



Observational learning with position uncertainty [☆]

Ignacio Monzón ^{a,*}, Michael Rapp ^b

^a *Collegio Carlo Alberto, Via Real Collegio 30, 10024 Moncalieri (TO), Italy*

^b *Research School of Economics, ANU College of Business and Economics, HW Arndt Building, The Australian National University, Canberra ACT 0200, Australia*

Received 25 January 2012; final version received 7 September 2013; accepted 15 September 2014

Available online 30 September 2014

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 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 provide a lower bound on information aggregation: individuals do at least as well as an agent with the strongest signal realizations would do in isolation. Finally, we show in a simple environment that position uncertainty slows down learning but not to a great extent.

© 2014 Elsevier Inc. All rights reserved.

JEL classification: C72; D83; D85

Keywords: Social learning; Complete learning; Information aggregation; Herds; Position uncertainty; Observational learning

[☆] We are grateful to Bill Sandholm for his advice, suggestions and encouragement. We also thank Dan Quint, Ricardo Serrano-Padial, Marzena Rostek, Marek Wernetka, Ben Cowan, María Eugenia Garibotti, Danqing Hu, Sang Yoon Lee, Fernando Louge and David Rivers for valuable comments, as well as the associate editor and three anonymous referees for their very helpful comments.

* Corresponding author.

E-mail addresses: ignacio@carloalberto.org (I. Monzón), michael.rapp@anu.edu.au (M. Rapp).

URLs: <http://www.carloalberto.org/people/monzon/> (I. Monzón), <http://rse.anu.edu.au/people/people.php?ID=2191> (M. Rapp).

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 smartphones 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 usually 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 smartphones 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 with a smartphone 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.

We take as our starting point the basic setup in the literature on social learning (Banerjee [2] and Bikhchandani et al. [4]) and introduce position uncertainty. In our setting, agents, exogenously ordered in a sequence, make a once-and-for-all decision between two competing technologies. The payoff from this decision is identical to all, but unknown. 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 [2], Bikhchandani et al. [4] and others 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 complete 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

¹ We study situations where network externalities do not play a salient role, including choosing between news sites, early search engines, computer software, smartphones, 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.

on the sample regardless of when they play. Heterogeneous play based on positions cannot drive complete learning in this setting.²

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 complete 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, complete learning occurs. This is due to the fact that as the number of agents in the game grows large, individuals rely less on the signal. In cases with bounded signal strength, we show agents achieve what we define as *bounded learning*: agents expect to do at least as well as somebody with the strongest signal structure consistent with the bound would do in isolation.

In the present setting, complete learning results from the combination of stationary sampling and a modified *improvement principle*. First introduced by Banerjee and Fudenberg [3], 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 [3] show complete learning must occur in their setting. Acemoglu et al. [1] and Smith and Sørensen [17] use similar arguments to show complete 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

² Position uncertainty leads to an additional difficulty. In the usual setup, the strategy of the agent in position 1 may affect the payoff of the agent in position 3, but the strategy of 3 does not affect the payoff of 1. With position uncertainty, 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.

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 complete learning if signal strength is unbounded and bounded learning if signal strength is bounded.

The result of bounded learning 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 complete learning, or perfect information aggregation. Complete learning becomes a limit result from bounded learning: as we relax the bounds on the signal strength, the lower bound on information aggregation approaches perfect information aggregation.

Once we establish that learning occurs, a question remains: how fast does information aggregate when agents do not know their positions? In general, the speed of convergence to the superior alternative depends on the signal and sampling structure, and on information about positions. We specify a simple signal structure and assume each agent observes the behavior of the preceding agent. In this simple environment, we show convergence is slower when agents do not know their positions relative to when they do, but in both cases it occurs at a polynomial rate.

1.1. Related literature

The seminal contributions to the social learning literature are Bikhchandani et al. [4] and Banerjee [2]. In these papers, 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. 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.

³ To see this, consider the following example. Signals are of unbounded strength. 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.

⁴ The 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).

Consequently, Banerjee [2] and Bikhchandani et al. [4] show that the optimal behavior of rational agents can prevent complete learning.

Smith and Sørensen [16] highlight that there is no information aggregation in the models of Banerjee [2] and Bikhchandani et al. [4] because these models assume signals are of *bounded strength*. In Smith and Sørensen [16], agents also know their positions in the adoption sequence and observe the behavior of all preceding agents. In contrast to Banerjee [2] and Bikhchandani et al. [4], 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 [16] show that if signals are of unbounded strength, complete learning occurs.⁵

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 [3], Çelen and Kariv [6], and Acemoglu et al. [1] are among them. Acemoglu et al. [1] 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, complete learning occurs.⁶

Several recent papers have highlighted the importance of position uncertainty. Smith and Sørensen [17] 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 complete learning occurs. In a context of observational learning, Costain [7] calls into question the uniqueness results of the global games literature. To do so, Costain [7] presents a model with position uncertainty but does not focus on the information aggregation dimension of the game. In Callander and Hörner [5], each agent observes the aggregate number of adopters, but does not know the order of the decisions. Callander and Hörner [5] show the counterintuitive result that it is sometimes optimal to follow the decision of a minority. Hendricks et al. [12] present a similar model and test its predictions using experimental data from an online music market. Herrera and Hörner [13] and Guarino et al. [10] 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 [13] propose a continuous time model, with agents' arrivals determined by a Poisson arrival process. They show that when signals are of unbounded strength, complete learning occurs. In Guarino et al. [10], time is discrete and agents make decisions sequentially. Guarino et al. [10] show that cascades cannot occur on the unobservable decision. Our paper differs from these recent papers in several dimensions. First, we focus on cases where both actions are observable.

⁵ Starting with Ellison and Fudenberg [8,9], other papers focus on boundedly rational agents. In Guarino and Jehiel [11] individuals observe all preceding agents and know the (expected) fraction of agents who choose the superior technology. They take their sample as coming from a binomial distribution with success probability equal to that fraction. Even with signals of unbounded strength, complete learning does not occur. Mistakes pile up fast because of agents' rule of behavior.

⁶ To see why position uncertainty matters in a setup like the one described in Acemoglu et al. [1], consider the following example. Each agent observes the agent before him and the very first agent. 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, complete learning would occur.

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

Larson [14] and Lobel et al. [15] focus on the speed of learning. Larson [14] suggests a model where agents observe a weighted average of past actions and make a choice in a continuous action space. Larson [14] 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. Lobel et al. [15] study the rate of convergence for two specific sampling rules: either agents sample their immediate predecessor or a random action from the past. Under a simple signal structure, we show that lack of information about positions slows down learning, but modestly relative to going from sampling the preceding agent to sampling a random action from the past.

In the next section we present the model. Section 3 defines complete learning, presents equilibrium existence and shows the intuition behind our main results in a simplified example. Section 4 presents our two building blocks. First, we show how under stationary sampling rules the expected utility of an observed agent must approach the average utility of agents over all positions. Second, we present our *ex-ante* improvement principle. In Section 5, using these two results, we show complete learning in the case that agents are *ex-ante* symmetric and play a symmetric equilibrium. Section 6 generalizes these findings to the asymmetric cases. Section 7 studies the speed of learning in our simplified example. Section 8 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. Assume first that each agent has equal probability of acting at any time: $\Pr(P = p) = \frac{1}{T!}$ for all $p \in \mathcal{P}$. We call this the case of *symmetric 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.

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 agents obtain a payoff of 1 when the action matches the state of the world, and of 0 otherwise. Moreover, we assume that $\Pr(\theta = 1) = \Pr(\theta = 0) = \frac{1}{2}$.⁷

⁷ Whenever the optimal technology depends on the state of the world, payoffs as presented are without loss of generality. We assume both states of the world are equally likely to simplify the exposition, but this assumption is not required for any of the results.

Agents may receive information about their position or about the position of those they observe through a (second) private signal $S_{P(i)}$, with realizations $s \in \mathcal{S}$. Initially, we assume players have symmetric position beliefs and no ex-ante information about positions. Thus, for now, $\mathcal{S} = \{\bar{s}\}$ and so $S_{P(i)} = \bar{s}$ for all i . We present our results in this environment. In Section 6, we show that our results hold also under general position beliefs and when agents receive information about positions. We describe next the conditions on the private signal Z and the sample ξ that guarantee complete learning.

2.1. Private signals about the state of the world

Agent i receives a private signal $Z_{P(i)}$, with realizations $z \in \mathcal{Z}$.⁸ 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 1 (*Unbounded Signal Strength*). Signal strength is **unbounded** if $0 < G_0(l) < 1$ for all likelihood ratios $l \in (0, \infty)$.

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 1 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 5.1.⁹

2.2. The sample

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

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

Let H_t be the (random) history at time t , with realizations $h_t \in \mathcal{H}_t$. Action a_0 is not chosen strategically but instead specified exogenously by an arbitrary distribution H_1 .¹⁰

⁸ Formally, (\mathcal{Z}, μ) is an arbitrary probability measure space.

⁹ We assume that $\underline{l} < 1 < \bar{l}$, 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.

¹⁰ 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 H_1 . Typically, one can think about agent 0 as receiving no sample. He knows he is the first agent, and so he follows his signal.

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 nonempty¹¹ and independent of other random variables. 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 determined by H_t and O_t . Formally, the sample $\xi : \mathcal{O}_t \times \mathcal{H}_t \rightarrow \mathcal{E} = \mathbb{N}^2$ is defined by

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

Therefore, $\xi(o_t, h_t)$ specifies the sample size and the number of observed agents who chose technology 1.¹² Agents may have more information about the positions of observed agents. This information is included in the second signal S , as explained in Section 6.

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 $w_{t,\tau} \equiv E[\frac{\mathbb{1}\{\tau \in O_t\}}{|O_t|}]$. These weights play a role in the present model because if agents 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, in the limit, weights only depend on relative positions.

Definition 2 (Stationary Sampling). A sampling rule is stationary if there exist limit weights $\{\bar{w}(i)\}_{i=1}^\infty$ with $\sum_{i=1}^\infty \bar{w}(i) = 1$ such that for all i : $\lim_{t \rightarrow \infty} w_{t,t-i} = \bar{w}(i)$.

The limit weights \bar{w} depend only on the distance between agents, and so no individual plays a predominant role in the game. 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.¹³ Another interesting stationary sampling rule is characterized by geometric weights: $w_{t,t-i} = \frac{\gamma-1}{\gamma^t-1} \gamma^{t-i}$ with $\gamma > 1$.

3. Existence and social learning in a nutshell

All information available to an agent is summarized by $I = (z, \xi, s)$, which is an element of $\mathcal{I} = \mathcal{X} \times \mathcal{E} \times \mathcal{S}$. 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 $\Omega(T) = (\theta, P, \{O_t\}_{t=1}^T, \{Z_t\}_{t=1}^T, \{S_t\}_{t=1}^T, H_1)$.

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.

¹¹ An agent who plays without observing others chooses the inferior technology with positive probability. Thus, whenever a positive fraction of agents observe empty samples, complete learning does not occur.

¹² For example, if $t = 5$, $h_5 = (0, 1, 1, 0, 0)$ and $o_5 = \{2, 3\}$ then $\xi_5(o_5, h_5) = (2, 1)$; that is, the agent in position 5 knows he observed two agents and that only one of them chose technology 1.

¹³ Of course, agents in positions $t < M$ sample differently, but the sampling rule is still stationary.

Definition 3 (*Ex-ante Utility*). The ex-ante utility u_i is given by

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

Profile $\sigma^* = (\sigma_i^*, \sigma_{-i}^*)$ is a Bayes–Nash equilibrium of the game if, for all σ_i and for all i , $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*)$. 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.2 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 equilibria.¹⁴ 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 $u_i(\sigma)$ 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 $u_i(\sigma) = \frac{1}{T} \sum_{t=1}^T u_t(\sigma)$.

For complete 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 4 (*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, 1. In principle, there can be multiple equilibria for each length T of the game. We say complete learning occurs in a particular sequence of equilibria, $\{\sigma^*(T)\}_{T=1}^{\infty}$, when the average utility approaches its maximum value.¹⁵

Definition 5 (*Complete Learning*). Complete learning occurs in a particular sequence of equilibria $\{\sigma^*(T)\}_{T=1}^{\infty}$ if $\lim_{T \rightarrow \infty} \bar{u}(\sigma^*(T)) = 1$.

¹⁴ In Section 6, investigating general position beliefs, we consider all equilibria and construct symmetric equilibria from possibly asymmetric equilibria. Thus, symmetric equilibria exist.

¹⁵ The sequence of games, each with length T , share weights $w_{t,\tau}$ for all $t \leq T$. Then, as $T \rightarrow \infty$, the sampling rule is fixed and we study the properties of equilibria under that sampling rule.

Our definition of complete learning also applies to the asymmetric cases studied in Section 6. 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, complete learning occurs if and only if every individual’s ex-ante utility converges to its maximum value.

We present an ex-ante improvement principle in Section 4.2: 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 6 (*Average Utility of Those Observed*). The average utility of those observed, denoted $\tilde{u}_i(\sigma_{-i})$, is given by

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

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 as*

$$\tilde{u}_i(\sigma_{-i}) = \frac{1}{T} \sum_{t=1}^T \sum_{\tau=0}^{t-1} w_{t,\tau} u_{\tau}(\sigma_{-i}).$$

See Appendix A.3 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, $\tilde{u}_i(\sigma_{-i}) = \tilde{u}(\sigma)$.

3.1. Social learning in a nutshell

Before going over our general results, we present a simple example which captures the main intuition behind these results.

Example. All T agents are equally likely to be placed in any position and have no ex-ante information about their position. Sampling follows a simple stationary rule: each agent observes the behavior of the preceding agent, except for agent 0 who observes no one. The signal structure is described by $\mu_1[(0, z)] = z^2$ and $\mu_0[(0, z)] = 2z - z^2$ with $z \in (0, 1)$. Then,

1. Given the sampling rule, the improvement vanishes: $\lim_{T \rightarrow \infty} |\bar{u}(\sigma) - \tilde{u}(\sigma)| = 0$.
2. In equilibrium agents imitate the sampled action if the signal is weak: $1 - \tilde{u}(\sigma^*) < z < \tilde{u}(\sigma^*)$. Otherwise, they follow their signal.

3. Agents improve upon those observed: $\bar{u}(\sigma^*) - \tilde{u}(\sigma^*) = (1 - \tilde{u}(\sigma^*))^2 > 0$.
4. Learning is complete: $\lim_{T \rightarrow \infty} \bar{u}(\sigma^*) = 1$.

First, note that each agent observes only his predecessor; thus logically, the average improvement is simply one in T parts of the total improvement from agent 0 to agent T : $|\bar{u}(\sigma) - \tilde{u}(\sigma)| = \frac{1}{T}|u_T(\sigma) - u_0(\sigma)|$. The improvement from agent 0 to agent T is bounded by one, thus $|\bar{u}(\sigma) - \tilde{u}(\sigma)| < \frac{1}{T}$, which leads to what we call *vanishing improvement* as $T \rightarrow \infty$. Second, because of the simple signal and sampling structure we can explicitly solve for the optimal behavior in equilibrium. The sample's informativeness is given by the likelihood of the action observed being the superior one. Moreover, since positions are unknown, the relevant likelihood is the average utility of those observed: $\tilde{u}(\sigma^*)$.¹⁶ That is why agents imitate what they observe if and only if $1 - \tilde{u}(\sigma^*) < z < \tilde{u}(\sigma^*)$. We can solve too for the expected improvement upon those observed, which is bounded away from zero as long as the average utility of those observed is not 1. Lastly, we put together the vanishing improvement and the improvement principle so we have that $(1 - \tilde{u}(\sigma^*))^2 < \frac{1}{T}$. As a result, both the average utility of those observed $\tilde{u}(\sigma^*)$ and the average utility $\bar{u}(\sigma^*)$ approach one as the number of players T grows.

4. Vanishing improvement and ex-ante improvement

Previous results on social learning rely on agents knowing their own position. If agents do not know their position, the improvement principle as developed by Banerjee and Fudenberg [3] cannot be used to show complete learning. Agents cannot base their behavior on the expected utility of those observed, conditional on their position. The example from the previous section highlights how the combination of vanishing improvement (for stationary sampling rules) and an improvement principle which does not rely on information on positions guarantee complete learning. In this section we present these two building blocks.

4.1. Vanishing improvement

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 Fig. 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 $w_{5,3}$ represents the weight agent 5 places on agent 3. As we move away from the first agents, weights $w_{t,\tau}$ approach limit weights $\bar{w}(t - \tau)$. Since the distance between agents 4 and 2 is equal to the distance between agents 5 and 3, then the limit weights are equal: $\bar{w}(2)$. The graph on the right side of Fig. 1 shows the sampling limit weights.

The horizontal ellipse in the left-hand side of Fig. 1 includes the weights subsequent agents place on agent 3. These weights approach the limit weights on the right side of Fig. 1. Since limit weights are only determined by the distance between agents, the sum of the horizontal limit weights adds up to 1. Thus, for any agent far enough from the beginning and 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.

¹⁶ Note that as $T \rightarrow \infty$ the equilibrium behavior of *all* agents limits to disregarding their private information. In contrast, if agents know their positions, the optimal strategy of an agent in a given position does not change with the total number of agents: early agents always put positive weight on their private signal.




Fig. 1. Stationary sampling.

Since the weights subsequent agents place on any one agent eventually 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 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 $\{y_t^T\}_{t=0}^T$ be an arbitrary sequence, with $0 \leq y_t^T \leq 1$. Let $\bar{y}(T) = \frac{1}{T} \sum_{t=1}^T y_t^T$ and $\tilde{y}(T) = \frac{1}{T} \sum_{t=1}^T \sum_{\tau=0}^{t-1} w_{t,\tau} y_\tau^T$. Then,*

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

See Appendix A.4 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 the sampling rule 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.




Fig. 2. Ex-ante improvement principle with unbounded signals.

4.2. The ex-ante improvement principle

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.

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 $\tilde{u}_i = \frac{1}{2}(\pi_0 + \pi_1)$. 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 Fig. 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 $C(U)$.




Fig. 3. Proof of complete learning.

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}*

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

See Appendix A.5 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 $\tilde{u}_i(\sigma^*(T)) = \tilde{u}(\sigma^*(T))$. With these two facts, we present the following corollary.

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

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

5. Social learning with symmetric strategies

We now present our main result.

Proposition 5 (*Complete Learning Under Symmetry With Unbounded Signals*). *If signal strength is unbounded, players’ beliefs are symmetric and sampling is stationary, then complete learning occurs in any sequence of symmetric equilibria.*

See Appendix A.6 for the proof.

Fig. 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 Fig. 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




Fig. 4. Ex-ante improvement principle with bounded signals.

the thick curve depicted in Fig. 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, complete learning occurs: $\bar{u}(\sigma^*(T))$ approaches 1.

This completes our discussion of the case of unbounded signal strength. As shown in Proposition 5, complete 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 complete learning occurs: position certainty is not a requirement for information aggregation.

5.1. Bounded signal strength: bounded 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, complete learning cannot be guaranteed. We can define, however, a bounded 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} .

In order to show that bounded 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 $l(Z)L(\tilde{\xi}) \geq 1$ 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 $L(1)\underline{l} \geq 1$ and $L(0)\bar{l} \leq 1$ both hold, all possible signals lead to the same outcome: copy the person observed, and no improvement can be obtained using \tilde{I} . The dotted lines in Fig. 4 correspond to combinations of π_0 and π_1 such that $L(1)\underline{l} = 1$ or $L(0)\bar{l} = 1$. Consequently, outside of the shaded area in Fig. 4, signals can be used to improve upon those observed, as Proposition 6 shows. We say that a model of observational learning satisfies bounded learning if agents receive at least u_{cl} in the limit.

Proposition 6 (*Learning Under Symmetry with Bounded Signals*). *If signal strength is bounded and players' beliefs are symmetric, then*

1. *if σ_{-i} is a symmetric strategy profile, there exists a strategy σ'_i and a positive-valued function C such that for all σ_{-i}*

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

2. *if sampling is stationary, bounded learning occurs in any sequence of symmetric equilibria.*

See Appendix A.7 for the proof.

The concept of bounded 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, Fig. 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 1.¹⁷

6. Social learning in asymmetric games

The model as presented in Section 2 provides a simple setup to analyze social learning. However, agents may possess more information than that contained in signal Z and the sample ξ . To allow for more general information frameworks, each agent receives a second private signal $S_{P(i)}$. 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 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, $S = \{1, \dots, T\}$ and $S_t = t$. This setting corresponds to the model of Smith and Sørensen [17] 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.¹⁸

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 fact, S plays a critical role when allowing for general position beliefs, as we explain next. To summarize, the signal S represents all information agents have that is not captured by Z or ξ .

As is clear by now, the results presented so far do not depend on the distribution of S . Our improvement principle provides a lower bound on the improvement, so if agents are given more information, they cannot do worse. In particular, complete learning does not depend on the information agents possess about their position.

¹⁷ The 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 $\bar{l} < \infty$, bounded 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.

¹⁸ Take the example 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.

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 (they are not equal to the average utilities \bar{u} and $\tilde{\bar{u}}$, respectively). 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 5 to show complete learning or bounded 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 the same 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.

6.1. 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 .

6.2. 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 7. 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, then

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

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, agents j and i get the same utility: $u_j(\sigma''_j, \sigma_{-j}^*, \Gamma_1) = u_i(\sigma'_i, \tilde{\sigma}_{-i}, \Gamma_2, j)$. 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 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 Proposition 7, 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 8 (*Learning in Asymmetric Games*). If sampling is stationary, then in any sequence of equilibria: 1. complete learning occurs if signal strength is unbounded and 2. bounded learning occurs if signal strength is bounded.

7. Speed of convergence under position uncertainty

Learning occurs even when agents do not know their positions, but how fast does it occur? In particular, is the speed of convergence to the superior technology affected by the lack of

T	\bar{u}^C	\bar{u}^U	\bar{u}^{RS}
10	0.894	0.885	0.846
30	0.938	0.927	0.864
100	0.971	0.955	0.882

Fig. 5. Average utilities under position certainty, position uncertainty, random sampling.

information about positions? The speed of learning depends on the sampling rule, signal distribution, and the information about positions. In this section, we fix the first two as in the example from Section 3.1: agents observe the action of the preceding agent, and receive signals from $\mu_1[(0, z)] = z^2$ and $\mu_0[(0, z)] = 2z - z^2$ with $z \in (0, 1)$. We compare how fast $\bar{u}(\sigma^*) \rightarrow 1$ when 1) agents have no information about their positions and when 2) agents know their positions perfectly. Let $\bar{u}^C(T)$ and $\bar{u}^U(T)$ denote the average utility with and without information about positions respectively. We show that convergence is slower when agents do not know their positions, but in both cases it occurs at a polynomial rate.

Proposition 9 (Rate of Convergence). $1 - \bar{u}^U(T) = \Theta(T^{-\frac{1}{2}})$ and $1 - \bar{u}^C(T) = \Theta(\log(T)T^{-1})$.

When agents do not know their positions, $(1 - \bar{u}^U(T))^2 = T^{-1}|u_T(\sigma^*) - u_0(\sigma^*)|$, as shown in Section 3.1. The utility of the last agent $u_T(\sigma^*) \rightarrow 1$, so $1 - \bar{u}^U(T) = \Theta(T^{-\frac{1}{2}})$. Additionally, $\bar{u}^U(T) - \tilde{u}^U(T) = (1 - \bar{u}^U(T))^2$, so also $1 - \bar{u}^U(T) = \Theta(T^{-\frac{1}{2}})$.

When agents have full information about positions, the informational content of the sample depends on the position. Still, the optimal decision in equilibrium is simple: imitate the sampled action if $1 - u_t(\sigma^*) < z < u_t(\sigma^*)$ and follow the signal otherwise. Thus, the improvement depends on the position: $u_{t+1}(\sigma^*) - u_t(\sigma^*) = (1 - u_t(\sigma^*))^2$. It is immediate that $1 - u_t(\sigma^*) = \Theta(t^{-1})$. Given this, it is also easy to show that $1 - \bar{u}^C(T) = \Theta(\log(T)T^{-1})$.

Lobel et al. [15] had already shown that $1 - u_t(\sigma^*) = \Theta(t^{-1})$ when agents know their positions. They study a second sampling rule: agents sample a random agent from the past. Given full information about positions and the same signal structure, they show $1 - u_t(\sigma^*) = \Theta((\log(t))^{-1})$, so one can show that $1 - \bar{u}(T)(\sigma^*) = \Theta((\log(T))^{-1})$ in that case.

Focusing only on agents' information can be misleading. Take an individual's ex-ante belief about the position of the sampled agent. First, when the predecessor is observed and agents know their positions, information is perfect. Second, with the same sampling rule but position uncertainty, all agents from 0 to $T - 1$ are equally likely to be observed. Third, with position certainty and random sampling, the last agent T has the same beliefs as all agents in the second model. Other agents actually have more precise information about who they sample. Naively, one may assume that agents in the third model do better than those in the second. However, learning in the second model is much faster than in the third. See Fig. 5 for an example.

To sum up, in the present setting no information about positions does slow down learning, but in a modest way relative to the extreme impact of going from sampling the preceding agent to sampling a random action from the past.

8. 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,

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, complete 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. *Bounded learning* holds; individuals do at least as well as somebody with the strongest signal structure consistent with the bound would do in isolation. This result is useful for two reasons. First, it describes a lower bound on information aggregation for all information structures. Second, complete learning becomes a limit result from bounded learning: as we relax the bounds on the signal strength, the lower bound on information aggregation approaches perfect information aggregation.

Then, 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.

Finally, we show in a simple environment that learning is slower when agents do not know their positions relative to when they do, but occurs in both cases at a polynomial rate.

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.

Appendix 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 three¹⁹ possible signals $Z \in \mathcal{Z} = \{0, \frac{1}{2}, 1\}$, distributed as follows, with $p > q$,

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

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 $z = \frac{1}{2}$.²⁰ 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 10 (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 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 $\xi = 1$. 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.

¹⁹ We include only three possible signals for simplicity. Signals are of bounded strength, since there are finitely many realizations. The equilibria we present are strict, so a small probability of receiving arbitrarily informative signals can be added without changing any of the analysis.

²⁰ Having 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.

A.2. Equilibrium existence

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

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

Then $\rho_i : \mathcal{E} \times \mathcal{S} \times \Theta \rightarrow [0, 1]$ 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 $R_i = \{\rho_i : \rho_i(\xi, s, \theta) = E[\sigma_i(Z_{P(i)}, \xi, s) \mid \theta]\}$ for some $\sigma_i \in \Sigma_i$.

The advantage of dealing with ρ_i is that R_i is a subset of an Euclidean space with finite dimension $|\mathcal{E}| \cdot |\mathcal{S}| \cdot |\Theta|$. 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,

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

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 $|\mathcal{E}| \cdot |\mathcal{S}| \cdot |\Theta| \cdot T$.

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

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

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.²¹ \square

A.3. Average utility of those observed

First, we show by induction that O_t is independent of the history of play H_τ for all $\tau \leq t$. 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

²¹ We focus on Bayes–Nash Equilibria (BNE), but for any BNE one can construct an outcome equivalent Perfect Bayesian Equilibrium (PBE). One player's action only affects another's payoffs through the distribution of histories. Histories leading to zero probability samples have probability zero. Therefore, a player's action in response to a sample with zero probability has no impact on the distribution of histories. Then, to construct a PBE, simply take a BNE, fix any beliefs for samples received with zero probability (i.e. agents think $\theta = 0$) and set an optimal strategy according to those beliefs (agents choose $a = 0$).

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 .

With this 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 $\tilde{\xi}_t$ denote the action of that agent, that is, $\tilde{\xi}_t = a_{\tilde{O}_t}$. Since beliefs and strategies are symmetric, then

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

Then, by definition,

$$\begin{aligned} \tilde{u}_i(\sigma_{-i}) &= \frac{1}{2} \sum_{\theta} \left[\frac{1}{T} \sum_{t=1}^T \Pr(\tilde{\xi}_t(\sigma_{-i}) = \theta \mid \theta) \right] \\ &= \frac{1}{2} \sum_{\theta} \left[\frac{1}{T} \sum_{t=1}^T \sum_{\tau=0}^{t-1} \Pr(\tilde{O}_t(\sigma_{-i}) = \tau \mid \theta) \Pr(a_\tau(\sigma_{-i}) = \theta \mid \theta, \tilde{O}_t = \tau) \right] \\ &= \frac{1}{2} \sum_{\theta} \left[\frac{1}{T} \sum_{t=1}^T \sum_{\tau=0}^{t-1} w_{t,\tau} \Pr(a_\tau(\sigma_{-i}) = \theta \mid \theta, \tilde{O}_t = \tau) \right]. \end{aligned}$$

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, $\Pr(\tilde{O}_t = \tau) = E[\frac{\mathbb{1}\{\tau \in O_t\}}{|O_t|}] = w_{t,\tau}$. 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, $\Pr(a_\tau(\sigma_{-i}) = \theta \mid \theta, \tilde{O}_t = \tau) = \Pr(a_\tau(\sigma_{-i}) = \theta \mid \theta)$. As a result,

$$\begin{aligned} \tilde{u}_i(\sigma_{-i}) &= \frac{1}{2} \sum_{\theta} \left[\frac{1}{T} \sum_{t=1}^T \sum_{\tau=0}^{t-1} w_{t,\tau} \Pr(a_\tau(\sigma_{-i}) = \theta \mid \theta) \right] \\ &= \frac{1}{T} \sum_{t=1}^T \sum_{\tau=0}^{t-1} w_{t,\tau} \left[\frac{1}{2} \sum_{\theta} \Pr(a_\tau(\sigma_{-i}) = \theta \mid \theta) \right] \\ &= \frac{1}{T} \sum_{t=1}^T \sum_{\tau=0}^{t-1} w_{t,\tau} u_\tau(\sigma_{-i}) \quad \square \end{aligned}$$

A.4. Vanishing improvement with stationary sampling

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

$$\begin{aligned} \bar{y} - \tilde{y} &= \frac{1}{T} \sum_{t=1}^T y_t - \frac{1}{T} \sum_{t=1}^T \sum_{\tau=0}^{t-1} w_{t,\tau} y_\tau = \frac{1}{T} (y_T - y_0) + \frac{1}{T} \sum_{t=0}^{T-1} y_t - \frac{1}{T} \sum_{t=0}^{T-1} \sum_{\tau=t+1}^T w_{\tau,t} y_t \\ &= \frac{1}{T} (y_T - y_0) + \frac{1}{T} \sum_{t=0}^{T-1} y_t \left(1 - \sum_{\tau=t+1}^T w_{\tau,t} \right) \leq \frac{1}{T} + \frac{1}{T} \sum_{t=0}^{T-1} |y_t| \left| 1 - \sum_{\tau=t+1}^T w_{\tau,t} \right| \end{aligned}$$

$$\begin{aligned}
 &\leq \frac{1}{T} + \frac{1}{T} \sum_{t=0}^{T-1} \left| 1 - \sum_{\tau=t+1}^T w_{\tau,t} \right| \\
 &\leq \frac{1}{T} + \frac{1}{T} \sum_{t=0}^{T-1} \left| 1 - \sum_{\tau=t+1}^T \bar{w}(\tau - t) \right| + \frac{1}{T} \sum_{t=0}^{T-1} \left| \sum_{\tau=t+1}^T w_{\tau,t} - \bar{w}(\tau - t) \right| \\
 &\leq \frac{1}{T} + \frac{1}{T} \sum_{i=1}^T \left(1 - \sum_{i=1}^t \bar{w}(i) \right) + \frac{1}{T} \sum_{i=1}^T \sum_{i=1}^t |w_{t,t-i} - \bar{w}(i)| \tag{1}
 \end{aligned}$$

Formally, we need to show that for all $\varepsilon > 0$ there exists a $T^* < \infty$ such that for all $T \geq T^*$, it is true that $\sup_{\{y^T\}_{T=0}^T} |\bar{y}(T) - \tilde{y}(T)| \leq \varepsilon$. First, we define j so most of the limit weight \bar{w} is placed on the first j agents: $1 - \sum_{i=1}^j \bar{w}(i) < \frac{\varepsilon}{10}$. Second, we define $t' \geq j$ so that the first j weights are close to the limit weights: $\sum_{i=1}^j |w_{t,t-i} - \bar{w}(i)| < \frac{\varepsilon}{10}$ for all $t \geq t'$. Since the previous expression is a finite sum of terms going to zero, t' must exist. Third, we note that the remaining terms of the sum are also as small as needed: $\sum_{i=j+1}^t |w_{t,t-i} - \bar{w}(i)| \leq \frac{3}{10}\varepsilon$ for all $t \geq t'$.²²

We return now to Eq. (1),

$$\begin{aligned}
 \bar{y} - \tilde{y} &\leq \frac{1}{T} + \frac{1}{T} \sum_{i=1}^{j-1} \left(1 - \sum_{i=1}^t \bar{w}(i) \right) + \frac{1}{T} \sum_{i=j}^T \left(1 - \sum_{i=1}^t \bar{w}(i) \right) \\
 &\quad + \frac{1}{T} \sum_{t=1}^{t'-1} \sum_{i=1}^t |w_{t,t-i} - \bar{w}(i)| + \frac{1}{T} \sum_{t=t'}^T \left(\sum_{i=1}^t |w_{t,t-i} - \bar{w}(i)| \right) \\
 &\leq \frac{1}{T} + \frac{j-1}{T} + \frac{1}{T} \sum_{i=j}^T \frac{\varepsilon}{10} + \frac{1}{T} \sum_{t=1}^{t'-1} \left(\sum_{i=1}^t |w_{t,t-i}| + \sum_{i=1}^t |\bar{w}(i)| \right) + \frac{1}{T} \sum_{t=t'}^T \frac{4}{10}\varepsilon \\
 &\leq \frac{j}{T} + \frac{\varepsilon}{10} + \frac{t'(t'+1)}{T} + \frac{4}{10}\varepsilon \leq \varepsilon \quad \text{for } T \geq T^* \equiv 2(j + t'(t'+1))\varepsilon^{-1}. \quad \square
 \end{aligned}$$

A.5. Ex-ante improvement principle with unbounded signals

We present first the following auxiliary proposition.

Proposition 11. For all $l \in (\underline{l}, \bar{l})$, $G_\theta(l)$ satisfies:

$$l > \frac{G_1(l)}{G_0(l)} \quad \text{and} \quad l < \frac{1 - G_1(l)}{1 - G_0(l)} \tag{2}$$

²² To see why, note that:

$$\begin{aligned}
 \sum_{i=j+1}^t |w_{t,t-i} - \bar{w}(i)| &\leq \sum_{i=j+1}^t w_{t,t-i} + \sum_{i=j+1}^t \bar{w}(i) \leq 1 - \sum_{i=1}^j w_{t,t-i} + \frac{\varepsilon}{10} \\
 \sum_{i=j+1}^t |w_{t,t-i} - \bar{w}(i)| &\leq 1 - \sum_{i=1}^j \bar{w}(i) - \sum_{i=1}^j (w_{t,t-i} - \bar{w}(i)) + \frac{\varepsilon}{10} \leq \sum_{i=1}^j |w_{t,t-i} - \bar{w}(i)| + \frac{2}{10}\varepsilon \leq \frac{3}{10}\varepsilon.
 \end{aligned}$$

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

$$[1 - G_1(k)] - k[1 - G_0(k)] \geq [1 - G_1(k')] - k'[1 - G_0(k')] \tag{3}$$

$$G_0(k') - G_1(k')(k')^{-1} \geq G_0(k) - G_1(k)(k)^{-1} \tag{4}$$

The proof of Eq. (2) follows Lemma A.1 of Smith and Sørensen [16]. Let $Z(L) = \{Z \in \mathcal{Z} : l(Z) \leq L\}$. By the definition of Radon–Nikodym derivative,

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

With respect to (3) and (4), $[1 - G_1(k)] - [1 - G_1(k')] = G_1(k') - G_1(k)$ and,

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

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, $\bar{k}(\pi_0, \pi_1) = \frac{\pi_0}{1-\pi_1}$ and $\underline{k}(\pi_0, \pi_1) = \frac{1-\pi_0}{\pi_1}$.

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

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

The lower bound on $\Delta(\pi_0, \pi_1)$ is constructed as follows. Pick any $U < 1$ and define $\pi_0^* = \pi_1^* = U$, 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:

$$C(U) \equiv \frac{1}{2} \min\{(1 - \pi_0^*)[G_0(\underline{k}^*) - (\underline{k}^*)^{-1}G_1(\underline{k}^*)]; (1 - \pi_1^*)[[1 - G_1(\bar{k}^*)] - \bar{k}^*[1 - G_0(\bar{k}^*)]]\} \tag{5}$$

As a direct implication of Proposition 11, $C(U)$ 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 Eqs. (3) and (4), $C(U)$ is decreasing in U .

Next, pick any (π_0, π_1) , corresponding to $\tilde{u}(\sigma_{-i}) = U$, that is $\frac{\pi_0 + \pi_1}{2} = U$. There are two possible cases. First, when $\pi_0 < \pi_0^*$, then $\underline{k}(\pi_0, \pi_1) > \underline{k}^*$.²³ Consequently, by (4),

$$(1 - \pi_0)[G_0(\underline{k}(\pi_0, \pi_1)) - (\underline{k}(\pi_0, \pi_1))^{-1}G_1(\underline{k}(\pi_0, \pi_1))] > (1 - \pi_0^*)[G_0(\underline{k}^*) - (\underline{k}^*)^{-1}G_1(\underline{k}^*)]$$

The second case occurs when $\pi_0 > \pi_0^*$. Then $\pi_1 < \pi_1^*$ and $\bar{k}(\pi_0, \pi_1) < \bar{k}^*$. Then, by (3),

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

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

A.6. Complete learning with unbounded signals

First, fix $U < 1$. 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, $\tilde{u}(\sigma^*(T)) \rightarrow 1$. Finally, since $\bar{u}(\sigma^*(T)) \geq \tilde{u}(\sigma^*(T))$, complete learning must occur. □

A.7. Bounded learning with bounded signals

Assume that the agent follows the smart strategy σ'_i defined in Appendix A.5 for the case of unbounded signal strength. Fig. 6 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 $(\hat{\pi}_0, \hat{\pi}_1) = (\bar{l} \frac{1-\underline{l}}{1-\underline{l}}, \frac{\underline{l}-1}{1-\underline{l}})$ 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 $(\pi_0^*, \pi_1^*) = (\hat{\pi}_0 \frac{U}{u_{cl}}, \hat{\pi}_1 \frac{U}{u_{cl}})$ be the only combination that 1) lies on the straight line that links $(0, 0)$ and $(\hat{\pi}_0, \hat{\pi}_1)$ and 2) yields an average expected utility of U . Let $\underline{k}^* = \underline{k}(\pi_0^*, \pi_1^*)$ and

²³ To see that $\underline{k}(\pi_0, \pi_1) > \underline{k}^*$, note that an isoutility line is characterized by $\pi_0 = 2U - \pi_1$. Then, along the isoutility line, $\underline{k}(\pi_0, \pi_1) = 1 - \frac{2U-1}{\pi_1}$. Finally, note $U > \frac{1}{2}$ (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 $\underline{l} < 1 < \bar{l}$, the minimum utility such that no improvement is possible is attained at the intersection of conditions $L(0)\bar{l} = 1$ and $L(1)\underline{l} = 1$.




Fig. 6. Bound on the improvement.

$\bar{k}^* = \bar{k}(\pi_0^*, \pi_1^*)$. We show next that $\underline{k}^* > \underline{l}$ and $\bar{k}^* < \bar{l}$. Since $U < u_{cl}$, then

$$\underline{k}^* = \frac{1 - \hat{\pi}_0 \frac{U}{u_{cl}}}{\hat{\pi}_1 \frac{U}{u_{cl}}} = \frac{1 - \hat{\pi}_0 \frac{1}{u_{cl}}}{\hat{\pi}_1 \frac{1}{u_{cl}}} > \frac{1 - \hat{\pi}_0 \frac{1}{u_{cl}}}{\hat{\pi}_1 \frac{1}{u_{cl}}} = \frac{1 - \hat{\pi}_0}{\hat{\pi}_1} = \underline{l} \quad \text{and}$$

$$\bar{k}^* = \frac{\hat{\pi}_0 \frac{U}{u_{cl}}}{1 - \hat{\pi}_1 \frac{U}{u_{cl}}} = \frac{\hat{\pi}_0 \frac{1}{u_{cl}}}{\frac{1}{U} - \hat{\pi}_1 \frac{1}{u_{cl}}} < \frac{\hat{\pi}_0 \frac{1}{u_{cl}}}{\frac{1}{u_{cl}} - \hat{\pi}_1 \frac{1}{u_{cl}}} = \frac{\hat{\pi}_0}{1 - \hat{\pi}_1} = \bar{l}.$$

Define, as before, the lower bound $C(U)$ on the improvement by Eq. (5). Since $\underline{k}^* > \underline{l}$ and $\bar{k}^* < \bar{l}$, Proposition 11 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, Eqs. (3) and (4) again guarantee that any sample distribution (π_0, π_1) corresponding to $\tilde{u}(\sigma_{-i}) = U$ leads to $\Delta(\pi_0, \pi_1) \equiv u_i(\sigma'_i, \sigma_{-i}) - \tilde{u}_i(\sigma_{-i}) \geq C(U)$.

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 $T > T(U)$. This implies $\tilde{u}(\sigma_{-i}^*(T)) > U$. Consequently, $\tilde{u}(\sigma_{-i}^*(T)) \rightarrow u_{cl}$. Finally, since $\bar{u}(\sigma^*(T)) \geq \tilde{u}(\sigma_{-i}^*(T))$, bounded learning must occur. \square

References

- [1] D. Acemoglu, M.A. Dahleh, I. Lobel, A. Ozdaglar, Bayesian learning in social networks, *Rev. Econ. Stud.* 78 (2011) 1201–1236.
- [2] A. Banerjee, A simple model of Herd behavior, *Quart. J. Econ.* 107 (1992) 797–817.
- [3] A. Banerjee, D. Fudenberg, Word-of-mouth learning, *Games Econ. Behav.* 46 (2004) 1–22.
- [4] S. Bikhchandani, D. Hirshleifer, I. Welch, A theory of fads, fashion, custom, and cultural change as informational cascades, *J. Polit. Economy* 100 (1992) 992–1026.
- [5] S. Callander, J. Hörner, The wisdom of the minority, *J. Econ. Theory* 144 (2009) 1421–1439.
- [6] B. Çelen, S. Kariv, Observational learning under imperfect information, *Games Econ. Behav.* 47 (2004) 72–86.
- [7] J.S. Costain, A Herding perspective on global games and multiplicity, *J. Inst. Theoretical Econ.* 7 (2007).
- [8] G. Ellison, D. Fudenberg, Rules of thumb for social learning, *J. Polit. Economy* 101 (1993) 612–643.
- [9] G. Ellison, D. Fudenberg, Word-of-mouth communication and social learning, *Quart. J. Econ.* 110 (1995) 93–125.
- [10] A. Guarino, H. Harmgart, S. Huck, Aggregate information cascades, *Games Econ. Behav.* 73 (2011) 167–185.
- [11] A. Guarino, P. Jehiel, Social learning with Coarse inference, *Amer. Econ. J. Microecon.* 5 (2013) 147–174.
- [12] K. Hendricks, A. Sorensen, T. Wiseman, Observational learning and demand for search goods, *Amer. Econ. J. Microecon.* 4 (2012) 1–31.
- [13] H. Herrera, J. Hörner, Biased social learning, *Games Econ. Behav.* 80 (2013) 131–146.

- [14] N. Larson, Inertia in social learning from a summary statistic, Working paper, 2011.
- [15] I. Lobel, D. Acemoglu, M.A. Dahleh, A. Ozdaglar, Rate of convergence of learning in social networks, in: *Proceedings of the American Control Conference*, 2009.
- [16] L. Smith, P. Sørensen, Pathological outcomes of observational learning, *Econometrica* 68 (2000) 371–398.
- [17] L. Smith, P. Sørensen, Rational social learning with random sampling, Working paper, 2008.