Preliminary Results on Social Learning with Partial Observations

Ilan Lobel, Daron Acemoglu, Munther Dahleh and Asuman Ozdaglar *†‡

ABSTRACT

We study a model of social learning with partial observations from the past. Each individual receives a private signal about the correct action he should take and also observes the action of his immediate neighbor. We show that in this model the behavior of asymptotic learning is characterized in terms of two threshold values that evolve deterministically. Individual actions are fully determined by the value of their signal relative to these two thresholds. We prove that asymptotic learning from an ex ante viewpoint applies if and only if individual beliefs are unbounded. We also show that symmetry between the states implies that the minimum possible amount of asymptotic learning occurs.

1. INTRODUCTION

Many important decision are taken by individuals under conditions of imperfect information. In such situations, it is natural for individuals to gather information in order to improve their decisions. A major source of information is the past actions of other individuals facing similar decision problems. This motivates the analysis of social learning problems, where a group of individuals are simultaneously learning from others and also taking important economic or social decisions. Examples of social learning problems include behavior in financial markets, where each trader may try to learn from the positions of other traders or from prices, consumer decisions in product markets, where purchases by other consumers are a key source of information, and political decision-making, where in voting or other political actions individuals typically learn from and condition on the behavior of others. A central question is therefore whether the equilibrium process of social learning will lead to the correct actions by groups.[1]

A large literature in game theory investigates the first question. A well-known result in this context, first derived by Banerjee [1] and Bikchandani et al. [3], establishes the possibility of a "pathological" result that features no learning and the possibility of incorrect actions by a large group of individuals. Consider N individuals ordered exogenously and choosing between two actions, say 0 and 1. Each individual receives a signal about which action is the right one and also observes the actions of all other agents that have moved before him. The signal received by each individual takes two possible values (one favoring 0 the other one favoring 1) and is identically and independently distributed across individuals. Banerjee [1] and Bikchandani et al. [3] show that the perfect Bayesian equilibrium of this game involves a particular type of "herding" in which following two consecutive actions in the same direction (for example, two individuals choosing 0), each subsequent individual ignores his own signal and follows the actions of these two individuals. Clearly, since two individuals choosing the action 0 is possible even when the right action is 1, this result illustrates a pathological form of non-learning and incorrect actions by individuals.

A more complete analysis of this model is provided by Smith and Sorensen [9], who analyze the case in which signals can also differ in their informativeness. Smith and Sorensen's main result can be summarized as follows. Let us refer to signals as *unbounded* if the likelihood ratio of a particular state can be arbitrarily large conditional on individual signals and as *bounded* otherwise. Smith and Sorensen show that with unbounded signals, there will be asymp-

*Ilan Lobel (lobel@mit.edu) is with the Operations Research Center at the Massachusetts Institute of Technology, Cambridge, MA.
†Daron Acemoglu (daron@mit.edu) is with the Department of Economics at the Massachusetts Institute of Technology, Cambridge, MA.
‡Munther Dahleh (dahleh@mit.edu) and Asuman Ozdaglar (asuman@mit.edu) are with Department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology, Cambridge, MA.

Inter-Perf '07, October 26, 2007, Nantes, France

[1] A related and equally important question concerns what types of communication and observation structures will facilitate learning. For example, is learning more or less likely when individuals observe actions and communicate within their narrow communities? More generally, what is the impact of the topology of a social network on the patterns of learning? We study this question in our companion paper [7].

totic learning, i.e., the probability of the correct action being chosen converges to 1.

This literature typically focuses on social learning environments in which individuals observe all previous actions. Consequently, the information set of individuals making decision later is necessarily finer than those moving earlier, which implies that Bayesian posteriors form *a martingale*. This property enables the use of the martingale convergence theorem and significantly simplifies the analysis. However, most relevant cases of social learning in practice do not feature this property. Often, each individual will have observed a different sample of actions than those who have acted before and will not necessarily have superior information relative to them. The existing literature, except for the more recent paper by Smith and Sorensen [8], has not studied the properties of equilibrium social learning in this more realistic environment. An investigation of the patterns of social learning in such an environment is not only important because of its greater realism, but also because it will enable us to address the second question posed above and study what types of social structures are more conducive to learning and information aggregation.

In this paper, we take a step in this direction by studying the simplest model of social learning without the martingale property. Each individual again receives a signal (with varying degree of informativeness) but only observes the action of the person who has moved before him. Despite the simplicity of this environment, existing results in the literature do not apply. Moreover, the mathematical structure of this simple case is very similar to the case in which each individual observes a uniformly random decision from the past and our result extend in a straightforward manner.

Our main results are as follows. First, we provide a recursive characterization of individual decisions in terms of two deterministic thresholds, such that the value of individual signals relative to these thresholds completely determines decisions. Second, as in Smith and Sorensen [9], unbounded signals ensure asymptotic learning. Third, when signals are bounded, there will never be asymptotic learning. Finally, we show that under a symmetry condition on the conditional signal distributions and with bounded signals, there will exist an equilibrium with the minimum amount of learning in the long-run. Under very mild conditions, this equilibrium is unique. In contrast, with asymmetry between the states, the amount of asymptotic learning can be quite high.

Our paper is related to the large and growing social learning literature (see [1], [3], [5], [4], [10]). Most closely related are the recent papers by Banerjee and Fudenberg [2] and Smith and Sorensen [8]. Banerjee and Fudenberg analyze a model of social learning in which individuals observe a random sample of past actions under the assumption that there is a continuum of agents, so that past actions reveal sufficient information about the underlying state. Smith and Sorensen study a related environment of social learning without

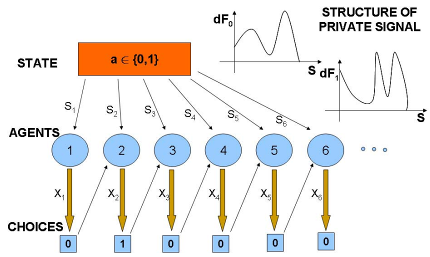

Figure 1: Model of Social Learning with Limited Information.

the martingale property. While their method of analysis is different from ours, a number of our results are present in their work. In particular, Smith and Sorensen also show that unbounded signals will lead to social learning. However, our results on the dynamics of beliefs, the limiting distribution of probabilities and the role that asymmetry plays in asymptotic learning are novel.

The rest of the paper is organized as follows. In Section 2 we present the model, followed by an analysis of the properties of private beliefs in Section 3. In Section 4, we characterize the evolution of ex ante probabilities of taking the correct action. Section 5 presents our main results on asymptotic learning under unbounded signals and characterizes the convergence behavior of actions under bounded signals.

2. THE MODEL

The game consists of a countably infinite number of agents indexed by $n \in \mathbb{N}$, acting sequentially. Each agent n has a single action $x_n \in \{0,1\}$. The underlying state of the world is $a \in \{0,1\}$. If $x_n = a$, then the payoff of agent n is given by $u_n = 1$, and otherwise, $u_n = 0$. A priori, both states of the world are equally likely.

Let the information set of agent n be Ω_n. We assume that $\Omega_n = \{s_n, x_{n-1}\}$, where s_n is the private signal of the individual drawn independently from the conditional distribution F_a given the underlying state $a \in \{0,1\}$, and x_{n-1} is the action of the previous agent.

Our goal is to understand the limiting properties of a perfect Bayes-Nash equilibrium in this model. In particular, we want to determine the level of learning that is achieved by the agents as measured by their ex ante probability of choosing the best decision, i.e., $P(x_n = a)$.

DEFINITION 1. *(Asymptotic Learning) There is asymptotic learning if x_n converges to a in probability, i.e.,* $\lim_{n\to\infty} P(x_n = a) = 1$.

3. PRIVATE BELIEFS

How the sequence of decisions $\{x_n\}$ evolves depends on inference based on individuals' signals regarding the underlying state. It is convenient to work with a transformation of these signals, which we refer to as *private beliefs* (see [9]).

DEFINITION 2. *(Private Belief) Agent n's private belief p_n is the probability that the state is equal to 1 conditional on his private signal s_n, i.e., $p_n = P(a=1|s_n)$.*

For a given signal s_n, by Bayes' rule, the private belief is

$$p_n = \frac{1}{1+\frac{dF_0(s_n)}{dF_1(s_n)}}, \qquad (1)$$

where dF_a reduces to the density of F_a if the distribution function has a density and the ratio in the denominator is the likelihood ratio.

Since p_n is a function of s_n only, the sequence of random variables $\{p_n\}$ is also independent and identically distributed. We will denote the cumulative distribution function for private beliefs given the true state a by G_a. That is,

$$G_a(x) = P(p_n \le x|a), \quad \text{for all } n \in \mathbb{N}. \qquad (2)$$

Because the private beliefs contain all the useful information about the signals, we will directly work with private beliefs, or equivalently we suppose that each agent n knows only x_{n-1} and p_n when making his decision.

DEFINITION 3. *(Bounded and Unbounded Private Beliefs) Let β and $1-\gamma$ be the infimum and the supremum of the support of the distribution function G_1, i.e.,*

$$\beta = \inf_{x\in[0,1]} \{x : G_1(x) > 0\}. \qquad (3)$$

$$\gamma = 1 - \sup_{x\in[0,1]} \{x : G_1(x) < 1\}. \qquad (4)$$

Then, private beliefs are unbounded if $\beta = \gamma = 0$. The beliefs are bounded if both $\beta > 0$ and $\gamma > 0$.

We ignore the possibility that only one of β and γ is strictly positive to simplify the presentation.[2]

Unbounded private beliefs correspond to the likelihood ratio in Eq. (1) being unbounded, which implies that an agent can receive an arbitrarily strong signal about the underlying state. As in the existing work on the social learning literature, this feature will have important implications for the limiting behavior of the sequence $\{x_n\}$.

4. EVOLUTION OF THE PROCESS

In this paper, we will characterize the limiting behavior of the agents by focusing on *ex ante* probabilities of correct decisions conditional on the true state a. These probabilities will be denoted

$$Y_n = P(x_n = 1|a=1) \text{ and } N_n = P(x_n = 0|a=0). \qquad (5)$$

The unconditional probability of a correct decision is then

$$P(x_n = a) = \frac{Y_n + N_n}{2}, \qquad (6)$$

[2]Note that β and γ can be alternatively defined in terms of G_0 since the two distributions have the same support.

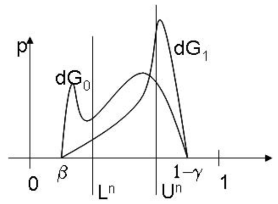

Figure 2: Equilibrium Decision Rule Depicted on the Private Belief Interval.

and therefore asymptotic learning (from an ex ante point of view) is equivalent to the convergence of the sequence $\{(Y_n, N_n)\}$.

Let us next define the thresholds

$$U_n = \frac{N_n}{1-Y_n+N_n} \text{ and } L_n = \frac{1-N_n}{1-N_n+Y_n}, \qquad (7)$$

which will fully characterize the decision rule as described by Lemma 1 below. Note that the sequence $\{(U_n, L_n)\}$ only depend on $\{(Y_n, N_n)\}$ and therefore are deterministic. This reflects the fact that each individual recognizes the amount of information that will be contained in the action of the previous agent, which determines his own decision thresholds. Individual actions are still stochastic since they are determined by whether the individual's private beliefs is below L_n, above U_n or in between.

DEFINITION 4. *Agent n's strategy σ_n is a mapping from his information set to his possible actions, i.e., $\sigma_n : \Omega_n \to \{0,1\}$. A* perfect Bayesian equilibrium *of the game is a sequence of strategies for the players $\{\sigma_n^*\}$ such that for each n, σ_n^* maximizes the agent's expected utility given $\{\sigma_1^*, \ldots, \sigma_{n-1}^*, \sigma_{n+1}^*, \ldots\}$.*

LEMMA 1. *Let U_n and N_n be given by Eq. (7). Then, in all perfect Bayesian equilibria agent n's decision rule satisfies:*

$$x_n = \begin{cases} 0, & \text{if } p_n < L_{n-1}, \\ x_{n-1}, & \text{if } p_n \in (L_{n-1}, U_{n-1}), \\ 1, & \text{if } p_n > U_{n-1}. \end{cases}$$

Using Eq. (7), it follows for any $\beta > 0$ that $L_n \ge \beta$ if and only if

$$N_n + \left(\frac{\beta}{1-\beta}\right) Y_n \le 1. \qquad (8)$$

Similarly, $U_n \le 1-\gamma$ if and only if

$$\left(\frac{\gamma}{1-\gamma}\right) N_n + Y_n \le 1. \qquad (9)$$

We obtain a stationary zone (the shaded area in Figure 3) such that once the sequence $\{(Y_n, N_n)\}$ enters this area, it remains constant. This region is the singleton (1,1) when beliefs are unbounded and is a non-degenerate quadrilateral when beliefs are bounded. Asymptotic learning is clearly equivalent to $\lim_{n\to\infty}\{(Y_n, N_n)\} = (1,1)$.

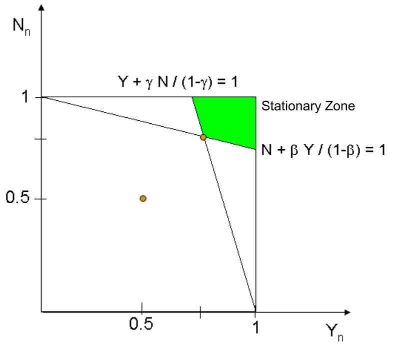

Figure 3: Stationary Zone on (Y_n, N_n) Graph.

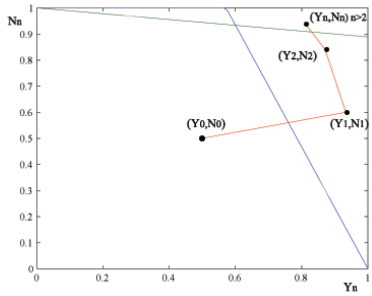

Figure 4: Example Showing Asymmetry Could Lead to More Learning.

5. CONVERGENCE ANALYSIS

The next proposition is one of the main results of our paper and shows that the sequence $\{(Y_n, N_n)\}$ asymptotically approaches the stationary zone given by the shaded area in Figure 3.

PROPOSITION 1. *Let L_n and U_n be as defined in Eq. (7) and β and γ as in Eq. (3). The sequences L_n and U_n satisfy*

$$\limsup_{n\to\infty} L_n \leq \beta, \quad \textit{and} \quad \liminf_{n\to\infty} U_n \geq 1-\gamma.$$

The proof of this and all other results are provided in [6].

An implication of Proposition 1 is that asymptotic learning occurs when the private beliefs are unbounded.

PROPOSITION 2. *When the private beliefs are unbounded, asymptotic learning occurs, i.e., $\lim_{n\to\infty} P(x_n = a) = 1$. When the beliefs are bounded, asymptotic learning does not occur, i.e., $\lim_{n\to\infty} P(x_n = a) < 1$.*

5.1 Learning under Symmetry

When beliefs are bounded, Proposition 2 does not specify whether and where the sequence $\{(Y_n, N_n)\}$ will converge. We will next establish that under a symmetry assumption there exists an equilibrium with the minimum amount of asymptotic learning possible.

ASSUMPTION 1. *(Symmetry) The states are symmetric if $G_0(r) = G_1(1-r)$ for all $r \in [0,1]$.*

ASSUMPTION 2. *G_0 and G_1 have densities.*

The next proposition contains the main convergence result of this subsection. In particular, we show that both sequences $\{N_n\}$ and $\{Y_n\}$ converge to the limit $(1-\beta)$.

PROPOSITION 3. *Assume that symmetry holds. Then, there exists an equilibrium where the sequences $\{N_n\}$ and $\{Y_n\}$ both converge to the limit $(1-\beta)$, i.e., $\lim_{n\to\infty} N_n = \lim_{n\to\infty} Y_n = (1-\beta)$. If Assumption 2 also holds, this equilibrium is unique.*

If symmetry does not hold, then the sequence $\{Y_n + N_n\}$ might converge to a value greater than $2(1-\beta)$, i.e., not to the edge of the shaded region in Figure 3.

As an example of the behavior of asymptotic learning without symmetry, Figure 4 represents the dynamics of $\{(Y_n, N_n)\}$ for the following the distribution $G_0(r) = \frac{18}{30}$, for $r \in [0.1, 1-0.7)$ and the cumulative distributions having value 0 if $r < 0.1$ and value 1 for $r \geq 0.7$ (there is a unique G_1 associated with this G_0). In this example, private beliefs can take only two values, 0.1 and 0.7. The private belief of 0.1 implies a strong likelihood that 0 is the true state, while a belief of 0.7 implies a much weaker likelihood in favor of state 1. In this example, the sequence $\{(Y_n, N_n)\}$ converges to a point in the interior of the stationary zone as can be seen in Figure 4. As noted above, this limit point involves a greater amount of asymptotic learning than in the case with symmetric pair.

6. CONCLUSIONS

In this paper, we presented an analysis of social learning when individuals only observe the action of their immediate neighbor. Despite the simplicity of this environment, the evolution of beliefs is substantially different than the typical models of social learning in the game theory literature. We characterized the behavior of asymptotic learning in terms of two threshold values that evolve deterministically. Individual actions are fully determined by the value of their signal relative to these two thresholds. We prove that asymptotic learning from an ex ante viewpoint applies if and only if individual beliefs are unbounded. We also show that for symmetric states bounded signals imply the minimum possible amount of asymptotic learning.

The tools introduced in this paper can be generalized to analyze social learning in environments in which individuals observe many samples of past actions and investigate how the topology of communication across agents affects information aggregation and the likelihood of asymptotic learning. This is an area we are investigating in [7].

7. REFERENCES

[1] Banerjee A., “A Simple Model of Herd Behavior," *Quarterly Journal of Economics,* vol. 107, pp. 797-817, 1992.

[2] Banerjee A. and Fudenberg D., “Word-of-mouth Learning," *Games and Economic Behavior*, vol. 46, pp. 1-22, 2004.

[3] Bikchandani S., Hirshleifer D., and Welch I., “A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy*, vol. 100, pp. 992-1026, 1992.

[4] Ellison G., Fudenberg D., “Rules of Thumb for Social Learning," *Journal of Political Economy*, vol. 101, no. 4, pp. 612-643, 1993.

[5] Ellison G. and Fudenberg, D., “Word-of-mouth communication and social learning," *The Quarterly Journal of Economics*, vol. 110, pp. 93-126, 1995.

[6] Lobel I., Acemoglu D., Dahleh M., Ozdaglar A., “Learning from Neighbors,", MIT Mimeo, 2007.

[7] Lobel I., Acemoglu D., Dahleh M., Ozdaglar A., “The Structure of Information and Limits of Learning,", MIT Mimeo, 2007.

[8] Smith L. and Sorensen P., “Rational Social Learning with Random Sampling," unpublished manuscript, 1998.

[9] Smith L. and Sorensen P., “Pathological Outcomes of Observational Learning," *Econometrica*, vol. 68, no. 2, pp. 371-398, 2000.

[10] Vives X., “Social Learning and Rational Expectations," *European Economic Review*, vol. 40, pp. 589-601, 1996.