

Collegio Carlo Alberto



Observational Learning with Position Uncertainty

Ignacio Monzón

Michael Rapp

No. 206

May 2011

Carlo Alberto Notebooks

www.carloalberto.org/working_papers

© 2011 by Ignacio Monzón and Michael Rapp. Any opinions expressed here are those of the authors and not those of the Collegio Carlo Alberto.

Observational Learning with Position Uncertainty*

Ignacio Monzón[†] and Michael Rapp[‡]

May 3, 2011

Abstract

Observational learning is typically examined when agents have precise information about their position in the sequence of play. We present a model in which agents are uncertain about their positions. Agents are allowed to have arbitrary ex-ante beliefs about their positions: they may observe their position perfectly, imperfectly, or not at all. Agents sample the decisions of past individuals and receive a private signal about the state of the world. We show that social learning is robust to position uncertainty. Under any sampling rule satisfying a stationarity assumption, learning is complete if signal strength is unbounded. In cases with bounded signal strength, we show that agents achieve what we define as constrained efficient learning: individuals do at least as well as the most informed agent would do in isolation.

JEL Classification: C72, D83, D85

*We are grateful to Bill Sandholm for his advice, suggestions and encouragement. We also thank Dan Quint, Ricardo Serrano-Padial, Marzena Rostek, Marek Weretka, Ben Cowan, María Eugenia Garibotti, Danqing Hu, Sang Yoon Lee, Fernando Louge and David Rivers for valuable comments. All remaining errors are ours.

[†]Collegio Carlo Alberto, Via Real Collegio 30, 10024 Moncalieri (TO), Italy. (ignacio@carloalberto.org, <http://www.carloalberto.org/people/monzon/>)

[‡]Department of Economics, University of Wisconsin–Madison, 1180 Observatory Drive, Madison, WI 53706, USA. (mrapp@wisc.edu, <https://mywebspace.wisc.edu/mrapp/web>)

1. Introduction

In a wide range of economic situations, agents possess private information regarding some shared uncertainty. These include choosing between competing technologies, deciding whether to invest in a new business, deciding whether to eat at a new restaurant in town, or selecting which new novel to read.

If actions are observable but private information is not, one agent's behavior provides useful information to others. Consider the example of choosing between two MP3 players recently released to the market. One is inherently better for all individuals, but agents do not know which.¹ Individuals form a personal impression about the quality of each device based on information they receive privately. They also observe the choices of others, which partly reveal the private information of those agents.

An important characteristic of this environment has been overlooked: when choosing between competing alternatives, an agent may not know how many individuals have faced the same decision before him. Moreover, when an agent observes someone else's decision, he also may not know when that decision was made. For example, even when individuals know when the competing MP3 players were released, they may be unaware of exactly how many individuals have already chosen between the devices. In addition, when an agent observes someone else listening to an MP3 player on the bus, he does not know whether the person using the device bought it that morning or on the day it was released. We study settings with position uncertainty, and ask whether individuals eventually learn and choose the superior technology by observing the behavior of others.

The seminal contributions to the social learning literature are [Banerjee \[1992\]](#) and [Bikhchandani, Hirshleifer, and Welch \[1992\]](#). In these papers, a set of rational agents choose sequentially between two technologies. The payoff from this decision is identical to all, but unknown. Agents know their position in the adoption sequence. In each period, one agent receives a signal and observes what all agents before him have chosen. Afterwards, this agent makes a once-and-for-all decision between the two technologies. Given that each agent knows that the signal he has received is no better than the signal other individuals have received, agents eventually follow the behavior of others and ignore their own signals. Consequently, [Banerjee and Bikhchandani et al.](#)

¹We study situations where network externalities do not play a salient role, including choosing between news sites, early search engines, computer software, MP3 players or computer brands. Our focus is on informational externalities. We are currently working on a setting with network externalities, for which we have partial results.

show that the optimal behavior of rational agents can prevent social learning.

Smith and Sørensen [2000] highlight that there is no information aggregation in the models of Banerjee [1992] and Bikhchandani et al. [1992] because these models assume signals are of *bounded strength*. In Smith and Sørensen [2000], agents also know their positions in the adoption sequence and observe the behavior of all preceding agents. In contrast to Banerjee [1992] and Bikhchandani et al. [1992], agents receive signals of *unbounded strength*: signals can get arbitrarily close to being perfectly informative. In such a context, the conventional wisdom represented by a long sequence of agents making the same decision can always be overturned by the action of an agent with a strong enough signal. As a result, individuals never fully disregard their own information. In fact, Smith and Sørensen [2000] show that if signals are of unbounded strength, social learning occurs.

We take as our starting point the basic setup in the literature on social learning and introduce position uncertainty. In our setting, agents, exogenously ordered in a sequence, choose between two competing technologies. Agents receive a noisy private signal about the quality of each technology. We study cases of both bounded and unbounded signal strength. We depart from the literature by assuming that agents may not know 1) their position in the sequence or 2) the position of those they observe. We refer to both phenomena as *position uncertainty*. We also depart from Banerjee [1992], Bikhchandani et al. [1992] and Smith and Sørensen [2000] in that agents observe the behavior of only a sample of preceding agents.

To understand the importance of position uncertainty, note that if an agent knows his position in the sequence, he can condition his behavior on *three* pieces of information: his private signal, his sample and *his position in the sequence*. In fact, he may weigh his sample and signal differently depending on his position. The typical story of social learning is one of learning over time. Early agents place a relatively larger weight on the signal than on the sample. As time progresses, information is aggregated and the behavior of agents becomes a better indicator of the true state of the world. Later agents place a relatively larger weight on the sample, which, in turn, can lead to precise information aggregation. In contrast, if agents have no information about their positions, such heterogeneity in play is impossible. Instead, agents place the same weight on the sample regardless of when they play. Heterogeneous play based on positions cannot drive social learning in this setting.²

²Position uncertainty leads to an additional difficulty. Agents A and B do not know if A precedes B or vice versa. This adds a strategic component to the game: A's strategy affects B's payoffs inasmuch as B's strategy affects A's payoffs. Thus, the game cannot be solved recursively and, in fact, there are cases with multiple equilibria.

This paper presents a flexible framework for studying observational learning under position uncertainty. Our framework is general in two ways. First, agents are placed in the adoption sequence according to an arbitrary distribution. For example, some agents may, *ex-ante*, be more likely to have an early position in the sequence, while others are more likely to be towards the end. Second, our setup allows for general specifications of the information agents obtain about their position, once it has been realized. Agents may observe their position, receive no information about it, or have imperfect information about it. For instance, agents may know if they are early or late adopters, even if they do not know their exact positions.

We focus on stationary sampling, which allows for a rich class of natural sampling rules. Sampling rules are *stationary* if the probability that an agent samples a particular set of individuals is only a function of the distance between the agent and those he observes (that is, how far back in time those decisions were made). This assumption implies that no individual plays a decisive role in everyone else's samples.

We say social learning occurs if the probability that a randomly selected agent chooses the superior technology approaches one as the number of agents grows large. We find that learning is robust to the introduction of position uncertainty. We specify weak requirements on the information of agents: they observe a private signal and at least an unordered sample of past play. Agents need not have *any* information on their position in the sequence. Under unbounded signal strength, social learning occurs. This is due to the fact that as the number of agents grows large, individuals rely less on the signal. In cases with bounded signal strength, we show agents achieve what we define as *constrained efficient learning*: agents expect to do at least as well as the most informed agent would do in isolation.

In the present setting, social learning results from the combination of stationary sampling and a modified *improvement principle*. First introduced by [Banerjee and Fudenberg \[2004\]](#), the improvement principle states that since agents can copy the decisions of others, they must do at least as well, in expectation, as those they observe. In addition, when an agent receives a very strong signal and follows it, he must do better than those he observes. As long as the agents observed choose the inferior technology with positive probability, the improvement is bounded away from zero. In [Banerjee and Fudenberg's](#) model, agents know their positions and observe others from the preceding period. If learning did not occur, this would mean that agents choose the inferior technology with positive probability in every period. Thus, there would be an improvement between any two consecutive periods. This would lead to an infinite sequence of improvements,

which cannot occur, since utility is bounded. In that way, Banerjee and Fudenberg show social learning must occur in their setting. Acemoglu, Dahleh, Lobel, and Ozdaglar [Forthcoming] and Smith and Sørensen [2008] use similar arguments to show social learning in their models.

We develop an *ex-ante* improvement principle, which allows for position uncertainty, and so places weaker requirements on the information that agents have. However, this *ex-ante* improvement principle does not guarantee learning by itself. Under position uncertainty, the improvement upon those observed is only true *ex-ante* (i.e. before the positions are realized). Conditional on his position, an agent may actually do worse, on average, than those he observes.³

Stationary sampling rules have the following implication: as the number of agents grows large, all agents become equally likely to be sampled.⁴ This leads to a useful accounting identity. Note that the *ex-ante* (expected) utility of an agent is the average over the utilities he would get in each possible position. As all agents become equally likely to be sampled, the expected utility of an *observed* agent approaches the average over all possible positions. Thus, for all stationary sampling rules, the difference between the *ex-ante* utility and the expected utility of those observed must vanish as the number of agents grows large.

To summarize, our *ex-ante* improvement principle states that the difference between the *ex-ante* utility and the expected utility of an observed agent only goes to zero if, in the limit, agents cannot improve upon those observed. This, combined with stationary sampling rules, implies there is social learning if signal strength is unbounded and constrained efficient learning if signal strength is bounded.

The result of constrained efficient learning is novel in the literature and is useful for two reasons. First, it describes a lower bound on information aggregation for all information structures. For any signal strength, agents do, in expectation, at least as well as if they were provided with the strongest signals available. Second, it highlights the continuity of social learning, or perfect information aggregation. Social learning becomes a limit result from constrained efficient learn-

³To see this, consider the following example. There is a large number of agents in a sequence. Everyone knows their position exactly except for the agents in the second and last positions, who have an equal chance of being in each of these positions. If each agent observes the decision of the agent that preceded him, then, by the standard improvement principle, the penultimate agent in the sequence makes the correct choice with a probability approaching one. The agent who plays last, on the other hand, is not sure that he is playing last. He does not know if he is observing the penultimate agent, who is almost always right, or the first agent, who often makes mistakes. As a result, the agent who plays last relies on his signal too often, causing him to make mistakes more often than the agent he observes.

⁴The following example illustrates this point. Let all agents be equally likely to be placed in any position in a finite sequence and let sampling follow a simple stationary rule: each agent observes the behavior of the preceding agent. An agent does not know his own position and wonders about the position of the individual he observes. Since the agent is equally likely to be in any position, the individual observed is also equally likely to be in any position (except for the last position in the sequence).

ing: as we relax the bounds on the signal strength, the lower bound on information aggregation approaches perfect information aggregation.

Before introducing the formal model, we describe the related literature on observational learning. Most of the literature has focused on cases in which agents know their own positions and the positions of those they observed, and sample from past behavior. The aforementioned Banerjee and Fudenberg [2004], Çelen and Kariv [2004], and Acemoglu et al. [Forthcoming] are among them. Acemoglu et al. [Forthcoming] present a model on social learning with stochastic sampling. In their model, agents know both their own position and the position of sampled individuals. Under mild conditions on sampling, social learning occurs.⁵ Larson [2008] focuses on the speed of learning. He suggests a model where agents observe a weighted average of past actions and make a choice in a continuous action space. Larson shows that if agents could choose weights, they would place higher weights on more recent actions. Moreover, learning is faster when the effects of early actions fade out quickly.

Several recent papers have highlighted the importance of position uncertainty. Smith and Sørensen [2008] is the first paper to allow for some uncertainty about positions. In their model, agents know their own position but do not know the positions of the individuals sampled. Since agents know their own positions, an improvement principle and a mild assumption on sampling ensures that social learning occurs. In Callander and Hörner [2009], each agent observes the aggregate number of adopters, but does not know the order of the decisions. Callander and Hörner show the counterintuitive result that it is sometimes optimal to follow the decision of a minority. Hendricks, Sorensen, and Wiseman [2010] present a similar model and test its predictions using experimental data on an online music market. Herrera and Hörner [2009] and Guarino, Harmgart, and Huck [Forthcoming] focus on the interesting case where only one decision (to invest) is observable whereas the other one (not to invest) is not. Herrera and Hörner propose a continuous time model, with agents' arrivals determined by a Poisson arrival process. They show that when signals are of unbounded strength, social learning occurs. In Guarino et al. [Forthcoming], time is discrete and agents make decisions sequentially. Guarino et al. show that cascades cannot occur on the unobservable decision. This paper differs from these recent papers in several dimensions. First, we focus on cases where both actions are observable. Second, we allow for agents to observe

⁵To see why position uncertainty matters in a setup like the one described in Acemoglu et al. [Forthcoming], consider the following example. Each agent observes the very first agent and the agent before him. However, they do not know which is which. Since the first agent may choose the inferior technology, his incorrect choice may carry over. This happens because individuals cannot identify who is the first agent in their sample. If agents knew the position of those observed, social learning would occur.

a sample from past behavior. Third, we present a model where information aggregation can be studied for any case of position uncertainty.

In the next section we explain the model and discuss the implications of stationary sampling. In Section 3, we present our results on social learning in the case that agents are ex-ante symmetric and play a symmetric equilibrium. Section 4 generalizes these findings to the asymmetric cases. We also extend our results to a non-stationary sampling rule: geometric sampling. Section 5 concludes.

2. Model

There are T players, indexed by i . Agents are exogenously ordered as follows. Let $p : \{1, \dots, T\} \rightarrow \{1, \dots, T\}$ be a permutation and let \mathcal{P} be the set of all possible permutations. Then, P is a random variable with realizations $p \in \mathcal{P}$. The random variable $P(i)$ specifies the period in which player i is asked to play.

DEFINITION 1. SYMMETRIC POSITION BELIEFS. *Players have **symmetric position beliefs** if P has a uniform random distribution, that is $\Pr(P = p) = \frac{1}{T!}$ for all $p \in \mathcal{P}$.*

Initially, we assume players have symmetric position beliefs. In Section 4 we show that our results also hold under general position beliefs.

Agents know the length T of the sequence but may not know their position in it. We are interested in how information is transmitted and aggregated in a large society with diffuse decentralized information, so we study the behavior of agents as T approaches infinity.

All agents choose one of two technologies: $a \in \mathcal{A} = \{0, 1\}$. There are two states of the world: $\theta \in \Theta = \{0, 1\}$. The timing of the game is as follows. First, $\theta \in \Theta$ and $p \in \mathcal{P}$ are chosen. The true state of the world and the agents' order in the sequence are not directly revealed to agents. Instead, each individual i receives a noisy signal $Z_{p(i)}$ about the true state of the world, a second signal $S_{p(i)}$ that may include information about his position, and a sample $\xi_{p(i)}$ of the decisions of agents before him. With these three sources of information, each agent decides between technologies 0 and 1, collects payoffs, and dies. Thus, the decision of each individual is once-and-for-all.

2.1 Payoffs

We study situations with no payoff externalities. Let $u(a, \theta)$ be the payoff from choosing technology a when the state of the world is θ . We assume that the optimal technology depends on the

state of the world; that is, $u(0,0) > u(1,0)$ and $u(1,1) > u(0,1)$. We show in Appendix A.2 that under these assumptions, any model is equivalent to a case with $\Pr(\theta = 1) = \Pr(\theta = 0) = \frac{1}{2}$ and payoffs as given in Table 1, for some $\lambda > 0$.

		True State (θ)	
		0	1
Technology (a)	0	λ	0
	1	0	1

Table 1: Payoff Function

2.2 Minimum Information Structure

We describe the information available to agents in two steps. First, we present some minimum restrictions on the private signal Z and the sample ξ that guarantee social learning. Next, we describe the additional signal S , which includes other information about the game that agents may have and that does not disrupt learning.

2.2.1 Private Signals about the State of the World

Agent i receives a private signal $Z_{P(i)}$, with realizations $z \in \mathcal{Z}$. Conditional on the true state of the world, signals are i.i.d. across individuals and distributed according to μ_1 if $\theta = 1$ or μ_0 if $\theta = 0$. We assume μ_0 and μ_1 are mutually absolutely continuous. Then, no perfectly-revealing signals occur with positive probability, and the following likelihood ratio (Radon-Nikodym derivative) exists:

$$l(z) \equiv \frac{d\mu_1}{d\mu_0}(z)$$

An agent's behavior depends on the signal Z only through the likelihood ratio l . Given ξ and S , individuals that choose technology 1 are those that receive a likelihood ratio greater than some threshold. For this reason, it is of special interest to define a distribution function G_θ for this likelihood ratio: $G_\theta(l) \equiv \Pr(l(Z) \leq l \mid \theta)$. We define signal strength as follows.

DEFINITION 2. UNBOUNDED SIGNAL STRENGTH. *Signal strength is **unbounded** if for all likelihood ratios $l \in (0, \infty)$,*

$$0 < G_0(l) < 1.$$

Since μ_0 and μ_1 are mutually absolutely continuous, the support of G_0 and G_1 has to coincide,

and so the previous definition also holds for G_1 . Let $\text{supp}(G)$ be the support of both G_0 and G_1 . If Definition 2 does not hold, we assume that the convex hull of $\text{supp}(G)$ is given by $\text{co}(\text{supp}(G)) = [\underline{l}, \bar{l}]$, with both $\underline{l} > 0$ and $\bar{l} < \infty$, and we call this the case of bounded signal strength. We study this case in Section 3.2.⁶

2.2.2 The Sample

Let $a_t \in A$ be the behavior of the agent playing in period t . The history of past behavior at period t is defined by

$$h_t = (\dots, a_{-1}, a_0, a_1, a_2, \dots, a_{t-1}).$$

Let H_t be the (random) history at time t , with realizations $h_t \in \mathcal{H}_t$. Actions a_t for $t \leq 0$ are not chosen strategically but are instead specified exogenously by an arbitrary distribution over a set \mathcal{H}_1 of initial histories.⁷ For example, the sample of the first agent in the sequence is always exogenously specified by this distribution. Conditional on θ , H_1 is independent of P . The results we present in this paper do not place any other restrictions on the distribution of the actions of early agents.

Agents observe the actions of a sample of individuals from previous periods. The random set O_t , which takes realizations $o_t \in \mathcal{O}_t$, lists the positions of the agents observed by the individual in position t . For example, if the individual in position $t = 5$ observes the actions of individuals in periods 2 and 3 then $o_5 = \{2, 3\}$. We assume O_t is nonempty⁸ and independent of the other random variables, except possibly Z_τ and S_τ for $\tau > t$.⁹ We also assume sampling to be stationary. We describe this assumption later.

At a minimum, we assume that each agent observes an unordered sample of past play, fully

⁶We assume that $\underline{l} < \lambda < \bar{l}$, and so we disregard cases where one action is dominant if the only source of information is the signal. Also, there are intermediate cases, where the bound is only in one side. They do not add much to the understanding of the problem, so we only mention them after presenting the results from the main two cases.

⁷Restaurants, hotels, clubs, and other establishments change hands occasionally but agents are not always aware of this. At the same time, changes in characteristics of the options may imply significant variations in consumers' utility but may not always be observed. Consequently, the arbitrary distribution over initial histories may result from agents $t \leq 0$ playing the same game when payoffs were different.

⁸An agent who plays without observing the actions of anyone else chooses the inferior technology with positive probability. Thus, if agents observe empty samples with positive probability, then there is a positive fraction of agents who choose the inferior technology and so social learning does not occur.

⁹In detail, O_t is independent of θ , P , and its concurrent signal Z_t . Moreover, the set O_t of sampled individuals must be independent of the actual decisions of those individuals. To accomplish this, we assume that O_t is independent of H_1 , independent of Z_τ and S_τ for all $\tau < t$, and independent of O_τ for all $\tau \neq t$.

determined by H_t and O_t . Formally, the sample $\zeta : \mathcal{O}_t \times \mathcal{H}_t \rightarrow \Xi = \mathbb{N}^2$ is defined by

$$\zeta(o_t, h_t) \equiv \left(|o_t|, \sum_{\tau \in o_t} a_\tau \right), \text{ where } |o_t| \text{ is the number of elements in } o_t.$$

Therefore, $\zeta(o_t, h_t)$ specifies the sample size and the number of observed agents that chose technology 1.¹⁰

Agents know the sample size and how many agents in the sample chose each technology. Choosing technology 1 with a probability equal to the proportion of agents choosing technology 1 in the sample is equivalent to picking an agent at random and following his behavior. Such a behavior leads, in expectation, to a utility equal to the expected utility of those observed. In other words, this minimum requirement on the information in the sample guarantees that agents can mimic those observed and obtain at least their average expected utility. It is worth noting, however, that agents may have more information about the positions of observed agents. This information is included in the second signal S , as explained in the next subsection.

We impose restrictions on the stochastic set of observed agents O_t . To do so, we first define the expected weight of agent τ on agent t 's sample by $w_{t,\tau} \equiv E \left[\frac{\mathbb{1}\{\tau \in O_t\}}{|O_t|} \right]$. These weights play a role in the present model because if agents indeed pick a random agent from the sample and follow his strategy, the ex-ante probability agent t picks agent τ is given by $w_{t,\tau}$. Now, different distributions for O_t induce different weights $w_{t,\tau}$. We restrict the sampling rule by imposing restrictions directly on the weights. We suppose that weights only depend on relative positions.

DEFINITION 3. STATIONARY SAMPLING. *A sampling rule is stationary if the weight of agent τ on agent t 's sample depends only on the distance between these agents:*

$$w_{t,\tau} = w(t - \tau) \quad \text{for all } t \text{ and all } \tau.$$

The weights $w_{t,\tau}$ depend only on the distance between agents, and so no individual plays a predominant role in the game. A rich family of such rules can be obtained by assuming that for all t, t' and (t_1, t_2, \dots, t_n) , the following condition holds:

$$\Pr(O_{t'} = \{t' - t_1, t' - t_2, \dots, t' - t_n\}) = \Pr(O_t = \{t - t_1, t - t_2, \dots, t - t_n\}). \quad (1)$$

¹⁰For example, if $t = 5, h_5 = (\dots, a_1 = 1, a_2 = 1, a_3 = 0, a_4 = 0)$ and $o_5 = \{2, 3\}$ then $\zeta_5(o_5, h_5) = (2, 1)$; that is, the agent in position 5 knows he observed two agents and that only one of them chose technology 1.

This condition imposes a stronger restriction than $w_{t,\tau} = w(t - \tau)$ since it requires that only relative positions determine the likelihood of samples.¹¹ Some examples of stationary sampling rules include when 1) agent t observes the set of his M predecessors, or when 2) agent t observes a uniformly random agent from the set of his M predecessors, or finally when 3) agent t observes a uniformly random subset K of the set of his M predecessors.

Section 4.3 studies an example of non-stationary sampling rules, geometric sampling, which includes the case when agents draw uniformly from all those who have already played.

2.3 Additional Information: The Signal S

The previous subsections presented the minimum information requirements of our model. Together with stationary sampling, the assumptions on the signal Z and the sample ξ guarantee social learning, as we show later. The random variables Z and ξ are observable, whereas P , O , H and of course θ are not. Our model, however, allows agents to possess more information about the game they play. We add a second observable signal. To allow for more general information frameworks, each agent receives a (second) private signal $S_{P(i)}$, with realizations $s \in \mathcal{S}$. Signal S may provide information about the unobserved random variables, so it may depend on θ , P , O_t and H_t . We require that conditional on θ , S_t be independent of H_1 and of Z_τ for all τ . To simplify our proof of existence, we also assume that \mathcal{S} is finite for any length T of the sequence.

To illustrate how S can provide additional information, assume that agents know their position. In such a case, $\mathcal{S} = \{1, \dots, T\}$ and $S_t = t$. This setting corresponds to the model of Smith and Sørensen [2008] where agents know their position and observe an unordered sample of past behavior. Our setup can also accommodate cases in which agents have no information about their position, but have some information about the sample, which may be ordered.¹²

If agents have imperfect information about their position in the sequence, then S specifies a distribution over positions. Moreover, S may provide information on several unobservable variables at the same time. In Section 4 we show that S plays a critical role when allowing for general position beliefs. To summarize, the signal S represents all information agents have that is not cap-

¹¹The following example shows a stationary sampling rule that violates equation (1). Every individual in an even position t observes two agents, those in positions $t - 1$ and $t - 2$. Every individual in an odd position t observes *one* of two agents, either the agent in position $t - 1$ or that in position $t - 2$. Weights are $w_{t,\tau} = \frac{1}{2}$ if $\tau \in \{t - 1, t - 2\}$ and zero otherwise, so sampling is stationary.

¹²Take the example presented in the previous subsection, with $t = 5$, $h_5 = (\dots, a_1 = 1, a_2 = 1, a_3 = 0, a_4 = 0)$ and $o_5 = \{2, 3\}$, which leads to $\xi_5(o_5, h_5) = (2, 1)$. If samples are ordered, then $s_5 = (1, 0)$, denoting that the first agent in the sample chose technology 1 and the second one chose technology 0.

tured by Z or ζ . The results we present in this paper do not depend on the distribution of S . In particular, social learning does not depend on the information agents possess about their position.

2.4 Strategies and Equilibrium

All information available to an agent is summarized by $I = (z, \zeta, s)$, which is an element of $\mathcal{I} = \mathcal{Z} \times \Xi \times \mathcal{S}$. Agent i 's strategy is a function $\sigma_i : \mathcal{I} \rightarrow [0, 1]$ that specifies a probability $\sigma_i(I)$ for choosing technology 1 given the information available. Let σ_{-i} be the strategies for all players other than i . Then the profile of play is given by $\sigma = (\sigma_i, \sigma_{-i})$.

The random variables in this model are

$$\Omega(T) = \left(\theta, P, \{O_t\}_{t=1}^T, \{Z_t\}_{t=1}^T, \{S_t\}_{t=1}^T, H_1 \right).$$

We focus on properties of the game *before* both the order of agents and the state of the world are determined. We define the *ex-ante utility* as the expected utility of a player.

DEFINITION 4. EX-ANTE UTILITY. *The ex-ante utility u_i is given by*

$$u_i(\sigma) \equiv \frac{1}{2} \sum_{\theta} \Pr \left(a_{P(i)}(\sigma) = \theta \mid \theta \right) (\lambda(1 - \theta) + \theta).$$

Profile $\sigma^* = (\sigma_i^*, \sigma_{-i}^*)$ is a Bayes-Nash equilibrium of the game if:

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*) \quad \text{for all } \sigma_i \text{ and for all } i.$$

Using a fixed point argument, we show that equilibria exist, for any length T of the game.

PROPOSITION 1. EXISTENCE. *For each T there exists an equilibrium $\sigma^*(T)$.*

See Appendix A.3 for the proof.

When beliefs are symmetric, agents are ex-ante homogeneous; the information they obtain depends on their position $P(i)$ but not on their identity i . To simplify the analysis of the model with symmetric beliefs, we focus on strategies where the agent in a given position behaves in the same way regardless of his identity.

A profile of play σ is symmetric if $\sigma_i = \sigma_j$ for all i, j . A symmetric equilibrium is one in which the profile of play is symmetric. In the model with symmetric beliefs, we focus on symmetric equi-

libria.¹³ In general, multiple equilibria may arise, and some may be asymmetric. Appendix A.1 presents an example of such equilibria.

Under symmetric profiles of play, an agent's utility is not affected by the order of the other agents. In other words, if σ is symmetric, the utility $u_t(\sigma)$ of any agent i in position t is the same for any permutation p with $p(i) = t$. Only the uncertainty over one's own position matters, and so when beliefs and strategies are symmetric the ex-ante utility of every player can be expressed as follows:

$$u_i(\sigma) = \frac{1}{T} \sum_{t=1}^T u_t(\sigma).$$

For social learning, we require that the expected proportion of agents choosing the right technology approaches 1 as the number of players grows large. We can infer this expected proportion through the *average utility* of the agents.

DEFINITION 5. AVERAGE UTILITY. *The average utility \bar{u} is given by*

$$\bar{u}(\sigma) = \frac{1}{T} \sum_{i=1}^T u_i(\sigma).$$

The expected proportion of agents choosing the right technology approaches 1 if and only if the average utility reaches its maximum possible value, $\frac{\lambda+1}{2}$. In principle, there can be multiple equilibria for each length T of the game. We say social learning occurs in a particular sequence of equilibria, $\{\sigma^*(T)\}_{T=1}^{\infty}$, when the average utility approaches its maximum value.

DEFINITION 6. SOCIAL LEARNING. *Social learning occurs in a sequence $\{\sigma^*(T)\}_{T=1}^{\infty}$ of equilibria if*

$$\lim_{T \rightarrow \infty} \bar{u}(\sigma^*(T)) = \frac{\lambda+1}{2}.$$

Our definition of social learning also applies to the asymmetric cases studied in Section 4. In the main part of the paper we focus on the symmetric case (all players are ex-ante identical and play symmetric strategies). In such a case, all agents have the same ex-ante utility, which of course coincides with the average utility: $u_i(\sigma) = \bar{u}(\sigma)$. In this symmetric setting, social learning occurs if and only if every individual's ex-ante utility converges to its maximum value.

¹³In Section 4, investigating general position beliefs, we consider all equilibria and construct symmetric equilibria from possibly asymmetric equilibria. Thus, symmetric equilibria exist.

2.5 Average Utility of Those Observed

We present an ex-ante improvement principle in Section 3.1: as long as those observed choose the inferior technology with positive probability, an agent's utility can always be strictly higher than the average utility of those observed. To accomplish this, we need to define the average utility \tilde{u}_i of those that agent i observes. Let $\tilde{\xi}_{P(i)}$ denote the action of a randomly chosen agent from those sampled by agent i . Then, we define \tilde{u}_i as follows.

DEFINITION 7. AVERAGE UTILITY OF THOSE OBSERVED. *The average utility of those observed, denoted $\tilde{u}_i(\sigma_{-i})$, is given by*

$$\tilde{u}_i(\sigma_{-i}) \equiv \frac{1}{2} \sum_{\theta} \Pr \left(\tilde{\xi}_{P(i)}(\sigma_{-i}) = \theta \mid \theta \right) (\lambda(1 - \theta) + \theta).$$

The average utility \tilde{u}_i of those that agent i observes is the expected utility of a randomly chosen agent from agent i 's sample.

When i is in position t the weight $w_{t,\tau}$ represents the probability that agent τ is selected at random from agent i 's sample. Thus, under symmetry, \tilde{u}_i can be reexpressed as follows.

PROPOSITION 2. *When beliefs and strategies are symmetric, the average utility of those observed can be rewritten as¹⁴*

$$\tilde{u}_i(\sigma_{-i}) = \frac{1}{T} \sum_{t=1}^T \sum_{\tau:\tau < t} w_{t,\tau} u_{\tau}(\sigma_{-i}).$$

See Appendix A.4 for the proof. From now on, we utilize the expression for \tilde{u}_i given by Proposition 2.

Finally, we also use $\tilde{u}(\sigma)$ to denote the average of $\tilde{u}_i(\sigma_{-i})$ taken across individuals. Notice that since players are ex-ante symmetric, when beliefs and the strategy profile are symmetric, $\tilde{u}_i(\sigma_{-i}) = \tilde{u}(\sigma)$.

2.6 Consequences of Stationary Sampling

Stationary sampling has the following useful implication: the weights that one agent places on previous individuals can be translated into weights that subsequent individuals place on him. The graph on the left side of Figure 1 represents the weights $w_{t,\tau}$ agent t place on agent τ . For example, $w_{4,2}$ represents the weight agent 4 places on agent 2 and $w_{2,0}$ represents the weight agent 2 places on agent 0. These weights are highlighted in the left-hand graph of Figure 1. Since the

¹⁴We allow agents to observe individuals from the initial history, so τ can be less than 1.

distance between agents 2 and 4 is equal to the distance between agents 2 and 0, then these weights must be equal: $w_{4,2} = w_{2,0} = w(2)$. The graph on the right side of Figure 1 shows the sampling weights only as a function of the distance between agents.




Figure 1: Stationary Sampling

The horizontal ellipse in the right-hand side of Figure 1 includes the weights subsequent agents place on agent 2. The vertical dashed ellipse shows the weights agent 2 places on previous agents. Since weights are only determined by the distance between agents, the sum of the horizontal weights must equal the sum of the vertical weights.¹⁵ By definition,¹⁶ the vertical sum of weights is given by $\sum_{\tau:\tau < t} w_{t,\tau} = 1$. Thus, for any agent far enough from the end of the sequence, the sum of horizontal weights — those subsequent agents place on him — is arbitrarily close to 1. Intuitively, this means that all individuals in the sequence are “equally important” in the samples.

Since the weights subsequent agents place on any one agent add up to 1, the average observed utility ultimately weights all agents equally, and the proportion of the total weight placed on any fixed group of agents vanishes. Then, as we show next, the homogeneous role of agents

¹⁵ Algebraically, for any $J > 0$, $\sum_{\tau=t-J}^{t-1} w_{t,\tau} = \sum_{\tau=t-J}^{t-1} w(t-\tau) = \sum_{j=1}^J w(j) = \sum_{\tau=t+1}^{t+J} w_{t,\tau}$

¹⁶ $\sum_{\tau:\tau < t} w_{t,\tau} = \sum_{\tau:\tau < t} E \left[\frac{\mathbb{1}_{\{\tau \in O_t\}}}{|O_t|} \right] = E \left[\frac{1}{|O_t|} E \left[\sum_{\tau:\tau < t} \mathbb{1}_{\{\tau \in O_t\}} \mid |O_t| \right] \right] = E \left[\frac{|O_t|}{|O_t|} \right] = 1$.

under stationary sampling imposes an accounting identity: average utilities and average observed utilities get closer as the number of agents grows large. This is not an equilibrium result but a direct consequence of stationary sampling, as the following proposition shows.

PROPOSITION 3. *Let sampling be stationary, and, for each T , let $\{y_t^T\}_{t=-\infty}^T$ be an arbitrary sequence, with $0 \leq y_t^T \leq \frac{\lambda+1}{2}$. Let $\bar{y}(T) = \frac{1}{T} \sum_{t=1}^T y_t^T$ and $\tilde{y}(T) = \frac{1}{T} \sum_{t=1}^T \sum_{\tau:\tau < t} w_{t,\tau} y_\tau^T$. Then,*

$$\lim_{T \rightarrow \infty} \sup_{\{y_t^T\}_{t=-\infty}^T} |\bar{y}(T) - \tilde{y}(T)| = 0.$$

See Appendix A.5 for the proof.

We use arbitrarily specified sequences to highlight that Proposition 3 holds for all possible payoff sequences. Turning to the case of interest, let $\{\sigma(T)\}_{T=1}^\infty$ be a sequence of symmetric strategy profiles that induces average utility $\bar{u}(\sigma(T))$ and average utility of those observed $\tilde{u}(\sigma(T))$. Define the ex-ante improvement $v(T) \equiv \bar{u}(\sigma(T)) - \tilde{u}(\sigma(T))$ to be the difference between the average utility and the average utility of those observed. Proposition 3 has the following immediate corollary.

COROLLARY 1. *If beliefs and strategies are symmetric and sampling is stationary, then $\lim_{T \rightarrow \infty} v(T) = 0$.*

The ex-ante improvement $v(T)$ represents how much an agent in an unknown position expects to improve upon an agent he selects at random from his sample. Corollary 1 shows that under stationary sampling rules, the ex-ante improvement $v(T)$ vanishes as the number of agents in the sequence grows larger.

3. Social Learning

Previous results on social learning rely on agents knowing their own position. A common way to show social learning is based on an improvement principle, first introduced in Banerjee and Fudenberg [2004]. Smith and Sørensen [2008] and Acemoglu et al. [Forthcoming] adopt a similar approach. The basic idea of the improvement principle is as follows. Agents can always copy the behavior of an agent picked at random from those observed. This behavior leads to an expected utility equal to that of who is observed. Agents know their position, and thus know the expected utility of those observed. Consequently, the expected utility of an agent cannot decrease with his position. As long as the signal structure is informative enough, some agents receive signals that

convey more information than the sample. By using this information, the expected utility of an agent strictly increases with his position. This, together with a lower bound on the improvement, determines that the expected utility of agents converges to the maximum possible value; that is, agents who come later in the sequence choose the right technology with probability approaching one.

If agents do not know their position, the improvement principle as developed by Banerjee and Fudenberg [2004] cannot be used to show social learning. Agents cannot base their behavior on the expected utility of those observed, conditional on their position. The expected utility of agents need not increase over time.

We develop an improvement principle for settings with position uncertainty. Even if an agent has no information about his own position or the position of those observed, he can still copy the behavior of somebody picked at random from those observed. Moreover, from an ex-ante perspective, he can obtain an expected utility higher than that of those observed by using information from the signal. The difference in our argument is that although the agent may not improve upon those observed, conditional on his position, he can improve upon those observed unconditional on his position.

In what follows, we first explain why an ex-ante improvement principle holds in the present setting. Then, we show how this ex-ante improvement principle combines with stationary sampling to show both social learning and constrained efficient learning.

3.1 The Ex-ante Improvement Principle

Because of position uncertainty, the ex-ante improvement principle we develop does not rely on any information on positions. Let us restrict the information set in the following way. First, disregard all information contained in S . Second, pick an agent at random from those agent i observes. Let $\tilde{\xi}$ denote the action of the selected agent. The restricted information set is defined by $\tilde{I} = (z, \tilde{\xi})$ with $\tilde{I} \in \tilde{\mathcal{I}} = \mathcal{Z} \times \{0, 1\}$.

The average utility \tilde{u}_i of those observed by agent i depends only on the likelihood that those observed choose the superior technology in each state of the world. Define those likelihoods by $\pi_0 \equiv \Pr(\tilde{\xi} = 0 \mid \theta = 0)$ and $\pi_1 \equiv \Pr(\tilde{\xi} = 1 \mid \theta = 1)$. Then, copying a random agent from the sample leads to a utility $\tilde{u}_i = \frac{1}{2}(\pi_0\lambda + \pi_1)$. Moreover, the information contained in the action $\tilde{\xi}$

can be summarized by the likelihood ratio $L : \{0, 1\} \rightarrow (0, \infty)$ given by

$$L(\tilde{\xi}) \equiv \frac{\Pr(\theta = 1 \mid \tilde{\xi})}{\Pr(\theta = 0 \mid \tilde{\xi})}.$$

Then, the information provided by $\tilde{\xi}$ is captured by $L(1) = \frac{\pi_1}{1-\pi_0}$ and $L(0) = \frac{1-\pi_1}{\pi_0}$.

Based on the restricted information set \tilde{I} , agents can actually do better in expectation than those observed. Basically, if the information from the signal is more powerful than the information contained in the action $\tilde{\xi}$, agents are better off following the signal than following the sample. We use this to define a “smarter” strategy σ'_i that utilizes both information from $\tilde{\xi}$ and from the signal. If signal strength is unbounded and observed agents do not always choose the superior technology, the “smarter” strategy does strictly better than simply copying a random agent from the sample. In fact, fix a level U for the utility of observed agents. The thick line in Figure 2 corresponds to combinations of π_0 and π_1 such that $\tilde{u}_i(\sigma_{-i}) = U$, and the shaded area represents combinations such that $\tilde{u}_i(\sigma_{-i}) > U$. Outside of the shaded area, the improvement upon those observed is bounded below by a positive-valued function $C(U)$.




Figure 2: Ex-ante Improvement Principle with Unbounded Signals

PROPOSITION 4. EX-ANTE IMPROVEMENT PRINCIPLE. *If signal strength is unbounded, players have symmetric beliefs and σ_{-i} is a symmetric strategy profile, then there exists a strategy σ'_i and a positive-valued function C such that for all σ_{-i}*

$$u_i(\sigma'_i, \sigma_{-i}) - \tilde{u}_i(\sigma_{-i}) \geq C(U) > 0 \quad \text{if } \tilde{u}_i(\sigma_{-i}) \leq U < \frac{\lambda + 1}{2}.$$

See Appendix A.6 for the proof.

Proposition 4 presents a strategy that allows the agent to improve upon those observed. Then, in any equilibrium $\sigma^*(T)$, he must improve at least that much. In addition, since we focus on symmetric equilibria, $u_i(\sigma^*(T)) = \bar{u}(\sigma^*(T))$ and $\tilde{u}_i(\sigma_{-i}^*(T)) = \tilde{u}(\sigma^*(T))$. With this two facts, we present the following corollary.

COROLLARY 2. *If signal strength is unbounded and beliefs are symmetric, then in any symmetric equilibrium $\sigma^*(T)$,*

$$\bar{u}(\sigma^*(T)) - \tilde{u}(\sigma^*(T)) \geq C(U) > 0 \quad \text{if } \tilde{u}(\sigma^*(T)) \leq U < \frac{\lambda + 1}{2}.$$

We now present our first result.

PROPOSITION 5. SOCIAL LEARNING UNDER SYMMETRY WITH UNBOUNDED SIGNALS. *If signal strength is unbounded, players' beliefs are symmetric and sampling is stationary, then social learning occurs in any sequence of symmetric equilibria.*

See Appendix A.7 for the proof.

Figure 3 depicts the proof of Proposition 5. On the horizontal axis is the average utility of those observed, while on the vertical axis is players' ex-ante utility. We focus on the behavior of the average utility $\bar{u}(\sigma^*(T))$ as the number of players T approaches infinity. First, agents do at least as well as those observed, so any equilibrium must correspond to a point above the 45° line. Second, because of stationary sampling, $v(T) = \bar{u}(\sigma^*(T)) - \tilde{u}(\sigma^*(T))$ must vanish as the number of players T grows large. The dashed lines in Figure 3 represent how as T grows large, equilibria must approach the 45° line. Finally, the improvement principle guarantees that $\bar{u}(\sigma^*(T)) \geq \tilde{u}(\sigma^*(T)) + C(\tilde{u}(\sigma^*(T)))$. In any equilibrium the value of $\bar{u}(\sigma^*(T))$ must be above the thick curve depicted in Figure 3. These conditions together imply that as the number of players T grows large, any sequence of equilibria must lead to a sequence of payoffs that approach the northeast corner of the square. In particular, social learning occurs: $\bar{u}(\sigma^*(T))$ approaches $\frac{\lambda+1}{2}$.

This completes our discussion of the case of unbounded signal strength. As shown in Proposition 5, social learning occurs even without any information on positions. Of course, if individuals do have some information on their own position or the position of others, they may use it. However, in any case, information gets aggregated and social learning occurs: position certainty is not a requirement for information aggregation.




Figure 3: Proof of Social Learning

3.2 Bounded Signals: Constrained Efficient Learning

So far, we have studied social learning when signal strength is unbounded. However, one can imagine settings in which this assumption may be too strong, that is, there may be a lower bound \underline{l} and an upper bound \bar{l} to the likelihood ratio $l(Z)$. If signal strength is bounded, social learning cannot be guaranteed. We can define, however, a constrained efficient learning concept that is satisfied in this setting.

Imagine there is an agent in isolation who only receives the strongest signals, those yielding likelihoods \underline{l} and \bar{l} . Under such a signal structure the likelihood ratios determine how often each signal occurs in each state. Let u_{cl} denote the *maximum utility in isolation*, that is the expected utility this agent would obtain by following his signal. We show that in the limit, the expected utility of a random agent cannot be lower than u_{cl} .¹⁷

In order to show that constrained efficient learning occurs, we present an ex-ante improvement principle when signal strength is bounded. Based on the restricted information set $\tilde{I} = \{z, \tilde{\xi}\}$ defined before, an agent follows his signal when its information is stronger than the information from $\tilde{\xi}$. An agent chooses technology 1 if $l(Z) L(\tilde{\xi}) \geq \lambda$ and technology 0 otherwise. Now, the agent improves in expectation upon those observed only if the informational content from the sample does not exceed the informational content of the *best possible* signals. In other words, if $L(1)\underline{l} \geq \lambda$ and $L(0)\bar{l} \leq \lambda$ both hold, all possible signals lead to the same outcome: copy the person

¹⁷A policy implication may be that if a regulator is no more informed than the most informed agents in the sequence, she would not be able to outperform the result of a social learning algorithm.

observed, and no improvement can be obtained using \tilde{I} . The dotted lines in Figure 4 correspond to combinations of π_0 and π_1 such that $L(1)\underline{l} = \lambda$ or $L(0)\bar{l} = \lambda$. Consequently, outside of the shaded area in Figure 4, signals can be used to improve upon those observed, as Proposition 6 shows.




Figure 4: Ex-ante Improvement Principle with Bounded Signals

PROPOSITION 6. EX-ANTE IMPROVEMENT PRINCIPLE WITH BOUNDED SIGNALS. *If signal strength is bounded, players have symmetric beliefs and σ_{-i} is a symmetric strategy profile, then there exists a strategy σ'_i and a positive-valued function C such that for all σ_{-i}*

$$u_i(\sigma'_i, \sigma_{-i}) - \tilde{u}_i(\sigma_{-i}) \geq C(U) > 0 \quad \text{if } \tilde{u}_i(\sigma_{-i}) \leq U < u_{cl}.$$

See Appendix A.8 for the proof.

We say that a model of observational learning satisfies *constrained efficient learning* if agents receive at least u_{cl} in the limit.

PROPOSITION 7. CONSTRAINED EFFICIENT LEARNING UNDER SYMMETRY. *If signal strength is bounded, players' beliefs are symmetric and sampling is stationary, then constrained efficient learning occurs in any sequence of symmetric equilibria.*

See Appendix A.9 for the proof.

The novel concept of constrained efficient learning provides a lower bound for information aggregation in contexts of bounded signal strength. This result highlights that learning is not restricted to the extreme case of unbounded signal strength. Moreover, the bound on information aggregation depends only on the most informative signals. Finally, Figure 4 can be used to explain

how the case of unbounded signal strength can be understood as a limit result of the case of bounded signal strength. If \underline{l} approaches 0 and \bar{l} approaches ∞ , the constraints represented by the dotted lines become less binding, and u_{cl} approaches $\frac{\lambda+1}{2}$.¹⁸

4. General Position Beliefs and Asymmetric Equilibria

This section extends our results to settings where individuals may have asymmetric position beliefs and to asymmetric equilibria. The generalization to asymmetric position beliefs is necessary to analyze position beliefs in their most general form. Naturally, when the players themselves are not symmetric one must also analyze asymmetric equilibria. In addition, as Appendix A.1 shows, asymmetric equilibria can arise even when players have symmetric beliefs. This section shows information aggregates in the most general setting, leading to learning.

Studying asymmetric settings creates a new difficulty: agents are not ex-ante identical. The ex-ante utility u_i and the utility of those observed \tilde{u}_i vary across individuals. The average utilities \bar{u} and \tilde{u} do not reflect the individual magnitudes. To address this difficulty, we construct an auxiliary game where agents behave in the same way as in the original game but are ex-ante identical. Since the auxiliary game is symmetric, we can use the tools developed in Section 3 to show social learning or constrained efficient learning, depending on the signal strength.

The intuition behind this construction is simple. For any game with an arbitrary distribution over permutations we can construct an analogous one in which players are ex-ante symmetric but have interim beliefs corresponding to those with an arbitrary distribution over permutations. We do this by adding information through the additional signal. Likewise, when agents are allowed to have extra information, the assumption of a symmetric equilibrium is without loss of generality. To see this, suppose that some players do not react in the same way to the same information. Then, we simply take them as different types. Since their type is part of the information they receive, they react in the same way to the same information. Before agents are assigned their type, they act symmetrically.

¹⁸The reasoning presented in this section can also be used to study cases where the signal strength is bounded only on one side. For example, if $\underline{l} = 0$ but $\bar{l} < \infty$, constrained efficient learning implies that agents always choose the right action in state of the world 0 but might not do so in state of the world 1.

4.1 Construction of Auxiliary Symmetric Game

From the primitives of the game with general position beliefs, we construct an auxiliary game with symmetric beliefs. Let Γ_1 denote the original game and Γ_2 denote the auxiliary game. We construct an alternative random ordering \tilde{P} and signal \tilde{S} for Γ_2 . The other primitives are identical across the two games.

To construct \tilde{P} we first shuffle the agents according to a uniform random ordering Q , and then order them according to P , the ordering of Γ_1 , so that $\tilde{P} = P \circ Q$. The additional set of signals in Γ_2 is formally given by $\tilde{S}_t = (P^{-1}(t), S_t)$.

As constructed, Γ_2 is an observational learning game with symmetric beliefs. First, players have symmetric position beliefs, $\Pr(\tilde{P} = \tilde{p}) = \frac{1}{T!}$. Second, conditional on θ , \tilde{S}_t is independent of Z_τ . Since all other primitives are identical to those in Γ_1 , Γ_2 is an observational learning game with symmetric beliefs, as described in Section 2.

The players of Γ_2 have the same information as those in Γ_1 . For any symmetric profile of play in Γ_2 , if agent i knows the realization of $Q(i)$, then the remaining uncertainty in the ordering is identical to that in Γ_1 . Since the profile of play is symmetric in Γ_2 , all other information contained in Q is irrelevant. Agent i in Γ_2 is told $\tilde{S}_{\tilde{p}(i)} = (Q(i), S_{\tilde{p}(i)})$, which is the same information player $Q(i)$ has in Γ_1 .

4.2 Relationship Between Γ_1 and Γ_2

The following proposition completes the description of the relationship between these games, showing that for any equilibrium of Γ_1 there is a corresponding symmetric equilibrium of Γ_2 that leads to the same average utility.

PROPOSITION 8. *For any equilibrium σ^* of Γ_1 , there exists a symmetric equilibrium $\tilde{\sigma}$ of Γ_2 such that*

$$\bar{u}(\sigma^*, \Gamma_1) = \bar{u}(\tilde{\sigma}, \Gamma_2).$$

Proof. We construct $\tilde{\sigma}$ first. Letting $j = q(i)$, notice that agent i plays in position $\tilde{p}(i) = p(q(i)) = p(j)$ and receives $p^{-1}(\tilde{p}(i)) = p^{-1}(p(j)) = j$ as part of his signal. The information that player i receives in Γ_2 is $I_i = (z_{\tilde{p}(i)}, \xi_{\tilde{p}(i)}, s_{\tilde{p}(i)}, p^{-1}(\tilde{p}(i))) = (z_{p(j)}, \xi_{p(j)}, s_{p(j)}, j)$ while a player j in Γ_1 receives $I_j = (z_{p(j)}, \xi_{p(j)}, s_{p(j)})$. Then we can construct $\tilde{\sigma}$ from σ^* by setting

$$\tilde{\sigma}_i(z, \xi, s, j) = \sigma_j^*(z, \xi, s)$$

for all i, j, z, ξ and s . Notice that $\tilde{\sigma}_i$ does not vary with i and so $\tilde{\sigma}$ is a symmetric profile of play.

Next, we show that when the two games are realized on the same probability space, the action sample paths always coincide. When realized on the same probability space, the two games have identical initial histories h_1 . Position by position, in every realization, the players use the same strategy in both games. Then, the action sample paths coincide and the sum of the utilities must also be the same: $\bar{u}(\sigma^*, \Gamma_1) = \bar{u}(\tilde{\sigma}, \Gamma_2)$.

We show next that $\tilde{\sigma}$ is an equilibrium. To do so, we argue that in Γ_2 a player cannot achieve a higher average utility than $\bar{u}(\sigma^*, \Gamma_1)$ when others follow $\tilde{\sigma}$. Suppose to the contrary that there is some σ'_i such that $u_i(\sigma'_i, \tilde{\sigma}_{-i}, \Gamma_2) > \bar{u}(\sigma^*, \Gamma_1)$ and let $u_i(\sigma'_i, \tilde{\sigma}_{-i}, \Gamma_2, j)$ be the expected payoff to this player when $q(i) = j$. Since all permutations of Q are equally likely,

$$u_i(\sigma'_i, \tilde{\sigma}_{-i}, \Gamma_2) = \frac{1}{T} \sum_{j=1}^T u_i(\sigma'_i, \tilde{\sigma}_{-i}, \Gamma_2, j) > \bar{u}(\sigma^*, \Gamma_1) = \frac{1}{T} \sum_{j=1}^T u_j(\sigma^*, \Gamma_1).$$

There must be some agent j in Γ_1 with $u_j(\sigma^*, \Gamma_1) < u_i(\sigma'_i, \tilde{\sigma}_{-i}, \Gamma_2, j)$. If agent i in Γ_2 can do better, there is some agent j in Γ_1 that is not playing a best response. Agent j in Γ_1 can copy the behavior of agent i in Γ_2 when $q(i) = j$ by playing σ''_j , given by $\sigma''_j(z, \xi, s) = \sigma'_i(z, \xi, s, j)$. In this way, agent j gets the same utility as agent i : $u_j(\sigma''_j, \sigma^*_{-j}, \Gamma_1) = u_i(\sigma'_i, \tilde{\sigma}_{-i}, \Gamma_2, j)$. To see this, notice that for each realization of p , player i of Γ_1 and player j of Γ_2 play in the same position, i.e. $\tilde{p}(i) = p(q(i)) = p(j)$. For reasons identical to those explained above, since the other players are still following $\tilde{\sigma}$ and σ^* in Γ_1 and Γ_2 respectively, the players are facing identical distributions over action histories, and since they are following the same strategy they must receive the same utility.

If $\tilde{\sigma}$ is not an equilibrium of Γ_2 , agent j has a profitable deviation in Γ_1 , so σ^* cannot be an equilibrium of Γ_1 . Equivalently, if σ^* is an equilibrium of Γ_1 then $\tilde{\sigma}$ must be an equilibrium of Γ_2 .

■

With this proposition, extending results from symmetric position uncertainty to general position uncertainty is straightforward. The average utilities under general position uncertainty must exhibit the same behavior as average utilities under symmetric position uncertainty.

PROPOSITION 9. SOCIAL LEARNING WITH UNBOUNDED SIGNALS. *If signal strength is unbounded and sampling is stationary, then social learning occurs in any sequence of equilibria.*

PROPOSITION 10. CONSTRAINED EFFICIENT LEARNING. *If signal strength is bounded and sampling is stationary, then constrained efficient learning occurs in any sequence of equilibria.*

4.3 Finite Geometric Sampling

An interesting family of non-stationary sampling rules is characterized by geometric weights, defined by

$$w_{t,\tau} = \begin{cases} \frac{\gamma-1}{\gamma^t-1} \gamma^\tau & \text{if } \tau \geq 0 \\ 0 & \text{if } \tau < 0 \end{cases} \quad \text{with } \gamma > 0 \text{ and } \gamma \neq 1.$$

In this case, agent 0 is the only non-strategic agent who is observed.¹⁹

Geometric sampling encompasses two interesting cases. When $\gamma > 1$, agents are more likely to sample from the recent play whereas when $\gamma < 1$ agents are more likely to sample from the distant past. The case $\gamma = 1$, which lies in between, corresponds to agents sampling uniformly from past play.

If $\gamma < 1$, early agents outweigh late ones, so learning cannot be guaranteed. But, when $\gamma > 1$, the sampling rule is approximately stationary and so an analog to Proposition 3 can be shown.

PROPOSITION 11. *Let sampling be stationary and let $\{y_t^T\}_{t=-\infty}^T$ be any arbitrarily specified sequence for each T , with $0 \leq y_t^T \leq \frac{\lambda+1}{2}$. Let $\bar{y}(T) = \frac{1}{T} \sum_{t=1}^T y_t^T$ and $\tilde{y}(T) = \frac{1}{T} \sum_{t=1}^T \sum_{\tau:\tau < t} w_{t,\tau} y_\tau^T$. Then,*

$$\lim_{T \rightarrow \infty} \sup_{\{y_t^T\}_{t=-\infty}^T} |\bar{y}(T) - \tilde{y}(T)| = 0.$$

See Appendix A.10 for the proof.

Proposition 11 guarantees that the results from Section 3 also hold when sampling is geometric and agents are more likely to sample from the recent past. These results hold for any value of $\gamma > 1$. Sampling uniformly from the past is a limit case; social learning holds for geometric sampling as it gets arbitrarily close to uniform, but it is not guaranteed to hold for uniform sampling. Finally, if agents are more likely to sample from the distant past, social learning need not occur.

5. Conclusion

In many real-world economic activities, each agent observes the behavior of others but does not know how many individuals have faced the same decision before him, or when those observed actually made their decisions. We present a model that allows for position uncertainty. Agents,

¹⁹For this sampling rule, we can think about agent 0 as receiving no sample. He knows he is the first agent, and so he follows his signal. Consequently, we do not need to specify a prehistory of play in this case.

exogenously ordered in a sequence, choose between two competing technologies. They receive a noisy private signal about the quality of each technology and observe a sample of past play.

We present a flexible framework for studying observational learning under position uncertainty: agents are placed in the adoption sequence according to an arbitrary distribution and receive information about their positions that can be arbitrarily specified. We focus on stationary sampling, which allows for a rich class of natural sampling rules and guarantees that no individual plays a decisive role in everyone else's sample.

We first show that even under complete position uncertainty social learning occurs. In fact, for any information on positions, under unbounded signal strength, the fraction of adopters of the superior technology goes to one as the number of agents grows.

Next, we show that information also aggregates in cases of bounded signal strength. *Constrained efficient learning* holds; individuals do at least as well as the most informed agent would do in isolation. This novel result is useful for two reasons. First, it describes a lower bound on information aggregation for all information structures. Second, social learning becomes a limit result from constrained efficient learning: as we relax the bounds on the signal strength, the lower bound on information aggregation approaches perfect information aggregation.

Finally, we discuss how position uncertainty may lead to asymmetric and multiple equilibria, even if agents are symmetric *ex-ante*. We then show how to translate any asymmetric case into a symmetric one, and so we are able to show that learning occurs for all equilibria.

Our results are driven by two factors. First, the homogeneous role of agents under stationary sampling results in a useful accounting identity: as the number of agents grows large, the difference between the *ex-ante* utility and the utility of observed agents must vanish. Second, our minimum information requirement guarantees that an *ex-ante improvement principle* holds: on average, agents must do better than those they observe.

Future work should address environments with network externalities. In some economic situations, payoffs depend both on an uncertain state of the world and on the proportion of agents choosing each technology. Agents are interested in learning about both the state of the world and the aggregate profile of play. In such situations, informational externalities get confounded with coordination motives. Agents do not know the true state of nature, so it is not obvious on what outcome they should coordinate. In addition, since they do not observe the aggregate play, even if they knew the state of nature, they would not know which action to choose. Finally, this environment is interesting because agents may take into consideration that their behavior provides others

with information. As a result, agents may change their behavior in order to influence others.

A. Proofs and Examples

A.1 Example of Multiple and Asymmetric Equilibria

We present an example with $T \geq 3$ agents. Each agent observes the behavior of his immediate predecessor. The agent in the first position knows his own position. There are two relevant agents in this example: John and Paul. They believe (correctly) that they are equally likely to be in positions 2 and 3 and know they are not placed elsewhere. The beliefs of agents in positions 4 to T do not play a role in this example and are therefore not specified.

The signal structure is simple: each agent receives one of three possible signals $Z \in \mathcal{Z} = \{0, \frac{1}{2}, 1\}$.²⁰ The distribution of the signals is given by

$$\mu_0(z) = \begin{cases} p & \text{if } z = 0 \\ q & \text{if } z = 1 \\ 1 - p - q & \text{if } z = \frac{1}{2} \end{cases} \quad \text{and} \quad \mu_1(z) = \begin{cases} p & \text{if } z = 1 \\ q & \text{if } z = 0 \\ 1 - p - q & \text{if } z = \frac{1}{2} \end{cases} \quad \text{with } p > q.$$

The behavior of the first agent does not depend on the strategies employed by other agents. He simply follows his signal when it is informative and randomizes with equal probability when he receives the uninformative signal $z = \frac{1}{2}$.²¹ The behavior of agents in positions 4 to T does not affect John and Paul. Before presenting the result, we define biased strategies σ_0 and σ_1 as follows:

$$\sigma_0(\xi, z) = \begin{cases} 0 & \text{if } \xi = 0 \\ 0 & \text{if } \xi = 1 \text{ and } z = 0 \\ 1 & \text{if } \xi = 1 \text{ and } z \in \{\frac{1}{2}, 1\} \end{cases} \quad \text{and} \quad \sigma_1(\xi, z) = \begin{cases} 1 & \text{if } \xi = 1 \\ 1 & \text{if } \xi = 0 \text{ and } z = 1 \\ 0 & \text{if } \xi = 0 \text{ and } z \in \{0, \frac{1}{2}\}. \end{cases}$$

PROPOSITION 12. MULTIPLE ASYMMETRIC EQUILIBRIA. *There are three possible equilibria among John and Paul for parameters $p > q$ and $(p - q)(2 + p^2 + q^2) - 3(p^2 + q^2) \leq 0$. First, there is a*

²⁰We include only three possible signals for simplicity. Signals are of bounded strength, since there are finitely many realizations of the signals. The equilibria we present are strict, so a small probability of receiving arbitrarily informative signals can be added without changing any of the analysis.

²¹Having the first player randomize symmetrically simplifies the example. However, this example does not rely on the first player being indifferent. We could instead assume that there are two approximately uninformative signals that occur with equal probability. This would make the first player strictly prefer to take each action.

symmetric equilibrium in which John and Paul follow the signal when it is informative and the sample otherwise. Also, there are two asymmetric equilibria: John playing σ_0 and Paul playing σ_1 , or John playing σ_1 and Paul playing σ_0 .

In the asymmetric equilibria, the biases reinforce one another. To see this, consider the equilibrium when John plays σ_0 and Paul plays σ_1 and assume uninformative signals $z = \frac{1}{2}$ are highly unlikely. When John observes somebody choosing technology 0, he does not know whether the observed agent is the first agent or Paul. If Paul is observed, Paul himself observed a sample $\xi = 0$ and a signal $z \in \{0, \frac{1}{2}\}$. Then, disregarding the unlikely cases of uninformative signals, *both* the first agent and Paul observed signals $z = 0$. In this way, when John observes somebody choosing 0, he chooses to disregard his own signal. Note that this in contrast to the symmetric equilibrium, in which a sample never overpowers an informative signal. It is the fact that John and Paul might observe each other that allows these biases to reinforce one another, and thus the asymmetric equilibria arise.

Formally, let Paul play σ_1 . In that case, it is straightforward to show that John follows an informative signal when he observes $\xi = 1$. The condition $(p - q)(2 + p^2 + q^2) - 3(p^2 + q^2) \leq 0$, which is satisfied, for example, by $p = \frac{1}{2}$ and $q = \frac{1}{3}$, guarantees that John chooses action 0 after observing $\xi = 0$, disregarding his own signal. These two facts imply that σ_1 is a best response to σ_0 . The fact that σ_0 is a best response to σ_1 can be seen in an analogous way. Thus there are two asymmetric equilibria in which each player is playing one of the biased strategies.

A.2 General Payoffs

Each agent chooses technology 1 as long as:

$$\begin{aligned}
 & u(a = 1 | I) > u(a = 0 | I) \\
 & u(1,1) \Pr(\theta = 1 | I) + u(1,0) \Pr(\theta = 0 | I) > u(0,1) \Pr(\theta = 1 | I) + u(0,0) \Pr(\theta = 0 | I) \\
 & \frac{\Pr(\theta = 1 | I)}{\Pr(\theta = 0 | I)} > \frac{u(0,0) - u(1,0)}{u(1,1) - u(0,1)}
 \end{aligned}$$

Then, the agent chooses technology 1 only if:

$$\frac{\Pr(I | \theta = 1)}{\Pr(I | \theta = 0)} > \frac{\Pr(\theta = 1) u(0,0) - u(1,0)}{\Pr(\theta = 0) u(1,1) - u(0,1)} \equiv \lambda.$$

Note that any combination of payoffs and probabilities that result in λ lead to the same behavior. In particular, the following combination works:

$$\lambda \equiv \frac{\Pr(\theta = 1) u(0,0) - u(1,0)}{\Pr(\theta = 0) u(1,1) - u(0,1)} \equiv \frac{\frac{1}{2} \lambda - 0}{\frac{1}{2} 1 - 0}. \blacksquare$$

A.3 Equilibrium Existence

Agent i 's strategy $\sigma_i \in \Sigma_i = \prod_{(z,\xi,s) \in \mathcal{Z} \times \Xi \times \mathcal{S}} [0,1]$. We can collapse the strategy σ_i into the probability of choosing technology 1, conditional on ξ , S and θ , as follows

$$\rho_i(\xi, s, \theta) \equiv E \left[\sigma_i \left(Z_{P(i)}, \xi, s \right) \mid \theta \right].$$

Then $\rho_i : \Xi \times \mathcal{S} \times \Theta \rightarrow [0,1]$ defines a many to one mapping $\sigma_i \mapsto \rho_i$. Let ρ be the profile of such functions. Any two strategy profiles σ that lead to the same ρ give the same probability distribution over histories and also lead to the same utilities. It is without loss of generality then to consider agents choosing ρ_i directly from the feasible set R_i defined by

$$R_i = \left\{ \rho_i : \rho_i(\xi, s, \theta) = E \left[\sigma_i \left(Z_{P(i)}, \xi, s \right) \mid \theta \right] \text{ for some } \sigma_i \in \Sigma_i \right\}.$$

The advantage of dealing with ρ_i is that R_i is a subset of an Euclidean space with finite dimension $|\Xi| \cdot |\mathcal{S}| \cdot |\Theta|$. The set R_i is bounded since all elements are positive and less than one. To see that R_i is also closed, consider a sequence $\rho_i^n \in R_i$ for each n , with $\rho_i^n \rightarrow \rho_i$. For each ρ_i^n , pick a σ_i^n which yields ρ_i^n . Since Σ_i is sequentially compact in the product topology, there is a subsequence $\sigma_i^{n_m}$ with $\sigma_i^{n_m} \rightarrow \sigma_i$ pointwise. By the dominated convergence theorem,

$$\rho_i(\xi, s, \theta) = \lim_{n_m \rightarrow \infty} \rho_i^{n_m}(\xi, s, \theta) = \lim_{n_m \rightarrow \infty} E \left[\sigma_i^{n_m} \left(Z_{P(i)}, \xi, s \right) \mid \theta \right] = E \left[\sigma_i \left(Z_{P(i)}, \xi, s \right) \mid \theta \right],$$

and so R_i is closed, which implies that it is compact. Let $R = \prod_{i=1}^T R_i$. Then R is a compact set in an Euclidean space of dimension $|\Xi| \cdot |\mathcal{S}| \cdot |\Theta| \cdot T$.

Next, we rewrite the ex-ante utility of agent i as follows,

$$\begin{aligned} u_i(\rho_i, \rho_{-i}) &= \frac{1}{2} \sum_{\theta \in \Theta} \sum_{s \in \mathcal{S}} \Pr(s \mid \theta) \sum_{t=1}^T \Pr(P(i) = t \mid s, \theta) \sum_{h_t \in \mathcal{H}_t} \Pr(H_t(\rho_{-i}) = h_t \mid s, \theta) \\ &\quad \times \sum_{\xi \in \Xi} \Pr(\xi \mid h_t, s) [\theta \rho_i(\xi, s, \theta) + \lambda (1 - \theta) (1 - \rho_i(\xi, s, \theta))]. \end{aligned}$$

Utility is continuous in one's own strategy ρ_i . Next, utility only depends on the strategies of others through the distribution over histories. This distribution is continuous in ρ_{-i} . Therefore payoffs are continuous in ρ .

We define $BR_i(\rho_{-i}) = \arg \max_{\rho_i \in R_i} u_i(\rho_i, \rho_{-i})$. Since payoffs are continuous, this correspondence is u.h.c. Next, let $BR(\rho) = \prod_{i=1}^T BR_i(\rho_{-i})$, and note that $BR(\rho)$ is also u.h.c. By Kakutani's fixed point theorem, there is a $\rho^* \in R$ such that $\rho^* \in BR(\rho^*)$. Thus if each player plays a strategy σ_i^* that maps to ρ_i^* they all play a best response. Then, there exists an equilibrium σ^* of the game. ■

A.4 Average Utility of Those Observed

We present first the following proposition, which verifies that the set of sampled individuals O_t is independent of the actual decisions H_t of those individuals.

PROPOSITION 13. *The random variable O_t is independent of the history of play H_τ for all $\tau \leq t$. Similarly, the random variable Z_t is independent of the history of play H_τ for all $\tau \leq t$, conditional on θ .*

Proof. We show this by induction. First, let $h_{t+1} = h_t \oplus 1$ if $a_t = 1$ and $h_{t+1} = h_t \oplus 0$ if $a_t = 0$. By assumption, O_t is independent of H_1 for all t . Then, it suffices to show that for all $\tau < t$, if H_τ is independent of O_t , then so is $H_{\tau+1}$. To see this, note that $H_{\tau+1} = h_\tau \oplus 1$ if and only if both $H_\tau = h_\tau$ and $a_\tau = 1$. Consequently, it suffices to show that $a_\tau = 1$ is independent of O_t for all $\tau < t$. Now, note that a_τ is a function of $\theta, P, \sigma, O_\tau, Z_\tau, S_\tau$ and H_τ . Since all of them are independent of O_t , then $H_{\tau+1}$ is independent of O_t .

Next, regarding Z_t , note that it is i.i.d., conditional on θ . Following the same argument as with O_t , note that conditional on θ , Z_t is independent of all the determinants of a_τ . Consequently, if Z_t is independent of H_τ , conditional on θ , then it is also independent of $H_{\tau+1}$, conditional on θ . ■

With Proposition 13 in hand, we can simplify \tilde{u}_i . Let \tilde{O}_t denote the position of a randomly chosen agent from O_t and let $\tilde{\zeta}_t$ denote the action of that agent, that is, $\tilde{\zeta}_t = a_{\tilde{O}_t}$. Since beliefs and strategies are symmetric, then

$$\Pr\left(\tilde{\zeta}_{P(i)}(\sigma_{-i}) = \theta \mid \theta\right) = \frac{1}{T} \sum_{t=1}^T \Pr\left(\tilde{\zeta}_t(\sigma_{-i}) = \theta \mid \theta\right).$$

Then, by definition,

$$\begin{aligned}
\tilde{u}_i(\sigma_{-i}) &= \frac{1}{2} \sum_{\theta} \left[\frac{1}{T} \sum_{t=1}^T \Pr \left(\tilde{\xi}_t(\sigma_{-i}) = \theta \mid \theta \right) (\lambda(1-\theta) + \theta) \right] \\
&= \frac{1}{2} \sum_{\theta} \left[\frac{1}{T} \sum_{t=1}^T \sum_{\tau:\tau < t} \Pr \left(\tilde{O}_t(\sigma_{-i}) = \tau \mid \theta \right) \Pr \left(a_{\tau}(\sigma_{-i}) = \theta \mid \theta, \tilde{O}_t = \tau \right) (\lambda(1-\theta) + \theta) \right] \\
&= \frac{1}{2} \sum_{\theta} \left[\frac{1}{T} \sum_{t=1}^T \sum_{\tau:\tau < t} w_{t,\tau} \Pr \left(a_{\tau}(\sigma_{-i}) = \theta \mid \theta, \tilde{O}_t = \tau \right) (\lambda(1-\theta) + \theta) \right].
\end{aligned}$$

The last step holds because the distribution of \tilde{O}_t is a function only of the distribution of O_t , and O_t is independent of the state of the world and the strategy profile. In fact, $\Pr \left(\tilde{O}_t = \tau \right) = E \left[\frac{\mathbb{1}_{\{\tau \in O_t\}}}{|O_t|} \right] = w_{t,\tau}$. Then, note that the action a_{τ} of the individual in position τ is a function of Z_{τ} , S_{τ} and ξ_{τ} , with ξ_{τ} itself being a function of O_{τ} and H_{τ} . Now, for all $\tau < t$, Z_{τ} , S_{τ} , O_{τ} and H_{τ} are independent of O_t and thus also independent of \tilde{O}_t . Consequently, $\Pr \left(a_{\tau}(\sigma_{-i}) = \theta \mid \theta, \tilde{O}_t = \tau \right) = \Pr \left(a_{\tau}(\sigma_{-i}) = \theta \mid \theta \right)$. As a result,

$$\begin{aligned}
\tilde{u}_i(\sigma_{-i}) &= \frac{1}{2} \sum_{\theta} \left[\frac{1}{T} \sum_{t=1}^T \sum_{\tau:\tau < t} w_{t,\tau} \Pr \left(a_{\tau}(\sigma_{-i}) = \theta \mid \theta \right) (\lambda(1-\theta) + \theta) \right] \\
&= \frac{1}{T} \sum_{t=1}^T \sum_{\tau:\tau < t} w_{t,\tau} \left[\frac{1}{2} \sum_{\theta} \Pr \left(a_{\tau}(\sigma_{-i}) = \theta \mid \theta \right) (\lambda(1-\theta) + \theta) \right] \\
&= \frac{1}{T} \sum_{t=1}^T \sum_{\tau:\tau < t} w_{t,\tau} u_{\tau}(\sigma_{-i}) \quad \blacksquare
\end{aligned}$$

A.5 Vanishing Improvement with Stationary Sampling

We need to show that for all $\tilde{\varepsilon} > 0$ there exists a $T^* < \infty$ such that for all $T \geq T^*$, it is true that

$$\sup_{\{y_t^T\}_{t=-\infty}^T} |\bar{y}(T) - \tilde{y}(T)| \leq \tilde{\varepsilon}.$$

We define T' so that at most ε of the weight is placed outside of a range of T' agents:

$$\sum_{\tau=t+1}^{t+T'(\varepsilon)} w_{t,\tau} > 1 - \varepsilon$$

In order to see that the difference must vanish, we split the expression of interest in three parts,

$$\begin{aligned}
|\bar{y} - \tilde{y}| &= \left| \frac{1}{T} \sum_{t=1}^T y_t - \frac{1}{T} \sum_{t=1}^T \sum_{\tau:\tau < t} w_{t,\tau} y_\tau \right| = \frac{1}{T} \left| \sum_{t=1}^T y_t - \sum_{\tau=-\infty}^T \sum_{t=\tau+1}^T w_{t,\tau} y_\tau \right| \\
&= \frac{1}{T} \left| \sum_{t=1}^T y_t - \sum_{t=-\infty}^T \left(\sum_{\tau=t+1}^T w_{t,\tau} \right) y_t \right| \\
&= \frac{1}{T} \left| \sum_{t=1}^T y_t - \sum_{t=1}^T \left(\sum_{\tau=t+1}^T w_{t,\tau} \right) y_t - \sum_{t=-\infty}^0 \left(\sum_{\tau=t+1}^T w_{t,\tau} \right) y_t \right| \\
&\leq \frac{1}{T} \frac{\lambda + 1}{2} \left[\left| \sum_{t=-\infty}^0 \left(\sum_{\tau=t+1}^T w_{t,\tau} \right) \right| + \left| \sum_{t=1}^{T-T'} \left(1 - \sum_{\tau=t+1}^T w_{t,\tau} \right) \right| + \left| \sum_{t=T-T'+1}^T \left(1 - \sum_{\tau=t+1}^T w_{t,\tau} \right) \right| \right] \tag{2}
\end{aligned}$$

The first part of equation (2) represents the difference between \bar{y} and \tilde{y} given by the fact that some observed agents are non-strategic (those indexed from $-\infty$ to 0). These non-strategic agents enter in \tilde{y} but are not part of \bar{y} . As the number of agents grows large, the non-strategic agents form a smaller part of \tilde{y} . Consequently, the first part can be as small as needed. The second part of equation (2) is comprised of agents that are strategic and whose sum of weights is close enough to 1, since the number of agents is large. Finally, the third part of equation (2) includes the last agents, those that have not been observed by many agents yet. We show next that these three parts can be as small as needed. Regarding the first part,²²

$$\sum_{t=-\infty}^0 \sum_{\tau=1}^T w_{t,\tau} = \sum_{t=1}^T \sum_{\tau=-\infty}^0 w_{t,\tau} = \sum_{t=1}^{T'} \sum_{\tau=-\infty}^0 w_{t,\tau} + \sum_{t=T'+1}^T \sum_{\tau=-\infty}^0 w_{t,\tau} \leq \sum_{t=1}^{T'} 1 + \sum_{t=T'+1}^T \varepsilon < T' + T\varepsilon \tag{3}$$

Next, regarding the second part, note that $\sum_{\tau=t+1}^T w_{t,\tau} \geq \sum_{\tau=t+1}^{t+T'} w_{t,\tau} > 1 - \varepsilon$ for all $t \leq T - T'$. Consequently,

$$\sum_{t=1}^{T-T'} \left(1 - \sum_{\tau=t+1}^T w_{t,\tau} \right) < \sum_{t=1}^{T-T'} \varepsilon < T\varepsilon \tag{4}$$

Finally, regarding the third part,

$$\sum_{t=T-T'+1}^T \left(1 - \sum_{\tau=1}^T w_{t,\tau} \right) < \sum_{t=T-T'+1}^T 1 < T' \tag{5}$$

²²Note that $\sum_{\tau=t-T'}^{t+T'} w_{t,\tau} = \sum_{\tau=t+1}^{t+T'} w_{t,\tau} > 1 - \varepsilon$ and so $\sum_{\tau=t-T'}^{t+T'} w_{t,\tau} > \sum_{\tau=t-T'}^{t+T'} w_{t,\tau} + \sum_{\tau:\tau < t-T'} w_{t,\tau} - \varepsilon$. Then, $\sum_{\tau:\tau < t-T'} w_{t,\tau} < \varepsilon$.

Then, given (3), (4) and (5), we can express (2) as follows,

$$\begin{aligned} |\bar{y} - \tilde{y}| &\leq \frac{1}{T} \frac{\lambda + 1}{2} (T' + T\varepsilon + T\varepsilon + T') \\ &= \frac{\lambda + 1}{2} \left(2\frac{T'}{T} + 2\varepsilon \right) = (\lambda + 1) \left(\frac{T'}{T} + \varepsilon \right) \end{aligned}$$

Now, for all $\tilde{\varepsilon} > 0$ we can define $\varepsilon \equiv \frac{1}{2(\lambda+1)}\tilde{\varepsilon}$ and $T'' \equiv \frac{T'(\varepsilon)}{\varepsilon}$. Consequently,

$$|\tilde{y} - \bar{y}| \leq (\lambda + 1) \left(\frac{T'(\varepsilon)}{T} + \varepsilon \right) = (\lambda + 1) \left(\frac{T''}{T} \varepsilon + \varepsilon \right) = (\lambda + 1) \varepsilon \left(1 + \frac{T''}{T} \right) = \frac{\tilde{\varepsilon}}{2} \left(1 + \frac{T''}{T} \right) \leq \tilde{\varepsilon}$$

for all $T \geq T''$. ■

A.6 Ex-ante Improvement Principle with Unbounded Signals

We present first the following auxiliary proposition.

PROPOSITION 14. For all $l \in (L, \bar{l})$, $G_\theta(l)$ satisfies:

$$l > \frac{G_1(l)}{G_0(l)} \quad \text{and} \quad l < \frac{1 - G_1(l)}{1 - G_0(l)} \quad (6)$$

Moreover, if $k' \geq k$ then,

$$[1 - G_1(k)] - k[1 - G_0(k)] \geq [1 - G_1(k')] - k'[1 - G_0(k')] \quad (7)$$

$$G_0(k') - G_1(k') (k')^{-1} \geq G_0(k) - G_1(k) (k)^{-1} \quad (8)$$

The proof of equation (6) follows Lemma A.1 of [Smith and Sørensen \[2000\]](#). Let $Z(L) = \{Z \in \mathcal{Z} : l(Z) \leq L\}$. By the definition of Radon-Nikodym derivative,

$$G_1(L) = \int_{Z(L)} d\mu_1 = \int_{Z(L)} l(Z) d\mu_0 < \int_{Z(L)} L d\mu_0 = LG_0(L).$$

With respect to (7) and (8), note that:

$$\begin{aligned} [1 - G_1(k)] - [1 - G_1(k')] &= G_1(k') - G_1(k) = \int_{\{Z \in \mathcal{Z} : k \leq l(Z) \leq k'\}} d\mu_1 = \int_{\{Z \in \mathcal{Z} : k \leq l(Z) \leq k'\}} l(Z) d\mu_0 \\ &\geq k [G_0(k') - G_0(k)] = k [1 - G_0(k)] - k [1 - G_0(k')] \\ &\geq k [1 - G_0(k)] - k' [1 - G_0(k')] \quad \text{and also,} \end{aligned}$$

$$\begin{aligned}
G_0(k') - G_0(k) &= \frac{1}{k'} [k' (G_0(k') - G_0(k))] \geq \frac{1}{k'} \left[\int_{\{Z \in \mathcal{Z}: k \leq l(Z) \leq k'\}} l(Z) d\mu_0 \right] \\
&= \frac{1}{k'} \left[\int_{\{Z \in \mathcal{Z}: k \leq l(Z) \leq k'\}} d\mu_1 \right] = \frac{1}{k'} [G_1(k') - G_1(k)] \geq \frac{G_1(k')}{k'} - \frac{G_1(k)}{k}. \quad \blacksquare
\end{aligned}$$

In order to show that an ex-ante improvement principle holds, we first define a *smart* strategy σ'_i : “Follow the behavior of a random agent in the sample if the likelihood ratio from the signal lies between $[\underline{k}, \bar{k}]$. Otherwise, follow the signal.” Cutoffs \underline{k} and \bar{k} are optimal given the information provided by the only agent picked at random from the sample. Thus, $\bar{k}(\pi_0, \pi_1) = \lambda \frac{\pi_0}{1-\pi_1}$ and $\underline{k}(\pi_0, \pi_1) = \lambda \frac{1-\pi_0}{\pi_1}$.

This smart strategy has an advantage over copying a random agent in that sometimes a strong signal overrides an incorrect sample. At the same time, this also represents a disadvantage, since sometimes a strong and incorrect signal overrides a correct sample. Since O_t and Z_t are independent, the improvement Δ from following strategy σ'_i can be expressed as follows

$$\begin{aligned}
\Delta(\pi_0, \pi_1) &\equiv u_i(\sigma'_i, \sigma_{-i}) - \tilde{u}_i(\sigma_{-i}) \\
&= \Pr(\theta = 1)(1 - \pi_1) \Pr(l \geq \bar{k} \mid \theta = 1) + \Pr(\theta = 0)(1 - \pi_0) \Pr(l \leq \underline{k} \mid \theta = 0) \lambda \\
&\quad - \Pr(\theta = 1) \pi_1 \Pr(l \leq \underline{k} \mid \theta = 1) - \Pr(\theta = 0) \pi_0 \Pr(l \geq \bar{k} \mid \theta = 0) \lambda \\
&= \frac{1}{2} \left[(1 - \pi_1) \left[1 - G_1(\bar{k}) \right] + (1 - \pi_0) G_0(\underline{k}) \lambda - \pi_1 G_1(\underline{k}) - \pi_0 \left[1 - G_0(\bar{k}) \right] \lambda \right] \\
&= \frac{1}{2} \left\{ (1 - \pi_0) \lambda \left[G_0(\underline{k}) - (\underline{k})^{-1} G_1(\underline{k}) \right] + (1 - \pi_1) \left[\left[1 - G_1(\bar{k}) \right] - \bar{k} \left[1 - G_0(\bar{k}) \right] \right] \right\}
\end{aligned}$$

The lower bound on $\Delta(\pi_0, \pi_1)$ is constructed as follows. Pick any $U < \frac{\lambda+1}{2}$ and define $\pi_0^* = \pi_1^* = \left(\frac{\lambda+1}{2}\right)^{-1} U$, the sample distribution where the other agents perform equally well in both states and corresponds to an average observed utility of U . Let $\bar{k}^* = \bar{k}(\pi_0^*, \pi_1^*)$ and $\underline{k}^* = \underline{k}(\pi_0^*, \pi_1^*)$ and define the lower bound on the improvement as follows:

$$\begin{aligned}
C(U) &\equiv \frac{1}{2} \min \left\{ (1 - \pi_0^*) \lambda \left[G_0(\underline{k}^*) - (\underline{k}^*)^{-1} G_1(\underline{k}^*) \right]; \right. \\
&\quad \left. (1 - \pi_1^*) \left[\left[1 - G_1(\bar{k}^*) \right] - \bar{k}^* \left[1 - G_0(\bar{k}^*) \right] \right] \right\} \quad (9)
\end{aligned}$$

As a direct implication of Proposition 14, $C(U)$ is strictly positive. Next, note that a bigger U leads to a higher π_0^* , a higher π_1^* , a higher \bar{k}^* and a lower \underline{k}^* . Consequently, by equations (7) and (8), $C(U)$ is decreasing in U .

Next, pick any sample distribution (π_0, π_1) , corresponding to $\tilde{u}(\sigma_{-i}) = U$, that is $\frac{\lambda\pi_0 + \pi_1}{2} = U$. There are two possible cases. The first is when $\pi_0 < \pi_0^*$. Then $\underline{k}(\pi_0, \pi_1) > \underline{k}^*$.²³ Consequently, by (8),

$$\begin{aligned} (1 - \pi_0) \lambda \left[G_0(\underline{k}(\pi_0, \pi_1)) - (\underline{k}(\pi_0, \pi_1))^{-1} G_1(\underline{k}(\pi_0, \pi_1)) \right] \\ > (1 - \pi_0^*) \lambda \left[G_0(\underline{k}^*) - (\underline{k}^*)^{-1} G_1(\underline{k}^*) \right] \end{aligned}$$

The second is when $\pi_0 > \pi_0^*$. Then $\pi_1 < \pi_1^*$ and $\bar{k}(\pi_0, \pi_1) < \bar{k}^*$. Consequently, by (7),

$$\begin{aligned} (1 - \pi_1) \left[1 - G_1(\bar{k}(\pi_0, \pi_1)) \right] - \bar{k}(\pi_0, \pi_1) \left[1 - G_0(\bar{k}(\pi_0, \pi_1)) \right] \\ > (1 - \pi_1^*) \left[1 - G_1(\bar{k}^*) \right] - \bar{k}^* \left[1 - G_0(\bar{k}^*) \right] \end{aligned}$$

Then, $\Delta(\pi_0, \pi_1) \equiv u_i(\sigma'_i, \sigma_{-i}) - \tilde{u}_i(\sigma_{-i}) \geq C(U)$.

A.7 Social Learning with Unbounded Signals

First, fix $U < \frac{\lambda+1}{2}$. Corollary 1 states that $\lim_{T \rightarrow \infty} v(T) = 0$. Let $T(U)$ be such that $v(T) < C(U)$ for all $T > T(U)$. By Corollary 2, this implies $\tilde{u}(\sigma^*(T)) > U$. Consequently, $\tilde{u}(\sigma^*(T)) \rightarrow \frac{\lambda+1}{2}$. Finally, since $\bar{u}(\sigma^*(T)) \geq \tilde{u}(\sigma^*(T))$, social learning must occur. ■

A.8 Ex-ante Improvement Principle with Bounded Signals

Assume that the agent follows the smart strategy σ'_i defined in Appendix A.6 for the case of unbounded signal strength. Figure 5 presents the case of bounded signal strength. The shaded area in the top-right corner corresponds to combinations such that no improvement is possible with strategy σ'_i . The combination $(\widehat{\pi}_0, \widehat{\pi}_1) = \left(\frac{\bar{l} \lambda - \bar{l}}{\lambda \bar{l} - \bar{l}}, \frac{\bar{l} - \lambda}{\bar{l} - \bar{l}} \right)$ yields the lowest possible utility in that area, u_{cl} .²⁴ To construct the lower bound on $\Delta(\pi_0, \pi_1)$, pick any $U < u_{cl}$. Let $(\pi_0^*, \pi_1^*) = \left(\widehat{\pi}_0 \frac{U}{u_{cl}}, \widehat{\pi}_1 \frac{U}{u_{cl}} \right)$ be the only combination that 1) lies on the straight line that links $(0, 0)$ and $(\widehat{\pi}_0, \widehat{\pi}_1)$ and 2) yields an average expected utility of U . Let $\underline{k}^* = \underline{k}(\pi_0^*, \pi_1^*)$ and $\bar{k}^* = \bar{k}(\pi_0^*, \pi_1^*)$. We show next that $\underline{k}^* > \underline{l}$

²³ To see that $\underline{k}(\pi_0, \pi_1) > \underline{k}^*$, note that an isutility line is characterized by $\pi_0 = \lambda^{-1}(2U - \pi_1)$. Then, along the isutility line, $\underline{k}(\pi_0, \pi_1) = \lambda \frac{1 - \lambda^{-1}(2U - \pi_1)}{\pi_1} = \frac{\lambda - (2U - \pi_1)}{\pi_1} = 1 - \frac{2U - \lambda}{\pi_1}$. Finally, note $U > \frac{\lambda}{2}$ (otherwise, those observed are doing worse than following no information at all) and so an increase in π_1 leads to an increase in \underline{k} .

²⁴ With $\bar{l} < \lambda < \bar{l}$, the minimum utility such that no improvement is possible is attained at the intersection of conditions $L(0)\bar{l} = \lambda$ and $L(1)\underline{l} = \lambda$.

and $\bar{k}^* < \bar{l}$. Since $U < u_{cl}$, then

$$\begin{aligned} \underline{k}^* &= \lambda \frac{1 - \widehat{\pi}_0 \frac{U}{u_{cl}}}{\widehat{\pi}_1 \frac{U}{u_{cl}}} = \lambda \frac{\frac{1}{U} - \widehat{\pi}_0 \frac{1}{u_{cl}}}{\widehat{\pi}_1 \frac{1}{u_{cl}}} > \lambda \frac{\frac{1}{u_{cl}} - \widehat{\pi}_0 \frac{1}{u_{cl}}}{\widehat{\pi}_1 \frac{1}{u_{cl}}} = \lambda \frac{1 - \widehat{\pi}_0}{\widehat{\pi}_1} = \underline{l} \quad \text{and} \\ \bar{k}^* &= \lambda \frac{\widehat{\pi}_0 \frac{U}{u_{cl}}}{1 - \widehat{\pi}_1 \frac{U}{u_{cl}}} = \lambda \frac{\widehat{\pi}_0 \frac{1}{u_{cl}}}{\frac{1}{U} - \widehat{\pi}_1 \frac{1}{u_{cl}}} < \lambda \frac{\widehat{\pi}_0 \frac{1}{u_{cl}}}{\frac{1}{u_{cl}} - \widehat{\pi}_1 \frac{1}{u_{cl}}} = \lambda \frac{\widehat{\pi}_0}{1 - \widehat{\pi}_1} = \bar{l}. \end{aligned}$$

Define, as before, the lower bound $C(U)$ on the improvement by equation (9). Since $\underline{k}^* > \underline{l}$ and $\bar{k}^* < \bar{l}$, Proposition 14 implies that $C(U)$ is strictly positive. By the same argument as in the case of unbounded signal strength, $C(U)$ is decreasing in U . Finally, equations (7) and (8) again guarantee that any sample distribution (π_0, π_1) corresponding to $\tilde{u}(\sigma_{-i}) = U$ leads to $\Delta(\pi_0, \pi_1) \equiv u_i(\sigma'_i, \sigma_{-i}) - \tilde{u}_i(\sigma_{-i}) \geq C(U)$.




Figure 5: Bound on the Improvement

A.9 Constrained Efficient Learning with Bounded Signals

First, fix $U < u_{cl}$. Corollary 1 states that $\lim_{T \rightarrow \infty} v(T) = 0$. Let $T(U)$ be such that $v(T) < C(U)$ for all $T > T(U)$. By Proposition 6, this implies $\tilde{u}(\sigma_{-i}^*(T)) > U$. Consequently, $\tilde{u}(\sigma_{-i}^*(T)) \rightarrow u_{cl}$. Finally, since $\bar{u}(\sigma^*(T)) \geq \tilde{u}(\sigma_{-i}^*(T))$, constrained efficient learning must occur. ■

A.10 Vanishing Improvement with Geometric Sampling

We want to show that for all $\varepsilon > 0$, there exists $T(\varepsilon)$ such that $|\bar{y} - \tilde{y}| < \varepsilon$ for all $T \geq T(\varepsilon)$.

Equation (2) in Appendix A.5 shows that,

$$|\bar{y} - \tilde{y}| \leq \frac{1}{T} \frac{\lambda + 1}{2} \left[\left| \sum_{t=-\infty}^0 \left(\sum_{\tau=t+1}^T w_{t,\tau} \right) \right| + \left| \sum_{t=1}^{T-T''-1} \left(1 - \sum_{\tau=t+1}^T w_{t,\tau} \right) \right| + \left| \sum_{t=T-T''}^T \left(1 - \sum_{\tau=t+1}^T w_{t,\tau} \right) \right| \right]$$

Now, rearranging and taking into account that $w_{t,\tau} = 0$ for all $\tau < 0$, then

$$|\bar{y} - \tilde{y}| \leq \frac{1}{T} \frac{\lambda + 1}{2} \left[T' + 1 + \left| \sum_{t=T'+1}^{T-T''-1} \left(1 - \sum_{\tau=t+1}^T w_{t,\tau} \right) \right| + T'' \right]$$

Next, let $T'(\varepsilon) \equiv 1 + \frac{1}{\log(\gamma)} \max \{ -\log(\gamma - 1), -\log(\varepsilon) \}$ and let $T''(\varepsilon) \equiv \frac{-\log(\varepsilon)}{\log(\gamma)}$. Then, it can be shown that

$$\left| \left(\sum_{t=\tau+1}^T w_{t,\tau} \right) - 1 \right| < \varepsilon \quad \text{if} \quad T'(\varepsilon) < \tau < T - T''(\varepsilon)$$

and so

$$\begin{aligned} |\bar{y} - \tilde{y}| &\leq \frac{1}{T} \frac{\lambda + 1}{2} \left[T' + 1 + \left| \sum_{t=T'+1}^{T-T''-1} \varepsilon \right| + T'' \right] \\ &\leq \frac{1}{T} \frac{\lambda + 1}{2} (T' + 1 + T'' + T\varepsilon) \leq \frac{\lambda + 1}{2} \left(\frac{T' + 1}{T} + \frac{T''}{T} + \varepsilon \right) \end{aligned}$$

Now, let \tilde{T} be such that $\tilde{T} > \max \left\{ \frac{T'+1}{\varepsilon}, \frac{T''}{\varepsilon} \right\}$. Then, for all $T > \tilde{T}$,

$$|\bar{y} - \tilde{y}| \leq \frac{\lambda + 1}{2} (\varepsilon + \varepsilon + \varepsilon) = 3 \frac{\lambda + 1}{2} \varepsilon \equiv \tilde{\varepsilon} \quad \blacksquare$$

References

- ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (Forthcoming): "Bayesian Learning in Social Networks," *Review of Economic Studies*.
- BANERJEE, A. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797–817.
- BANERJEE, A. AND D. FUDENBERG (2004): "Word-of-mouth Learning," *Games and Economic Behavior*, 46, 1–22.

- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992–1026.
- CALLANDER, S. AND J. HÖRNER (2009): "The wisdom of the minority," *Journal of Economic Theory*, 144, 1421–1439.
- ÇELEN, B. AND S. KARIV (2004): "Observational Learning Under Imperfect Information," *Games and Economic Behavior*, 47, 72 – 86.
- GUARINO, A., H. HARMGART, AND S. HUCK (Forthcoming): "Aggregate information cascades," *Games and Economic Behavior*.
- HENDRICKS, K., A. SORENSEN, AND T. WISEMAN (2010): "Observational Learning and Demand for Search Goods," Working Paper.
- HERRERA, H. AND J. HÖRNER (2009): "Biased Social Learning," Working Paper.
- LARSON, N. (2008): "In with the New Or Out with the Old? Bottlenecks in Social Learning," Working Paper.
- SMITH, L. AND P. SØRENSEN (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398.
- (2008): "Rational Social Learning with Random Sampling," Working Paper.