Observational Learning in an Uncertain World

Daron Acemoglu, Munther Dahleh, Asuman Ozdaglar, and Alireza Tahbaz-Salehi

Abstract— **We study a model of observational learning in social networks in the presence of uncertainty about agents' type distributions. Each individual receives a private noisy signal about a payoff-relevant state of the world, and can observe the actions of other agents who have made a decision before her. We assume that agents do not observe the signals and types of others in the society, and are also uncertain about the type distributions. We show that information is correctly aggregated when preferences of different types are closely aligned. On the other hand, if there is sufficient heterogeneity in preferences, uncertainty about type distributions leads to potential identification problems, preventing asymptotic learning. We also show that even though learning is guaranteed to be incomplete ex ante, there are sample paths over which agents become certain about the underlying state of the world.**

I. Introduction

Since Savage's seminal work [1] in 1954, it has been well understood that under some regularity conditions, a Bayesian agent with access to a large collection of data can eventually learn an unknown parameter of interest. However, in many real world scenarios, the relevant information for learning the unknown state is not concentrated at the disposal of any single individual, and instead, is spread among a large collection of agents. For example, the information about the quality of a product is shared among all individuals who have purchased that product. Such observations have motivated a fairly large literature investigating the problems of information aggregation and learning in social and economic networks.

One of the most prominent frameworks for studying the problem of information aggregation in social networks has been the observational learning framework. In this framework, a collection of individuals with limited private information about a payoff-relevant state of the world make decisions sequentially. In addition to the private information available to the decision makers, they can also observe the *actions* of a subset of other agents who have already made their decisions. The central question in such models is whether eventually the true underlying state of the world is revealed as more agents make decisions.

The observational learning framework was independently introduced by Banerjee [2] and Bikhchandani, Hirshleifer, and Welch [3]. They showed that a bound on the informativeness of agents' private signals leads to an incorrect herd with positive probability, where individuals cannot properly infer the underlying state of the world from the actions of their predecessors. These works were followed by Smith and Sørensen [4], who showed that in the presence of unbounded private signals and with the entire history of actions observable, asymptotic learning is achieved in all equilibria. In a related work, Acemoglu, Dahleh, Lobel, and Ozdaglar [5] extend this result to general network topologies, where instead of observing the entire history, individuals can only observe the actions of a subset of agents. Observational learning models allowing for private, heterogeneous preferences have been studied by Smith and Sørensen [4], and more recently, Acemoglu, Dahleh, Lobel, and Ozdaglar [6].

Even though existing models, like [4] and [6], allow for private information on payoffs, they all include common knowledge about the distribution of preferences among their standing assumptions. However, it is natural to assume that there exists some uncertainty about the distribution of preferences of other individuals in the society. For example, consumers of a certain product might be uncertain about the exact distribution of other consumers' tastes.

In this paper, we study the evolution of beliefs and actions in the observational learning framework, when individuals are uncertain about the incentives of other agents in the society. To investigate this problem, we consider a model in which the dependence of the payoffs on the unknown state of the world varies across agents; that is, we consider a model where individuals are of different *types*. Moreover, we assume that not only each agent is unaware of her predecessors' types, but also she is uncertain about the probability distribution from which they were sampled. Such uncertainties require the agents to hold and update beliefs about the distribution of types as observations accumulate. It is the absence of common knowledge about the distribution of preferences that distinguishes our model of observational learning in an "uncertain world" from those already studied in the literature.

We show that if preferences of different types are sufficiently aligned with one another, uncertainties about type distributions have no effect on asymptotic learning, with information eventually aggregated through individuals' actions. On the other hand, we establish that if there is sufficient heterogeneity in preferences, then type distribution uncertainties lead to asymptotic identification problems with positive probability, and as a result, agents cannot infer any further information from their predecessors' actions. Our key observation is that with incentives of different types

This research is supported in parts by the following grants: AFOSR grant FA9550-09-1-0420, ARO grant W911NF-09-1-0556, NSF grant SES-0729361.

Daron Acemoglu is is with the Department of Economics, Massachusetts Institute of Technology, Cambridge, MA 02139 (e-mail: daron@mit.edu).

Munther Dahleh, Asuman Ozdaglar, and Alireza Tahbaz-Salehi are with Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139 (e-mails: dahleh@mit.edu, asuman@mit.edu, alirezat@mit.edu).

sufficiently apart, such identification problems arise even if the amount of uncertainty is arbitrarily small.

The failure of information aggregation in environments with type uncertainties is different from the herding behavior observed by Banerjee [2] and Bikhchandani *et al.* [3]. Whereas in herding outcomes agents discard their private information in favor of the public history, they continue to use their private information in the presence of type distribution uncertainties. However, the pattern of their decisions are such that an identification problem arises in the public history, and therefore, limits its informativeness. Following Smith and Sørensen [4], we refer to such outcomes as *confounded learning*.

Another key implication of our results is that in observational learning models with type distribution uncertainties, it is possible for agents to asymptotically become certain about the true state of the world, despite the fact that learning is guaranteed to be incomplete *ex ante*. In other words, depending on the signals observed by the agents, it is possible for the posterior beliefs to converge to the true state, even though ex ante, individuals do not believe that they will learn the state. Such a phenomenon does not occur in models consisting of a single type of agents (e.g. [3] and [5]), as ex post and ex ante asymptotic learning coincide.

In addition to the works in the social learning literature mentioned above, our paper is also related to a recent work by Acemoglu, Chernozhukov, and Yildiz [7] which studies the effect of uncertainty about the signal distributions on asymptotic learning and agreement of Bayesian agents. It shows that a vanishingly small individual uncertainty about the signal distributions can lead to substantial differences in asymptotic beliefs. Whereas in [7] the identification problems are inherent to the structure of the signals observed by the agents, the identification problem that arises in our model is an equilibrium phenomenon – a consequence of the rational decision making of agents who have already made their decisions.

The rest of the paper is organized as follows. In Section II, we present our model and define asymptotic learning over the social network. Section III contains our main results, where we show that uncertainties about type distributions result in incomplete learning if the incentives of different types are not closely aligned. In Section IV, we show that there are sample paths over which agents eventually become certain about the true state, despite the fact that learning is guaranteed to be incomplete. Section V concludes.

II. The Model

A. Agents and Observations

Consider a group of countably infinite agents indexed by $n \in \mathbb{N}$, making decisions sequentially. The payoff of each agent depends on her type, an unknown underlying state of the world $\theta \in \{0,1\}$, and the agent's decision. More specifically, we assume that agents are of two possible types: the normal type N and a biased type B, with payoffs

$$u_N(x,\theta) = \frac{1}{2} + \mathbb{I}\{x=\theta\}$$
$$u_B(x,\theta) = \begin{cases} \mathbb{I}\{\theta=1\}+1-h & \text{if } x=1 \\ \mathbb{I}\{\theta=0\}+h & \text{if } x=0 \end{cases}$$

where $x \in \{0,1\}$ is the action taken by the agent, $h \in [0,1]$ is a parameter characterizing the bias of agent B towards action $x=0$, and $\mathbb{I}$ is the indicator function.[1] Note that the payoff of the biased type reduces to that of the normal type for $h=1/2$. With some abuse of notation, we sometimes represent the type of agent n by $t_n \in \{h, \frac{1}{2}\}$ instead of $t_n \in \{N,B\}$.

Each agent $n \in \mathbb{N}$ forms beliefs about the unknown payoff-relevant state of the world θ after observing a private noisy signal $s_n \in S$ and the actions of all agents who have moved before her. The key assumption is that even though agent n can observe the actions of individuals in the set $\{1,2,\ldots,n-1\}$, she can observe neither their private signals nor their types.

Conditional on the state θ, agents' private signals are independently and identically distributed according to distribution $\mathbb{F}_\theta$. We assume that the pair $(\mathbb{F}_0,\mathbb{F}_1)$, known as the individuals' *signal structure*, are absolutely continuous with respect to one another; implying that no signal can fully reveal the state. Moreover, we assume that $\mathbb{F}_0$ and $\mathbb{F}_1$ are nonidentical, which guarantees that the private signals are informative about the state. Throughout the paper, the distribution functions $\mathbb{F}_\theta$ are assumed to be continuous. We also assume that *ex ante*, both states are equally likely; that is, $\mathbb{P}(\theta=1)=\mathbb{P}(\theta=0)=1/2$.

Agent n's type, t_n, is drawn randomly and independently from the set of possible types $\{N,B\}$ with probabilities λ and $1-\lambda$, respectively. However, these probabilities are not a priori known by the agents. Instead, they hold a common prior belief over the set of possible type distributions in the society.[2] We assume that agents' prior beliefs about λ has a binary support $\{\lambda_1,\lambda_2\}$, where $\lambda_2 > \lambda_1$. We denote the prior belief of the agents assigned to the event $\{\lambda=\lambda_1\}$ by $\mathbb{H}(\lambda_1)$.

Due to the incomplete information, agents need to form and update beliefs about parameter λ, in addition to the unknown payoff-relevant state θ. In the next sections, we show how this uncertainty about type distributions affects asymptotic aggregation of information over the network.

B. Solution Concept and Learning

Before making decision, agent n knows her type $t_n \in \{N,B\}$, has access to her private signal $s_n \in S$, and can observe actions of agents 1 through $n-1$. Therefore, the information set of agent n is $I_n = \{t_n, s_n, x_k, k<n\}$. We

[1]We have assumed binary state and action spaces to simplify notation and exposition. The main results presented in the paper hold for more general state and action spaces.

[2]Note that parameter h, characterizing the level of preference heterogeneity of the two types, is common knowledge among all individuals. What they are uncertain about is the probability according to which normal or biased agents appear in the sequence.

denote the set of all information sets of agent n with $\mathcal{I}_n$. A pure strategy for agent n is a mapping $\sigma_n : \mathcal{I}_n \longrightarrow \{0,1\}$, which maps her information sets to actions. A sequence of strategies $\sigma = \{\sigma_n\}_{n\in\mathbb{N}}$ defines a strategy profile.

Definition 1: A strategy profile σ^* is a *Perfect Bayesian Equilibrium* if for every $n \in \mathbb{N}$, σ_n^* maximizes the expected payoff of agent n, given the strategies of other agents, $\sigma_{-n}^* = \{\sigma_1^*, \ldots, \sigma_{n-1}^*, \sigma_{n+1}^*, \ldots\}$.

It is easy to verify that a Perfect Bayesian Equilibrium in pure strategies always exists. Given any equilibrium σ^*, the strategy of agent n is given by

$$\sigma_n^*(I_n) = \begin{cases} 1 & \text{if } \mathbb{P}(\theta = 1|I_n) > t_n \\ 0 & \text{if } \mathbb{P}(\theta = 1|I_n) < t_n \end{cases}$$

where $t_n \in \{h, \frac{1}{2}\}$ corresponds to her type.

Our main focus is on whether equilibrium behavior will lead to information aggregation. This is captured by the notion of asymptotic learning, defined below.

Definition 2: Asymptotic learning occurs in a Perfect Bayesian Equilibrium, if the posterior beliefs about the payoff-relevant state θ converge to the truth as more agents make decisions; that is, if

$$\lim_{n\to\infty} \mathbb{P}(\theta = j|I_n) = \mathbb{I}\{\theta = j\} \quad \text{for } j \in \{0,1\}.$$

Note that our notion of asymptotic learning does not require the individuals to learn the entire parameter vector (θ, λ), as only its first component is payoff-relevant. However, they still form and update beliefs about both components. In the next section, we show that individuals' beliefs about type distributions plays a key role in the asymptotic revelation of θ.

We also remark that the concept of asymptotic learning is defined in terms of the *ex post* probability assessments of the individuals. In other words, asymptotic learning occurs on a given sample path if the beliefs assigned to the true state converge to one *on that path*. This notion is distinct from the situation in which *ex ante*, agents believe that the sequence of actions will eventually reveal the payoff-relevant state θ. Such an ex ante assessment requires that asymptotic learning occurs for almost all sequences of information sets, $\{I_n\}_{n\in\mathbb{N}}$. More formally:

Definition 3: Asymptotic learning is *complete* in a Perfect Bayesian Equilibrium, if

$$\mathbb{E}\left[\lim_{n\to\infty} \mathbb{P}(\theta = j|I_n)\right] = \mathbb{I}\{\theta = j\}$$

for $j \in \{0,1\}$.

C. Bounded vs. Unbounded Private Beliefs

Smith and Sørensen [4] show that the key property of the signal structures that leads to learning in a world consisting of a single type is the level of information revealed to agents through their private signals. We define *private beliefs* of agent n as $p(s_n) = \mathbb{P}(\theta = 1|s_n)$ and remark that

$$p(s_n) = \left(1 + \frac{d\mathbb{F}_0}{d\mathbb{F}_1}(s_n)\right)^{-1}$$

where $d\mathbb{F}_0/d\mathbb{F}_1$ is the Radon-Nikodym derivative of $\mathbb{F}_0$ with respect to $\mathbb{F}_1$. We define the support of the private beliefs as $B = \{r \in [0,1] : 0 < \mathbb{P}(p(s_1) \le r) < 1\}$, and say private beliefs are *unbounded* if $\sup B = 1$ and $\inf B = 0$. In other words, under unbounded private beliefs, agents may receive arbitrarily strong signals favoring either state with positive probability, whereas with bounded private beliefs there is a maximum level of information in any signal. Throughout this paper, we assume that the private beliefs have an unbounded support.

Finally, for notational simplicity, we represent the distribution of agents' private beliefs conditional on the state $j \in \{0,1\}$ by $\mathbb{G}_j$:

$$\mathbb{G}_j(r) = \mathbb{P}(p(s_1) \le r|\theta = j).$$

We assume that $\mathbb{G}_j(r)$ is continuously differentiable and strictly increasing in r. In the next sections, we will use the following lemma, the proof of which can be found in [5].

Lemma 1: For any private belief distributions $(\mathbb{G}_0, \mathbb{G}_1)$ and for any $r \in (0,1)$, we have

$$\frac{d\mathbb{G}_0}{d\mathbb{G}_1}(r) = \frac{1-r}{r}.$$

Moreover, when private beliefs are unbounded, the ratio $\mathbb{G}_0(r)/\mathbb{G}_1(r)$ is non-increasing in r, and $\mathbb{G}_0(r) > \mathbb{G}_1(r)$ for all $0 < r < 1$.

Note that by definition, we have $\mathbb{G}_0(0) = \mathbb{G}_1(0) = 0$ and $\mathbb{G}_0(1) = \mathbb{G}_1(1) = 1$.

III. Main Results: Type Uncertainties and Social Learning

In this section, we study information aggregation and asymptotic learning when individuals are of different types. First, we compute the individuals' equilibrium strategies, and then, derive the social belief process dictated by the equilibrium.

A. Equilibrium Strategies

Given that in equilibrium she maximizes her expected payoff, an agent with index n takes action $x_n = 1$, if and only if

$$\mathbb{P}(\theta = 1|s_n, x^{n-1}) > t_n,$$

where $t_n \in \{h, \frac{1}{2}\}$ is her type, and $x^{n-1} = (x_1, \ldots, x_{n-1})$ is the public history of actions. On the other hand, by Bayes' rule

$$\begin{aligned} \frac{\mathbb{P}(\theta = 1|x^{n-1}, s_n)}{\mathbb{P}(\theta = 0|x^{n-1}, s_n)} &= \frac{\mathbb{P}(x^{n-1}|\theta = 1)d\mathbb{P}(s_n|\theta = 1)}{\mathbb{P}(x^{n-1}|\theta = 0)d\mathbb{P}(s_n|\theta = 0)} \\ &= \frac{\mathbb{P}(\theta = 1|s_n)\mathbb{P}(\theta = 1|x^{n-1})}{\mathbb{P}(\theta = 0|s_n)\mathbb{P}(\theta = 0|x^{n-1})}, \end{aligned}$$

where we have used the fact that, ex ante, all agents consider both states equally likely. Thus, the equilibrium decision rule of agent n is given by

$$x_n(s_n, x^{n-1}, t_n) = \begin{cases} 1 & \text{if } p(s_n) > \frac{t_n(1-q_n)}{t_n(1-2q_n)+q_n} \\ 0 & \text{if } p(s_n) < \frac{t_n(1-q_n)}{t_n(1-2q_n)+q_n} \end{cases} \tag{1}$$

where $p(s_n)$ is her private belief, and q_{n-1} is the *social belief*, defined as

$$q_n(x^{n-1}) = \mathbb{P}(\theta = 1 | x_1, \ldots, x_{n-1}).$$

The equilibrium strategies derived in (1) are such that all agents, except for the extreme scenarios of $h \in \{0, 1\}$, take both their private signals and the public history into account when making decisions. Note that this is a consequence of unbounded private beliefs. With bounded private beliefs, once social belief q_n passes a certain threshold, all agents discard their private signals; leading to the herding behavior observed by Banerjee [2] and Bikhchandani *et al.* [3].

B. Social Belief Process

Recall that besides the payoff-relevant state of the world, individuals are also uncertain about the probability distribution of types. Thus, they need to form and update beliefs on the pair (θ, λ). We define the set of social beliefs

$$q_n^{jk}(x^{n-1}) = \mathbb{P}(\theta = j, \lambda = \lambda_k | x^{n-1})$$

for $j \in \{0, 1\}$ and $k \in \{1, 2\}$. Note that the above definition immediately implies that $q_n = q_n^{11} + q_n^{12}$.

A simple application of Bayes' rule implies that the ratio of the social beliefs must satisfy

$$\frac{q_{n+1}^{il}(x^{n-1}, z)}{q_{n+1}^{jk}(x^{n-1}, z)} = \frac{q_n^{il}(x^{n-1})}{q_n^{jk}(x^{n-1})} \frac{\mathbb{P}(x_n = z | x^{n-1}, \lambda = \lambda_l, \theta = i)}{\mathbb{P}(x_n = z | x^{n-1}, \lambda = \lambda_k, \theta = j)}$$

where $z \in \{0, 1\}$. Setting $x_n = 0$ and applying Bayes' rule once again lead to

$$\frac{q_{n+1}^{il}(0)}{q_{n+1}^{jk}(0)} = \frac{q_n^{il}\left[\lambda_l \mathbb{G}_i(1 - q_n) + (1 - \lambda_l)\mathbb{G}_i\left(\frac{h(1-q_n)}{h(1-2q_n)+q_n}\right)\right]}{q_n^{jk}\left[\lambda_k \mathbb{G}_j(1 - q_n) + (1 - \lambda_k)\mathbb{G}_j\left(\frac{h(1-q_n)}{h(1-2q_n)+q_n}\right)\right]} \quad (2)$$

where we have dropped the dependence of the social beliefs on x^{n-1} for notational simplicity. Equation (2) which captures the evolution of the social belief ratios is determined by the equilibrium strategies. Note that according to the equilibrium decision rule, agent n takes action $x_n = 0$, whenever $p(s_n) < t_n(1 - q_n)/[t_n(1 - 2q_n) + q_n]$; an event which happens with probability $\mathbb{G}_j\left(\frac{t_n(1-q_n)}{t_n(1-2q_n)+q_n}\right)$ if the underlying parameters are $(\theta, \lambda) = (j, \lambda_k)$. A similar expression can be derived for the case that agent n takes action $x_n = 1$.

C. Asymptotic Beliefs

Our next observation is that social beliefs q_n^{jk} are bounded martingales with respect to the filtration generated by the public histories, and therefore, converge on almost all sample paths to some limit q_∞^{jk}. The central question we are interested in is whether the limiting social beliefs reveal the payoff-relevant state θ. Before addressing this question, we state and prove a few lemmas.

Lemma 2: Suppose that the underlying state of the world is $(\theta, \lambda) = (j, \lambda_k)$. Then, the belief q_n^{jk} almost surely converges to a limit which is not equal to zero.

Proof: Suppose that $A_{jk} = \{(\theta, \lambda) = (j, \lambda_k)\}$ holds. Then, conditional on this event, a simple application of Bayes' rule implies that

$$\begin{aligned}
\mathbb{E}\left[\left.\frac{1 - q_{n+1}^{jk}}{q_{n+1}^{jk}}\right| x^{n-1}, A_{jk}\right] &= \\
&= \sum_{z \in \{0,1\}} \mathbb{P}(x_n = z | x^{n-1}, A_{jk}) \frac{1 - q_{n+1}^{jk}(x^{n-1}, z)}{q_{n+1}^{jk}(x^{n-1}, z)} \\
&= \frac{1 - q_n^{jk}(x^{n-1})}{q_n^{jk}(x^{n-1})}.
\end{aligned}$$

As a result, conditional on A_{jk}, the likelihood ratio $\ell_n^{jk} \triangleq (1 - q_n^{jk})/q_n^{jk}$ forms a martingale, and by Doob's martingale convergence theorem [8], its limit must satisfy $\mathbb{E}(\ell_\infty^{jk}) = \ell_0^{jk} < \infty$. The fact that $\ell_\infty^{jk} = (1 - q_\infty^{jk})/q_\infty^{jk}$ has a bounded expectation implies that $q_\infty^{jk} > 0$ with probability one. ■

The above lemma states that fully incorrect learning is almost surely impossible. Our next lemma further characterizes the limiting beliefs.

Lemma 3: Suppose that $i \neq j$ and $l \neq k$. Then, conditional on the event $A_{jk} = \{(\theta, \lambda) = (j, \lambda_k)\}$, $q_n^{ik} \to 0$ and $q_n^{jl} \to 0$ with probability one.

Proof: First consider the sequence $\{q_n^{jl}\}$. Conditional on the event A_{jk}, we know that q_n^{jk} does not converge to zero almost surely. Therefore, the likelihood ratio q_n^{jl}/q_n^{jk} converges to some finite limit with probability one. By the social belief update (2), this limit is equal to zero almost surely, unless the equation

$$(\lambda_k - \lambda_l)\left[\mathbb{G}_j(1 - q) - \mathbb{G}_j\left(\frac{h(1 - q)}{h(1 - 2q) + q}\right)\right] = 0$$

has a solution $q \in (0, 1)$. However, no such q exists if $h \neq \frac{1}{2}$. Thus, conditional on A_{jk}, $q_n^{jl} \to 0$ with probability one.

Now consider the sequence $\{q_n^{ik}\}$. A similar argument shows that $q_n^{ik} \to 0$ on almost all sample paths, unless there exists a $q \in (0, 1)$ that solves

$$\begin{aligned}
\lambda_k \mathbb{G}_j(1 - q) + (1 - \lambda_k)\mathbb{G}_j\left(\frac{h(1 - q)}{h(1 - 2q) + q}\right) &= \\
\lambda_k \mathbb{G}_i(1 - q) + (1 - \lambda_k)\mathbb{G}_i\left(\frac{h(1 - q)}{h(1 - 2q) + q}\right)&.
\end{aligned}$$

However, recall that by Lemma 1, $\mathbb{G}_0(r) > \mathbb{G}_1(r)$ for all $r \in (0, 1)$. Therefore, there are no interior solutions to the above equation. This completes the proof. ■

The important consequence of Lemma 3 is that if agents asymptotically learn either component of the state vector (θ, λ), then they necessarily learn the other as well. In other words, there is enough information in the public histories so that agents can distinguish, say $(\theta = 0, \lambda_1)$ from $(\theta = 1, \lambda_1)$. However, Lemma 3 does not guarantee asymptotic learning. It is possible for identification problems to arise in the public histories in such a way that agents cannot distinguish, say $(\theta = 0, \lambda_1)$ from $(\theta = 1, \lambda_2)$.

The next result shows that if the incentives of the two types are not too far apart, then asymptotic learning occurs in equilibrium.

Proposition 1: Suppose $\lim_{r\uparrow 1}\mathbb{G}_1'(r)$ and $\lim_{r\downarrow 0}\mathbb{G}_0'(r)$ are non-zero. Then, there exists $\epsilon > 0$, such that complete asymptotic learning occurs in equilibrium for all $h \in (\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon)$.

Proof: Condition on event $A_{jk} = \{(\theta, \lambda) = (j, \lambda_k)\}$, and suppose that $i \neq j$ and $l \neq k$. An argument similar to the one made in the proof of Lemma 3 implies that asymptotic learning is complete if equation

$$\lambda_k \mathbb{G}_j(1-q) + (1-\lambda_k)\mathbb{G}_j\left(\frac{h(1-q)}{h(1-2q)+q}\right) = \lambda_l \mathbb{G}_i(1-q) + (1-\lambda_l)\mathbb{G}_i\left(\frac{h(1-q)}{h(1-2q)+q}\right) \quad (3)$$

has no solution in $(0,1)$. Since $|h - \frac{1}{2}| < \epsilon$ and ϵ is a small number, we can expand the above equation around $h_0 = \frac{1}{2}$ and get

$$q(1-q)\left[\frac{\mathbb{G}_j(1-q) - \mathbb{G}_i(1-q)}{q(1-q)} + 4\epsilon f(q) + O(\epsilon^2)\right] = 0,$$

where $f(q) = (1-\lambda_k)\mathbb{G}_j'(1-q) - (1-\lambda_l)\mathbb{G}_i'(1-q)$. Lemma 1 guarantees that the first term in the braces is sign-definite and uniformly bounded away from zero for all $q \in [0,1]$. Therefore, for small enough ϵ, the term in the braces is non-zero for all $0 \leq q \leq 1$, which completes the proof. ■

Proposition 1 establishes that complete asymptotic learning in the observational learning model of Smith and Sørensen [4] is robust to uncertainties in type distributions, as long as the bias in the incentive of the types are not too far apart. Note that this result is independent of the agents' prior beliefs and the level of uncertainty in the distributions; i.e., λ_1 and λ_2. However, the complete learning results of Proposition 1 do not hold for larger incentive biases. In the next proposition, we show that if there is sufficient heterogeneity in preference types, then certain identification problems can arise in public histories and hence, beliefs do not converge to the true underlying state of the world.

Proposition 2: Suppose that for all $r \in (0,1)$, we have $(\mathbb{G}_0/\mathbb{G}_1)(r) > r(1-r)(\mathbb{G}_0/\mathbb{G}_1)'(r)$. Then, there exists $\epsilon > 0$ such that for all $h \in [0, \epsilon)$, asymptotic learning is generically incomplete.

Proof: Due to the continuity of $\mathbb{G}_0$ and $\mathbb{G}_1$, it is sufficient to show that asymptotic learning is incomplete for $h = 0$.

Suppose that $h = 0$. By Lemma 3, conditional on the event A_{jk}, we have $q_\infty^{jl} = q_\infty^{ik} = 0$ with probability one. Therefore, complete asymptotic learning is obtained if and only if $q_n^{il} \to 0$ almost surely. Similar to the proof of Lemma 3, the limit of sequence $\{q_n^{il}\}$ depends on the number of interior solutions of equation (3), which for $h = 0$, reduces to

$$\lambda_k \mathbb{G}_j(1-q) = \lambda_l \mathbb{G}_i(1-q). \quad (4)$$

Note that we have $\lambda_2 > \lambda_1$ and $\mathbb{G}_0(r) > \mathbb{G}_1(r)$ for all $r \in (0,1)$, as shown by Lemma 1. Therefore, if $(j, \lambda_k) = (1, \lambda_1)$ or $(j, \lambda_k) = (0, \lambda_2)$, then (4) has no solution except for $q \in \{0,1\}$. Thus, conditional on the event A_{11} or A_{02}, the social belief assigned to the true state converges to one. However, if $(j, \lambda_k) = (1, \lambda_2)$, then there exists a $q^* \in (0,1)$, such that

$$\lambda_2 \mathbb{G}_1(1-q^*) = \lambda_1 \mathbb{G}_0(1-q^*).$$

Note that since $\mathbb{G}_1(r)/\mathbb{G}_0(r)$ is monotone in r, there is exactly one interior solution to the above equation. Similarly, an interior social belief exists when $(j,k) = (0,1)$.

In order to complete the proof, we need to show that the interior solution q^* is locally stable; that is, starting from a neighborhood of q^*, the social beliefs converge to q^* with positive probability. The key observation here is that that social beliefs q_n^{jk} are Markov-martingale stochastic processes, and therefore, we can use the first-order stability criterion for Markov-martingales to establish local stability of q^*. This stability criterion is stated in the Appendix.

For $z \in \{0,1\}$, we define the functions $\phi_z : \mathbb{R}^3 \longrightarrow \mathbb{R}^3$ as

$$(q_{n+1}^{12}, q_{n+1}^{11}, q_{n+1}^{02})(x^{n-1}, z) = \phi_z(q_n^{12}, q_n^{11}, q_n^{02}).$$

which express social beliefs of agent $n+1$ in terms of social beliefs of agent n and her action $x_n = z$. Note that it is sufficient to express the social beliefs only as a function of q_n^{12}, q_n^{11}, and q_n^{02}, as the remaining belief q_n^{01} linearly depends on these terms. Also note that $(q^*, 0, 0)$ is a fixed point of functions ϕ_0 and ϕ_1.

Using straightforward algebraic computations one can show that the eigenvalues of $\nabla\phi_1|_{(q^*,0,0)}$ are given by

$$\begin{aligned} v_1 &= \frac{1-\lambda_2\mathbb{G}_0(1-q^*)}{1-\lambda_2\mathbb{G}_1(1-q^*)} \\ v_2 &= \frac{1-\lambda_1\mathbb{G}_1(1-q^*)}{1-\lambda_2\mathbb{G}_1(1-q^*)} \\ v_3 &= 1 + \frac{q^*(1-q^*)}{1-\lambda_2\mathbb{G}_1(1-q^*)}\left[\lambda_2\mathbb{G}_1'(1-q^*) - \lambda_1\mathbb{G}_0'(1-q^*)\right] \end{aligned}$$

which are real, positive, non-unit, and generically distinct. Similarly, one can show that the eigenvalues of $\nabla\phi_0|_{(q^*,0,0)}$ are equal to

$$\begin{aligned} w_1 &= \lambda_2/\lambda_1 \\ w_2 &= \lambda_1/\lambda_2 \\ w_3 &= 1 - \frac{r^*(1-r^*)}{\mathbb{G}_0(r^*)\mathbb{G}_1(r^*)}\left[\mathbb{G}_0(r^*)\mathbb{G}_1'(r^*) - \mathbb{G}_1(r^*)\mathbb{G}_0'(r^*)\right] \end{aligned}$$

where $r^* = 1 - q^*$. Thus, eigenvalues of $\nabla\phi_0|_{(q^*,0,0)}$ are also real, positive, non-unit, and generically distinct. Therefore, all the conditions of Theorem 4 in the Appendix are satisfied, and as a result, conditional on the event A_{12}, the social belief q_n^{12} converges to q^* with positive probability. This completes the proof. ■

Proposition 2 establishes that for values of h close to zero, there are sample paths over which individuals forever remain uncertain about the true underlying state of the world. One can show that a similar result holds for h close to one. The main reason for the failure of learning in such cases is that after some point in time, agents' actions lead to an identification problem, where it is impossible for others to infer the truth. This phenomenon, which Smith and Sørensen [4] refer to as *confounded learning*, is different from the herding behavior. Notice that in confounded outcomes agents continue to use the information in their private signals for making decisions. However, their equilibrium decision making does not convey any further information to agents who appear later in the sequence. On the other hand, herding behavior is a consequence of bounded private beliefs and happens when agents discard their private signals completely, and solely base their decisions on the public histories. Hence, no new information is encoded in the social beliefs as time progresses.

We also remark that under the conditions of Proposition 2, asymptotic learning is incomplete even if λ_1 is arbitrarily close to λ_2. In other words, when incentive biases of the two types are far enough, even the slightest uncertainty in their distributions leads to a confounding outcome with a small positive probability. This is in contrast to Proposition 1, which shows that when incentives are sufficiently aligned, even high levels of uncertainty about type distributions do not preclude learning.

IV. Ex Post Learning

The main conclusion of Proposition 2 is that uncertainty about the distribution of preferences may lead to incomplete learning of the payoff-relevant state θ. However, the proposition does not rule out the possibility of ex post learning on some sample paths. In other words, it is possible for the posterior beliefs to converge to the true state on some sample path, even though *ex ante*, individuals do not believe that they will learn θ. Our next result shows that asymptotic learning occurs with positive probability even when there is significant heterogeneity in the preferences. Note that this is in contrast with the asymptotic learning results in the absence of type uncertainties, where ex ante and ex post learning coincide.

Corollary 3: Suppose that $h = 0$. Then, for $j \in \{0, 1\}$,

$$\lim_{n\to\infty} \mathbb{P}(\theta = j | I_n) = \mathbb{I}\{\theta = j\}$$

with positive probability.

Proof: The result follows from the proof of Proposition 2. Recall that equation (4) has no interior solutions when $(j, \lambda_k) = (1, \lambda_1)$ or $(j, \lambda_k) = (0, \lambda_2)$. Therefore, conditional on the event A_{11} (resp. A_{02}), we have $q_n \to 1$ (resp. $q_n \to 0$), with probability one. Since A_{11} and A_{02} have positive probabilities, the social beliefs converge to the truth with some non-zero probability. ■

V. Conclusions

In this paper, we studied the problem of information aggregation in Bayesian observational learning models with type uncertainties. We assumed that individuals have different incentives and are uncertain about the type distributions in the society. We showed that complete asymptotic learning occurs in the Perfect Bayesian Equilibrium of the game as long as the incentives of different types are not too far apart. On the other hand, our main result established that learning is not robust to type distribution uncertainties when different types have significantly different incentives. We showed that in the presence of such uncertainties, the equilibrium strategies lead to identification problems in the public histories, and as a result, agents cannot extract enough information from the actions of their predecessors to learn the state. We remarked that this confounded learning which is a consequence of type uncertainties is fundamentally different from the herding behavior, which is due to the boundedness of private signals. Finally, we showed that even though learning is guaranteed to be incomplete ex ante, ex post learning is still possible: there exists a positive measure set of sample paths over which agents become certain about the true underlying state of the world.

Appendix: Local Stability of Markov-Martingale Processes

The Appendix contains the first-order stability criterion for Markov-martingale stochastic processes. A more thorough treatment of the subject can be found in Smith and Sørensen [4] and Sørensen [9].

Consider the Markov process

$$y_{n+1} = \phi_z(y_n) \quad \text{with probability } \psi_z(y_n). \tag{5}$$

where $z \in \{0, 1\}$. We say the Markov process above is a *Markov-martingale* process if it satisfies

$$\psi_0(y)\phi_0(y) + \psi_1(y)\phi_1(y) = y,$$

for all y. Note that the above condition guarantees that $\mathbb{E}[y_{n+1}|y_1, \ldots, y_n] = y_n$. We have the following theorem.

Theorem 4: Let y^* be a fixed point of the Markov-Martingale process (5). Assume that $\psi_0(\cdot)$ and $\psi_1(\cdot)$ are continuous at y^*, with $0 < \psi_z(y^*) < 1$ for $z \in \{0, 1\}$, and that each $\phi_z(\cdot)$ is continuously differentiable at y^*. Also assume that each $\nabla_y \phi_z(y^*)$ has distinct, real, positive, non-unit eigenvalues. Then for any open ball $\mathcal{B}$ around y^*, there exists a $\theta < 1$ and an open ball $\mathcal{N} \subset \mathcal{B}$ around y^* such that

$$y_0 \in \mathcal{N} \Longrightarrow \mathbb{P}(\theta^{-n}\|y_n - y^*\| \to 0) > 0.$$

The proof can be found in [4].

References

[1] L. J. Savage, *The Foundations of Statistics*. New York: Wiley, 1954.
[2] A. Banerjee, "A simple model of herd behavior," *The Quarterly Journal of Economics*, vol. 107, no. 3, pp. 797–817, 1992.
[3] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as information cascades," *Journal of Political Economy*, vol. 100, no. 5, pp. 992–1026, 1992.
[4] L. Smith and P. N. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, Mar. 2000.
[5] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," 2008, LIDS Working Paper ♯2780, Massachusetts Institute of Technology.
[6] ——, "Heterogeneity and social learning," 2009, Working Paper.
[7] D. Acemoglu, V. Chernozhukov, and M. Yildiz, "Fragility of asymptotic agreement under Bayesian learning," Feb. 2009, Working paper.
[8] R. Durrett, *Probability: Theory and Examples*, 3rd ed. Belmont, CA: Duxbury Press, 2005.
[9] P. N. Sørensen, "Rational social learning," Ph.D. dissertation, Massachusetts Institute of Technology, 1996.