left(\frac{1}{s} - \frac{1}{x}\right) dG_1(s) \geq \int_{s=0}^{z} \left(\frac{1}{z} - \frac{2}{z+x}\right) dG_1(s) \geq \frac{x-z}{2} G_1(z)$$

establishing the first inequality. Similarly, the second inequality can be shown by

$$\begin{aligned} 1 - G_1(x) = \int_{s=x}^{1} dG_1(s) &= \int_{s=x}^{1} \frac{s}{1-s} dG_0(s) \\ &= (1 - G_0(x)) \left(\frac{x}{1-x}\right) + \int_{s=x}^{1} \left(\frac{s}{1-s} - \frac{x}{1-x}\right) dG_0(s) \end{aligned}$$

The last term can be bounded as follows

$$\begin{aligned} \int_{s=x}^{1} \left(\frac{s}{1-s} - \frac{x}{1-x}\right) dG_0(s) &\geq \int_{s=w}^{1} \left(\frac{s}{1-s} - \frac{x}{1-x}\right) dG_0(s) \\ &\geq \int_{s=w}^{1} \left(\frac{w}{1-w} - \frac{x+w}{2-x-w}\right) dG_0(s) \geq \frac{w-x}{2}(1 - G_0(w)) \end{aligned}$$

$\square$

To simplify exposition, we introduce some notations before proceeding. Let $p_n(s_n) := \mathbb{P}(\theta = 1|s_n)$ denote the private belief induced by s_n, and let $q_n^\sigma(m_n) := \mathbb{P}_\sigma(\theta = 1|m_n)$ denote the social belief induced by m_n in equilibrium σ. Using these notations, we make it clear that the private beliefs depend on the state-signal realization s_n, and the social beliefs depend on the action-signal realization m_n and the strategy profile σ. For each $m \in \text{supp}(\tilde{m}_n)$ in equilibrium σ, let

$$\begin{aligned} dY_n^\sigma(m) &:= d\mathbb{P}_\sigma(m_n = m|\theta = 1) = 2q_n^\sigma(m)d\mathbb{P}_\sigma(m_n = m) \\ dN_n^\sigma(m) &:= d\mathbb{P}_\sigma(m_n = m|\theta = 0) = 2(1 - q_n^\sigma(m))d\mathbb{P}_\sigma(m_n = m). \end{aligned}$$

where the last step in each of the two equations follows from Lemma 1 and the fact that

$$d\mathbb{P}_\sigma(m_n = m) = \frac{1}{2} \cdot (dY_n^\sigma(m) + dN_n^\sigma(m)). \tag{11}$$

Hence, for any m on the support of m_n, we have

$$\frac{dY_n^\sigma(m)}{dN_n^\sigma(m)} = \frac{q_n^\sigma(m)}{1 - q_n^\sigma(m)} \tag{12}$$

Lemma 3. *The probability that agent n's action match the correct state in equilibrium σ is such that*

$$\mathbb{P}_\sigma(a_n = \theta) \geq \frac{1}{2}\left(\int_{m_n \in \mathcal{M}} (1 - G_1^-(1 - q_n^\sigma(m_n)))dY_n^\sigma(m_n) + \int_{m_n \in \mathcal{M}} G_0(1 - q_n^\sigma(m_n))dN_n^\sigma(m_n)\right)$$

Proof. According to Proposition 1, we know that given a particular realization of m_n that induces a social belief $q_n^\sigma(m_n)$, agent n chooses $a_n = 1$ if $p_n(s_n) > 1 - q_n^\sigma(m_n)$ and chooses $a_n = 0$ if $p_n(s_n) < 1 - q_n^\sigma(m_n)$. Thus,

$$\begin{aligned} \mathbb{P}_\sigma(a_n = \theta) \geq & \int_{m_n \in \mathcal{M}} \mathbb{P}(p_n(s_n) > 1 - q_n^\sigma(m_n)|\theta = 1)d\mathbb{P}_\sigma(m_n|\theta = 1)\mathbb{P}(\theta = 1) \\ & + \int_{m_n \in \mathcal{M}} \mathbb{P}(p_n(s_n) < 1 - q_n^\sigma(m_n)|\theta = 0)d\mathbb{P}_\sigma(m_n|\theta = 0)\mathbb{P}(\theta = 0) \end{aligned}$$

where the inequality follows from the fact that the signals s_n and m_n are conditionally independent given θ, and the fact that we omitted the case where $p_n(s_n) = 1 - q_n^\sigma(m_n)$

under which the agent's action may also match the state with positive probability (when the distribution of the state-signals have mass points). The result follows by substituting the expressions of

$$1 - G_1^-(1 - q_n^\sigma(m_n)) = \mathbb{P}(p_n(s_n) > 1 - q_n^\sigma(m_n)|\theta = 1)$$

and

$$G_0(1 - q_n^\sigma(m_n)) = \mathbb{P}(p_n(s_n) < 1 - q_n^\sigma(m_n)|\theta = 0).$$

□

For any $n \in \mathbb{N}$ and $J \subset \mathbb{N}$, define:

$$\tilde{M}_n^1(J;\sigma) = \{m \in \mathcal{M} : \mathbb{P}_\sigma(\theta = 1|m_n = m, \sum_{j\in J} a_j = k) \geq 1/2 \text{ for every } k \tag{13}$$
$$\text{such that the posterior is well-defined}\}. \tag{14}$$

and

$$\tilde{M}_n^0(J;\sigma) = \{m \in \mathcal{M} : \mathbb{P}_\sigma(\theta = 1|m_n = m, \sum_{j\in J} a_j = k) \leq 1/2 \text{ for every } k \tag{15}$$
$$\text{such that the posterior is well-defined}\}. \tag{16}$$

By definition, we have

$$M_n^w(J;\sigma) = \tilde{M}_n^0(J;\sigma) \cup \tilde{M}_n^1(J;\sigma).$$

where $M_n^w(J,\sigma)$ is the set of action-signals used in the definition of weak separation.

Intuitively, upon observing an action-signal realization in $\tilde{M}_n^0(J,\sigma)$ or $\tilde{M}_n^1(J,\sigma)$, agent n can guarantee to get at least the expected utility of an average agent in J across histories that lead to $m_n = m$, by simply choosing action 0 or 1 respectively.

This idea can be formalized by considering the following *naive* strategy for agent n which only depends on his action-signal m_n:

$$\sigma_n^a(m_n, s_n) = \begin{cases} 1, & \text{if } m_n \in \tilde{M}_n^1(J_n,\sigma) \\ 0, & \text{if } m_n \in \tilde{M}_n^0(J_n,\sigma) \\ 0, & \text{if } m_n \notin M_n^w(J_n,\sigma) \end{cases} \tag{17}$$

Given an equilibrium σ, for any $J_n \subseteq \{1, ..., n-1\}$, let $\alpha_\sigma(J_n)$ be the average utility of agents in J_n, i.e.,

$$\alpha_\sigma(J_n) = \frac{\sum_{j \in J_n} \mathbb{P}_\sigma(a_j = \theta)}{|J_n|}.$$

The following lemma provides a lower bound of the expected utility that agent n can get by adopting the naive strategy.

Lemma 4. *Suppose for agent n and $J_n \subseteq \{1, ..., n-1\}$, we have*

$$\min_{a_{<n}} \mu_n(M_n^w(J_n; \sigma)|a_{<n}) > 1 - \epsilon. \tag{18}$$

Then the expected utility that agent n gets from adopting the naive strategy σ_n^a is such that

$$\mathbb{P}_{\sigma_n^a, \sigma_{-n}}(a_n = \theta) \geq \alpha_\sigma(J_n) - \epsilon \tag{19}$$

Proof. According to the definition of $\alpha_\sigma(J_n)$, we have

$$\begin{aligned} \alpha_\sigma(J_n) \cdot |J_n| \quad = \quad & \mathbb{P}(\theta = 1) \int_{a_{<n}} \left(\sum_{i \in J_n} a_i \right) d\mathbb{P}_\sigma(a_{<n}|\theta = 1) \\ & + \mathbb{P}(\theta = 0) \int_{a_{<n}} \left(|J_n| - \sum_{i \in J_n} a_i \right) d\mathbb{P}_\sigma(a_{<n}|\theta = 0) \end{aligned}$$

Given $\min_{a_{<n}} \mu_n(M_n^w(J_n, \sigma)|a_{<n}) > 1 - \epsilon$, we have

$$\begin{aligned} \alpha_\sigma(J_n) \cdot |J_n| \quad \leq \quad & \mathbb{P}(\theta = 1) \int_{a_{<n}} \sum_{i \in J_n} a_i \left[\epsilon + \mu_n(M_n^w(J_n, \sigma)|a_{<n}) \right] d\mathbb{P}_\sigma(a_{<n}|\theta = 1) \\ & + \mathbb{P}(\theta = 0) \int_{a_{<n}} \left(|J_n| - \sum_{i \in J_n} a_i \right) \cdot \left[\epsilon + \mu_n(M_n^w(J_n, \sigma)|a_{<n}) \right] d\mathbb{P}_\sigma(a_{<n}|\theta = 0) \\ \leq \quad & \mathbb{P}(\theta = 1) \int_{a_{<n}} \sum_{i \in J_n} a_i \mu_n(M_n^w(J_n, \sigma)|a_{<n}) d\mathbb{P}_\sigma(a_{<n}|\theta = 1) \\ & + \mathbb{P}(\theta = 0) \int_{a_{<n}} \left(|J_n| - \sum_{i \in J_n} a_i \right) \cdot \mu_n(M_n^w(J_n, \sigma)|a_{<n}) d\mathbb{P}_\sigma(a_{<n}|\theta = 0) \\ & + \epsilon \cdot |J_n| \end{aligned} \tag{20}$$

where the last inequality follows from the fact that $0 \leq \sum_{i \in J_n} a_i \leq |J_n|$.

Let the *date-n extended history* $h_{<n}$ be the ordered sequence of of the actions and the action-signal realizations before agent n takes his own action, i.e.

$$h_{<n} = (m_1, a_1, ..., m_{n-1}, a_{n-1}, m_n).$$

We use $\mathcal{H}_{<n}$ to denote the set of all date-n extended histories.

For any $n \in \mathbb{N}$, any $m \in \mathcal{M}$, any $J \subset \{1, ..., n-1\}$, let $\mu_n^{-1}(m, J, k)$ denote the set of date-n extended histories which specifies that exactly k agents in J choose action 1, and is such that $m_n = m$, i.e.,

$$\mu_n^{-1}(m, J, k) = \left\{ h_{<n} \in \mathcal{H}_{<n} : m_n = m \text{ and } \sum_{i \in J} a_i = k \right\}.$$

We refer to $\mu_n^{-1}(m, J, k)$ as the k^{th} *history-set of J contained in $m_n = m$.*

We note that the first term on the R.H.S. of Inequality (20) can be rewritten in terms of the extended histories $\mathcal{H}_{<n}$ as follows:

$$\begin{aligned}
& \mathbb{P}(\theta = 1) \int_{a_{<n}} \sum_{i \in J_n} a_i \mu_n(M_n^w(J_n, \sigma) | a_{<n}) d\mathbb{P}_\sigma(a_{<n} | \theta = 1) \\
= \;& \mathbb{P}(\theta = 1) \int_{h_{<n} : m_n \in M_n^w(J_n, \sigma)} \sum_{i \in J_n} a_i d\mathbb{P}_\sigma(h_{<n} | \theta = 1) \\
= \;& \mathbb{P}(\theta = 1) \left[\int_{h_{<n} : m_n \in \tilde{M}_n^1(J_n; \sigma)} \sum_{i \in J_n} a_i d\mathbb{P}_\sigma(h_{<n} | \theta = 1) + \int_{h_{<n} : m_n \in \tilde{M}_n^0(J_n; \sigma)} \sum_{i \in J_n} a_i d\mathbb{P}_\sigma(h_{<n} | \theta = 1) \right] \\
= \;& \mathbb{P}(\theta = 1) \sum_{k=1}^{|J_n|} \int_{m \in \tilde{M}_n^1(J_n; \sigma)} \int_{h_{<n} \in \mu_n^{-1}(m, J_n, k)} k d\mathbb{P}_\sigma(h_{<n} | \theta = 1, m_n = m) d\mathbb{P}_\sigma(m_n = m | \theta = 1) \\
& + \mathbb{P}(\theta = 1) \sum_{k=1}^{|J_n|} \int_{m \in \tilde{M}_n^0(J_n; \sigma)} \int_{h_{<n} \in \mu_n^{-1}(m, J_n, k)} k d\mathbb{P}_\sigma(h_{<n} | \theta = 1, m_n = m) d\mathbb{P}_\sigma(m_n = m | \theta = 1) \\
= \;& \mathbb{P}(\theta = 1) \sum_{k=1}^{|J_n|} \int_{m \in \tilde{M}_n^1(J_n; \sigma)} \left[k \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k) | \theta = 1, m_n = m) \right] d\mathbb{P}_\sigma(m_n = m | \theta = 1) \\
& + \mathbb{P}(\theta = 1) \sum_{k=1}^{|J_n|} \int_{m \in \tilde{M}_n^0(J_n; \sigma)} \left[k \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k) | \theta = 1, m_n = m) \right] d\mathbb{P}_\sigma(m_n = m | \theta = 1)
\end{aligned} \tag{21}$$

Similarly, the second term on the R.H.S. of Inequality (20) can be rewritten as

$$\begin{aligned}
&\mathbb{P}(\theta = 0)\int_{a_{<n}}\left(|J_n| - \sum_{i\in J_n} a_i\right)\cdot\mu_n(M_n^w(J_n,\sigma)|a_{<n})d\mathbb{P}_\sigma(a_{<n}|\theta = 0)\\
=\ &\mathbb{P}(\theta = 0)\sum_{k=1}^{|J_n|}\int_{m\in\tilde{M}_n^1(J_n;\sigma)}\left[(|J_n| - k)\cdot\mathbb{P}_\sigma(h_{<n}\in\mu_n^{-1}(m,J_n,k)|\theta = 0, m_n = m)\right]d\mathbb{P}_\sigma(m_n = m|\theta = 0)\\
&+\mathbb{P}(\theta = 0)\sum_{k=1}^{|J_n|}\int_{m\in\tilde{M}_n^0(J_n;\sigma)}\left[(|J_n| - k)\cdot\mathbb{P}_\sigma(h_{<n}\in\mu_n^{-1}(m,J_n,k)|\theta = 0, m_n = m)\right]d\mathbb{P}_\sigma(m_n = m|\theta = 0)
\end{aligned}\tag{22}$$

Combining Equations (21) and (22), we can rewrite Inequality (20) as

$$\begin{aligned}
&\alpha_\sigma(J_n)\cdot|J_n|\\
\le\ &\sum_{k=1}^{|J_n|}\int_{m\in\tilde{M}_n^1(J_n;\sigma)}\left[k\cdot\mathbb{P}_\sigma(h_{<n}\in\mu_n^{-1}(m,J_n,k),\theta = 1|m_n = m)\right]d\mathbb{P}_\sigma(m_n = m)\\
&+\sum_{k=1}^{|J_n|}\int_{m\in\tilde{M}_n^0(J_n;\sigma)}\left[k\cdot\mathbb{P}_\sigma(h_{<n}\in\mu_n^{-1}(m,J_n,k),\theta = 1|m_n = m)\right]d\mathbb{P}_\sigma(m_n = m)\\
&+\sum_{k=1}^{|J_n|}\int_{m\in\tilde{M}_n^1(J_n;\sigma)}\left[(|J_n| - k)\cdot\mathbb{P}_\sigma(h_{<n}\in\mu_n^{-1}(m,J_n,k),\theta = 0|m_n = m)\right]d\mathbb{P}_\sigma(m_n = m)\\
&+\sum_{k=1}^{|J_n|}\int_{m\in\tilde{M}_n^0(J_n;\sigma)}\left[(|J_n| - k)\cdot\mathbb{P}_\sigma(h_{<n}\in\mu_n^{-1}(m,J_n,k),\theta = 0|m_n = m)\right]d\mathbb{P}_\sigma(m_n = m)\\
&+\epsilon\cdot|J_n|
\end{aligned}\tag{23}$$

According to the definition of $\tilde{M}_n^1(J_n;\sigma)$, we know that for any $m\in\tilde{M}_n^1(J_n;\sigma)$, $J_n\in\{1,2,...,n-1\}$, and $k\in\{0,...,|J_n|\}$, we have

$$\begin{aligned}
&\mathbb{P}_\sigma(\theta = 1|h_{<n}\in\mu_n^{-1}(m,J_n,k))\ge\frac{1}{2}\\
\Leftrightarrow\ &\mathbb{P}_\sigma(\theta = 1|h_{<n}\in\mu_n^{-1}(m,J_n,k), m_n = m)\ge\frac{1}{2}\\
\Leftrightarrow\ &\mathbb{P}_\sigma(\theta = 1, h_{<n}\in\mu_n^{-1}(m,J_n,k)|m_n = m)\ge\mathbb{P}_\sigma(\theta = 0, h_{<n}\in\mu_n^{-1}(m,J_n,k)|m_n = m)
\end{aligned}\tag{24}$$

where the first equivalence follows from the fact that $m_n = m$ is satisfied for any $h_{<n}\in\mu_n^{-1}(m,J_n,k)$ according to the definition of $\mu_n^{-1}(m,J_n,k)$, and the second equivalence follows

from Bayes' rule. Therefore, according to Inequality (24), for any $m \in \tilde{M}_n^1(J_n;\sigma)$, we have

$$\begin{aligned}
& k \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 1 | m_n = m) + (|J_n| - k) \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 0 | m_n = m) \\
\leq\ & |J_n| \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 1 | m_n = m) + 0 \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 0 | m_n = m) \\
\leq\ & |J_n| \cdot \sigma_n^a(m, J_n) \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 1 | m_n = m) \\
& + |J_n| \cdot (1 - \sigma_n^a(m, J_n)) \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 0 | m_n = m) \qquad (25)
\end{aligned}$$

where the last step follows from the fact that $\sigma_n^a(m, J_n) = 1$ for any $m \in \tilde{M}_n^1(J_n;\sigma)$.

Similarly, for any $m \in \tilde{M}_n^0(J_n;\sigma)$, we have

$$\begin{aligned}
& k \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 1 | m_n = m) + (|J_n| - k) \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 0 | m_n = m) \\
\leq\ & 0 \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 1 | m_n = m) + |J_n| \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 0 | m_n = m) \\
\leq\ & |J_n| \cdot \sigma_n^a(m, J_n) \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 1 | m_n = m) \\
& + |J_n| \cdot (1 - \sigma_n^a(m, J_n)) \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 0 | m_n = m) \qquad (26)
\end{aligned}$$

Dividing $|J_n|$ on both sides of Inequality (23), and using inequalities (25) and (26), we get

$$\begin{aligned}
\alpha_\sigma(J_n) \leq\ & \sum_{k=1}^{|J_n|} \int_{m \in \tilde{M}_n^1(J_n;\sigma)} \left[\sigma_n^a(m, J_n) \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 1 | m_n = m)\right] d\mathbb{P}_\sigma(m_n = m) \\
& + \sum_{k=1}^{|J_n|} \int_{m \in \tilde{M}_n^0(J_n;\sigma)} \left[\sigma_n^a(m, J_n) \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 1 | m_n = m)\right] d\mathbb{P}_\sigma(m_n = m) \\
& + \sum_{k=1}^{|J_n|} \int_{m \in \tilde{M}_n^1(J_n;\sigma)} \left[(1 - \sigma_n^a(m, J_n)) \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 0 | m_n = m)\right] d\mathbb{P}_\sigma(m_n = m) \\
& + \sum_{k=1}^{|J_n|} \int_{m \in \tilde{M}_n^0(J_n;\sigma)} \left[(1 - \sigma_n^a(m, J_n)) \cdot \mathbb{P}_\sigma(h_{<n} \in \mu_n^{-1}(m, J_n, k), \theta = 0 | m_n = m)\right] d\mathbb{P}_\sigma(m_n = m) \\
& + \epsilon \qquad (27)
\end{aligned}$$

which is equivalent to

$$\begin{aligned}\alpha_\sigma(J_n) \leq & \int_{m\in\tilde{M}_n^1(J_n;\sigma)} \left[\sigma_n^a(m,J_n)\cdot \mathbb{P}_\sigma(\theta=1|m_n=m)\right] d\mathbb{P}_\sigma(m_n=m) \\ & + \int_{m\in\tilde{M}_n^0(J_n;\sigma)} \left[\sigma_n^a(m,J_n)\cdot \mathbb{P}_\sigma(\theta=1|m_n=m)\right] d\mathbb{P}_\sigma(m_n=m) \\ & + \int_{m\in\tilde{M}_n^1(J_n;\sigma)} \left[(1-\sigma_n^a(m,J_n))\cdot \mathbb{P}_\sigma(\theta=0|m_n=m)\right] d\mathbb{P}_\sigma(m_n=m) \\ & + \int_{m\in\tilde{M}_n^0(J_n;\sigma)} \left[(1-\sigma_n^a(m,J_n))\cdot \mathbb{P}_\sigma(\theta=0|m_n=m)\right] d\mathbb{P}_\sigma(m_n=m) \\ & +\epsilon \end{aligned} \tag{28}$$

Thus,

$$\begin{aligned}\alpha_\sigma(J_n) \leq & \int_{m\in M_n^w(J_n,\sigma)} \left[\sigma_n^a(m,J_n)\cdot \mathbb{P}_\sigma(\theta=1|m_n=m)\right] d\mathbb{P}_\sigma(m_n=m) \\ & + \int_{m\in M_n^w(J_n,\sigma)} \left[(1-\sigma_n^a(m,J_n))\cdot \mathbb{P}_\sigma(\theta=0|m_n=m)\right] d\mathbb{P}_\sigma(m_n=m)+\epsilon \\ \leq & \int_{m\in M_n^w(J_n,\sigma)} \left[\sigma_n^a(m,J_n)\cdot d\mathbb{P}_\sigma(m_n=m|\theta=1)\right] \mathbb{P}(\theta=1) \\ & + \int_{m\in M_n^w(J_n,\sigma)} \left[(1-\sigma_n^a(m,J_n))\cdot d\mathbb{P}_\sigma(m_n=m|\theta=0)\right] \mathbb{P}(\theta=0)+\epsilon \end{aligned} \tag{29}$$

Notice that the first two terms on the R.H.S. of the inequality (29) is agent n's expected utility in the event that $m_n \in M_n^w(J_n,\sigma)$, under the naive strategy σ_n^a. Therefore it is smaller than the ex ante expected utility $\mathbb{P}_{\sigma_n^a,\sigma_{-n}}(a_n=\theta)$ that agent n gets under σ_n^a (because there may be positive probability that $m_n \notin M_n^w(J_n,\sigma)$), i.e.,

$$\begin{aligned}\mathbb{P}_{\sigma_n^a,\sigma_{-n}}(a_n=\theta) \geq & \int_{m\in M_n^w(J_n,\sigma)} \left[\sigma_n^a(m,J_n)\cdot d\mathbb{P}_\sigma(m_n=m|\theta=1)\right] \mathbb{P}(\theta=1) \\ & + \int_{m\in M_n^w(J_n;\sigma)} \left[(1-\sigma_n^a(m,J_n))\cdot d\mathbb{P}_\sigma(m_n=m|\theta=0)\right] \mathbb{P}(\theta=0) \end{aligned} \tag{30}$$

Finally, inequalities (29) and (30) imply that

$$\mathbb{P}_{\sigma_n^a,\sigma_{-n}}(a_n=\theta) \geq \alpha_\sigma(J_n)-\epsilon$$

which concludes the proof of this lemma. □

Lemma 5. *Suppose for agent n and $J_n \subseteq \{1,...,n-1\}$, we have $\alpha_\sigma(J_n) < 1-\epsilon$ and*

$$\min_{a_{<n}} \mu_n(M_n^w(J_n;\sigma)|a_{<n}) > 1-\epsilon. \tag{31}$$

Then in any equilibrium σ, either

$$\mathbb{P}_\sigma(a_n = \theta) \geq \alpha_\sigma(J_n) - \epsilon + \frac{1}{256}(1 - \alpha_\sigma(J_n))^2 \min\left\{G_1\left(\frac{1 - \alpha_\sigma(J_n)}{16}\right), 1 - G_0\left(1 - \frac{1 - \alpha_\sigma(J_n)}{16}\right)\right\} \tag{32}$$

or

$$\mathbb{P}_\sigma(a_n = \theta) \geq \frac{1 + \alpha_\sigma(J_n)}{2} - \epsilon \tag{33}$$

Proof. The proof makes use of feasible strategies of agent n to derive lower bounds of his expected utility, i.e., $\mathbb{P}_\sigma(a_n = \theta)$. Define $\mathcal{E}_1$ and $\mathcal{E}_0$ as follows (for simplicity of exposition, I suppress their dependence on n, J_n, and σ):

$$\mathcal{E}_1 = \mathbb{P}_\sigma(\theta = 1, m_n \in \tilde{M}_n^0(J_n, \sigma)) \text{ and } \mathcal{E}_0 = \mathbb{P}_\sigma(\theta = 0, m_n \in \tilde{M}_n^1(J_n, \sigma)). \tag{34}$$

$\mathcal{E}_1$ is the probability that the true state is 1 but an action-signal in $\tilde{M}_n^0(J_n, \sigma)$ (which indicates that the state is more likely to be 0) is sent to agent n. Similarly, $\mathcal{E}_0$ is the probability that the true state is 0 but an action-signal in $\tilde{M}_n^1(J_n, \sigma)$ is sent to agent n. We should note that m_n is generated independently conditional on action-histories before agent n. Hence, $\mathcal{E}_1$ and $\mathcal{E}_0$ only depend on the strategy profile of the first $n-1$ agents and do not depend on agent n's strategy.

We consider three cases (which cover all possibilities):

- $\mathcal{E}_0 + \mathcal{E}_1 \leq (1 - \alpha_\sigma(J_n))/2$,
- $\mathcal{E}_0 \geq (1 - \alpha_\sigma(J_n))/4$,
- $\mathcal{E}_1 \geq (1 - \alpha_\sigma(J_n))/4$.

Case 1: $\mathcal{E}_1 + \mathcal{E}_0 \leq (1 - \alpha_\sigma(J_n))/2$

In this case, consider the following *naive* strategy for agent n which only depends on his action-signal m_n:

$$\sigma_n^a(m_n, s_n) = \begin{cases} 1, & \text{if } m_n \in \tilde{M}_n^1(J_n, \sigma) \\ 0, & \text{if } m_n \in \tilde{M}_n^0(J_n, \sigma) \\ 0, & \text{if } m_n \notin M_n^w(J_n, \sigma) \end{cases} \tag{35}$$

Now, we can interpret the term $\mathcal{E}_1+\mathcal{E}_0$ as the probability of agent n observing an action-signal in $M_n^w(J_n,\sigma)$ that specifies the "wrong" action under the naive strategy σ_n^a. In this case (case 1: $\mathcal{E}_1+\mathcal{E}_0 \le (1-\alpha_\sigma(J_n))/2$), the probability of such "wrong" specification is assumed to be no larger than $(1-\alpha_\sigma(J_n))/2$.

The expected utility that agent n gets under σ_n^a is:

$$\begin{aligned}
\mathbb{P}_{\sigma_n^a,\sigma_{-n}}(a_n=\theta) &= \mathbb{P}_{\sigma_n^a,\sigma_{-n}}(\theta=1,a_n=1)+\mathbb{P}_{\sigma_n^a,\sigma_{-n}}(\theta=0,a_n=0)\\
&\ge 1-\mathbb{P}_{\sigma_n^a,\sigma_{-n}}(\theta=1,m_n\in\tilde{M}_n^0(J_n,\sigma))-\mathbb{P}_{\sigma_n^a,\sigma_{-n}}(\theta=0,m_n\in\tilde{M}_n^1(J_n,\sigma))\\
&\quad -\mathbb{P}_\sigma(m_n\notin\tilde{M}_n^w(J_n,\sigma))\\
&\ge 1-\mathcal{E}_1-\mathcal{E}_0-\epsilon\\
&\ge \frac{1+\alpha_{\sigma(J_n)}}{2}-\epsilon
\end{aligned}$$

The first inequality follows from the specification of the strategy in (35) and the fact that $\mathbb{P}_\sigma(m_n\notin\tilde{M}_n^w(J_n,\sigma))\le\epsilon$ (implied by (31)); the second and third inequalities follow from the fact that the distribution of m_n is independent of agent n's strategy profile and the assumption that $\mathcal{E}_1+\mathcal{E}_0\le(1-\alpha_\sigma(J_n))/2$.

Note that σ_n^a is a feasible strategy for agent n because $\tilde{M}_n^1(J_n,\sigma)$ and $\tilde{M}_n^0(J_n,\sigma)$ are disjoint sets. So we conclude that $\frac{1+\alpha_{\sigma(J_n)}}{2}-\epsilon$ is a lower bound of agent n's equilibrium payoff $\mathbb{P}_\sigma(a_n=\theta)$, resulting in (33).

Case 2: $\mathcal{E}_0\ge(1-\alpha_\sigma(J_n))/4$

In this case, consider the following *fine-tuned naive strategy* for agent n which depends on his action-signal m_n and some *extreme values* of s_n:

$$\sigma_n^f(m_n,s_n)=\begin{cases}0, & \text{if } m_n\in\tilde{M}_n^1(J_n,\sigma),\ 1-q_n^\sigma(m_n)>\epsilon_1, \text{ and } p_n(s_n)<\epsilon_1\\ 1, & \text{if } m_n\in\tilde{M}_n^1(J_n,\sigma) \text{ and either } 1-q_n^\sigma(m_n)<\epsilon_1 \text{ or } p_n(s_n)>\epsilon_1\\ 0, & \text{if } m_n\in\tilde{M}_n^0(J_n,\sigma)\\ 0, & \text{if } m_n\notin M_n^w(J_n,\sigma)\end{cases} \tag{36}$$

where $p_n(s_n)=\mathbb{P}(\theta=1|s_n)$, $q_n^\sigma(m_n)=\mathbb{P}_\sigma(\theta=1|m_n)$, and $\epsilon_1=(1-\alpha_\sigma(J_n))/8$. Notice that the strategy σ_n^f differs from the naive strategy σ_n^a only in events where $m_n\in\tilde{M}_n^1(J_n,\sigma)$, $1-q_n^\sigma(m_n)>\epsilon_1$, and $p_n(s_n)<\epsilon_1$. These are events where the naive strategy specifies action 1, while the state-signal provides strong evidence that the state is actually $\theta=0$, i.e., the naive strategy is recommending the wrong action. The strategy σ_n^f, therefore, "reverses" the actions specified by the naive strategy in these events wherein it is very likely that the naive strategy is wrong. σ_n^f is again a feasible strategy for agent n.

Note that we assumed $\mathcal{E}_0 \geq (1 - \alpha_\sigma(J_n))/4$. This assumption indicates that, when the state is 0, there is positive probability (bounded away from zero) that the naive strategy specifies the wrong action (action 1). Therefore, the fine-tuned strategy σ_n^f (by reversing these actions) should (on average) strictly improve upon σ_n^a.

Following this intuition, I will show that the expected utility that agent n gets under σ_n^f is strictly higher than the utility he gets under σ_n^a, which in turn is higher than $\alpha_\sigma(J_n) - \epsilon$. In particular:

$$\mathbb{P}_{\sigma_n^f,\sigma_{-n}}(a_n = \theta) \geq \mathbb{P}_{\sigma_n^a,\sigma_{-n}}(a_n = \theta) + \frac{\epsilon_1}{2} G_1(\frac{\epsilon_1}{2}) \frac{1 - \alpha_\sigma(J_n)}{16} \tag{37}$$

and

$$\mathbb{P}_{\sigma_n^a,\sigma_{-n}}(a_n = \theta) \geq \alpha_\sigma(J_n) - \epsilon. \tag{38}$$

These two inequalities together implies (32), and concludes the proof of case 2.

The proof of (38) follows from Lemma 4. Just to reiterate the intuition: if $m \in \tilde{M}_n^0(J_n, \sigma)$ or $m \in \tilde{M}_n^1(J_n, \sigma)$, agent n can weakly improve upon an average agent in J_n, by choosing action 0 or 1. Since it is assumed that the probability that $m \in \tilde{M}_n^0(J_n, \sigma)$ or $m \in \tilde{M}_n^1(J_n, \sigma)$ is greater than $1 - \epsilon$ (implied by (31)), agent n's utility by following the naive strategy must be higher than the average utility of agents in J_n minus ϵ.

Now, I explain why (37) is correct: Notice that σ_n^f and σ_n^a only differs in the event when $m_n \in \tilde{M}_n^1(J_n, \sigma)$, $1 - q_n^\sigma(m_n) > \epsilon_1$, and $p_n(s_n) < \epsilon_1$. Thus,

$$\begin{aligned}
& \mathbb{P}_{\sigma_n^f,\sigma_{-n}}(a_n = \theta) - \mathbb{P}_{\sigma_n^a,\sigma_{-n}}(a_n = \theta) \\
= \; & \int_{m \in \tilde{M}_n^1(J_n \sigma), q_n^\sigma(m) < 1-\epsilon_1} (-1) \mathbb{P}(p_n(s_n) < \epsilon_1 | \theta = 1) d\mathbb{P}_\sigma(m_n = m | \theta = 1) \mathbb{P}(\theta = 1) \\
& + \int_{m \in \tilde{M}_n^1(J_n \sigma), q_n^\sigma(m) < 1-\epsilon_1} 1 \cdot \mathbb{P}(p_n(s_n) < \epsilon_1 | \theta = 0) d\mathbb{P}_\sigma(m_n = m | \theta = 0) \mathbb{P}(\theta = 0) \\
= \; & \int_{m \in \tilde{M}_n^1(J_n \sigma), q_n^\sigma(m) < 1-\epsilon_1} (-1) G_1(\epsilon_1) d\mathbb{P}_\sigma(m_n = m | \theta = 1) \mathbb{P}(\theta = 1) \\
& + \int_{m \in \tilde{M}_n^1(J_n \sigma), q_n^\sigma(m) < 1-\epsilon_1} G_0(\epsilon_1) d\mathbb{P}_\sigma(m_n = m | \theta = 0) \mathbb{P}(\theta = 0)
\end{aligned} \tag{39}$$

where the first term is the utility loss for σ_n^f from reversing the action of the naive strategy, and the second term is the utility gain from reversing the action. What we will show is that the gain is higher than the loss. According to Lemma 2 in the original paper, we have the following inequality: for any $0 < z < \epsilon_1$,

$$G_0(\epsilon_1) \geq \frac{1-\epsilon_1}{\epsilon_1} G_1(\epsilon_1) + \frac{\epsilon_1 - z}{2} G_1(z) \tag{40}$$

Take $z = \epsilon_1/2$ and apply inequality (40) to (39):

$$\begin{aligned}
& \mathbb{P}_{\sigma_n^f, \sigma_{-n}}(a_n = \theta) - \mathbb{P}_{\sigma_n^a, \sigma_{-n}}(a_n = \theta) \\
\geq \; & \int_{m \in \tilde{M}_n^1(J_n\sigma), q_n^\sigma(m) < 1-\epsilon_1} (-1) G_1(\epsilon_1) d\mathbb{P}_\sigma(m_n = m | \theta = 1) \mathbb{P}(\theta = 1) \\
& + \int_{m \in \tilde{M}_n^1(J_n\sigma), q_n^\sigma(m) < 1-\epsilon_1} \frac{1-\epsilon_1}{\epsilon_1} G_1(\epsilon_1) + \frac{\epsilon_1}{4} G_1(\frac{\epsilon_1}{2}) d\mathbb{P}_\sigma(m_n = m | \theta = 0) \mathbb{P}(\theta = 0)
\end{aligned} \tag{41}$$

Now since $q_n^\sigma(m_n) = \mathbb{P}_\sigma(\theta = 1 | m_n)$, according to Bayes' rule, we must have that

$$\frac{\mathbb{P}_\sigma(m_n | \theta = 1) \mathbb{P}(\theta = 1)}{\mathbb{P}_\sigma(m_n | \theta = 0) \mathbb{P}(\theta = 0)} = \frac{q_n^\sigma(m_n)}{1 - q_n^\sigma(m_n)} \tag{42}$$

and

$$\mathbb{P}_\sigma(m_n | \theta = 1) = 2 q_n^\sigma(m_n) \mathbb{P}_\sigma(m_n) \tag{43}$$

$$\mathbb{P}_\sigma(m_n | \theta = 0) = (1 - q_n^\sigma(m_n)) \mathbb{P}_\sigma(m_n) \tag{44}$$

Thus, using (42), we can rewrite (41) as

$$\begin{aligned}
& \mathbb{P}_{\sigma_n^f, \sigma_{-n}}(a_n = \theta) - \mathbb{P}_{\sigma_n^a, \sigma_{-n}}(a_n = \theta) \\
\geq \; & \int_{m \in \tilde{M}_n^1(J_n\sigma), q_n^\sigma(m) < 1-\epsilon_1} \left[\frac{1-\epsilon_1}{\epsilon_1} - \frac{1 - q_n^\sigma(m)}{q_n^\sigma(m)} \right] G_1(\epsilon_1) + \frac{\epsilon_1}{4} G_1(\frac{\epsilon_1}{2}) d\mathbb{P}_\sigma(m_n = m | \theta = 0) \mathbb{P}(\theta = 0)
\end{aligned}$$

Note that for any $q_n^\sigma(m) < 1 - \epsilon_1$, we have

$$\left[\frac{1-\epsilon_1}{\epsilon_1} - \frac{1 - q_n^\sigma(m)}{q_n^\sigma(m)} \right] \geq 0 \tag{45}$$

Therefore,

$$
\begin{aligned}
& \mathbb{P}_{\sigma_n^f,\sigma_{-n}}(a_n = \theta) - \mathbb{P}_{\sigma_n^a,\sigma_{-n}}(a_n = \theta) \\
\geq & \int_{m \in \tilde{M}_n^1(J_n\sigma), q_n^\sigma(m) < 1-\epsilon_1} \frac{\epsilon_1}{4} G_1(\frac{\epsilon_1}{2}) d\mathbb{P}_\sigma(m_n = m|\theta = 0)\mathbb{P}(\theta = 0) \\
\geq & \frac{\epsilon_1}{4} G_1(\frac{\epsilon_1}{2}) \mathbb{P}_\sigma(m \in \tilde{M}_n^1(J_n\sigma), q_n^\sigma(m) < 1-\epsilon_1 | \theta = 0)\mathbb{P}(\theta = 0) \qquad (46)
\end{aligned}
$$

Note that except the term $\mathbb{P}_\sigma(m \in \tilde{M}_n^1(J_n\sigma), q_n^\sigma(m) < 1-\epsilon_1|\theta = 0)$, all other terms on the R.H.S. of the above inequality are positive constants. We now use the assumption that $\mathcal{E}_0 > (1-\alpha_\sigma(J_n))/4$ to bound the term $\mathbb{P}_\sigma(m \in \tilde{M}_n^1(J_n\sigma), q_n^\sigma(m) < 1-\epsilon_1|\theta = 0)$.

Recall that $\mathcal{E}_0 = \mathbb{P}_\sigma(m_n \in \tilde{M}_n^1(J_n,\sigma), \theta = 0)$, so $\mathcal{E}_0 > (1-\alpha_\sigma(J_n))/4$ is equivalent to

$$
\mathbb{P}_\sigma(m_n \in \tilde{M}_n^1(J_n,\sigma)|\theta = 0)\mathbb{P}(\theta = 0) \geq \frac{1-\alpha_\sigma(J_n)}{4} \qquad (47)
$$

Note that

$$
\begin{aligned}
& \mathbb{P}_\sigma(m_n \in \tilde{M}_n^1(J_n,\sigma)|\theta = 0) \\
= & \mathbb{P}_\sigma(m_n \in \tilde{M}_n^1(J_n,\sigma), q_n^\sigma(m) < 1-\epsilon_1|\theta = 0) + \mathbb{P}_\sigma(m_n \in \tilde{M}_n^1(J_n,\sigma), q_n^\sigma(m) > 1-\epsilon_1|\theta = 0) \\
\leq & \mathbb{P}_\sigma(m_n \in \tilde{M}_n^1(J_n,\sigma), q_n^\sigma(m) < 1-\epsilon_1|\theta = 0) + \int_{m_n \in \tilde{M}_n^1(J_n,\sigma), q_n^\sigma(m) > 1-\epsilon_1} d\mathbb{P}_\sigma(m_n = m|\theta = 0) \\
= & \mathbb{P}_\sigma(m_n \in \tilde{M}_n^1(J_n,\sigma), q_n^\sigma(m) < 1-\epsilon_1|\theta = 0) + \int_{m_n \in \tilde{M}_n^1(J_n,\sigma), q_n^\sigma(m) > 1-\epsilon_1} 2(1-q_n^\sigma(m)) d\mathbb{P}_\sigma(m_n = m) \\
\leq & \mathbb{P}_\sigma(m_n \in \tilde{M}_n^1(J_n,\sigma), q_n^\sigma(m) < 1-\epsilon_1|\theta = 0) + \int_{m_n \in \tilde{M}_n^1(J_n,\sigma), q_n^\sigma(m) > 1-\epsilon_1} 2\epsilon_1 d\mathbb{P}_\sigma(m_n = m) \\
\leq & \mathbb{P}_\sigma(m_n \in \tilde{M}_n^1(J_n,\sigma), q_n^\sigma(m) < 1-\epsilon_1|\theta = 0) + 2\epsilon_1 \qquad (48)
\end{aligned}
$$

where the second equality follows from (44). Therefore, (48) implies that

$$
\begin{aligned}
& \mathbb{P}_\sigma(m_n \in \tilde{M}_n^1(J_n,\sigma), q_n^\sigma(m) < 1-\epsilon_1|\theta = 0) \\
\geq & \mathbb{P}_\sigma(m_n \in \tilde{M}_n^1(J_n,\sigma)|\theta = 0) - 2\epsilon_1 \\
\geq & \mathcal{E}_0/\mathbb{P}(\theta = 0) - 2\epsilon_1 \\
\geq & \frac{1-\alpha_\sigma(J_n)}{2} - 2\epsilon_1 \qquad (49)
\end{aligned}
$$

Substitute inequality (49) back to (46), we have

$$\begin{aligned}
& \mathbb{P}_{\sigma_n^f,\sigma_{-n}}(a_n = \theta) - \mathbb{P}_{\sigma_n^a,\sigma_{-n}}(a_n = \theta) \\
\geq\ & \frac{\epsilon_1}{4} G_1(\frac{\epsilon_1}{2}) \mathbb{P}_\sigma(m \in \tilde{M}_n^1(J_n\sigma), q_n^\sigma(m) < 1-\epsilon_1 | \theta = 0) \mathbb{P}(\theta = 0) \\
\geq\ & \frac{\epsilon_1}{4} G_1(\frac{\epsilon_1}{2})(\mathcal{E}_0 - \epsilon_1) \\
\geq\ & \frac{1-\alpha_\sigma(J_n)}{32} G_1(\frac{1-\alpha_\sigma(J_n)}{16})(\frac{1-\alpha_\sigma(J_n)}{8}) \\
\geq\ & \frac{(1-\alpha_\sigma(J_n))^2}{256} G_1(\frac{1-\alpha_\sigma(J_n)}{16})
\end{aligned} \tag{50}$$

where we used the fact that $\epsilon_1 = (1-\alpha_\sigma(J_n))/8$ and $\mathcal{E}_0 > (1-\alpha_\sigma(J_n))/4$.

Combining inequalities (50) with (38), we conclude that (32) is correct in case 2.

Case 3: $\mathcal{E}_1 \geq (1-\alpha_\sigma(J_n))/4$

In this case, consider the following *fine-tuned naive strategy* for agent n which depends on his action-signal m_n and some *extreme values* of s_n:

$$\sigma_n^{f'}(m_n, s_n) = \begin{cases} 1, & \text{if } m_n \in \tilde{M}_n^0(J_n,\sigma),\ 1 - q_n^\sigma(m_n) < 1-\epsilon_1, \text{ and } p_n(s_n) > 1-\epsilon_1 \\ 0, & \text{if } m_n \in \tilde{M}_n^0(J_n,\sigma) \text{ and either } 1 - q_n^\sigma(m_n) > 1-\epsilon_1 \text{ or } p_n(s_n) < 1-\epsilon_1 \\ 1, & \text{if } m_n \in \tilde{M}_n^1(J_n,\sigma) \\ 0, & \text{if } m_n \notin M_n^w(J_n,\sigma) \end{cases}$$

An analogous argument as in case 2 shows that (32) is again correct. □

Proof of Theorem 2. Suppose, by way of contradiction, that information aggregation does not occur, i.e.,

$$\liminf_{n\to\infty} \mathbb{P}_\sigma(a_n = \theta) = \hat{p} < 1.$$

This implies that for any $\epsilon_1 > 0$, there exists $n_1(\epsilon_1) \in \mathbb{N}$ such that for all $n > n_1(\epsilon_1)$, we have

$$\mathbb{P}_\sigma(a_n = \theta) \geq \hat{p} - \epsilon_1$$

Since $\alpha_\sigma(J_n) = \sum_{i\in J_n} \mathbb{P}_\sigma(a_i = \theta)/|J_n|$, for any J_n such that $\min_{i\in J_n} i \geq n_1(\epsilon_1)$ we have

$$\alpha_\sigma(J_n) \geq \hat{p} - \epsilon_1$$

Since the action-signals are separating, for any $\epsilon > 0$, there exists $n_2(\epsilon, \epsilon_1) > n_1(\epsilon_1)$ such that for any $n > n_2(\epsilon, \epsilon_1)$, there exists $J_n \in \{n_1(\epsilon_1), ..., n-1\}$ such that for all $a_{<n}$, we have

$$\mu_n(M_n(J_n, \sigma)|a_{<n}) > 1 - \epsilon$$

By Lemma 5, this implies that for any $n > n_2(\epsilon, \epsilon_1)$, there exists $J_n \in \{n_1(\epsilon_1), ..., n-1\}$ such that

$$\mathbb{P}_\sigma(a_n = \theta) \geq \alpha_\sigma(J_n) - \epsilon + \frac{1}{256}(1 - \alpha_\sigma(J_n))^2 \min\left\{G_1\left(\frac{1-\alpha_\sigma(J_n)}{16}\right), 1 - G_0\left(1 - \frac{1-\alpha_\sigma(J_n)}{16}\right)\right\}$$

or

$$\mathbb{P}_\sigma(a_n = \theta) \geq \frac{1 + \alpha_\sigma(J_n)}{2} - \epsilon$$

Since $\liminf_{n\to\infty} \mathbb{P}_\sigma(a_n = \theta) = \hat{p} < 1$, there exists a subsequence $\{n^i\}_{i=1}^\infty$ of agents such that

$$\lim_{i\to\infty} \mathbb{P}_\sigma(a_{n^i} = \theta) = \hat{p} < 1 \tag{51}$$

Therefore, for any $\epsilon_3 > 0$, there exists $n_3(\epsilon_3) \in \mathbb{N}$ such that for all $n^i > n_3(\epsilon_3)$, we have

$$\mathbb{P}_\sigma(a_{n^i} = \theta) \leq \hat{p} + \epsilon_3 \tag{52}$$

Therefore, for all $n^i > \max\{n_2(\epsilon, \epsilon_1), n_3(\epsilon_3)\}$, we have

$$\hat{p} + \epsilon_3 \quad \geq \quad \alpha_\sigma(J_{n^i}) - \epsilon + \frac{1}{256}(1 - \alpha_\sigma(J_{n^i}))^2 \min\left\{G_1\left(\frac{1-\alpha_\sigma(J_{n^i})}{16}\right), 1 - G_0\left(1 - \frac{1-\alpha_\sigma(J_{n^i})}{16}\right)\right\} \tag{53}$$

or

$$\mathbb{P}_\sigma(a_{n^i} = \theta) \geq \frac{1 + \alpha_\sigma(J_{n^i})}{2} - \epsilon$$

Consider the first case, for the inequality to hold for all n^i, we must have

$$\lim_{n\to\infty} \alpha_\sigma(J_{n^i}) = \hat{p}.$$

However, in this case, as $\epsilon, \epsilon_1, \epsilon_3$ goes to zero (i.e., n^i goes to infinity), the inequality features

$$\hat{p} \quad \geq \quad \hat{p} + \frac{1}{256}(1-\hat{p})^2 \min\left\{G_1\left(\frac{1-\hat{p}}{16}\right), 1 - G_0\left(1 - \frac{1-\hat{p}}{16}\right)\right\}$$

The R.H.S. is strictly larger than $\hat{p}$ because $\hat{p} \in [1/2, 1)$. Therefore, there exists $n(\epsilon, \epsilon_1, \epsilon_3)$ large enough such that for all $n^i > n(\epsilon, \epsilon_1, \epsilon_3)$, inequality (54) must be violated. This results in a contradiction.

Similarly, as n goes to infinity, the case

$$\mathbb{P}_\sigma(a_n = \theta) \geq \frac{1 + \alpha_\sigma(J_n)}{2} - \epsilon$$

will feature

$$\hat{p} \geq \frac{1 + \hat{p}}{2}$$

which also results in a contradiction. Thus, our assumption that information aggregation fails must be wrong. So information aggregation occurs. □

B3 Proof of Proposition 3

Proof. First, suppose an action-signal technology μ is deterministically strongly separating. Then there exist non-empty sets $\{J_n\}_{n=1}^{\infty}$ s.t. $J_n \subseteq \{1, ..., n-1\}$, $\limsup_{n\to\infty} J_n = \varnothing$ and

$$\lim_{n\to\infty} \min_{a_{<n}} \mu_n \left(M_n^s(J_n, \sum_{j\in J_n} a_j) \middle| a_{<n} \right) = 1$$

Then for any ϵ, there exists $n_\epsilon \in \mathbb{N}$ such that for all $n > n_\epsilon$, $\forall a_{<n} \in \mathcal{A}_{<n}$

$$\mu_n \left(M_n^s(J_n, \sum_{j\in J_n} a_j) \middle| a_{<n} \right) > 1 - \epsilon$$

According to the definition of $M_n^s(J_n, \sum_{j\in J_n} a_j)$ and $\mu_n^{-1}(m, J_n, k)$, we know that for any $m \in M_n^s(J_n, \sum_{j\in J_n} a_j)$ and any $k \neq \sum_{j\in J_n} a_j$, we have

$$\mu_n^{-1}(m, J_n, k) = \varnothing.$$

Therefore, given any equilibrium strategy profile σ and any $m \in M_n^s(J_n, \sum_{j\in J_n} a_j)$, since $\mathbb{P}\left(\theta = 1 | \mu_n^{-1}(m, J_n, \sum_{j\in J_n} a_j\right)$ must be either weakly larger than 1/2 or weakly smaller than

1/2 and $\mathbb{P}(\theta = 1|\mu_n^{-1}(m, J_n, k)$ is undefined for all $k \neq \sum_{j\in J_n} a_j$ (because $\mu_n^{-1}(m, J_n, k) = \varnothing$), we must have $m \in M_n^w(J_n, \sigma)$. This implies that

$$M_n(J_n, \sum_{j\in J_n} a_j) \subseteq M_n^w(J_n, \sigma).$$

Hence, for all $n > n_\epsilon$, $\forall a_{<n} \in \mathcal{A}_{<n}$,

$$\mu_n\left(M_n^w(J_n, \sigma)|a_{<n}\right) \geq \mu_n\left(M_n^s(J_n, \sum_{j\in J_n} a_j)\middle|a_{<n}\right) > 1 - \epsilon.$$

Thus,

$$\lim_{n\to\infty} \min_{a_{<n}} \mu_n\left(M_n^w(J_n; \sigma)|a_{<n}\right) = 1.$$

This implies that the action-signals are weakly separating.

□

B4 Explanation of the Independence Condition in Stochastic Separation

The assumption that γ_n are independent is important. Lobel and Sadler (2014) provides an interesting example: consider a stochastic social network in which agents' neighborhoods are correlatedly generated. Let E_n denote the (realized) set of agents before n whose observes an completely uninformative action-signal, and define the action-signal m_n to be uninformative with probability $1/(2^{|E_n|})$ and to reveal the action and identity of agent $\max_{i\in E_n} i$ otherwise.

In this example, all agents in the network either have an empty neighborhood or observe one agent whose neighborhood is empty in turn, so information aggregation fails. It satisfies all conditions in the definition stochastic strong separation except the independence condition.

B5 Explanation of Example 1 in Section 5.3

To see why the example doe not satisfy strong separation, note that for each agent n, the action-signal m_n has four possible realizations

$$\mathcal{M} = \{(\{n-1\}, 1), (\{n-1\}, 0), (\{n-2\}, 1), (\{n-2\}, 0)\} \tag{54}$$

We consider different choices of J_n:

Case 1: $n-1 \in J_n$

For any action-history $a_{<n}$ such that $a_i = 1$ for all $i \in J_n \bigcup \{n-2\}$, we have

$$\mu_n((\{n-1\},1)|a_{<n}) = 1/2 \text{ and } \mu_n((\{n-2\},1)|a_{<n}) = 1/2$$

For any action-history $a'_{<n}$ such that $a'_{n-2} = 1, a'_{n-1} = 0$, and $a'_i = 1$ for all $i \in J_n/\{n-1\}$, we have

$$\mu_n((\{n-1\},0)|a'_{<n}) = 1/2 \text{ and } \mu_n((\{n-2\},1)|a'_{<n}) = 1/2$$

Notice that the action-signal realization $(\{n-2\},1)$ are possible given both action-histories $a_{<n}$ and $a'_{<n}$. However,

$$\sum_{j \in J_n} a_j = |J_n| \neq \sum_{j \in J_n} a'_j = |J_n| - 1.$$

Thus, according to the definition of the set $M_n(J_n, k)$ in the section on strong separation, we must have

$$(\{n-2\},1) \notin M_n(J_n, |J_n|)$$

This implies that, for the action-history $a_{<n}$ introduced above such that $a_i = 1$ for all $i \in J_n \bigcup \{n-2\}$, we have

$$\mu_n\left(M_n\left(J_n, \sum_{j \in J_n} a_j\right)\middle| a_{<n}\right) \leq \frac{1}{2}.$$

This implies that

$$\min_{a_{<n}} \mu_n\left(M_n\left(J_n, \sum_{j \in J_n} a_j\right)\middle| a_{<n}\right) \leq \frac{1}{2}.$$

Case 2: $n-2 \in J_n$ Similar argument shows that

$$\min_{a_{<n}} \mu_n\left(M_n\left(J_n, \sum_{j \in J_n} a_j\right)\middle| a_{<n}\right) \leq \frac{1}{2}.$$

Since the argument above is correct for any $n \in \mathbb{N}$, we have

$$\lim_{n\to\infty} \min_{a_{<n}} \mu_n \left(M_n \left(J_n, \sum_{j \in J_n} a_j \right) \middle| a_{<n} \right) \leq \frac{1}{2}.$$

So the action-signals fails to satisfy strong separation. The argument for why the action-signals do not satisfy weak separation is similar.

B6 Relationship between Strong Separation and Networks with Expanding Observations

In this subsection, I show how the result on strong separation implies a classic result in the literature, established by Acemoglu et al. (2011). They study social learning in networks as follows: each agent n observes the identities and the actions of agents in a stochastic neighborhood $B(n) \subseteq \{1, 2, ..., n-1\}$ where $B(n)$ is generated according to some probability distribution $\mathbb{Q}_n$ over all subsets of $\{1, 2, ..., n-1\}$, and draws from each $\mathbb{Q}_n$ are independent from each other for all n and from the realizations of the state-signals.

Definition 13. *The network has expanding observations if for all $K \in \mathbb{N}$, we have*

$$\lim_{n\to\infty} \mathbb{Q}_n \left(\max_{b \in B(n)} b < K \right) = 0 \tag{55}$$

It is shown that if the network has expanding observations and the state-signals are unboundedly informative, information aggregation occurs in all equilibria. This classic result is a special case of the result on strong separation in my paper: for any social network that has expanding observations, there is an outcome-equivalent way to represent the information environment as one in which agents observe action-signals that are strongly separating.

To see this, consider the following environment: each agent n observes an action-signal that specifies the identities and the actions of the agent in his stochastic neighborhood, i.e.,

$$m_n = (B(n), a_k, \forall k \in B(n))$$

and is generated according to the information structure $\mu_n(\cdot|a_{<n})$ such that

$$\mu_n((B(n), a_{B(n)})|a_{<n}) = \mathbb{Q}_n(B(n)) \tag{56}$$

Let $M_n(J,k) := \{(B(n), a_{B(n)}) : B(n) = J, \sum_{j \in B(n)} a_j = k\}$. Since $\mathbb{Q}_n$ is independent of the action-history $a_{<n}$, condition (b) in the definition of strong separation is satisfied. Since $\mathbb{Q}_n$ satisfies expanding observations, condition (a) in the definition of strong separation is satisfied. Finally, it is easy to see that $M_n(J,k) \subseteq M_n^s(J,k)$ for every $n \in \mathbb{N}$, $J \subseteq \{1,...,n-1\}$, and $k \in \{0,...,|J|\}$, and that the sets $\{M_n(J,k)\}_{J,k}$ are mutually disjoint. Therefore the action-signal technology is strongly-separating, and leads to information aggregation when the state-signals are unboundedly informative.

B7 Proof of Theorem 3

Proof. Since stochastic strong separation implies stochastic weak separation, we only need to prove that stochastic weak separation give rise to information aggregation when the state-signals are unboundedly informative.

The proof parallels that of Theorem 2. The critical difference is in the proof of Lemma 4. Instead of using the $\alpha_\sigma(J_n)$ function, we define a function

$$\gamma_n(J) := \mu_n(M_n(J,\sigma)|a_{<n}) \text{ for all } a_{<n}$$

The function γ_n is well-defined due to the condition (b) in the definition of weak separation (Definition 7). Then we define the average utility of agents in sets generated according to γ_n to be

$$\alpha_\sigma(\gamma_n) := \sum_{J \in \mathcal{N}_1^{n-1}} \left(\frac{\sum_{j \in J} \mathbb{P}_\sigma(a_j = \theta)}{|J|} \right) \cdot \gamma_n(J)$$

Then the same proof technique as in Lemma 4 and 5 shows that for any n, if $\alpha_\sigma(\gamma_n)$ is not 1 and $\gamma_n(\mathcal{N}_k^{n-1})$ is close to 1, then agent n's ex ante expected utility will be strictly higher than $\alpha_\sigma(\gamma_n)$ (all the derivations in Lemma 4 and 5 will now involve a weighted average over all $J \in \mathcal{N}_k^{n-1}$ with weights equal to $\gamma_n(J)$). Then the same argument as in the proof of Theorem 2 shows that information aggregation must occur. $\square$

B8 Proof of Proposition 4

The first part of the statement follows from the fact that $M_n^s(J,k) \subseteq M_n^w(J,\sigma)$. The second part of the statement can be easily shown by taking γ_n (in the definitions of stochastic strong and weak separation) to be dirac on the J_n set in the definition of strong and weak separation.

B9 Proof of Theorem 4

Proof. If the action-signals are fullly informative, then herding and information cascade will occur as studied in Smith and Sorensen (2000). ◻

B10 Proof of Theorem 5

Proof. When there is an infinite sequence of fully ignorant agents, the collection of their actions will fully reveal the state due to law of large numbers.

Since $\lim_{n\to\infty}\min_{a_{<n}}\mu_n(M_n^s(\{n_i\},a_{n_i})) = 1$ for all i, the infinite sequence of ignorant agents are perfectly observed by the entire population (in the limit), so all agents will asymptotically learn the true state. ◻

B11 Proof of Theorem 6

I start by showing that upon observing any finite action-history, an agents' posterior belief about the state is always interior. We use $h_{<n}$ to denote a history of actions and action-signals that occurred before agent n takes his action. That is,

$$h_{<n} = (m_1, a_1, m_2, a_2, ..., a_{n-1}, m_n).$$

To distinguish from the action-histories, we will refer to such histories as *action-signal-histories*, and use $\mathcal{H}_{<n}$ to denote the set of all possible date-n action-signal-histories.

Lemma 6. *Given any strategy profile σ, any finite integer k, and any date-k action-history $a_{<k} \in \mathcal{A}_{<k}$, we have $\mathbb{P}_\sigma(\theta = 1|a_{<k}) \in (0,1)$.*

Proof. We prove the statement by induction. First, we establish that the statement is correct when $k = 1$. The distribution of the state-signals $(\mathbb{F}_0, \mathbb{F}_1)$ are assumed to not fully reveal the states. Thus,

$$\mathbb{P}_\sigma(\theta = 1|a_1 = 1) \in (0,1) \text{ and } \mathbb{P}_\sigma(\theta = 1|a_1 = 0) \in (0,1).$$

For any $h_{<2} = (m_1, a_1, m_2) \in \mathcal{H}_{<2}$, we must have

$$\mathbb{P}_\sigma(\theta = 1|h_{<2}) = \mathbb{P}_\sigma(\theta = 1|a_1, m_2) = \mathbb{P}_\sigma(\theta = 1|a_1) \in (0,1),$$

where the first equality follows from the fact that m_1 is uninformative about the state, and the second equality follows because the action-signal m_2 is independent of θ conditional on a_1.

Now, as an inductive step, suppose for $k \in \mathbb{N}^+$, we have

$$\mathbb{P}_\sigma(\theta = 1|a_{<k}) \in (0,1), \qquad \forall a_{<k} \in \mathcal{A}_{<k}, \tag{57}$$

$$\mathbb{P}_\sigma(\theta = 1|h_{<k}) \in (0,1), \qquad \forall h_{<k} \in \mathcal{H}_{<k}. \tag{58}$$

We want to show that, for any $a_{<k+1} \in \mathcal{A}_{<k+1}$ and any $h_{<k+1} \in \mathcal{H}_{<k+1}$, we have $\mathbb{P}_\sigma(\theta = 1|a_{<k+1}) \in (0,1)$ and $\mathbb{P}_\sigma(\theta = 1|h_{<k+1}) \in (0,1)$.

Since $|\mathcal{A}_{<k}| = 2^{k-1}$ is finite, $t_k = \max_{a_{<k}} \mathbb{P}_\sigma(\theta = 1|a_{<k})$ and $b_k = \min_{a_{<k}} \mathbb{P}_\sigma(\theta = 1|a_{<k})$ exist, and $t_k < 1$ and $b_k > 0$. According to Bayes' rule, we have, for any $m_k \in \bigcup_{a_{<k} \in \mathcal{A}_{<k}} \operatorname{supp}(\mu_n(\cdot|a_{<k}))$,

$$\mathbb{P}_\sigma(\theta = 1|m_k) = \frac{\int_{a_{<k}} \mathbb{P}_\sigma(\theta = 1|\tilde{a}_{<k} = a_{<k}, m_k)\mu_k(m_k|a_{<k})\mathbb{P}_\sigma(\tilde{a}_{<k} = a_{<k})}{\int_{a_{<k}} (m_k|a_{<k})\mathbb{P}_\sigma(\tilde{a}_{<k} = a_{<k})} \in (b_k, t_k) \subset (0,1), \tag{59}$$

because $\mathbb{P}_\sigma(\theta = 1|\tilde{a}_{<k} = a_{<k}, m_k) = \mathbb{P}_\sigma(\theta = 1|\tilde{a}_{<k} = a_{<k}) \in (b_k, t_k)$ for all $a_{<k} \in \mathcal{A}_{<k}$ and all $m_k \in \bigcup_{a_{<k} \in \mathcal{A}_{<k}} \operatorname{supp}(\mu_n(\cdot|a_{<k}))$.

Next, according to agent k's optimal decision rule, we know that

$$\frac{\mathbb{P}_\sigma(\theta = 1|h_{<k}, a_k = 1)}{\mathbb{P}_\sigma(\theta = 0|h_{<k}, a_k = 1)} = \frac{\mathbb{P}_\sigma(\theta = 1|h_{<k})(1 - F_1(1 - \mathbb{P}_\sigma(\theta = 1|m_k)))}{(1 - \mathbb{P}_\sigma(\theta = 1|h_{<k}))(1 - F_0(1 - \mathbb{P}_\sigma(\theta = 1|m_k)))} \tag{60}$$

We know that $\mathbb{P}_\sigma(\theta = 1|h_{<k}) > b_k$, $\mathbb{P}_\sigma(\theta = 0|h_{<k}) = 1 - \mathbb{P}_\sigma(\theta = 1|h_{<k}) > 1 - t_k$. Since F_1 and F_0 are mutually absolutely continuous, it must be that for any a_k that is on the equilibrium path (i.e. herding have not occurred yet, or the action is one that agent k herd on), we have

$$\frac{\mathbb{P}_\sigma(\theta = 1|h_{<k}, a_k = 1)}{\mathbb{P}_\sigma(\theta = 0|h_{<k}, a_k = 1)} \in (0, \infty) \tag{61}$$

Thus, $\mathbb{P}_\sigma(\theta = 1|h_{<k}, a_k = 1) \in (0,1)$. Taking expectation over the realizations of action-signals $(m_1, ..., m_k)$, we have

$$\mathbb{P}_\sigma(\theta = 1|a_{<k+1}) \in (0,1) \tag{62}$$

Finally, we also know that

$$\mathbb{P}_\sigma(\theta = 1|h_{<k+1}) = \mathbb{P}_\sigma(\theta = 1|h_{<k}, a_k, m_{k+1}) = \mathbb{P}_\sigma(\theta = 1|h_{<k}, a_k) \in (0,1). \tag{63}$$

$\square$

Proof of Proposition 6. Since the action-signals are not news-permitting, there must exist some $\epsilon > 0, k \in \mathbb{N}$, and a sequence of agents $\{n_i\}_{i=1}^{\infty}$ such that for all $i \in \mathbb{N}$, there exists some $m^i \in \mathcal{M}$ and $a^i_{<k} \in \mathcal{A}_{<k}$ such that $\mathbb{P}_\sigma(a^i_{<k}) > 0$,[20] and

$$\mu_{n_i}(m^i | a_{<n_i}) > \epsilon,\ \forall a_{n_i} \in \mathcal{A}_{n_i} \text{ with } a_{<k} = a^i_{<k}.$$

According to Lemma 6, for any equilibrium, there exists some $\epsilon_1 > 0$ such that

$$\mathbb{P}_\sigma(\theta = 1 | a^i_{<k}) \in (\epsilon_1, 1 - \epsilon_1)$$

which implies that $\mathbb{P}_\sigma(\theta = 0 | a^i_{<k}) = 1 - \mathbb{P}_\sigma(\theta = 0 | a^i_{<k}) \in (\epsilon_1, 1 - \epsilon_1)$. Since $\mathcal{A}_{<k}$ has finite elements, there exists $\epsilon_2 > 0$ such that

$$\mathbb{P}_\sigma(a^i_{<k}) > \epsilon_2,\ \forall i \in \mathbb{N}^+. \tag{64}$$

This implies that

$$\begin{aligned}
\mathbb{P}_\sigma(\theta = 1, m_{n_i} = m^i) &\geq \mathbb{P}_\sigma(\theta = 1, m_{n_i} = m^i, a_{<k} = a^i_{<k}) &(65)\\
&= \int_{a_{<n_i}: a_{<k} = a^i_{<k}} \mu_{n_i}(m | a_{<n_i}) \mathbb{P}_\sigma(\theta = 1 | a_{<n_i}) d\mathbb{P}_\sigma(a_{<n_i}) &(66)\\
&> \epsilon \cdot \int_{a_{<n_i}: a_{<k} = a^i_{<n_i}} \mathbb{P}_\sigma(\theta = 1 | a_{<n_i}) d\mathbb{P}_\sigma(a_{<n_i}) &(67)\\
&\geq \epsilon \cdot \epsilon_1 \cdot \epsilon_2 &(68)
\end{aligned}$$

Similarly, we have

$$\mathbb{P}_\sigma(\theta = 0, m_{n_i} = m^i) > \epsilon \cdot \epsilon_1 \cdot \epsilon_2 \tag{69}$$

We know that

$$\begin{aligned}
\mathbb{P}_\sigma(m_{n_i} = m^i, \theta = 0, a_{n_i} = 1) &> \epsilon \cdot \epsilon_1 \cdot \epsilon_2 \cdot (1 - G_0(1 - \mathbb{P}_\sigma(\theta = 1 | m_{n_i} = m^i))) &(70)\\
\mathbb{P}_\sigma(m_{n_i} = m^i, \theta = 1, a_{n_i} = 0) &> \epsilon \cdot \epsilon_1 \cdot \epsilon_2 \cdot G_1(1 - \mathbb{P}_\sigma(\theta = 1 | m_{n_i} = m^i)) &(71)
\end{aligned}$$

[20] For the action-signals to be not news-permitting (uniformly or for a strategy profile σ), we either assumed that the state-signals are unboundedly informative, which implies $\mathbb{P}_\sigma(a^i_{<k}) > 0$ when k is a fixed finite number, or we directly assumed $\mathbb{P}_\sigma(a^i_{<k}) > 0$ under state-signals that are boundedly informative.

Therefore, the probability that agent n_i takes the wrong action is such that

$$\mathbb{P}_\sigma(a_{n_i} \neq \theta) > \epsilon \cdot \epsilon_1 \cdot \epsilon_2 \cdot \max[1 - G_0(1 - \mathbb{P}_\sigma(\theta = 1 | m_{n_i} = m^i)), G_1(1 - \mathbb{P}_\sigma(\theta = 1 | m_{n_i} = m^i))] \tag{72}$$

I claim that the R.H.S. of inequality (72) is strictly bounded above zero. To see this, consider the function $f(x) = \max[1 - G_0(x), G_1(x)]$ for $x \in [0,1]$. Suppose that there is a sequence of real numbers $\{x^i\}_{i=1}^{\infty}$ in the interval $[0,1]$ such that $\lim_{i\to\infty} f(x^i) = 0$. Then, we must have

$$\lim_{i\to\infty} 1 - G_0(x^i) = 0 \text{ and } \lim_{i\to\infty} G_1(x^i) = 0 \tag{73}$$

Since G_0 and G_1 are mutually absolutely continuous, when $\lim_{i\to\infty} 1 - G_0(x^i) = 0$, we must have $\lim_{i\to\infty} 1 - G_1(x^i) = 0$ which contradicts with $\lim_{i\to\infty} G_1(x^i) = 0$. Thus, there exists $\epsilon_3 > 0$ such that $f(x) > \epsilon_3$ for $x \in [0,1]$. Hence,

$$\mathbb{P}_\sigma(a_{n_i} \neq \theta) > \epsilon \cdot \epsilon_1 \cdot \epsilon_2 \cdot \epsilon_3 > 0 \tag{74}$$

So agent n_i takes the correct action with probability strictly less than $1 - \epsilon \cdot \epsilon_1 \cdot \epsilon_2 \cdot \epsilon_3$. With infinite number of agents $\{n_i\}_{i=1}^{\infty}$ taking the correct action with probability strictly bounded below 1, information aggregation is guaranteed to fail.

□

C12 Proof of Proposition 5

Proof. As shown in the proof of Theorem 6, for any finite k and any history $a_{<k}$, the posterior belief $\mathbb{P}_\sigma(\theta = 1 | a_{<k})$ and $\mathbb{P}_\sigma(\theta = 1 | m_k)$ must be interior for all $m_k \in \mathcal{M}$. Then according to Proposition 7, the equilibrium strategy is such that both actions are chosen with positive probabilities (because $p_n > 1 - \mathbb{P}_\sigma(\theta = 1 | m_k)$ and $p_n < 1 - \mathbb{P}_\sigma(\theta = 1 | m_k)$ both occur with strictly positive probabilities. Therefore, by an inductive argument, all finite histories are reached with positive probabilities.

□

C13 Proof of Corollary 1

Follows from Proposition 5 and Theorem 6.

C Proofs and supplementary materials for Section 6

C1 Proof of Proposition 6

Lemma 7. *For any private belief distributions (G_0, G_1), the following relations hold.*

- *The ratio $G_0(x)/G_1(x)$ is non-increasing in x.*
- *The ratio $(1-G_0(x))/(1-G_1(x))$ is non-decreasing in x.*
- *$G_0(x) > G_1(x)$ for all $x \in (0,1)$.*

Proof. By Lemma 2, we have for any $x \in (0,1)$,

$$G_0(x) = \int_{s=0}^{x} dG_0(s) = \int_{s=0}^{x} \left(\frac{1-s}{s}\right) dG_1(s) \geq \int_{s=0}^{x} \left(\frac{1-x}{x}\right) dG_1(s) = \left(\frac{1-x}{x}\right) G_1(x) \tag{75}$$

We also have

$$\begin{aligned} d\left(\frac{G_0(x)}{G_1(x)}\right) &= \frac{dG_0(x)G_1(x) - G_0(x)dG_1(x)}{(G_1(x))^2} \\ &= \frac{dG_1(x)}{(G_1(x))^2}\left[\left(\frac{1-x}{x}\right)G_1(x) - G_0(x)\right]. \end{aligned}$$

We know that $dG_1(x) \geq 0$, $G_1(x) > 0$, and the term in brackets above is non-positive by Equation (75). Therefore,

$$d\left(\frac{G_0(x)}{G_1(x)}\right) \leq 0 \tag{76}$$

which implies that the ratio $G_0(x)/G_1(x)$ is non-increasing. A similar argument proves the second part of the lemma.

Finally, using the fact that $G_0(1) = G_1(1) = 1$, we must have for any $x \in (0,1)$, we have

- $G_0(x) > G_1(x)$
- $1 - G_0(x) < 1 - G_1(x)$

□

Proof of Proposition 6. Consider an agent n in week c, and take $J_n = \{(c-1)t+1, ..., ct\}$ to be the agents in week $c-1$. For any action-signal realization $m_n = x_{<c} = (x_0, ..., x_{c-1}) \in \{0,1\}^c$,

if $x_{c-1} = 1$, then I know for sure that

$$\sum_{j=(c-1)t+1}^{ct} a_j > \phi_c(x_{<c-1}) \tag{77}$$

We want to show that for all $k > \phi_c(x_{<(c-1)})$,

$$\mathbb{P}_\sigma\left(\theta = 1 \left| \sum_{j=(c-1)t+1}^{ct} a_j = k, \mathbf{1}\left\{\sum_{j=rt+1}^{(r+1)t} a_j > \phi_k(x_{<r})\right\} = x_r \text{ for all } r = 0, 1, ..., c-2\right.\right) \geq 1/2.$$

To see this, first note that in the event

$$\mathbf{1}\left\{\sum_{j=rt+1}^{(r+1)t} a_j > \phi_k(x_{<r})\right\} = x_r \text{ for all } r = 0, 1, ..., c-2$$

All agents in week $(c-1)$ observes exactly the same action-signal realization that is equal to $x_{<(c-1)}$, and have the same posterior belief about the state

$$\mu^* = \mathbb{P}_\sigma\left(\theta = 1 \left| \mathbf{1}\left\{\sum_{j=rt+1}^{(r+1)t} a_j > \phi_k(x_{<r})\right\} = x_r \text{ for all } r = 0, 1, ..., c-2\right.\right).$$

In this case, given state θ, each agent will take action $a_n = 1$ if the private belief (induced by s_n) is such that

$$p_n > 1 - G_\theta(1 - \mu^*),$$

and he or she will take action $a_n = 0$ if the private belief is such that

$$p_n < 1 - G_\theta(1 - \mu^*).$$

Therefore, denote

$$\hat{\mu} = \mathbb{P}_\sigma\left(\theta = 1 \left| \sum_{j=(c-1)t+1}^{ct} a_j = k, \mathbf{1}\left\{\sum_{j=rt+1}^{(r+1)t} a_j > \phi_k(x_{<r})\right\} = x_r \text{ for all } r = 0, 1, ..., c-2\right.\right)$$

Then according to Bayes' rule and the conditional independence of the state-signals (or the private belief), we must have

$$\frac{\hat{\mu}}{1-\hat{\mu}} = \frac{\mu^*}{1-\mu^*}\left(\frac{1-G_1(1-\mu^*)}{1-G_0(1-\mu^*)}\right)^k \left(\frac{G_1(1-\mu^*)}{G_0(1-\mu^*)}\right)^{t-k} \tag{78}$$

Note that, according to Lemma 7, the R.H.S. is increasing in k. By definition of $\phi_c(x_{(c-1)})$, the R.H.S. is larger than 1 if $k = \lceil \phi_c(x_{<(c-1)}) \rceil$, and thus is also larger than 1 for all $k \geq \phi_c(x_{<(c-1)})$. These are precisely the values of $\sum_{J_n} a_j$ that can occur, given the action-signal $m_n = x_{<c}$. Since the posterior belief minus 1/2 all have the same (positive) sign, it satisfies weak separation. Similar argument works for the case when $x_{c-1} = 0$. □

C2 Proof of Proposition 7

Proof. When the state-signals are unboundedly informative, we know that all strongly separating action-signals give rise to information aggregation. It is obvious from the definition of strong separation that if an additional signal is disclosed the resulting information environment will still be strongly separating. This is because the additional signal can only make the environment more informative about past action-histories. Thus, information aggregation is robust to third party manipulation.

□

C3 Proof of Proposition 8

Proof. Since the state-signals are boundedly informative, a manipulation strategy $\mu^\dagger$ which fully disclose the action-histories will block information aggregation by inducing herding behavior. □

C4 Proof of Theorem 9

Proof. Notice that in the definition of strong separation. Only the aggregate numbers of actions in J_n matters, not the individual actions. Therefore, when identity information in m_n is lost, it will make the action-signals less informative, but will not break the strong separation condition. Thus, information aggregation should still occur. □

D Information Diffusion

In this section, we analyze the long-run efficiency of social learning using the alternative metric of information diffusion.

Since $\tilde{s}_n$ are identically independently distributed across n, the distribution of the private beliefs can be summarized by the distribution of $\tilde{p}_1$. Let the convex hall of the support of the private beliefs be the region $[\underline{\beta}, \bar{\beta}]$, where

$$\underline{\beta} = \inf\{x \in [0,1] \,|\, \mathbb{P}(\tilde{p}_1 \leq x) > 0\} \text{ and } \bar{\beta} = \sup\{x \in [0,1] \,|\, \mathbb{P}(\tilde{p}_1 \leq x) < 1\}.$$

There is a unique binary signal $s^* \in \{0,1\}$, a random variable such that

$$\mathbb{P}(\tilde{\theta} = 1 | s^* = 0) = \underline{\beta} \text{ and } \mathbb{P}(\tilde{\theta} = 1 | s^* = 1) = \bar{\beta}.$$

We shall call s^* the *expert signal.*

Definition 14. *We say that **information diffusion** occurs if we have*

$$\lim_{n\to\infty} \mathbb{E}_\sigma[u(\tilde{a}_n, \tilde{\theta})] \geq \mathbb{E}[u(\tilde{s}^*, \tilde{\theta})] \equiv u^*.$$

Intuitively, we have diffusion if agents perform as though they were guaranteed to receive one of the strongest possible signals. Note that aggregation is generally a stronger criterion than diffusion; if state-signals are unboundedly informative (i.e., $1 - \underline{\beta} = \bar{\beta} = 1$), the two metrics coincide. The metric of information diffusion provides an alternative perspective that emphasizes the role of the social environment π in social learning. Often, when aggregation turns on whether the state-signals are unboundedly or boundedly informative, information diffuses according to Definition 14.

With the definition of information diffusion at hand, I will now discuss how the main characterization results in this paper should be interpreted if we use information diffusion as the metric for long-run efficiency. First, if the action-signals are not news-permitting, then information diffusion always fails. So the necessary condition for long-run efficiency remains the same. Second, if the action-signals are strongly or weakly separating, then information diffusion occurs regardless of whether the state-signals are boundedly or unboundedly informative. Thus, the sufficient condition for information diffusion is only dependent on the information structure of the action-signals, not the state-signals. Finally, in terms of the robustness properties in this paper, information diffusion under strongly separating action-signals are robust to manipulation by third-parties who can provide additional information, and it is robust to distortion caused by agents' limited cognitive capacity to remember sources of prior actions in the separating sequence.

In conclusion, if we adopt information diffusion as the metric for long-run learning, the characterization results are solely on the action-signals. On the downside, the economic meaning of information diffusion is less-appealing because agents may still choose incorrect actions when information diffusion occurs, and the metric of information diffusion itself depends on the information structure of the state-signal.

E Welfare Non-Monotonicity

By focusing on information aggregation, I have offered a simple method to determine whether an information environment is consistent with the *long-run* efficiency of social learning. A normative interpretation of the result is that if individual privacy is preserved in a way that ensures that a minimum amount of information is revealed in action-signals that are separating, then long-run efficiency will not be impacted. When short- and medium-run efficiency are taken into account, one might expect an even simpler relationship between efficiency and privacy: as more information about past actions is disclosed, agents become better informed and so make wiser decisions. Thus, efficiency can only be achieved at the cost of privacy, and it is maximized under full transparency. Indeed, a natural analogy is with a *single* agent who receives an informative signal about other agents' past actions: the agent's expected utility increases if his signal becomes more informative about others' actions, and it is highest if he perfectly learns other individual's actions.

Proposition 10 offers a sharp contrast to the single-agent benchmark: It shows that there is a non-monotonic relationship between the informativeness of the action-signals and the welfare of the agents. In particular, increasing the informativeness of the action-signals may increase the welfare of earlier agents *at the cost of later agents*. This result suggests that a fully transparent environment, which has been the main focus of the prior literature, may *not* be socially optimal. From a policy perspective, the result implies that the commonly perceived "either/or" dichotomy between efficiency and privacy is a "false binary" in social learning.

Formally, consider a social planner with a discount factor $\delta \in (0,1)$. Suppose the social planner can flexibly design the action-signals by choosing any action-signal technology $\mu = (\{\mu_n\}_{n=1}^{\infty}, \mathcal{M})$ where $\mu_n(\cdot|h_{<n}) \in \triangle(\mathcal{M})$ specifies the distribution of agent n's action signal

given the *social history* $h_{<n} = (m_1, a_1, ..., m_{n-1}, a_{n-1})$ that specifies the first $n-1$ actions *and realizations of the first* $n-1$ *action-signals*.[21] We use $\mathcal{H}_{<n}$ to denote the set of date-n social histories. The social planner chooses μ to maximize the discounted social welfare:

$$\sum_{n=1}^{\infty} \delta^n \mathbb{E}_{\sigma,\mu}[u(a_n, \theta)] \tag{79}$$

where σ is an equilibrium strategy profile under μ. We refer to the optimal action-signal technology for the social planer as the *socially optimal action-signal technology*. It may seem that the socially optimal action-signal technology should be fully informative. However, as the following proposition shows, this intuition is incorrect.

Proposition 10. *There exists* $[a, b] \subset (0, 1)$ *such that whenever* μ_n^* *garbles* μ_n *for signal realizations that has posterior outside* $[a, b]$ *under* μ_n, $h_{<n+1}$ *always become Blackwell more informative under* μ_n^* *than under* μ_n.

Therefore, the social environment can be optimally opaque when short to medium-run welfare is taken into account. In practice, we may want to design information environments that protect individual privacy, which requires making the action-signals less informative about past actions. Theorem 10 implies that such practices may not be socially inefficient: unlike single-agent settings (the special case of $\delta = 0$ in the social planer's problem), efficiency and privacy do not have a simple "either/or" trade-off.

E1 Proof of Proposition 10

Proof of Theorem 10. Throughout this proof, we assume that the information structures $\{\mu_i\}_{i=1}^{n-1}$ of the first $n-1$ action-signals and the equilibrium strategy profiles $\{\sigma\}_{i=1}^{n-1}$ of the first $n-1$ agents are fixed. Therefore, for simplicity of exposition, we will not explicitly specify the dependence of probabilities measures on $\{\mu_i, \sigma_i\}_{i=1}^{n-1}$.

First, note that when we fix $\{\mu_i, \sigma_i\}_{i=1}^{n-1}$, the distribution of the histories $h_{<n}$ is fixed. We compare two information structures of agent n's action-signal: μ_n and μ_n^*, with measurable signal spaces $\mathcal{M}$ and $\mathcal{M}^*$ respectively. Suppose μ_n^* is a garbling of μ_n, i.e., there exists a

[21] We allow the information structure to depend on past action-signals so that the social planer has full flexibility to design the environment. This can also be done in the main characterization theorems.

transition probability $\gamma(\cdot|m) \in \triangle(\mathcal{M}^*)$ for each $m \in \mathcal{M}$ such that:

$$\mu_n^*(M^*|h_{<n}) = \int_{m\in\mathcal{M}} \gamma(M^*|m)\mu_n(dm|h_{<n}) \text{ for every measurable set } M^* \subseteq \mathcal{M}^* \tag{80}$$

We want to show that there exist equilibrium strategies σ_n and σ_n^* that are optimal for agent n under μ_n and μ_n^* respectively, such that the history $h_{<n+1}$ is Blackwell more informative about the state θ under (μ_n^*, σ_n^*) than it is under (μ_n, σ_n). According to the Blackwell Theorem, this is equivalent to the statement that $\mathbb{P}_{\mu_n^*,\sigma_n^*}(\theta = 1|h_{<n+1})$ is a mean-preserving spread of $\mathbb{P}_{\mu_n,\sigma_n}(\theta = 1|h_{<n+1})$. We will prove this equivalent statement.

Take any measurable set $M^* \subseteq \mathcal{M}^*$, according to Bayes' rule, we must have

$$\begin{aligned}
\frac{\mathbb{P}_{\mu_n^*}(\theta = 1|m_n \in M^*)}{\mathbb{P}_{\mu_n^*}(\theta = 0|m_n \in M^*)} &= \frac{\sum_{h_{<n}} \pi_1 \mathbb{P}(h_{<n}|\theta = 1)\mu_n^*(M^*|h_{<n})}{\sum_{h_{<n}} \pi_0 \mathbb{P}(h_{<n}|\theta = 0)\mu_n^*(M^*|h_{<n})} \\
&= \frac{\sum_{h_{<n}} \pi_1 \mathbb{P}(h_{<n}|\theta = 1) \int_m \gamma(M^*|m) d\mu_n(m|h_{<n})}{\sum_{h_{<n}} \pi_0 \mathbb{P}(h_{<n}|\theta = 0) \int_m \gamma(M^*|m) d\mu_n(m|h_{<n})} \\
&= \frac{\sum_{h_{<n}} \int_m \gamma(M^*|m) d\mathbb{P}_{\mu_n}(\theta = 1, h_{<n}, m_n = m)}{\sum_{h_{<n}} \int_m \gamma(M^*|m) d\mathbb{P}_{\mu_n}(\theta = 0, h_{<n}, m_n = m)} \\
&= \frac{\int_m \gamma(M^*|m) \sum_{h_{<n}} d\mathbb{P}_{\mu_n}(\theta = 1, h_{<n}, m_n = m)}{\int_m \gamma(M^*|m) \sum_{h_{<n}} d\mathbb{P}_{\mu_n}(\theta = 0, h_{<n}, m_n = m)} \\
&= \frac{\int_m \gamma(M^*|m) d\mathbb{P}_{\mu_n}(\theta = 1, m_n = m)}{\int_m \gamma(M^*|m) d\mathbb{P}_{\mu_n}(\theta = 0, m_n = m)} \\
&= \frac{\int_m \gamma(M^*|m) d\mathbb{P}_{\mu_n}(\theta = 1|m_n = m)}{\int_m \gamma(M^*|m) d\mathbb{P}_{\mu_n}(\theta = 0|m_n = m)}
\end{aligned} \tag{81}$$

where the second equality follows from Equation (80).

Equation (81) implies that

$$\mathbb{P}_{\mu_n^*}(\theta = 1|m_n \in M^*) = \frac{\int_m \gamma(M^*|m)\mathbb{P}_{\mu_n}(\theta = 1|m_n = m) d\mathbb{P}_{\mu_n}(m_n = m)}{\int_m \gamma(M^*|m) d\mathbb{P}_{\mu_n}(m_n = m)} \tag{82}$$

Therefore, the posterior belief $\mathbb{P}_{\mu_n^*}(\theta = 1|m_n \in M^*)$ is a weighted average of the posterior beliefs $\mathbb{P}_{\mu_n}(\theta = 1|m_n = m)$ with weights $\gamma(M^*|m)$. Since this is true for any measurable set M^* and γ is a well-defined transition probability, there must exist a version of the condition probability $\mathbb{P}_{\mu_n^*}(\theta = 1|m_n = m^*)$ and a function $\gamma' : \mathcal{M} \times \mathcal{M}^* \to \mathbb{R}_+$ such that for all $m^* \in \mathcal{M}^*$,

$$\mathbb{P}_{\mu_n^*}(\theta = 1|m_n = m^*) = \frac{\int_m \gamma(m, m^*)\mathbb{P}_{\mu_n}(\theta = 1|m_n = m) d\mathbb{P}_{\mu_n}(m_n = m)}{\int_m \gamma(m, m^*) d\mathbb{P}_{\mu_n}(m_n = m)} \tag{83}$$

This implies that for all $m^* \in \mathcal{M}^*$,

$$\min_{m:\gamma(m,m^*)>0} \mathbb{P}_{\mu_n}(\theta = 1|m_n = m) \leq \mathbb{P}_{\mu_n^*}(\theta = 1|m_n = m^*) \leq \max_{m:\gamma(m,m^*)>0} \mathbb{P}_{\mu_n}(\theta = 1|m_n = m). \quad (84)$$

If μ_n^* is not identical to μ_n, there must exist a set $\hat{M}^*$ in $\mathcal{M}^*$ with positive measure such that for all $m^* \in \hat{M}^*$, the inequalities are strict:

$$\min_{m:\gamma(m,m^*)>0} \mathbb{P}_{\mu_n}(\theta = 1|m_n = m) < \mathbb{P}_{\mu_n^*}(\theta = 1|m_n = m^*) < \max_{m:\gamma(m,m^*)>0} \mathbb{P}_{\mu_n}(\theta = 1|m_n = m) \quad (85)$$

Let $\bar{M}(m^*)$ and $\underline{M}(m^*)$ denote the sets of signal realizations in $\mathcal{M}$ such that the induced posterior beliefs under μ_n is strictly larger or smaller than the posterior belief induced by $m_n = m^*$ under μ_n^*, respectively: i.e.,

$$\begin{aligned}
\bar{M}(m^*) &= \{m \in \mathcal{M} : \mathbb{P}_{\mu_n}(\theta = 1|m_n = m) > \mathbb{P}_{\mu_n^*}(\theta = 1|m_n = m^*)\} \\
\underline{M}(m^*) &= \{m \in \mathcal{M} : \mathbb{P}_{\mu_n}(\theta = 1|m_n = m) < \mathbb{P}_{\mu_n^*}(\theta = 1|m_n = m^*)\}
\end{aligned}$$

Fix any $m^* \in \hat{M}^*$ and any $h_{<n} \in \mathcal{H}_{<n}$, we adopt the following notations:

$$\begin{aligned}
\nu_n^* &= \mathbb{P}_{\mu_n^*}(\theta = 1|m_n = m^*) \\
\xi_n &= \mathbb{P}(\theta = 1|h_{<n})
\end{aligned}$$

Then we can calculate the posterior beliefs of the state induced by the history $h_{<n+1} = (h_{<n}, m_n = m^*, a_n)$ under μ_n^* and agent n's optimal strategy σ_n^*, for $a_n = 0$ and 1 respectively:

$$\begin{aligned}
P_{\mu_n^*,\sigma_n^*}(\theta = 1|h_{<n}, m_n = m^*, a_n = 1) &= \frac{\xi_n(1 - G_1(1 - \nu_n^*))}{\xi_n(1 - G_1(1 - \nu_n^*)) + (1 - \xi_n)(1 - G_0(1 - \nu_n^*))} \\
P_{\mu_n^*,\sigma_n^*}(\theta = 1|h_{<n}, m_n = m^*, a_n = 0) &= \frac{\xi_n G_1(1 - \nu_n^*)}{\xi_n G_1(1 - \nu_n^*) + (1 - \xi_n) G_0(1 - \nu_n^*)}
\end{aligned}$$

In comparison, we can also calculate the posterior beliefs about the state induced by the history $h_{<n+1} = (h_{<n}, m_n = m, a_n)$ under μ_n and agent n's corresponding optimal strategy σ_n, for $a_n = 0$ and 1 respectively:

$$\begin{aligned}
P_{\mu_n,\sigma_n}(\theta = 1|h_{<n}, m_n = m, a_n = 1) &= \frac{\xi_n(1 - G_1(1 - \nu_n(m)))}{\xi_n(1 - G_1(1 - \nu_n(m))) + (1 - \xi_n)(1 - G_0(1 - \nu_n(m)))} \\
P_{\mu_n,\sigma_n}(\theta = 1|h_{<n}, m_n = m, a_n = 0) &= \frac{\xi_n G_1(1 - \nu_n(m))}{\xi_n G_1(1 - \nu_n(m)) + (1 - \xi_n) G_0(1 - \nu_n(m))}
\end{aligned}$$

where $\nu_n(m) := \mathbb{P}_{\mu_n}(\theta = 1|m_n = m)$ denote the posterior belief about the state induced by the signal $m_n = m$ under μ_n.

Since $m^* \in \hat{M}^*$, the sets $\bar{M}(m^*)$ and $\underline{M}(m^*)$ must have positive measures under μ_n. For any $m \in \bar{M}(m^*)$, we have $\nu_n(m) > \nu_n^*$. Thus, by Lemma 7, we must have

$$P_{\mu_n,\sigma_n}(\theta = 1|h_{<n}, m_n = m, a_n = 1) < P_{\mu_n^*,\sigma_n^*}(\theta = 1|h_{<n}, m_n = m^*, a_n = 1)$$
$$P_{\mu_n,\sigma_n}(\theta = 1|h_{<n}, m_n = m, a_n = 0) < P_{\mu_n^*,\sigma_n^*}(\theta = 1|h_{<n}, m_n = m^*, a_n = 0)$$

Similarly, for any $m \in \underline{M}(m^*)$, we must have

$$P_{\mu_n,\sigma_n}(\theta = 1|h_{<n}, m_n = m, a_n = 1) > P_{\mu_n^*,\sigma_n^*}(\theta = 1|h_{<n}, m_n = m^*, a_n = 1)$$
$$P_{\mu_n,\sigma_n}(\theta = 1|h_{<n}, m_n = m, a_n = 0) > P_{\mu_n^*,\sigma_n^*}(\theta = 1|h_{<n}, m_n = m^*, a_n = 0)$$

Since the histories that lead to $m_n = m \in \bar{M}(m^*)$ have higher average posteriors than $m_n = m^*$, and those that lead to $m_n = m \in \underline{M}(m^*)$ have lower average posteriors than $m_n = m^*$ and that the probability of the two intermediate realizations will go to zero as the two posteriors become extreme, the above inequalities must imply that the distribution of the posterior beliefs under μ_n^* is a mean-preserving spread of the distribution of posteriors under μ_n when the realizations under μ_n that are garbled involve extreme histories.

□

F Equilibrium Information Manipulation

Propositions 7 and 8 studies whether there exists a manipulation strategy that blocks information aggregation; but will such manipulation strategies be adopted by the third party in some equilibrium? How does the preference of the third-party affects information manipulation? I focus on a (single) third party who cares about long-run average actions, i.e., the third party's utility function $V(a, \theta)$ is

$$V(a, \theta) = \lim_{T\to\infty} \frac{1}{T} \sum_{n=1}^{T} v(a_n, \theta)$$

for some $v : \{0, 1\}^2 \to \mathbb{R}$.

Definition 15. *The third party's preference and the agents' preference are* ***aligned*** *if* $v(\theta, \theta) \geq$

*$v(1-\theta,\theta)$ for $\theta \in \{0,1\}$, and they are **misaligned** near state θ if $v(\theta,\theta) < v(1-\theta,\theta)$.*

I assume the third party has full commitment power and can flexibly implement any manipulation technology $\mu^\dagger$. In the appendix, I characterize the third party's problem given any μ. The problem can be formulated as a variant of a Bayesian persuasion problem with a lower bound on the receiver's information. Then I study the optimal information design for the third party.

As we already know, when the state-signals are unboundedly informative, information aggregation in AS environments cannot be blocked by the third party. The third party's payoff will be $\pi_1 v(1,1) + \pi_0 v(0,0)$, regardless of his strategy $\mu^\dagger$.

Now suppose the action-signal technology is AS and information aggregation occurs, but the state-signals are boundedly informative near a state θ (i.e., $\exists \epsilon > 0$ s.t. $\mathbb{P}(p_n \in [\theta-\epsilon, \theta+\epsilon]) = 0$). We know from Proposition 8 that information aggregation *can* be blocked. Whether or not it will be optimally blocked by the third party depends on his preference.

Proposition 11. *Suppose the state-signals are boundedly informative only near state θ and information aggregation occurs under μ. Then information aggregation fails under any optimal $\mu^\dagger$ for the third party if and only if the third party's preference is misaligned with that of the agents' at θ.*

Proof. If the third-party's preference is aligned with the agents', inducing information aggregation will maximize the third-parties expected utility because the third party also prefer the agents to take the correct action given each state.

On the other hand, if the third party's preference is misaligned with the agents' at θ, then the third party can provide fully informative $m_n^\dagger$ to agents that are "ignorant" (there has to be an infinite number of such ignorant agents for information aggregation to occur under μ), this will induce herding on action θ which is preferred by the third party. However, if the third party's preference is misaligned only at $1-\theta$, since the state-signals are not boundedly informative near $1-\theta$, blocking information aggregation (by inducing herding) will not benefit the third party, thus, information aggregation occurs. □