

Observational learning in large anonymous games

IGNACIO MONZÓN

Collegio Carlo Alberto, Università degli Studi di Torino

I present a model of observational learning with payoff interdependence. Agents, ordered in a sequence, receive private signals about an uncertain state of the world and sample previous actions. Unlike in standard models of observational learning, an agent's payoff depends both on the state and on the actions of others. Agents want both to learn the state and to anticipate others' play. As the sample of previous actions provides information on both dimensions, standard informational externalities are confounded with payoff externalities. I show that in spite of these confounding factors, when signals are of unbounded strength, there is learning in a strong sense: agents' actions are *ex post* optimal given both the state of the world and others' actions. With bounded signals, actions approach *ex post* optimality as the signal structure becomes more informative.

KEYWORDS. Observational learning, payoff interdependence, information aggregation, position uncertainty.

JEL CLASSIFICATION. C72, D83, D85.

1. INTRODUCTION

In several economic environments, the utility of an agent is affected both by some uncertain state of the world and by the actions of others. Consider a brand new operating system, of unknown quality. Each consumer cares not only about its quality, but also about whether others will adopt it. A consumer who considers buying the new system does not know how many after him will also adopt it, and may not know exactly how many before him have already adopted it. Alternatively, consider a farmer who must decide to plant either corn or soybeans at the start of the season, but is uncertain about their relative demand by the end of the season. Even though he does not know what other farmers will choose, the choices of others affect the relative profitability of each crop: if most farmers plant corn, then the price of corn will be lower and so it is more profitable to plant soybeans. Similar stories apply to investment in assets with unknown fundamentals, voting, contributions to public goods of uncertain quality, network congestion, and many other environments.

Ignacio Monzón: ignacio@carloalberto.org

I am grateful to Bill Sandholm for his advice, suggestions, and encouragement. I thank Nageeb Ali, Andrea Gallice, Leandro Gorno, Daniel Hauser, Toomas Hinnosaar, Eeva Muring, Giorgio Martini, Alexei Parakhonyak, Lones Smith, and Aleksey Tetenov for valuable comments and suggestions. I also thank two anonymous referees for their very thoughtful reports. Some early ideas that led to this paper were present in "Social Learning in Games," the third chapter of my dissertation, which was joint work with Michael Rapp. I especially thank him.

By observing a sample of the actions of others, an agent obtains information both about the state of the world and about how others behave. The farmer deciding between crops may have private information on how the demand will be at the end of the season. He may also observe the decisions of some of his neighbors. With these two sources of information, he must form beliefs about both the future demand and about the actions of those farmers he does not observe. When a farmer observes that one of his neighbors plants corn, this might mean that his neighbor believes corn to be in high demand, as in standard models of observational learning. It may also mean that most farmers are planting corn.

I study the outcomes of observational learning in large games. In the standard setup of observational learning, (complete) learning has a simple definition: the fraction of adopters of the superior action must approach 1. When payoffs depend on others' actions, the *right* action depends not only on the state of the world, but also on what others do. I focus then on whether realized actions are ex post optimal. I say that *strategic learning* occurs when agents' actions are ex post optimal given both the state of the world and the realized actions of others. The main message of this paper, as stated formally in [Proposition 2](#), is simple: Strategic learning occurs, provided that the signal structure is sufficiently informative.

The notion of strategic learning is demanding: it requires that agents not only learn about the state of the world, but also that they correctly anticipate others' actions and best respond to them. In what follows I describe the framework and I present the intuition behind this result.

Agents are exogenously ordered in a sequence and are uncertain about their position in it. There are two, a priori equally likely, states of the world. Each agent receives a private signal about the underlying state of the world and observes the actions of some of his predecessors. Then he makes a once-and-for-all decision between two actions (0 and 1). The main innovation with respect to the standard setup is that an agent's payoff depends not only on his own action and the unknown state of the world, but also on the proportion X of agents who choose action 1. My framework applies to the examples described before (coordination games, such as the adoption of a new operating system, and anti-coordination games, such as the example of farmers). I do not impose any particular functional form on how payoffs depend on X .

Payoff interdependence adds a strategic consideration to observational learning: each agent understands that since his own action is observed by some of his successors, it partly determines *their* decisions. An agent who can affect aggregate outcomes needs to take into account the effect of his decision on others' actions. [Gallice and Monzón \(forthcoming\)](#) show that this strategic component can have a strong effect on the aggregate play when there is a finite number of agents who never make mistakes.¹ However, this should intuitively be less relevant in large games. Individual farmers do not expect to be able to affect aggregate supply, and individual consumers typically do not believe

¹In [Gallice and Monzón \(forthcoming\)](#), a finite number of agents must decide sequentially whether to contribute to a public good. Since there are no mistakes, each individual agent can determine the aggregate outcome. This allows for full contribution in equilibrium (see [Example 3](#) in [Section 4.1](#) for a detailed explanation).

that they can determine the overall adoption rate of a new operating system. In this paper, I assume that agents make mistakes with arbitrarily small probability. I show that this implies that agents cannot individually determine the aggregate play.

The intuition behind strategic learning is simple and has two components. First, although each agent could, in principle, affect aggregate outcomes, in practice there are no butterfly effects. As the number of agents grows large, each individual's action has a smaller effect on the proportion X . Second, as each agent foresees that each action has a small effect on X , he can treat the proportion X as given. Realized payoffs depend (approximately) only on the state of the world and his own action. In this sense, I translate a game of observational learning *with* payoff externalities into a game of observational learning *without* them. Then I use tools of standard observational learning to show that strategic learning occurs. I develop this intuition in detail in what follows.

The first main result ([Proposition 1](#)) shows that as the number of agents grows large, the proportion X converges to its expectation in each state of the world. This proposition addresses two challenges that result from the additional strategic factors associated with payoff externalities. First, each agent needs to anticipate how others will behave. Second, each agent may need to account for the effect of his own action on others' decisions.

I develop a novel approach to show convergence of the proportion X . If the equilibrium strategy profile were the same regardless of the number of agents, [Proposition 1](#) would be straightforward. Agents make mistakes with positive probability, so a fixed strategy profile would create an irreducible and aperiodic Markov chain over actions. Thus, a standard ergodic argument would lead to this result. However, as the number of agents grows, the game changes, so the equilibrium strategy profile varies with the number of agents. I use a coupling argument to show that any Markov chain induced by a strategy profile converges to its stationary distribution. The speed of convergence has a geometric lower bound that is independent of the particular equilibrium strategy profile. Thus, the effect of one individual's action on the proportion X wanes as the number of agents grows, even with strategy profiles that change with the number of agents. I show through this argument that the proportion X converges to its expectation. As a direct consequence, no individual agent can affect the aggregate outcome. This result holds true for all payoff specifications.

The second main result ([Proposition 2](#)) explains why strategic learning must occur in equilibrium when signals are of unbounded strength. Since the proportion X converges to its expectation in each state of the world, each agent can anticipate the payoffs he would get from each action in each state of the world. Optimality considerations limit the possible combinations of proportions X and payoffs that can occur in equilibrium. To see this, consider first a long-run outcome where in both states of the world, the payoff from choosing action 1 exceeds that from choosing action 0. Any agent who chooses action 0 regrets it *ex post*. Intuitively, an agent could instead choose action 1 always and obtain higher payoffs. It follows that no positive proportion of agents can choose a dominated action.

The final step in [Proposition 2](#) deals with long-run outcomes where agents want to choose different actions in different states of the world. I provide an improvement principle that applies to environments with payoff externalities. An individual can always

copy a random action from the sample he observes. Moreover, when his private signal is strong enough, he can go against the observed action and do (in expected terms) strictly better than the observed agent. Then, as the number of agents grows large, it must be the case that either (i) the fraction of agents who choose the superior action approaches 1 or (ii) the extra payoff from choosing the right action approaches 0. In either case, there is strategic learning.

[Proposition 2](#) provides a unique prediction of play for games with only one Nash equilibrium (e.g., an anti-coordination game). In the farmers' example, the proportion of crops planted correctly matches the demand. If instead there are several equilibria in each state of the world, [Proposition 2](#) does not select among them. I illustrate this point through a coordination game ([Example 8](#)).

Finally, I show that some degree of information aggregation also occurs with signals of bounded strength. [Lemma 7](#) presents a notion of bounded strategic learning. Although actions may be ex post suboptimal with bounded signals, there is a bound on how far actions can be from optimality. This bound depends on the information structure, and approaches zero as signals' informativeness increases.

1.1 *Related literature*

There is a large literature that studies observational learning, starting from the seminal contributions of [Bikhchandani et al. \(1992\)](#) and [Banerjee \(1992\)](#). In these papers, a set of rational agents choose sequentially between two actions. An agent's payoff depends on whether his action matches the unknown state of the world, but not on others' actions. The actions of others are relevant only because of their informational content. In [Bikhchandani et al. \(1992\)](#) and [Banerjee \(1992\)](#), each agent knows that his own signal is not better than the signals others have received. Agents eventually follow others' behavior and disregard their own signals. Then the optimal behavior of rational agents can prevent complete learning. [Smith and Sørensen \(2000\)](#) show that when signals are of unbounded strength, individuals never fully disregard their own information and complete learning occurs. [Monzón and Rapp \(2014\)](#) present conditions for information aggregation when agents are uncertain both about their own position in the sequence and about the positions of those they observe.

Starting with [Dekel and Piccione \(2000\)](#), a line of research focuses on the outcomes of sequential voting. In [Dekel and Piccione \(2000\)](#), a finite sequence of agents cast votes between two alternatives. Their focus is on the comparison between simultaneous and sequential voting. [Dekel and Piccione](#) show that any equilibrium of a simultaneous voting game is also an equilibrium when voting is sequential. In [Callander \(2007\)](#), agents vote sequentially and care not only about electing the superior candidate, but also about voting for the winning candidate. [Callander](#) shows that a *bandwagon* eventually starts: voters ignore their private information and vote for the leading candidate. [Ali and Kartik \(2012\)](#) present a model motivated by sequential voting, but that encompasses the class of *collective preferences*: an agent's utility increases when others choose an action that matches the unknown state. [Ali and Kartik](#) show how herds can arise. My paper differs from this line of research in several dimensions. First, I allow for payoff externalities that

can be either positive or negative. My model can accommodate both incentives to conform and incentives to go against the crowd. Second, agents observe a sample of past behavior instead of the whole history of play. Together with position uncertainty and a positive probability of mistakes, this implies that agents cannot individually determine the aggregate outcome. Third, my focus is not on herds, but rather on whether agents are ex post satisfied with their action.

Several recent papers have highlighted the importance of payoff externalities in other environments. [Eyster et al. \(2014\)](#) present a model of observational learning with congestion. As usual, agents want to match their action to the state of the world. But when previous agents in the sequence choose an action, they make it less attractive for those coming later. [Eyster et al.](#) study whether learning occurs as a function of congestion costs. [Cripps and Thomas \(forthcoming\)](#) present a model of (possibly informative) queues. Service to those in the queue is provided only in the good state of the world, but at a stochastic rate. [Cripps and Thomas](#) study the dynamics of the queue. [Arieli \(2017\)](#) focuses on recurring games: successive generations of agents play the same game. As in my paper, payoffs depend on the unknown state of the world and also on the actions of others. However, payoff externalities are only local: an agent's utility is affected by the actions of others in the same generation. [Arieli](#) studies when complete learning occurs. In addition to the points already mentioned, my paper differs from [Eyster et al. \(2014\)](#), [Cripps and Thomas \(forthcoming\)](#), and [Arieli \(2017\)](#) in that an agent's payoff depends on the actions of those before and also after him in the sequence. This adds a strategic consideration to the analysis, as agents may affect future decisions.

2. MODEL

Let $\mathcal{I} = \{1, \dots, T\}$ be a set of agents, indexed by i . Agents are exogenously placed in a sequence in positions indexed by $t \in \{1, \dots, T\}$. The random variable Q assigns a position $Q(i)$ to each agent i . Let $q : \{1, \dots, T\} \rightarrow \{1, \dots, T\}$ be a permutation and let \mathcal{Q} be the set of all possible permutations. All permutations are ex ante equally likely: $\Pr(Q = q) = \frac{1}{T!}$ for all $q \in \mathcal{Q}$. Each individual has no ex ante information about his position in the sequence.²

There are two equally likely states of the world, $\theta \in \Theta = \{0, 1\}$. Agents must choose between two possible actions $a \in \mathcal{A} = \{0, 1\}$.³ The timing of the game is as follows. First, nature chooses the state of the world θ and the order of the sequence q . Agents do not observe these directly. Instead, each agent i receives a noisy signal about the state of the world and a sample of past actions. Then he makes a once-and-for-all choice.

Payoffs may depend on the actions of others. Let $X \equiv \frac{1}{T} \sum_{j \in \mathcal{I}} a_j$ denote the proportion of agents who choose action 1, with realizations $x \in [0, 1]$. Agent i obtains utility $u(a_i, X, \theta) : \mathcal{A} \times [0, 1] \times \Theta \rightarrow \mathbb{R}$, where $u(a_i, X, \theta)$ is a continuous function in X .⁴

²This setup corresponds to the case of symmetric position beliefs as defined in [Monzón and Rapp \(2014\)](#).

³In the [Supplemental Material](#), which is available in a supplementary file on the journal website, <http://econtheory.org/supp/3014/supplement.pdf>, I extend the results to any finite number of states and actions.

⁴Note that an agent's payoff depends on the actions of both those who came before him and those who come after him in the sequence.

2.1 Private signals

Each agent i receives a private signal $S_{Q(i)}$, with realizations $s \in \mathcal{S}$. Conditional on the true state of the world, signals are independent and identically distributed (i.i.d.) across individuals and distributed according to F_0 if $\theta = 0$ or F_1 if $\theta = 1$. I assume that F_0 and F_1 are mutually absolutely continuous. Then, no perfectly revealing signals occur with positive probability, and the likelihood ratio (Radon–Nikodym derivative) $l(s) \equiv \frac{dF_1}{dF_0}(s)$ exists. Let G_θ be the distribution function for this likelihood ratio: $G_\theta(l) \equiv \Pr(l(S) \leq l \mid \theta)$. Since F_0 and F_1 are mutually absolutely continuous, the support $\text{supp}(G)$ of G_0 coincides with the support of G_1 . I define signal strength as follows.

DEFINITION (Signal strength). Signal strength is *unbounded* if $0 < G_0(l) < 1$ for all likelihood ratios $l \in (0, \infty)$; signal strength is *bounded* if the convex hull of $\text{supp}(G)$ is given by $\text{co}(\text{supp}(G)) = [\underline{l}, \bar{l}]$, both with $0 < \underline{l} < 1 < \bar{l} < \infty$.⁵

2.2 The sample of past actions

Agents observe others' actions through a simple sampling rule. Let $h_t = (a_1, a_2, \dots, a_{t-1})$ denote a possible history of actions up to period $t - 1$. Let H_t be the (random) history at time t , with realizations $h_t \in \mathcal{H}_t$. Agent i in position $q(i) = t$ receives a sample $\xi_t : \mathcal{H}_t \rightarrow \Xi$ containing the ordered choices of his M predecessors (if available):

$$\xi_t = \begin{cases} \emptyset & \text{if } t = 1 \\ (a_1, \dots, a_{t-1}) & \text{if } 1 < t \leq M \\ (a_{t-M}, \dots, a_{t-1}) & \text{if } t > M. \end{cases}$$

The first agent observes nobody's action, so he receives an empty sample. Agents in positions $t \in \{2, \dots, M\}$ observe the actions of all their predecessors. Subsequent agents observe the actions of their M immediate predecessors.⁶

2.3 Strategies, mistakes, and equilibrium existence

All information available to an agent is summarized by $\{s, \xi\}$, which is an element of $\mathcal{S} \times \Xi$. I assume that individuals make mistakes with small probability $\varepsilon > 0$, so their strategies are ε -constrained. Formally, agent i 's strategy is a function $\sigma_i : \mathcal{S} \times \Xi \rightarrow [\varepsilon, 1 - \varepsilon]$ that specifies a probability $\sigma_i(s, \xi)$ for choosing action 1 given the information available. The set of ε -constrained strategies is denoted by Σ . Let σ_{-i} denote the strategies for all players other than i . Then the profile of play is given by $\sigma = (\sigma_i, \sigma_{-i})$.⁷

⁵I disregard intermediate cases, since they do not add much to the understanding of observational strategic learning.

⁶I present alternative sampling rules in [Appendix A.9](#). First, I consider the case where agents sample from *all* predecessors. I show that Propositions 1 and 2 also hold whenever closer predecessors are more likely to be sampled. Second, I allow for the sample size M to grow as the number of agents T increases. I show that as long as M does not grow too fast, Propositions 1 and 2 also hold.

⁷Mistakes are rationally anticipated. This model is equivalent to one in which agents choose from $[0, 1]$, but they know in advance that there is a 2ε chance that their decision will be overruled by a coin flip. An

Every profile σ induces a probability distribution \mathbf{P}_σ over histories H_t and, consequently, over proportions X . Profile $\sigma^* = (\sigma_i^*, \sigma_{-i}^*)$ is a Bayes–Nash equilibrium of the game if

$$E_{\sigma^*}[u(a_i, X, \theta)] \geq E_{(\sigma_i, \sigma_{-i}^*)}[u(a_i, X, \theta)] \quad \text{for all } \sigma_i \in \Sigma \text{ and for all } i.$$

A profile of play is symmetric if $\sigma_i = \sigma_j$ for all $i, j \in \mathcal{I}$.

LEMMA 1. *For each T there exists a symmetric equilibrium $\sigma^{*,T}$.*

The [Appendix](#) contains most of the proofs.

2.4 Definition of strategic learning

I study the outcomes of large anonymous games, so I let the number of agents grow large and study symmetric equilibria. Agents face a different stage game in each state of the world. Ex ante, each agent is uncertain not only about the state of the world θ , but also about the realization of the proportion X . An agent receives his private signal and observes the actions of some predecessors. Given this information, he forms beliefs both about the underlying state of the world and about the possible realizations of the proportion X . Then he chooses an action.

I study whether agents can successfully learn both about the state of the world and about the proportion X . In standard observational learning models, complete learning occurs when the fraction of adopters of the superior action approaches 1. When payoff externalities exist, I say *strategic learning* occurs whenever agents' actions are ex post optimal given both the state of the world and the realization of the proportion X . I first present two simple examples that illustrate when agents will be ex post satisfied with their actions. I then introduce the formal definition of strategic learning.

EXAMPLE 1 (Anti-coordination). Let $u(1, X, 0) = \frac{1}{5} - X$, $u(1, X, 1) = \frac{4}{5} - X$, and $u(0, X, \theta) = 0$.

Example 1 presents an environment where choosing action 1 becomes less attractive as more agents also choose it. In state $\theta = 0$, action 1 is preferred as long as $X \leq \frac{1}{5}$, while in state $\theta = 1$, action 1 is preferred whenever $X \leq \frac{4}{5}$. Let x_θ be the realized proportion in state θ , so $x = (x_0, x_1)$ is the vector of realized proportions in each state. When $(x_0, x_1) = (\frac{1}{5}, \frac{4}{5})$ agents are ex post satisfied with their choices. For example, if $(x_0, x_1) = (0, \frac{4}{5})$ instead, an agent would prefer to (ex post) switch his action from 0 to 1 in state $\theta = 0$. In fact, $(\frac{1}{5}, \frac{4}{5})$ is the only vector of realized proportion that makes all agents ex post satisfied with their actions in both states of the world.

Formally, define the *excess utility* from choosing action 1 in state θ given X as $v_\theta(X) \equiv u(1, X, \theta) - u(0, X, \theta)$. I say that x_θ corresponds to a Nash equilibrium (NE)

alternative interpretation of this model is as follows. With probability $1 - 2\varepsilon$, an agent chooses rationally from $[0, 1]$. With probability 2ε , the agent is a “behavioral” type. Half of behavioral types always choose action 0, while the others always choose action 1.

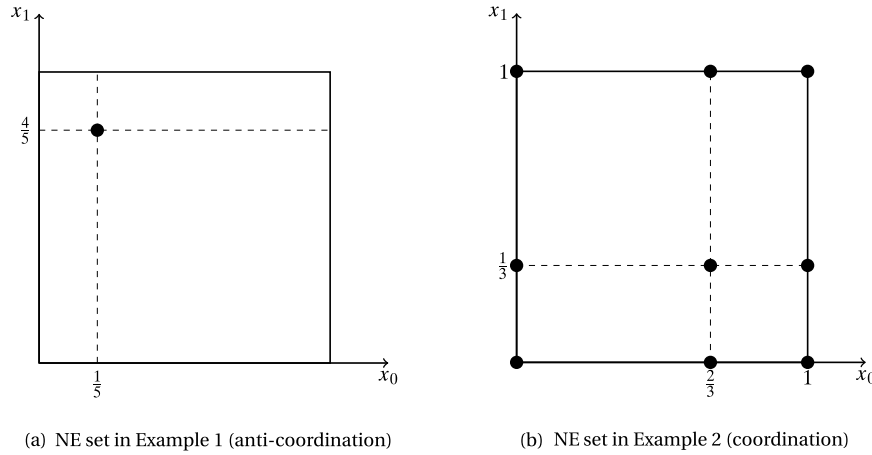


FIGURE 1. The NE sets in Examples 1 and 2.

of the stage game θ (and denote it by $x_\theta \in \text{NE}^\theta$) whenever $v_\theta(x_\theta) > 0 \Rightarrow x_\theta = 1$ and $v_\theta(x_\theta) < 0 \Rightarrow x_\theta = 0$. Similarly, $x \in \text{NE}$ whenever $x_\theta \in \text{NE}^\theta$ for $\theta \in \{0, 1\}$.

The circle in Figure 1(a) depicts the set NE for Example 1. There is a unique $x_\theta \in \text{NE}^\theta$ for each $\theta \in \{0, 1\}$, so NE is the singleton $\{(\frac{1}{5}, \frac{4}{5})\}$. Other games can have multiple elements in NE. Consider, for example, the following simple coordination game.

EXAMPLE 2 (Coordination). Let $u(1, X, 0) = X - \frac{2}{3}$, $u(1, X, 1) = X - \frac{1}{3}$, and $u(0, X, \theta) = 0$.

In Example 2, $\text{NE}^0 = \{0, \frac{2}{3}, 1\}$ and $\text{NE}^1 = \{0, \frac{1}{3}, 1\}$. Then there are nine elements in NE, which are depicted in Figure 1(b) with circles.

It is not obvious a priori whether the realized proportion will be close to elements of NE. The main result in this paper (Proposition 2) shows that this is, in fact, the case. Intuitively, there is strategic learning when, as the number of agents grows large, the (random) proportion X gets close to NE. Because mistakes occur with positive probability, the proportion X may not get arbitrarily close to elements in NE. This is why I first take the number of agents to infinity and then take the probability of mistakes to zero. Let the distance between the realized proportion x and the set NE be defined by $d(x, \text{NE}) \equiv \min_{y \in \text{NE}} |x - y|$.

DEFINITION (Strategic learning). There is *strategic learning* when for all $\delta > 0$, there exists $\tilde{\varepsilon} > 0$ such that

$$\lim_{T \rightarrow \infty} \mathbf{P}_{\sigma^{*,T}}(d(X, \text{NE}) < \delta) = 1$$

for all sequences of symmetric equilibria $\{\sigma^{*,T}\}_{T=1}^\infty$ in games with probability of mistakes $\varepsilon < \tilde{\varepsilon}$.

The main result in the paper (Proposition 2) shows that strategic learning occurs when signals are of unbounded strength.

3. RESULTS

3.1 Average action convergence

The (random) proportion X converges to its expectation in both states of the world. Let the random variable $X_\theta|\sigma$ represent the proportion of agents who choose action 1, conditional on the state of the world θ , and given the strategy σ .⁸ The random vector $X|\sigma = (X_0|\sigma, X_1|\sigma)$ has realizations $x = (x_0, x_1)$. A sequence of symmetric strategy profiles $\{\sigma^T\}_{T=1}^\infty$ induces a sequence of proportions $\{X|\sigma^T\}_{T=1}^\infty$ and a sequence of expected proportions $\{(E[X_0|\sigma^T], E[X_1|\sigma^T])\}_{T=1}^\infty$. As highlighted by the notation, the expected proportions may change with T . I show that in spite of this, $X|\sigma^T$ converges in probability to its expectation.

PROPOSITION 1 (Average action converges in probability). *Take any sequence of symmetric strategy profiles $\{\sigma^T\}_{T=1}^\infty$. Then $X_\theta|\sigma^T - E[X_\theta|\sigma^T] \xrightarrow{P} 0$. More generally, take any sequence $\{\tilde{\sigma}_i^T\}_{T=1}^\infty$ of alternative strategies for agent i . Let the profile of play $\tilde{\sigma}^T = (\tilde{\sigma}_i^T, \sigma_{-i}^T)$ include i 's alternative strategy. Then $X_\theta|\tilde{\sigma}^T - E[X_\theta|\sigma^T] \xrightarrow{P} 0$.*

A symmetric strategy profile σ^T induces a Markov chain over M -period histories of play. Any agent in positions $t > M$ observes the actions of his M immediate predecessors. As σ^T is symmetric, the likelihood that agent i in position $Q(i) > M$ chooses action 1 given sample ξ is independent of both his identity and his position. Then σ^T induces a Markov chain $\{Y_t\}_{M < t \leq T}$ over M -period histories. Moreover, as agents make mistakes with probability $\varepsilon > 0$, the Markov chain is irreducible and aperiodic. If σ^T was fixed for all T , a standard argument (ergodic theorem) would suffice to show **Proposition 1**. However, there is no guarantee that the equilibrium play is independent of the number of agents. In fact, it is easy to find examples where this is not the case.

A Markov chain induced by an ε -constrained symmetric strategy profile σ^T converges to its unique stationary distribution geometrically. Fix a strategy profile σ^T and its induced Markov chain $\{Y_t\}_{t > M}$, but let $t \rightarrow \infty$. A coupling argument provides a geometric lower bound on the speed of convergence to the stationary distribution. Moreover, for any $\varepsilon > 0$, this lower bound is independent of the particular strategy profile σ^T . As a result, although $\{Y_t\}_{M < t \leq T}$ depends on a particular σ^T , it must approach its expectation as T grows. In fact, $X_\theta|\sigma^T - E[X_\theta|\sigma^T]$ converges in the L^2 norm, so it also converges in probability.

Finally, the long-run behavior of the proportion X does not change when one agent deviates and picks a different strategy. To see this, compare the random proportion $X|\sigma^T$ induced by the symmetric profile to the random proportion $X|\tilde{\sigma}^T$ induced by one agent deviating. Let i be the agent who deviates and chooses strategy $\tilde{\sigma}_i$. Agents in positions earlier than $Q(i)$ are not affected by agent i 's strategy. Agents in positions

⁸The underlying random variables in this model are the state of the world θ , the positions Q , the private signals $\{S_t\}_{t=1}^T$, and draws from a randomizing device. The action a_1 is determined by $\sigma_{Q^{-1}(1)}$, the signal s_1 , and a draw from the randomizing device. This in turn determines the sample ξ_2 and affects action a_2 . Recursively, this affects all actions a_t and the proportion $X = \frac{1}{T} \sum_{t=1}^T a_t$. The notation $X_\theta|\sigma$ highlights that the distribution of X is for a given strategy profile σ and depends on the state of the world θ .

right after $Q(i)$ are directly affected. Because of mistakes, the effect that agent i 's action has on subsequent actions $t > Q(i)$ vanishes (geometrically) as t increases. So as the total number of agents T increases, the fraction of agents who are directly affected by i 's action goes to zero. Then also $X_{\theta|\sigma^T} - X_{\theta|\tilde{\sigma}^T} \xrightarrow{P} 0$.

The ε -mistakes are a necessary ingredient for **Proposition 1**: with $\varepsilon = 0$, an individual's action can determine the proportion X .⁹ As I show next, **Proposition 1** allows for a simple approximation to the utility agents obtain from playing this game.

3.2 Utility convergence

Agents' expected utility under symmetric profile σ^T is simply

$$u(\sigma^T) \equiv E_{\sigma^T}[u(a_i, X, \theta)] = \frac{1}{2} \sum_{\theta \in \{0,1\}} E_{\sigma^T}[X_{\theta} \cdot u(1, X_{\theta}, \theta) + (1 - X_{\theta}) \cdot u(0, X_{\theta}, \theta)].$$

Define the *utility of the expected average action* \bar{u}^T by

$$\bar{u}^T \equiv \frac{1}{2} \sum_{\theta \in \{0,1\}} E_{\sigma^T}[X_{\theta}] \cdot u(1, E_{\sigma^T}[X_{\theta}], \theta) + (1 - E_{\sigma^T}[X_{\theta}]) \cdot u(0, E_{\sigma^T}[X_{\theta}], \theta).$$

Agents' expected utility converges to the utility of the expected average action.

LEMMA 2 (Expected utility convergence). *Take any sequence of symmetric strategy profiles $\{\sigma^T\}_{T=1}^{\infty}$. Then $\lim_{T \rightarrow \infty}[u(\sigma^T) - \bar{u}^T] = 0$.*

PROOF. By **Proposition 1**, $X_{\theta|\sigma^T} - E[X_{\theta|\sigma^T}] \xrightarrow{P} 0$. The function $u(a_i, X, \theta)$ is continuous in X . Then $X_{\theta|\sigma^T} \cdot u(a_i, X_{\theta|\sigma^T}, \theta) \xrightarrow{P} E_{\sigma^T}[X_{\theta}] \cdot u(a_i, E_{\sigma^T}[X_{\theta}], \theta)$ because of the continuous mapping theorem. Moreover, $u(a_i, X, \theta)$ is bounded, so $X_{\theta|\sigma^T} \cdot u(a_i, X_{\theta|\sigma^T}, \theta)$ is also bounded. Then $\lim_{T \rightarrow \infty} E_{\sigma^T}[X_{\theta} \cdot u(a_i, X_{\theta}, \theta)] = \lim_{T \rightarrow \infty} E_{\sigma^T}[X_{\theta}] \cdot u(a_i, E_{\sigma^T}[X_{\theta}], \theta)$ by the portmanteau theorem. This leads directly to $\lim_{T \rightarrow \infty}[u(\sigma^T) - \bar{u}^T] = 0$. \square

Proposition 1 also allows for a simple approximation of the expected utility of deviations. Suppose that agent i chooses an alternative strategy $\tilde{\sigma}_i$ and let $u(\tilde{\sigma}_i^T, \sigma_{-i}^T)$ denote the resulting expected utility from this deviation. Define the *approximate utility of the deviation* \tilde{u}^T as

$$\tilde{u}^T \equiv \frac{1}{2} \sum_{\theta \in \{0,1\}} \sum_{a \in \mathcal{A}} \mathbf{P}_{\tilde{\sigma}_i^T}(a_i = a \mid \theta) \cdot u(a, E_{\sigma^T}[X_{\theta}], \theta).$$

⁹A simple example illustrates this. To fix ideas, consider the standard herding model. Let $M = 1$, so each agent observes his immediate predecessor. There are two signals $S = \{0, 1\}$, with $dF_1(1) = dF_0(0) = 0.8$. For any T , let all agents follow $\sigma_i^T(s, \xi) = s$ if $\xi = \emptyset$ and $\sigma_i^T(s, \xi) = \xi$ if $\xi \neq \emptyset$. The first agent observes an empty sample and follows his signal. Agents later in the sequence copy the action they observe. It is simple to see that the action of the first agent determines X . In this example $\mathbf{P}_{\sigma^T}(X_1 = 1) = 0.8$ and $\mathbf{P}_{\sigma^T}(X_1 = 0) = 0.2$ for all T . The random proportion $X_1|\sigma^T$ does not converge in probability. A similar example can be constructed to show that the second part of **Proposition 1** can fail with $\varepsilon = 0$, even if the first part holds (that is, even if $X|\sigma^T$ converges in probability to its expectation).

LEMMA 3 (Expected utility of deviations). *Take any sequence of symmetric strategy profiles $\{\sigma^T\}_{T=1}^\infty$ and a sequence of alternative strategies for agent i : $\{\tilde{\sigma}_i^T\}_{T=1}^\infty$. Then $\lim_{T \rightarrow \infty} [u(\tilde{\sigma}_i^T, \sigma_{-i}^T) - \tilde{u}^T] = 0$.*

The proof closely follows that of Lemma 2; see Appendix A.4 for the details.

3.3 The set of limit points

Different profiles of play σ^T induce different distributions over X . I study the set L of limit points for sequences of equilibrium strategies $\{\sigma^T\}_{T=1}^\infty$. Let $E[X|\sigma] = (E_\sigma[X_0], E_\sigma[X_1])$ denote the vector of expected proportions.

DEFINITION (Limit points). The vector $x = (x_0, x_1)$ is a limit point if there exists a sequence of symmetric equilibrium strategy profiles $\{\sigma^T\}_{T=1}^\infty$ such that for some subsequence $\{\sigma^{T_\tau}\}_{\tau=1}^\infty$, $\lim_{\tau \rightarrow \infty} E[X|\sigma^{T_\tau}] = x$.

The following corollary, which is an immediate consequence of Proposition 1, shows why one should focus on the set L of limit points. As the number of agents grows large, only proportions X close to L occur with positive probability

COROLLARY 1. *Take any sequence of symmetric equilibrium strategy profiles $\{\sigma^T\}_{T=1}^\infty$ and any $\delta > 0$. Then $\lim_{T \rightarrow \infty} \mathbf{P}_{\sigma^T}(d(X, L) < \delta) = 1$.*

The set of limit points L is generated by equilibrium strategies. Optimality considerations allow for a partial characterization of L . Pick a sequence of symmetric equilibria $\{\sigma^T\}_{T=1}^\infty$ and also a sequence of (alternative) ε -constrained strategies for agent i : $\{\tilde{\sigma}_i^T\}_{T=1}^\infty$. Since σ^T are equilibrium strategies, $u(\tilde{\sigma}_i^T, \sigma_{-i}^T) - u(\sigma^T) \leq 0$ for all σ_{-i}^T and for all T . Computing exactly $u(\tilde{\sigma}_i^T, \sigma_{-i}^T)$ and $u(\sigma^T)$ is not possible in general. It requires specifying payoffs, the signal structure, and the number M of agents sampled, and then also computing the equilibrium play. Fortunately, Lemmas 2 and 3 together make it easy to work with alternative strategies. Let the *approximate improvement* Δ^T be given by

$$\Delta^T \equiv \tilde{u}^T - \bar{u}^T = \frac{1}{2} \sum_{\theta \in \{0,1\}} [\mathbf{P}_{\tilde{\sigma}^T}(a_i = 1 | \theta) - E_{\sigma^T}[X_\theta]] \cdot v_\theta(E_{\sigma^T}[X_\theta]).$$

No alternative strategy can do better than the equilibrium strategy. The following corollary formalizes this intuition for the approximate improvement Δ^T .¹⁰

COROLLARY 2. *Take any sequence of symmetric equilibrium strategy profiles $\{\sigma^T\}_{T=1}^\infty$ and a sequence of ε -constrained strategies $\{\tilde{\sigma}_i^T\}_{T=1}^\infty$ for agent i . Then $\limsup_{T \rightarrow \infty} \Delta^T \leq 0$.*

I present two simple alternative strategies that restrict the possible elements of the set L of limit points. The first strategy is to always choose a particular action, regardless of the information received. This strategy proves useful when one action dominates

¹⁰Since agents make mistakes with positive probability, alternative strategies must be ε -constrained.

the other in the limit. The second strategy is to copy the action of one of the observed agents, unless the signal received is extremely informative. This strategy resembles the standard improvement principle in observational learning and is useful when no action strictly dominates the other in the limit.

3.4 *Alternative Strategy 1: Always follow a given action*

The first alternative strategy is simple: follow a given action, regardless of the information received. Lemma 4 shows how this strategy imposes restrictions on the elements of L .

LEMMA 4 (Dominance). *Any limit points $(x_0, x_1) \in L$ must satisfy*

$$(x_0 - \varepsilon)v_0(x_0) + (x_1 - \varepsilon)v_1(x_1) \geq 0 \tag{1}$$

$$(1 - \varepsilon - x_0)v_0(x_0) + (1 - \varepsilon - x_1)v_1(x_1) \leq 0. \tag{2}$$

Moreover, let one action be weakly dominant in the limit: $v_0(x_0)v_1(x_1) \geq 0$. Then $v_\theta(x_\theta) > 0$ implies $(x_0, x_1) = (1 - \varepsilon, 1 - \varepsilon)$ and $v_\theta(x_\theta) < 0$ implies $(x_0, x_1) = (\varepsilon, \varepsilon)$.

To illustrate how Lemma 4 partially characterizes the long-run outcomes of large games, consider first (2). When (2) is not satisfied, always playing action 1 leads to a utility that is strictly higher than the expected utility of the game. Then points that do not satisfy (2) cannot be limit points. Take again the simple anti-coordination game presented in Example 1. The shaded area in Figure 2(a) shows all points that satisfy (2).¹¹ Take, for example $(\frac{4}{5}, \frac{1}{5})$. For a large enough number of players, agents' expected payoffs become arbitrarily close to $\frac{1}{2}[\frac{4}{5}u(1, \frac{4}{5}, 0) + \frac{1}{5}u(1, \frac{1}{5}, 1)] = -\frac{1}{2}(\frac{3}{5})^2$. An agent who *always*

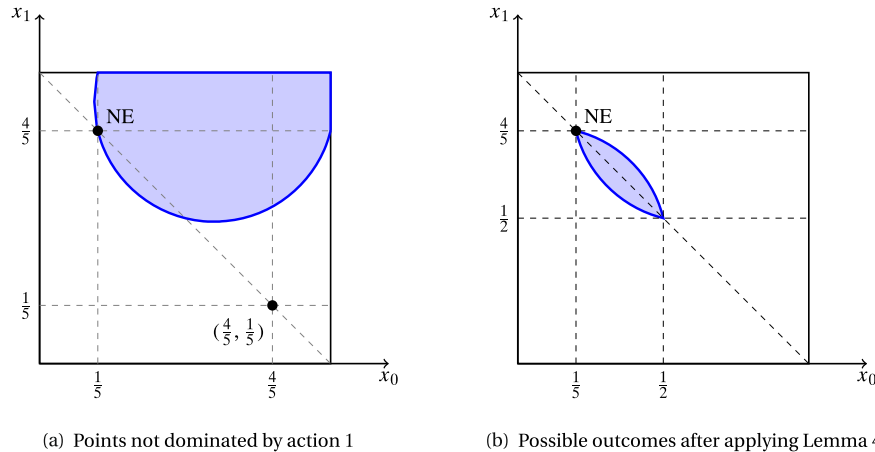


FIGURE 2. Applying Lemma 4 to Example 1 (anti-coordination).

¹¹The exact shape of the sets depicted in Figures 2 and 3 depend on the value of ε . I present them with $\varepsilon = 0$.

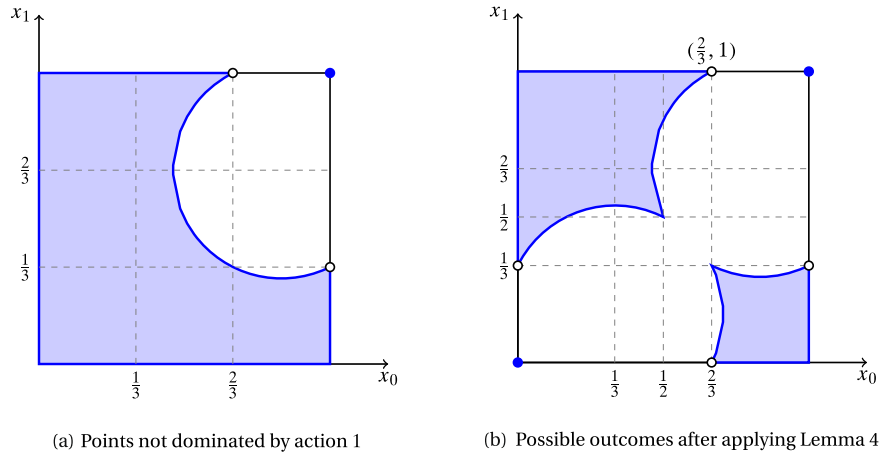


FIGURE 3. Applying Lemma 4 to Example 2 (coordination).

chooses action 1 obtains instead payoffs arbitrarily close to $\frac{1}{2}[u(1, \frac{4}{5}, 0) + u(1, \frac{1}{5}, 1)] = 0$. Then there cannot be a sequence of equilibria that induces $(\frac{4}{5}, \frac{1}{5})$ as a limit point.

Equation (1) describes the outcomes not dominated by action 0 instead. In the case of Example 1, (1) generates an area symmetric to that presented in Figure 2(a). In fact, it is easy to see that $(\frac{4}{5}, \frac{1}{5})$ is also dominated by always playing action 0. The shaded area in Figure 2(b) represents the possible outcomes that remain after applying Lemma 4 in Example 1.

Outcomes that make agents indifferent between actions in one state but not in the other can only be in L if all agents choose the nondominated action in both states, so either $x = (1 - \varepsilon, 1 - \varepsilon)$ or $x = (\varepsilon, \varepsilon)$. Figure 3 illustrates this for the coordination game from Example 2. The shaded area in Figure 3(a) shows the points that are not dominated by action 1. The shaded area in Figure 3(b) depicts the outcomes that remain after applying Lemma 4. The nonshaded circles in Figure 3(b) such as $(\frac{2}{3}, 1)$ cannot be limit points. These points are not (strictly) dominated by always playing some action. For example, as Figure 3(a) shows, $(\frac{2}{3}, 1)$ is not (strictly) worse than always choosing action 1. However, $(\frac{2}{3}, 1)$ cannot be a limit point because of the last result in Lemma 4.

3.5 Alternative Strategy 2: Improve upon a sampled agent

The second alternative strategy deals with the most interesting case: nondominated actions. Consider a general $x = (x_0, x_1)$ with $v_0(x_0)v_1(x_1) < 0$. At such a point, agents want to choose different actions in different states of the world. For simplicity, assume first that $v_0(x_0) < 0$ and $v_1(x_1) > 0$, so agents want their action to match the state of the world. In what follows, I show that unless $x = (0, 1)$, an improvement holds for sufficiently informative signals and small probability of mistakes ε . This environment resembles one from observational learning without payoff externalities, so the approach that I follow takes advantage of tools from those environments.

I introduce an improvement principle based on a simple strategy. Each individual selects *one* individual at random from his sample. Let $\tilde{\xi} = 1$ if the action of the selected

individual is $a = 1$ and let $\tilde{\xi} = 0$ otherwise. The simple strategy mandates that the sampled action must be copied unless a strong enough signal is received. Formally, focus on T big enough so that $v_0(E_{\sigma^T}[X_0]) < 0$ and $v_1(E_{\sigma^T}[X_1]) > 0$. The simple strategy $\tilde{\sigma}^T$ is

$$\tilde{\sigma}^T(\tilde{\xi}, s) = \begin{cases} 1 & \text{if } \tilde{\xi} = 1 \text{ and } l(s) \geq \underline{k}^T \equiv \frac{-v_0(E_{\sigma^T}[X_0]) \mathbf{P}_{\sigma^T}(\tilde{\xi} = 1 | \theta = 0)}{v_1(E_{\sigma^T}[X_1]) \mathbf{P}_{\sigma^T}(\tilde{\xi} = 1 | \theta = 1)} \\ 1 & \text{if } \tilde{\xi} = 0 \text{ and } l(s) \geq \bar{k}^T \equiv \frac{-v_0(E_{\sigma^T}[X_0]) \mathbf{P}_{\sigma^T}(\tilde{\xi} = 0 | \theta = 0)}{v_1(E_{\sigma^T}[X_1]) \mathbf{P}_{\sigma^T}(\tilde{\xi} = 0 | \theta = 1)} \\ 0 & \text{otherwise.} \end{cases}$$

This simple strategy improves upon the average utility \bar{u}^T whenever signals are sufficiently informative and mistakes are not that common. This is derived from two intuitive reasons. First, as long as signals more informative than the observed action $\tilde{\xi}$ exist, the strategy $\tilde{\sigma}$ is strictly better than just imitating $\tilde{\xi}$. Second, without mistakes, the utility of imitating the observed action $\tilde{\xi}$ approaches the average utility \bar{u}^T as the number of agents grows large.

LEMMA 5 (Improvement principle). *Any limit point $(x_0, x_1) \in L$ with $v_0(x_0) < 0$ and $v_1(x_1) > 0$ must satisfy*

$$\begin{aligned} & -v_0(x_0)[(1 - 2\varepsilon)x_0[G_0(\underline{k}) - (\underline{k})^{-1}G_1(\underline{k})] - \varepsilon(1 - 2x_0)] \\ & + v_1(x_1)[(1 - 2\varepsilon)(1 - x_1)[[1 - G_1(\bar{k})] - \bar{k}[1 - G_0(\bar{k})]] - \varepsilon(2x_1 - 1)] \leq 0 \end{aligned} \quad (3)$$

with $\underline{k} = \frac{-v_0(x_0)}{v_1(x_1)} \frac{x_0}{x_1}$ and $\bar{k} = \frac{-v_0(x_0)}{v_1(x_1)} \frac{1-x_0}{1-x_1}$.

Whenever a point x does not satisfy (3), agents can profit from following the simple strategy $\tilde{\sigma}$, so such an x cannot be a limit point. The term $G_0(\underline{k}) - \underline{k}^{-1}G_1(\underline{k}) \geq 0$ in (3) increases in \underline{k} . Symmetrically, the term $[1 - G_1(\bar{k})] - \bar{k}[1 - G_0(\bar{k})] \geq 0$ in (3) decreases in \bar{k} . With signals of unbounded strength, both terms are strictly positive. Then, whenever a positive proportion of agents choose an action that does not match the state of the world (that is, whenever $x_0 > 0$ or $x_1 < 1$), an agent can do better than a random sampled individual. The existence of mistakes may prevent such an improvement.¹² With signals of bounded strength, the term $G_0(\underline{k}) - \underline{k}^{-1}G_1(\underline{k}) \geq 0$ is strictly positive whenever $\underline{k} > \underline{l}$ and the term $[1 - G_1(\bar{k})] - \bar{k}[1 - G_0(\bar{k})]$ is strictly positive whenever $\bar{k} < \bar{l}$. Then, as long as $\underline{k} > \underline{l}$ or $\bar{k} < \bar{l}$, there is potential for improvement upon those observed.

To illustrate how Lemma 5 provides a partial characterization of the outcomes of large games, consider first the anti-coordination game presented in Example 1. Lemma 5 applies when $v_0(x_0) < 0$ and $v_1(x_1) > 0$, which holds whenever $x_0 > \frac{1}{5}$ and $x_1 < \frac{4}{5}$. Take a signal structure with $\bar{l}^{-1} = \underline{l} = \frac{1}{2}$. Points outside of the lightly shaded area in Figure 4(a) have $\underline{k} > \underline{l}$. The term $G_0(\underline{k}) - \underline{k}^{-1}G_1(\underline{k})$ is strictly positive there, so for ε small enough, (3) cannot hold. Next, take a more informative signal structure:

¹²Because of mistakes, it can happen that $v_0(x_0) < 0$ and $v_1(x_1) > 0$ but $(x_0, x_1) \neq (\varepsilon, 1 - \varepsilon)$. Example 4 in the next section shows how this can happen in the standard observational learning setup.

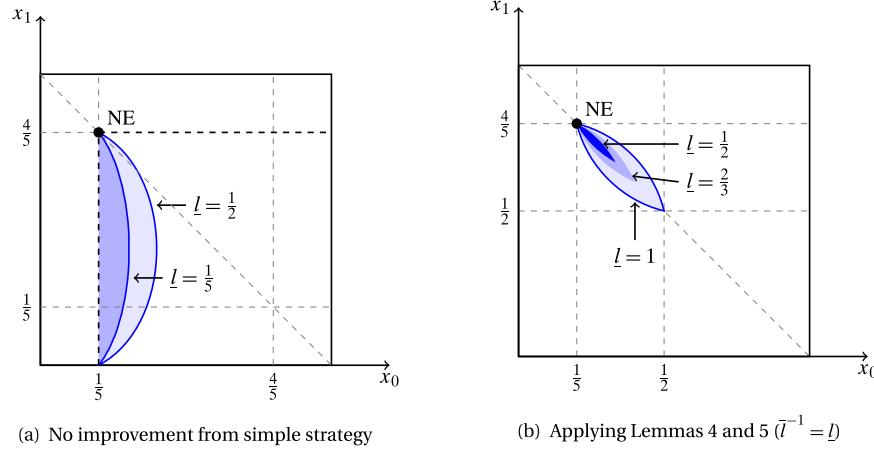


FIGURE 4. Applying Lemma 5 to Example 1 (anti-coordination).

$\bar{l}^{-1} = \underline{l} = \frac{1}{5}$. Points outside of the dark shaded area have $\underline{k} > \underline{l}$. As the bounds on the informativeness of the signal become less restrictive, the shaded area becomes smaller.¹³

Symmetrically, whenever $\bar{k} < \bar{l}$, then $[1 - G_1(\bar{k})] - \bar{k}[1 - G_0(\bar{k})]$ is strictly positive. The area determined by this condition is not depicted in Figure 4(a), but is symmetric to those depicted therein. The shaded areas in Figure 4(b) depict outcomes (x_0, x_1) that satisfy both conditions $\underline{k} \leq \underline{l}$ and $\bar{l} \leq \bar{k}$. As $\bar{l}^{-1} = \underline{l}$ gets smaller, the area satisfying $\underline{k} \leq \underline{l} < \bar{l} \leq \bar{k}$ shrinks. Figure 4(b) provides a preview of the main result of this paper. Only outcomes close to NE remain after applying Lemmas 4 and 5.

For simplicity, I have so far discussed only the case with $v_0(x_0) < 0$ and $v_1(x_1) > 0$. Lemma 6 presents an improvement principle that applies when $v_0(x_0) > 0$ and $v_1(x_1) < 0$, so agents want their action to match the opposite state of the world. The argument behind Lemma 6 is symmetric to that of Lemma 5. See the Supplemental Material for details.

LEMMA 6. Take a limit point $(x_0, x_1) \in L$ with $v_0(x_0) > 0$ and $v_1(x_1) < 0$. Then

$$v_0(x_0)[(1 - 2\varepsilon)(1 - x_0)[G_0(\bar{k}) - (\bar{k})^{-1}G_1(\bar{k})] - \varepsilon(2x_0 - 1)] - v_1(x_1)[(1 - 2\varepsilon)x_1[[1 - G_1(\underline{k})] - \underline{k}[1 - G_0(\underline{k})]] - \varepsilon(1 - 2x_1)] \leq 0.$$

3.6 Strategic learning

Lemmas 4, 5, and 6 jointly lead to the main result of this paper: there is strategic learning. I illustrate this result with the coordination game presented in Example 2. Consider first signals of bounded strength. The shaded areas in Figure 5(a) depict the possible outcomes that satisfy (3) in Lemma 5 for different values of $\bar{l}^{-1} = \underline{l}$. Lemma 5 applies to outcomes with $v_0(x_0) < 0$ and $v_1(x_1) > 0$, which correspond to $x_0 < \frac{2}{3}$ and $x_1 > \frac{1}{3}$. Lemma 6

¹³As before, the exact shape of the sets shown in Figures 4 and 5 depends on the value of ε . I present them with $\varepsilon = 0$.

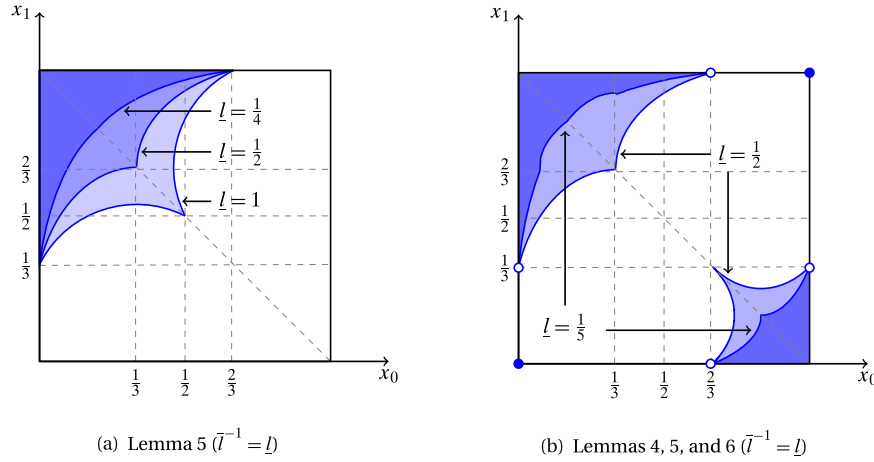


FIGURE 5. Observational learning in games. Example 2 (coordination).

applies when $v_0(x_0) > 0$ and $v_1(x_1) < 0$, which correspond to $x_0 > \frac{2}{3}$ and $x_1 < \frac{1}{3}$. Figure 5(b) shows the set of possible limit points that remain after applying Lemmas 4, 5, and 6. The shaded areas shrink when signals become more informative. Then, with signals of unbounded strength, only points arbitrarily close to NE can be limit points. Proposition 2 formalizes this intuition.

PROPOSITION 2. *Assume signals are of unbounded strength. Then there is strategic learning.*

3.7 Signals of bounded strength

With signals of bounded strength, agents' play need not become arbitrarily close to elements of NE. I show, however, that there must be some degree of learning through the observations of others. I provide a bound on how far long-run outcomes can be from elements of NE. This result is a direct consequence of Lemmas 4, 5, and 6. Intuitively, whenever an agent's choice is ex post suboptimal, it is because he was wrong about the state. To fix ideas, let there be a positive proportion of agents who choose action 1 in state 0 ($x_0 > 0$), but who are ex post dissatisfied ($v_0(x_0) < 0$). Those agents would have preferred choosing action 0. The loss in the population is approximately $-v_0(x_0)x_0$. Instead, the gain in the population from choosing action 1 in state 1 is $v_1(x_1)x_1$. I show that the ratio between the loss and the gain must be bounded above by the informativeness of signals. This ratio is given by $\underline{k} = [-v_0(x_0)x_0]/[v_1(x_1)x_1]$. It must happen that $\underline{k} \leq \underline{l}$. Similarly, the ratio between the gain and the loss from choosing action 0 is given by $\bar{k} = [-v_0(x_0)(1-x_0)]/[v_1(x_1)(1-x_1)]$, and it must happen that $\bar{k} \geq \bar{l}$.

In general, let the set $NE_{(\underline{l}, \bar{l})}$ contain all outcomes with ratios bounded by (\underline{l}, \bar{l}) :

$$x \in NE_{(\underline{l}, \bar{l})} \quad \text{if} \quad \begin{cases} v_0(x_0)v_1(x_1) \geq 0 \Rightarrow x \in NE \\ v_0(x_0) < 0 \text{ and } v_1(x_1) > 0 \Rightarrow \underline{k} \leq \underline{l} < \bar{l} \leq \bar{k} \\ v_0(x_0) > 0 \text{ and } v_1(x_1) < 0 \Rightarrow \bar{k} \leq \underline{l} < \bar{l} \leq \underline{k}. \end{cases}$$

The following result shows that *bounded strategic learning* must occur. Its definition is analogous to the definition of strategic learning, with NE replaced by $\text{NE}_{(\underline{L}, \bar{I})}$.

LEMMA 7 (Bounded strategic learning). *For all $\delta > 0$, there exists $\tilde{\varepsilon} > 0$ such that*

$$\lim_{T \rightarrow \infty} \mathbf{P}_{\sigma^{*,T}}(d(X, \text{NE}_{(\underline{L}, \bar{I})}) < \delta) = 1$$

for all sequences of symmetric equilibria $\{\sigma^{,T}\}_{T=1}^{\infty}$ in games with probability of mistakes $\varepsilon < \tilde{\varepsilon}$.*

The argument behind Lemma 7 is similar to that of Proposition 2. See the Supplemental Material for details.

Lemma 7 highlights the continuity of information aggregation. Strategic learning is a limit result of bounded strategic learning: as I relax the bounds on the signal strength, $\text{NE}_{(\underline{L}, \bar{I})}$ approaches NE.

4. EXAMPLES AND APPLICATIONS

This paper studies the long-run outcomes of observational learning in games. The examples that follow shed further light in this direction. First, I show how strategic learning need not occur without mistakes. Second, I explore the role of mistakes in an example of pure observational learning (without payoff externalities). The third example illustrates the key role of the observation of others to attain strategic learning. Fourth, I provide an example of a coordination game with multiple equilibria. In one equilibrium, agents coordinate on the superior technology, but in a different equilibrium, agents coordinate on a given technology, regardless of its inherent quality. Finally, I illustrate the long-run outcomes of games with preferences like those from Callander (2007) and Eyster et al. (2014).

4.1 Strategic learning can fail with $\varepsilon = 0$

EXAMPLE 3 (Gallice and Monzón (forthcoming)). Each agent observes his immediate predecessor: $M = 1$. Let $u(1, X, \theta) = 3(\frac{1}{T} \sum_{j \neq i} a_j + 1) - 1$ and $u(0, X, \theta) = 3(\frac{1}{T} \sum_{j \neq i} a_j)$. Individuals make no mistakes: $\varepsilon = 0$.

In this example, agents must sequentially decide whether to contribute to a public good. They observe the action of their immediate predecessor but they do not know their position in the sequence. Each agent must choose whether to contribute a fixed amount 1 to a common pool or to refrain from contributing. After all agents choose an action, the total amount invested gets multiplied by 3 and equally shared among all agents. There is only one state of the world in Gallice and Monzón (forthcoming), so signals play no role. Contribution by all agents is socially optimal. However, if an agent takes the actions of others as given, then he strictly prefers not to contribute. As a result, $\text{NE} = \{(0, 0)\}$ is a singleton in Example 3. In stark contrast to Proposition 2 in this paper, Gallice and Monzón show that full contribution can occur in equilibrium. Since there are no mistakes, each individual agent can determine the aggregate outcome. Each individual contributes to induce his potential successors to do the same.

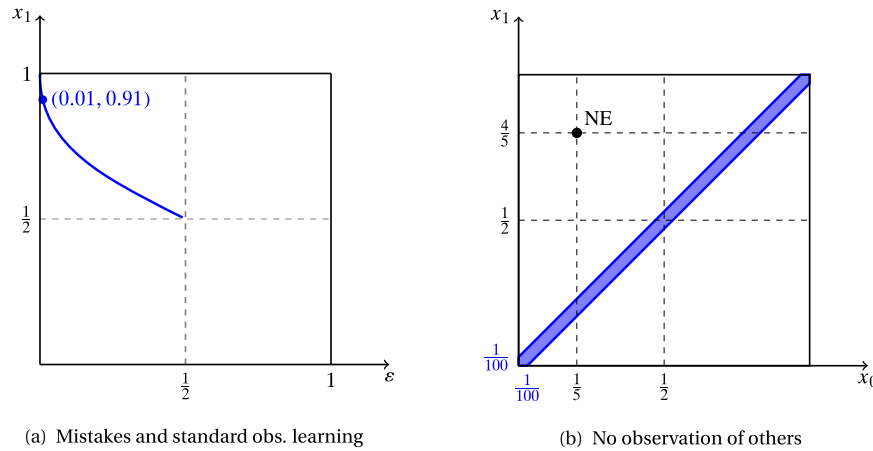


FIGURE 6. Examples 4 and 5.

4.2 Mistakes in observational learning without payoff externalities

EXAMPLE 4 (Standard observational learning). Let $u(1, X, 1) = u(0, X, 0) = 1$ and $u(1, X, 0) = u(0, X, 1) = 0$. Each agent observes his immediate predecessor: $M = 1$. The signal structure is described by $F_1[(0, s)] = s^2$ and $F_0[(0, s)] = 2s - s^2$ with $s \in (0, 1)$.

In this symmetric example, the average action X_1 represents the fraction of agents who choose the right action. Signals are of unbounded strength and the set $NE = (0, 1)$ is a singleton. Then [Proposition 2](#) guarantees that X_1 will be δ -close to 1 for low enough ε . This example provides a simple environment to illustrate what happens when ε is positive. What is the link between δ and ε ? Is it true that (as the number of agents grows large) X_1 must approach $1 - \varepsilon$? This example shows that this is not the case.

The simple signal and observational structure in [Example 4](#) allow for an analytical solution. As the number of agents grows large, the fraction of adopters of the superior technology approaches $\bar{x}_1 \equiv \frac{1-\varepsilon}{1-2\varepsilon} \left(1 - \sqrt{\frac{\varepsilon}{1-\varepsilon}}\right)$. See the [Supplemental Material](#) for details. [Figure 6\(a\)](#) shows the long-run fraction of adopters of the superior technology \bar{x}_1 as a function of the probability of mistakes ε . For example, when $\varepsilon = 0.01$, $\bar{x}_1 \approx 0.91 < 1 - \varepsilon$.¹⁴

4.3 No observation of others' actions

Consider next an anti-coordination game like that presented in [Example 1](#), but with agents who do not observe others' actions.

¹⁴Why can mistakes lead to such a low fraction X_1 ? Intuitively, agents learn whenever they can improve upon those they observe. Let π denote the likelihood that an observed agent chooses the right technology. Without mistakes and with signals of unbounded strength, agents can always improve (as long as $\pi < 1$). Let $I(\pi)$ denote the improvement without mistakes. With mistakes, the improvement becomes $\Pr(\text{mistakes})(1/2 - \pi) + [1 - \Pr(\text{mistakes})]I(\pi) = -2\varepsilon(\pi - 1/2) + (1 - 2\varepsilon)I(\pi)$. Although the second term is always positive, the first term is negative. Then the magnitude of $I(\pi)$ becomes relevant: two signal structures, both with unbounded strength, will lead to different fractions X_1 .

EXAMPLE 5 (No observation of others). Let $u(1, X, 0) = \frac{1}{5} - X$, $u(1, X, 1) = \frac{4}{5} - X$, and $u(0, X, \theta) = 0$. Agents do not observe others' actions. The signal structure is as follows. Let $S = \{0, \frac{1}{2}, 1\}$, with $dF_1(\frac{1}{2}) = dF_0(\frac{1}{2}) = 99/100$, $dF_1(1) = dF_0(0) = (1 - \gamma)/100$, and $dF_1(0) = dF_0(1) = \gamma/100$. Let $\gamma < 1/2$.

A signal $s = \frac{1}{2}$ is uninformative about the state of the world. Signals $s = 0$ and $s = 1$ are instead informative. As γ gets smaller, signals become closer to being of unbounded strength. So Lemma 7 guarantees that the lower the γ , the closer one gets to strategic learning if agents observe the actions of some predecessors. When there is no observation of others, information cannot get transmitted through actions. It is easy to see that at most $|E[X_1] - E[X_0]| \leq 1/100$. Outcomes outside of the shaded area in Figure 6(b) can never be attained without observing others.

4.4 Application to common payoff functions in the literature

EXAMPLE 6 (Congestion. Example 1 in Eyster et al. (2014)). Payoffs are given by $u(1, X, \theta) = \theta - kX$ and $u(0, X, \theta) = 1 - \theta - k(1 - X)$. Signals are of unbounded strength.

An agent obtains a utility of 1 when he chooses the superior technology. On top of it, others who choose the same action as him exert a congestion effect of amount k .¹⁵ The excess utility function is $v_\theta(X) = 2\theta - 2kX - (1 - k)$. When $k < 1$, $v_0(X) < 0$ and $v_1(X) > 0$ for all X . Then NE = $\{(0, 1)\}$. If instead $k \geq 1$, NE = $\{(\frac{1}{2} - \frac{1}{2k}, \frac{1}{2} + \frac{1}{2k})\}$. Signals are of unbounded strength, so Proposition 2 guarantees that there is strategic learning. The long-run outcome will be the unique element of NE. The analysis is analogous to that for the anti-coordination game presented in Example 1.

EXAMPLE 7 (Desire to conform with the majority. Callander (2007)). Payoffs are given by $u(a_i, X, \theta) = \theta f(X) + (1 - \theta)(1 - f(X)) + k[a_i f(X) + (1 - a_i)(1 - f(X))]$. The continuous and monotonically increasing function $f(X)$ has $f(0) = 0$ and $f(1) = 1$. Signals are of unbounded strength.

There is an election with two candidates: 0 and 1. The function $f(X)$ denotes the probability that candidate 1 wins the election given that a fraction X choose him.¹⁶ Each voter obtains a payoff of 1 if the better candidate gets elected. On top of it, he obtains a payoff of k if he votes for the better candidate. The excess utility function is $v_\theta(X) = k(2f(X) - 1)$. An individual cannot affect the result of the election. Then only the cooperation component remains. The possible long-run outcomes are analogous to those in Example 2.

¹⁵In Eyster et al. (2014), only predecessors' actions have a negative effect. Instead, in this paper, it is both predecessors and successors. I have adapted the payoff function to account for this.

¹⁶In Callander (2007), $f(X) = 1$ if $X < \frac{1}{2}$, $f(X) = 0$ if $X > \frac{1}{2}$, and $f(X) = \frac{1}{2}$ if $X = \frac{1}{2}$. Instead, in this paper, payoffs are continuous, so $f(X)$ must be continuous.

4.5 Multiple equilibria in coordination games

EXAMPLE 8 (Coordination. No selection of equilibria). Payoffs are as in Example 2: $u(1, X, 0) = X - \frac{2}{3}$, $u(1, X, 1) = X - \frac{1}{3}$, and $u(0, X, \theta) = 0$. The signal structure is as follows. Let $S = \{0, \frac{1}{2}, 1\}$, with $dF_1(\frac{1}{2}) = dF_0(\frac{1}{2}) = 99/100$, $dF_1(1) = dF_0(0) = (1 - \gamma)/100$, and $dF_1(0) = dF_0(1) = \gamma/100$. Let $\gamma < 1/2$. Each agent observes his immediate predecessor: $M = 1$.

It is easy to show that there is an equilibrium where all agents choose action 1, regardless of what they observe. Under such strategy of play, when the number of agents grows large, the proportion X is close to $1 - \varepsilon$ in both states of the world. Then it is always optimal to choose action 1.¹⁷

Interestingly, there is another equilibrium where agents coordinate on the superior technology. This equilibrium has a simple form. Take a sequence of symmetric strategy profiles where $\sigma^T(s, \xi) = \sigma(s, \xi)$ does not change with T and is given by

$$\sigma(s, \xi) = \begin{cases} s & \text{if } s = \{0, 1\} \\ \xi & \text{if } s = 1/2. \end{cases}$$

Agents follow an informative signal and mimic their predecessor if the signal is uninformative. Under this profile of play, the proportion X is close to γ in state 0 and to $1 - \gamma$ in state 1 (for T large and ε small). This implies that an agent wants his action to match the state of the world. Moreover, the sample is informative about the state of the world. So indeed an agent who receives an uninformative signal copies the action of his predecessor. To sum up, for large enough T , strategy σ is an equilibrium.

5. DISCUSSION

I study the long-run outcomes of observational learning with payoff externalities. In several economic situations, payoffs depend both on an uncertain state of the world and on others' actions. Individuals obtain information about their environment from private signals and also by observing others. Agents need to learn both about the state of the world and about the play of others. Thus, the actions of others exert a direct payoff externality on top of the standard informational externality from observational learning. Even if agents knew the state, they would still not observe the aggregate play, so it would not be obvious which action to choose. Finally, a new strategic consideration arises with payoff externalities: agents may change their behavior so as to influence others.

I show that in spite of these confounding factors, there is strategic learning: agents' actions are ex post optimal given the state of the world and the actions of others. As long as the number of agents grows large and they sometimes make mistakes, each agent's individual influence on the aggregate outcome becomes negligible. Individuals are aware

¹⁷This example illustrates how the proportion X can converge to a Nash equilibrium of the true state without agents learning the state. In this equilibrium, past actions do not provide any information about the state of the world. In spite of this, agents are ex post satisfied by choosing action 1 in both states, and so strategic learning occurs.

of this and so they act as if they could not influence the aggregate play. In large games, the aggregate behavior becomes almost deterministic. I can then translate an environment with payoff externalities into one without them. I then use standard arguments in observational learning to show that information aggregates. Agents are ex post satisfied with their actions in both states of the world.

APPENDIX: PROOFS AND ALTERNATIVE SAMPLING RULES

A.1 Proof of Lemma 1

The proof of existence of a symmetric equilibrium builds upon Theorem 3 in Cheng et al. (2004). Cheng et al. (2004) show that a pure strategy symmetric equilibrium exists in symmetric infinite games with compact, convex strategy sets and continuous and quasiconcave utility functions. I first present Theorem 3 in Cheng et al. (2004) and then show how it applies to the environment in the present paper.

For each player $i \in \mathcal{I}$, let R_i be a set of strategies (with $\rho_i \in R_i$). Agent i 's payoffs from profile (ρ_1, \dots, ρ_T) are denoted by $u_i(\rho_1, \dots, \rho_T)$. The tuple $[\mathcal{I}, \{R_i\}_{i=1}^T, \{u_i\}_{i=1}^T]$ denotes a game.

DEFINITION (Symmetric games (Definition 2 in Cheng et al. (2004))). A normal-form game is symmetric if the players have identical strategy spaces ($R_i = R$ for all $i \in \mathcal{I}$) and $u_i(\rho_i, \rho_{-i}) = u_j(\rho_j, \rho_{-j})$ for $\rho_i = \rho_j$ and $\rho_{-i} = \rho_{-j}$ for all $i, j \in \mathcal{I}$. Thus we can write $u(\rho_i, \rho_{-i})$ for the utility to any player playing strategy ρ_i in profile ρ . Then the tuple $[\mathcal{I}, R, u(\cdot)]$ denotes a symmetric game.

THEOREM 1 (Theorem 3 in Cheng et al. (2004)). A symmetric game $[\mathcal{I}, R, u(\cdot)]$ with R a nonempty, convex, and compact subset of some Euclidean space and $u(\rho_i, \rho_{-i})$ continuous in (ρ_i, ρ_{-i}) and quasiconcave in ρ_i has a symmetric pure-strategy equilibrium.

In the current paper, agent i 's strategy is a function $\sigma_i : S \times \Xi \rightarrow [\varepsilon, 1 - \varepsilon]$, with $\sigma_i \in \Sigma$. I collapse the strategy σ_i into the likelihood $\rho_i(\xi, \theta)$ of choosing action 1 given the sample received and the state of the world. Formally, define $\rho_i(\xi, \theta) \equiv \mathbf{P}_{\sigma_i}(a_i = 1 \mid \theta, \xi)$. There is a many-to-one mapping $\sigma_i \mapsto \rho_i$. It is without loss of generality to work directly with agents who choose ρ_i from the feasible set

$$R_i = \{\rho_i : \rho_i(\xi, \theta) = E[\sigma_i(S_{Q(i)}, \xi) \mid \theta] \text{ for some } \sigma_i \in \Sigma_i\}.$$

The set of strategies Σ is the same for all agents, so $R_i = R$ for all $i \in \mathcal{I}$. Conveniently, R is a subset of a Euclidean space of dimension $|\Xi| \cdot |\Theta|$, and is nonempty and compact (see Appendix A.2 in Monzón and Rapp (2014) for the proof). Next, take $\rho_i \in R$ and $\rho'_i \in R$, with ρ_i derived from σ_i and ρ'_i derived from σ'_i . Then

$$\begin{aligned} \alpha \rho_i(\xi, \theta) + (1 - \alpha) \rho'_i(\xi, \theta) &= \alpha E[\sigma_i(S_{Q(i)}, \xi) \mid \theta] + (1 - \alpha) E[\sigma'_i(S_{Q(i)}, \xi) \mid \theta] \\ &= E[\alpha \sigma_i(S_{Q(i)}, \xi) + (1 - \alpha) \sigma'_i(S_{Q(i)}, \xi) \mid \theta]. \end{aligned}$$

As Σ is convex, then $\alpha\rho_i(\xi, \theta) + (1 - \alpha)\rho'_i(\xi, \theta) \in R$, so R is convex. Agent i 's expected utility as a function of ρ becomes

$$\begin{aligned} u_i(\rho_i, \rho_{-i}) &= \frac{1}{2} \sum_{\theta \in \Theta} \frac{1}{T} \sum_{t=1}^T \sum_{h_t \in \mathcal{H}_t} \mathbf{P}_{\rho_{-i}}(H_t = h_t | \theta) \sum_{\xi \in \Xi} \Pr(\xi | h_t) \\ &\quad \times [\rho_i(\xi, \theta) E_{\rho_{-i}}[u(1, X, \theta) | \theta, h_t, a_t = 1] \\ &\quad + (1 - \rho_i(\xi, \theta)) E_{\rho_{-i}}[u(0, X, \theta) | \theta, h_t, a_t = 0]]. \end{aligned}$$

It is simple to see that $u_i(\rho_i, \rho_{-i})$ is continuous in ρ_i . Others' ρ_{-i} affect $u_i(\rho_i, \rho_{-i})$ through two channels: first, they affect the distribution of H_t ; second, they affect the distribution of X_θ . Utility $u_i(\rho_i, \rho_{-i})$ is continuous in ρ_{-i} through both channels (note that $u(a_i, X, \theta)$ is continuous in X). Therefore, payoffs $u_i(\rho_i, \rho_{-i})$ are continuous in ρ . Finally, note that $u_i(\rho_i, \rho_{-i})$ is linear in $\rho_i(\xi, \theta)$, so $u(\rho_i, \rho_{-i})$ is quasiconcave in ρ_i . Then, by Theorem 3 in Cheng et al. (2004), there exists $\rho^* \in R$ such that ρ^* is a best response to $\rho_{-i} = (\rho^*, \dots, \rho^*)$. Thus, if each agent plays a strategy σ^* that maps to ρ^* , all play a best response. As a result, there exists a symmetric equilibrium σ^* of the game.

A.2 Proof of Proposition 1

I present first an intermediate lemma. Let $P = (p_{ij})$ be a transition matrix on a finite state space \mathcal{Y} . Assume that the Markov chain $Y = (Y_n)_{n=0}^\infty$ associated with P is aperiodic and irreducible. Let μ denote the unique stationary distribution of Y .

LEMMA 8. *Let $\mathcal{Y}^1 \subseteq \mathcal{Y}$ be a nonempty subset of the state space and let $\mu^1 \equiv \sum_{y \in \mathcal{Y}^1} \mu_y$. Then there exists $K < \infty$ such that $\rho = \min_{i,j} p_{ij}^{(K)} > 0$. Moreover, for any distribution over states in period t ,*

$$|\Pr(Y_{t+n} \in \mathcal{Y}^1) - \mu^1| \leq (1 - 2\rho)^{(n-K)/K}.$$

This lemma follows directly from Corollary 4.1.5 in Kemeny and Snell (1983).

With Lemma 8 in hand, I turn to the proof of Proposition 1. Let $\{\sigma^\tau\}_{\tau=1}^\infty$ be a sequence of symmetric strategy profiles. After the first M periods, all samples are of size M . Let $\mathcal{Y} = \{0, 1\}^M$ be the set of all possible histories of length M . Each symmetric strategy profile σ^τ induces a Markov chain $Y^\tau = (Y_t)_{t \geq M}$ over \mathcal{Y} . Since mistakes occur with positive probability, these Markov chains are irreducible and aperiodic. Then each Y^τ has a unique stationary distribution, which I denote by μ^τ . After exactly M periods, transition probabilities are bounded below: $\min_{y, y' \in \mathcal{Y} \times \mathcal{Y}} \Pr(Y_{n+M} = y' | Y_n = y) \geq \varepsilon^M$. The lower bound ε^M is independent of the strategy profile σ^τ .

Let \mathcal{Y}^1 be all histories where the last agent chose action $a = 1$ and let $\bar{\mu}^\tau \equiv \sum_{y \in \mathcal{Y}^1} \mu_y^\tau$. Then Lemma 8 guarantees that for any distribution over states in period t ,

$$|\Pr(Y_{t+n}^\tau \in \mathcal{Y}^1) - \bar{\mu}^\tau| \leq (1 - 2\varepsilon^M)^{\frac{n-M}{M}} = (1 - 2\varepsilon^M)^{-1} [(1 - 2\varepsilon^M)^{\frac{1}{M}}]^n \equiv c\delta^n. \quad (4)$$

This bound holds for any symmetric strategy profile σ^τ .

In what follows, I fix a state of the world θ , so from now on I drop the subindex θ . Also, I fix a strategy profile σ^τ . I use τ to index strategy profiles and use T to index the number of agents. Let $V(\sigma^\tau)$ denote the variance of $X|\sigma^\tau$ for any number of players T : $V(\sigma^\tau) \equiv E_{\sigma^\tau}[(X - E_{\sigma^\tau}[X | \theta])^2 | \theta]$. I show that for any $\tilde{\delta} > 0$, there exists $\tilde{T} < \infty$ such that $E_{\sigma^\tau}[X^2 | \theta] - (E_{\sigma^\tau}[X | \theta])^2 < \tilde{\delta}$ for all $T > \tilde{T}$ and for all τ . This implies that

$$\lim_{T \rightarrow \infty} E_{\sigma^\tau}[X^2 | \theta] - (E_{\sigma^\tau}[X | \theta])^2 = 0,$$

that is, $X_\theta|\sigma^T - E[X_\theta|\sigma^T]$ converges to zero in the L^2 norm, which implies convergence in probability.

Fix a strategy profile σ^τ and define

$$\begin{aligned} V(\sigma^\tau) &\equiv E_{\sigma^\tau} \left[\left(\frac{1}{T} \sum_{t=1}^T a_t \right)^2 \right] - \left(E_{\sigma^\tau} \left[\frac{1}{T} \sum_{t=1}^T a_t \right] \right)^2 \\ &= \frac{1}{T^2} \left[\sum_{t=1}^T (E_{\sigma^\tau}[a_t^2] - E_{\sigma^\tau}[a_t]^2) \right. \\ &\quad \left. + 2 \sum_{t=1}^T \sum_{n=1}^{T-t} (E_{\sigma^\tau}[a_t a_{t+n}] - E_{\sigma^\tau}[a_t] E_{\sigma^\tau}[a_{t+n}]) \right]. \end{aligned} \quad (5)$$

It is easy to see that $\sum_{t=1}^T (E_{\sigma^\tau}[a_t^2] - E_{\sigma^\tau}[a_t]^2) \leq T$. Regarding the remaining terms, note that

$$\begin{aligned} E_{\sigma^\tau}[a_t a_{t+n}] - E_{\sigma^\tau}[a_t] E_{\sigma^\tau}[a_{t+n}] &= \mathbf{P}_{\sigma^\tau}(a_t = 1) \mathbf{P}_{\sigma^\tau}(a_{t+n} = 1 | a_t = 1) \\ &\quad - \mathbf{P}_{\sigma^\tau}(a_t = 1) \mathbf{P}_{\sigma^\tau}(a_{t+n} = 1) \\ &= \mathbf{P}_{\sigma^\tau}(a_t = 1) [\mathbf{P}_{\sigma^\tau}(a_{t+n} = 1 | a_t = 1) - \mathbf{P}_{\sigma^\tau}(a_{t+n} = 1)] \\ &\leq |\mathbf{P}_{\sigma^\tau}(a_{t+n} = 1 | a_t = 1) - \mathbf{P}_{\sigma^\tau}(a_{t+n} = 1)|. \end{aligned}$$

Given (4), $|\mathbf{P}_{\sigma^\tau}(a_{t+n} = 1 | a_t = 1) - \bar{\mu}| < c\delta^n$ and $|\mathbf{P}_{\sigma^\tau}(a_{t+n} = 1) - \bar{\mu}| < c\delta^{t+n}$ for any σ^τ . Then

$$|\mathbf{P}_{\sigma^\tau}(a_{t+n} = 1 | a_t = 1) - \mathbf{P}_{\sigma^\tau}(a_{t+n} = 1)| < c\delta^n + c\delta^{t+n} \leq 2c\delta^n.$$

So the second term in (5) becomes

$$\begin{aligned} 2 \sum_{t=1}^T \sum_{n=1}^{T-t} (E_{\sigma^\tau}[a_t a_{t+n}] - E_{\sigma^\tau}[a_t] E_{\sigma^\tau}[a_{t+n}]) &\leq 2 \sum_{t=1}^T \sum_{n=1}^{T-t} 2c\delta^n = 4c \sum_{t=1}^T \sum_{n=1}^{T-t} \delta^n \\ &\leq 4c \sum_{t=1}^T \frac{\delta(1 - \delta^{T-t})}{1 - \delta} \leq 4c \sum_{t=1}^T \frac{\delta}{1 - \delta} \\ &= 4c \frac{\delta}{1 - \delta} T. \end{aligned}$$

Then, for all σ^τ ,

$$V(\sigma^\tau) \leq \frac{1}{T} \left(1 + 4c \frac{\delta}{1-\delta} \right),$$

where $(1 + 4c \frac{\delta}{1-\delta})$ is independent of σ .

Then pick any $b > 0$. There exists \tilde{T} such that for all $T > \tilde{T}$ and for all σ^τ , $V(\sigma^\tau) < b$. So in particular, for all $b > 0$, there exists \tilde{T} such that for all $T > \tilde{T}$, $V(\sigma^T) < b$; that is, $V(\sigma^T) \rightarrow 0$.

The proof of the second part of [Proposition 1](#) is as follows. Let agent i be in position $t = Q(i)$. Define two Markov chains, both with the same transition matrix P . These chains start right after agent i plays. Their only difference is the starting distribution over states. First, $(Y_n)_{n \geq t+1}$ has agent i following strategy σ_i . Second, $(\tilde{Y}_n)_{n \geq t+1}$ has agent i following strategy $\tilde{\sigma}_i$. Let N be the first period in which these two chains meet. By (4), $\Pr(N > n) \leq c\delta^n$. Note that for any $N = n$,

$$\begin{aligned} X|\sigma^T - X|\tilde{\sigma}^T &= \frac{1}{T} \left[\sum_{t=1}^{Q(i)-1} (a_t|\sigma^T - a_t|\tilde{\sigma}^T) + \sum_{t=Q(i)}^{Q(i)+n-1} (a_t|\sigma^T - a_t|\tilde{\sigma}^T) \right. \\ &\quad \left. + \sum_{t=Q(i)+n}^T (a_t|\sigma^T - a_t|\tilde{\sigma}^T) \right]. \end{aligned}$$

But $a_t|\sigma^T = a_t|\tilde{\sigma}^T$ for $t \in \{1, Q(i) - 1\}$ and for $t \in \{Q(i) + n, T\}$. Then

$$\left| X|\sigma^T - X|\tilde{\sigma}^T \right| = \left| \frac{1}{T} \sum_{t=Q(i)}^{Q(i)+n-1} (a_t|\sigma^T - a_t|\tilde{\sigma}^T) \right| \leq \frac{n}{T}.$$

To sum up, for any strategy profile σ^T ,

$$\Pr\left(|X|\sigma^T - X|\tilde{\sigma}^T| \geq \frac{n}{T} \right) \leq c\delta^n. \tag{6}$$

Then for all $b > 0$, there exists n such that $b \geq c\delta^n$. Fix b and n . There is always a T , so that $n/T < b$. Then

$$\Pr(|X|\sigma^T - X|\tilde{\sigma}^T| \geq b) \leq \Pr\left(|X|\sigma^T - X|\tilde{\sigma}^T| \geq \frac{n}{T} \right) \leq c\delta^n \leq b.$$

Finally, note that both $X_\theta|\sigma^T - X_\theta|\tilde{\sigma}^T \xrightarrow{P} 0$ and $X_\theta|\sigma^T - E[X_\theta|\sigma^T] \xrightarrow{P} 0$. Then also $X_\theta|\tilde{\sigma}^T - E[X_\theta|\sigma^T] \xrightarrow{P} 0$.

A.3 Proof of Corollary 1

The distance $d(X, L)$ can be bounded above as

$$\begin{aligned} d(X, L) &= \min_{y \in L} |X - y| \leq \min_{y \in L} [|X - E_{\sigma^T}[X]| + |E_{\sigma^T}[X] - y|] \\ &\leq |X - E_{\sigma^T}[X]| + \min_{y \in L} |E_{\sigma^T}[X] - y| = |X - E_{\sigma^T}[X]| + d(E_{\sigma^T}[X], L). \end{aligned}$$

The set L includes all limit points for convergent subsequences of $\{E_{\sigma^T}[X]\}_{T=1}^{\infty}$. Then $\lim_{T \rightarrow \infty} d(E_{\sigma^T}[X], L) = 0$. For some \tilde{T} large enough, $d(E_{\sigma^T}[X], L) < \delta/2$ for all $T > \tilde{T}$. Then $\mathbf{P}_{\sigma^T}(d(X, L) < \delta) \geq \mathbf{P}_{\sigma^T}(|X - E_{\sigma^T}[X]| < \delta/2)$. Finally, [Proposition 1](#) guarantees that $\lim_{T \rightarrow \infty} \mathbf{P}_{\sigma^T}(|X - E_{\sigma^T}[X]| < \delta/2) = 1$.

A.4 Proof of Lemma 3

Agent i 's expected utility $u(\tilde{\sigma}_i^T, \sigma_{-i}^T)$ is given by

$$\begin{aligned} u(\tilde{\sigma}_i^T, \sigma_{-i}^T) &= E_{\tilde{\sigma}^T}[u(a_i, X, \theta)] = \frac{1}{2} \sum_{\theta \in \{0,1\}} \sum_{a \in \mathcal{A}} E_{\tilde{\sigma}^T}[u(a, X, \theta) \mathbb{1}\{a_i = a\} | \theta] \\ &= \frac{1}{2} \sum_{\theta \in \{0,1\}} \sum_{a \in \mathcal{A}} E_{\tilde{\sigma}^T}[u(a, X_\theta, \theta) | a_i = a] \mathbf{P}_{\tilde{\sigma}^T}(a_i = a | \theta). \end{aligned}$$

Then

$$\begin{aligned} u(\tilde{\sigma}_i^T, \sigma_{-i}^T) - \tilde{u}^T &= \frac{1}{2} \sum_{\theta \in \{0,1\}} \sum_{a \in \mathcal{A}} \mathbf{P}_{\tilde{\sigma}^T}(a_i = a | \theta) \\ &\quad \times [E_{\tilde{\sigma}^T}[u(a, X_\theta, \theta) | a_i = a] - u(a, E_{\sigma^T}[X_\theta], \theta)]. \end{aligned}$$

If $\lim_{T \rightarrow \infty} \mathbf{P}_{\tilde{\sigma}^T}(a_i = a | \theta) = 0$, then trivially

$$\lim_{T \rightarrow \infty} \mathbf{P}_{\tilde{\sigma}^T}(a_i = a | \theta) [E_{\tilde{\sigma}^T}[u(a, X_\theta, \theta) | a_i = a] - u(a, E_{\sigma^T}[X_\theta], \theta)] = 0.$$

Assume instead that there exists $\delta > 0$ such that $\mathbf{P}_{\tilde{\sigma}^T}(a_i = a | \theta) \geq \delta$ infinitely often. By [Proposition 1](#), for any $\delta > 0$, $\lim_{T \rightarrow \infty} \mathbf{P}_{\tilde{\sigma}^T}(|X_\theta - E_{\sigma^T}[X_\theta]| \geq \delta) = 0$. Then it is also true that for any $\delta > 0$, $\lim_{T \rightarrow \infty} \mathbf{P}_{\tilde{\sigma}^T}(|X_\theta - E_{\sigma^T}[X_\theta]| \geq \delta | a_i = a) = 0$.¹⁸ So by the portmanteau theorem, $\lim_{T \rightarrow \infty} E_{\tilde{\sigma}^T}[u(a, X_\theta, \theta) | a_i = a] = \lim_{T \rightarrow \infty} u(a, E_{\sigma^T}[X_\theta], \theta)$. This leads directly to $\lim_{T \rightarrow \infty} [u(\tilde{\sigma}_i^T, \sigma_{-i}^T) - \tilde{u}^T] = 0$.

A.5 Proof of Corollary 2

We have

$$\begin{aligned} \limsup_{T \rightarrow \infty} \Delta^T &= \limsup_{T \rightarrow \infty} [\tilde{u}^T - u(\tilde{\sigma}_i^T, \sigma_{-i}^T) + u(\tilde{\sigma}_i^T, \sigma_{-i}^T) - u(\sigma^T) + u(\sigma^T) - \bar{u}^T] \\ &\leq \limsup_{T \rightarrow \infty} [\tilde{u}^T - u(\tilde{\sigma}_i^T, \sigma_{-i}^T)] + \limsup_{T \rightarrow \infty} [u(\tilde{\sigma}_i^T, \sigma_{-i}^T) - u(\sigma^T)] \\ &\quad + \limsup_{T \rightarrow \infty} [u(\sigma^T) - \bar{u}^T]. \end{aligned}$$

¹⁸To see this, note that $\mathbf{P}_{\tilde{\sigma}^T}(|X_\theta - E_{\sigma^T}[X_\theta]| > \delta) = \sum_{a \in \mathcal{A}} \mathbf{P}_{\tilde{\sigma}^T}(|X_\theta - E_{\sigma^T}[X_\theta]| > \delta | \theta, a_i = a) \mathbf{P}_{\tilde{\sigma}^T}(a_i = a | \theta)$. By [Proposition 1](#), $\lim_{T \rightarrow \infty} \mathbf{P}_{\tilde{\sigma}^T}(|X_\theta - E_{\sigma^T}[X_\theta]| > \delta) = 0$. Then if $\mathbf{P}_{\tilde{\sigma}^T}(a_i = a | \theta) \geq \delta$ infinitely often, it must be the case that $\lim_{T \rightarrow \infty} \mathbf{P}_{\tilde{\sigma}^T}(|X_\theta - E_{\sigma^T}[X_\theta]| > \delta | a_i = a) = 0$.

Lemmas 2 and 3 imply that $\lim_{T \rightarrow \infty} [u(\sigma^T) - \bar{u}^T]$ and $\lim_{T \rightarrow \infty} [\bar{u}^T - u(\tilde{\sigma}_i^T, \sigma_{-i}^T)] = 0$, respectively. Next, σ^T are equilibrium strategies, so $u(\tilde{\sigma}_i^T, \sigma_{-i}^T) - u(\sigma^T) \leq 0$ for all σ_{-i}^T and for all T . These two facts together imply that

$$\limsup_{T \rightarrow \infty} \Delta^T \leq \limsup_{T \rightarrow \infty} [u(\tilde{\sigma}_i^T, \sigma_{-i}^T) - u(\sigma^T)] \leq 0.$$

A.6 Proof of Lemma 4

Lemma 4 deals with the case in which an action is dominant (either weakly or strictly) in the limit. Consider two alternative strategies: $\tilde{\sigma}^0$, “always play action 0,” and $\tilde{\sigma}^1$, “always play action 1.” Define accordingly $\Delta^{0,T} \equiv \frac{1}{2} \sum_{\theta \in \{0,1\}} (\varepsilon - E_{\sigma^T}[X_\theta]) v_\theta(E_{\sigma^T}[X_\theta])$ and $\Delta^{1,T} \equiv \frac{1}{2} \sum_{\theta \in \{0,1\}} (1 - \varepsilon - E_{\sigma^T}[X_\theta]) v_\theta(E_{\sigma^T}[X_\theta])$. Then, by Corollary 2,

$$\begin{aligned} \limsup_{T \rightarrow \infty} \Delta^{0,T} &= \frac{1}{2} \sum_{\theta \in \{0,1\}} (\varepsilon - x_\theta) v_\theta(x_\theta) \leq 0 \\ \limsup_{T \rightarrow \infty} \Delta^{1,T} &= \frac{1}{2} \sum_{\theta \in \{0,1\}} (1 - \varepsilon - x_\theta) v_\theta(x_\theta) \leq 0. \end{aligned}$$

Next assume $v_0(x_0)v_1(x_1) \geq 0$. Then if $v_\theta(x_\theta) < 0$, (1) requires $x_\theta = \varepsilon$. If, alternatively, $v_\theta(x_\theta) > 0$, (2) requires $x_\theta = 1 - \varepsilon$.

The rest of the proof is a direct result of the following lemma.

LEMMA 9. *If $x_\theta = \varepsilon$ for some $\theta \in \{0, 1\}$, then $x = (\varepsilon, \varepsilon)$. Similarly, if $x_\theta = 1 - \varepsilon$ for some $\theta \in \{0, 1\}$, then $x = (1 - \varepsilon, 1 - \varepsilon)$.*

PROOF. Assume that $x_1 = 1 - \varepsilon$, but $x_0 \neq 1 - \varepsilon$. The proof is analogous for all other cases. The expected proportion $E_{\sigma^T}[X_\theta]$ can be expressed as

$$\begin{aligned} E_{\sigma^T}[X_\theta] &= E_{\sigma^T} \left[\frac{1}{T} \sum_{t=1}^T a_t \mid \theta \right] = \frac{1}{T} \sum_{t=1}^T E_{\sigma^T}[a_t \mid \theta] = \frac{1}{T} \sum_{t=1}^T \mathbf{P}_{\sigma^T}(a_t = 1 \mid \theta) \\ &= \mathbf{P}_{\sigma^T}(a_i = 1 \mid \theta) = \sum_{\xi \in \Xi} \mathbf{P}_{\sigma^T}(\xi \mid \theta) \int_{s \in S} \sigma^T(s, \xi) dF_\theta(s). \end{aligned}$$

Let $\Xi^M \subset \Xi$ be the set of all samples with exactly M actions. All agents in positions $M < t \leq T$ receive samples $\xi_t \in \Xi^M$. Since mistakes occur with positive probability $\varepsilon > 0$, all samples $\xi \in \Xi^M$ occur with positive probability: $\mathbf{P}_{\sigma^T}(\xi \mid \theta) \geq \varepsilon^M$ for any strategy profile σ^T . Then $\lim_{T \rightarrow \infty} \int_{s \in S} \sigma^T(s, \xi) dF_1(s) = 1 - \varepsilon$ for all $\xi \in \Xi^M$. Since $\sigma^T(s, \xi) \leq 1 - \varepsilon$, then, for any $\tilde{c} > 0$,

$$\lim_{T \rightarrow \infty} \int_{s \in S} \mathbb{1}\{\sigma^T(s, \xi) \geq 1 - \varepsilon - \tilde{c}\} dF_1(s) = 1.$$

I show next that the previous equation must also hold for measure F_0 . That is, for all $\tilde{c} > 0$,

$$\lim_{T \rightarrow \infty} \int_{s \in S} \mathbb{1}\{\sigma^T(s, \xi) \geq 1 - \varepsilon - \tilde{c}\} dF_0(s) = 1, \tag{7}$$

which implies that $\lim_{T \rightarrow \infty} \int_{S \in \mathcal{S}} \sigma^T(s, \xi) dF_0(s) = 1 - \varepsilon$ for all $\xi \in \Xi^M$ and so $x_0 = 1 - \varepsilon$.

To see why (7) must hold for measure F_0 , consider the sequence of sets $\{S^t\}_{t=1}^\infty$ with $S^t = \mathbb{1}\{\sigma^T(s, \xi) < 1 - \varepsilon - \tilde{c}\}$. We know that $\lim_{T \rightarrow \infty} \int_{S \in S^t} dF_1(s) = 0$. Assume that for some $c > 0$, $\int_{S \in S^t} dF_0(s) \geq c > 0$ for all t . Pick $l \in (0, \infty)$ such that¹⁹

$$0 < \int_{\{s: l(s) \leq l\}} dF_0(s) \leq c.$$

Then

$$\begin{aligned} \int_{\{s: l(s) \leq l\}} dF_0(s) &\leq c \leq \int_{S \in S^t} dF_0(s) \\ \int_{\{s: l(s) \leq l, s \notin S^t\}} dF_0(s) &\leq \int_{\{s: l(s) > l, s \in S^t\}} dF_0(s) \\ \int_{\{s: l(s) \leq l, s \notin S^t\}} l(s)^{-1} dF_1(s) &\leq \int_{\{s: l(s) > l, s \in S^t\}} l(s)^{-1} dF_1(s) \\ l^{-1} \int_{\{s: l(s) \leq l, s \notin S^t\}} dF_1(s) &\leq l^{-1} \int_{\{s: l(s) > l, s \in S^t\}} dF_1(s) \\ G_1(l) = \int_{\{s: l(s) \leq l\}} dF_1(s) &\leq \int_{\{s \in S^t\}} dF_1(s). \end{aligned}$$

Because of absolute continuity, since $G_0(l) > 0$, then $G_1(l) > 0$. So for all elements of $\{S^t\}_{t=1}^\infty$, $\int_{S \in S^t} dF_1(s) \geq G_1(l) > 0$. Then $\int_{S \in S^t} dF_1(s)$ cannot converge to zero. \square

A.7 Proof of Lemma 5

Let $\pi_\theta^T \equiv \mathbf{P}_{\sigma^T}(\tilde{\xi} = 1 \mid \theta)$. I show first the following intermediate lemma.

LEMMA 10. *For any sequence of strategy profiles $\{\sigma^T\}_{T=1}^\infty$, $\lim_{T \rightarrow \infty} \pi_\theta^T - E_{\sigma^T}[X_\theta] = 0$ for $\theta \in \{0, 1\}$.*

PROOF. Fix the state of the world θ and consider $T \geq 2M$:

$$\begin{aligned} \pi_\theta^T &= \frac{1}{T} \sum_{t=1}^T \mathbf{P}_{\sigma^T}(\tilde{\xi}_t = 1) \\ &= \frac{1}{T} \left[\sum_{t=2}^M \frac{1}{t-1} \sum_{\tau=1}^{t-1} \mathbf{P}_{\sigma^T}(a_\tau = 1) + \sum_{t=M+1}^T \frac{1}{M} \sum_{\tau=t-M}^{t-1} \mathbf{P}_{\sigma^T}(a_\tau = 1) \right] \\ &= \frac{1}{T} \left[\sum_{\tau=1}^{T-1} \mathbf{P}_{\sigma^T}(a_\tau = 1) \sum_{t=\tau+1}^{\min\{\tau+M, T\}} [\min\{t-1, M\}]^{-1} \right] \end{aligned}$$

¹⁹It may happen that the lowest possible interval $\{s : l(s) \leq l\}$ with positive mass starts with a mass point (say at \tilde{l}). If so, its mass may be $\int_{\{s: l(s) \leq \tilde{l}\}} dF_0(s) > c$. In such a case, consider $\alpha \in (0, 1)$ with $\alpha \int_{\{s: l(s) \leq \tilde{l}\}} dF_0(s) = c$. The same argument holds.

$$\begin{aligned}
&= \frac{1}{T} \left[\sum_{\tau=1}^{M-1} \mathbf{P}_{\sigma^T}(a_\tau = 1) \sum_{t=\tau+1}^{\tau+M} (t-1)^{-1} + \sum_{\tau=M}^{T-M} \mathbf{P}_{\sigma^T}(a_\tau = 1) \sum_{t=\tau+1}^{\tau+M} M^{-1} \right. \\
&\quad \left. + \sum_{\tau=T-M+1}^{T-1} \mathbf{P}_{\sigma^T}(a_\tau = 1) \sum_{t=\tau+1}^T M^{-1} \right] \\
&= \frac{1}{T} \sum_{\tau=1}^T \mathbf{P}_{\sigma^T}(a_\tau = 1) \\
&\quad + \frac{1}{T} \left[\sum_{\tau=1}^{M-1} \mathbf{P}_{\sigma^T}(a_\tau = 1) \left(\sum_{t=\tau}^{\tau+M-1} t^{-1} - 1 \right) - \sum_{\tau=T-M+1}^T \mathbf{P}_{\sigma^T}(a_\tau = 1) \left(1 - \frac{T-\tau}{M} \right) \right].
\end{aligned}$$

So

$$\begin{aligned}
\pi_\theta^T - E_{\sigma^T}[X_\theta] &= \frac{1}{T} \left[\sum_{\tau=1}^{M-1} \mathbf{P}_{\sigma^T}(a_\tau = 1) \left(\sum_{t=\tau}^{\tau+M-1} t^{-1} - 1 \right) \right. \\
&\quad \left. - \sum_{\tau=T-M+1}^T \mathbf{P}_{\sigma^T}(a_\tau = 1) \left(1 - \frac{T-\tau}{M} \right) \right].
\end{aligned} \tag{8}$$

Then it follows directly that $\lim_{T \rightarrow \infty} \pi_\theta^T - E_{\sigma^T}[X_\theta] = 0$. \square

With Lemma 10 in hand, I turn to the proof of Lemma 5. Given the simple strategy, the approximate improvement is given by

$$\begin{aligned}
\Delta^T &= \frac{1}{2} \sum_{\theta \in \{0,1\}} v_\theta(E_{\sigma^T}[X_\theta]) [\varepsilon + (1-2\varepsilon) [\pi_\theta^T [1 - G_\theta(\underline{k}^T)] + (1 - \pi_\theta^T) [1 - G_\theta(\bar{k}^T)]] \\
&\quad - E_{\sigma^T}[X_\theta]] \\
&= \frac{1}{2} \sum_{\theta \in \{0,1\}} v_\theta(E_{\sigma^T}[X_\theta]) [\pi_\theta^T - E_{\sigma^T}[X_\theta] + (1 - 2\pi_\theta^T) \varepsilon] \\
&\quad + \frac{1-2\varepsilon}{2} \sum_{\theta \in \{0,1\}} v_\theta(E_{\sigma^T}[X_\theta]) [-\pi_\theta^T G_\theta(\underline{k}^T) + (1 - \pi_\theta^T) [1 - G_\theta(\bar{k}^T)]] \\
&= \frac{1}{2} \sum_{\theta \in \{0,1\}} v_\theta(E_{\sigma^T}[X_\theta]) [\pi_\theta^T - E_{\sigma^T}[X_\theta] + (1 - 2\pi_\theta^T) \varepsilon] \\
&\quad + \frac{1-2\varepsilon}{2} \left[(-v_0(E_{\sigma^T}[X_0])) \pi_0^T \left[G_0(\underline{k}^T) - \frac{v_1(E_{\sigma^T}[X_1])}{-v_0(E_{\sigma^T}[X_0])} \frac{\pi_1^T}{\pi_0^T} G_1(\underline{k}^T) \right] \right. \\
&\quad \left. + v_1(E_{\sigma^T}[X_1]) (1 - \pi_1^T) \left[[1 - G_1(\bar{k}^T)] - \frac{-v_0(E_{\sigma^T}[X_0])}{v_1(E_{\sigma^T}[X_1])} \frac{1 - \pi_0^T}{1 - \pi_1^T} [1 - G_0(\bar{k}^T)] \right] \right] \\
&= \frac{1}{2} \sum_{\theta \in \{0,1\}} v_\theta(E_{\sigma^T}[X_\theta]) [\pi_\theta^T - E_{\sigma^T}[X_\theta]]
\end{aligned}$$

$$\begin{aligned}
& + \frac{1-2\varepsilon}{2} \left[(-v_0(E_{\sigma^T}[X_0])) \left[\pi_0^T [G_0(\underline{k}^T) - (\underline{k}^T)^{-1} G_1(\underline{k}^T)] - \frac{\varepsilon}{1-2\varepsilon} (1-2\pi_0^T) \right] \right. \\
& \left. + v_1(E_{\sigma^T}[X_1]) \left[(1-\pi_1^T) [[1-G_1(\bar{k}^T)] - \bar{k}^T [1-G_0(\bar{k}^T)]] - \frac{\varepsilon(2\pi_1^T-1)}{1-2\varepsilon} \right] \right].
\end{aligned}$$

Let $\underline{k} \equiv \frac{-v_0(x_0)}{v_1(x_1)} \frac{x_0}{x_1}$ and $\bar{k} \equiv \frac{-v_0(x_0)}{v_1(x_1)} \frac{1-x_0}{1-x_1}$. Note that $\lim_{T \rightarrow \infty} \underline{k}^T = \underline{k}$ and $\lim_{T \rightarrow \infty} \bar{k}^T = \bar{k}$. However, $G_\theta(l)$ may be discontinuous if there are mass points. In spite of this,

$$\lim_{T \rightarrow \infty} G_0(\underline{k}^T) - (\underline{k}^T)^{-1} G_1(\underline{k}^T) = G_0(\underline{k}) - (\underline{k})^{-1} G_1(\underline{k}). \quad (9)$$

To see this, first let $\lim_{l \searrow \underline{k}} G_\theta(l)$ denote the limit when l approaches \underline{k} from the right. Since $G_\theta(l)$ is always right-continuous, then (9) holds. Next, let $\lim_{l \nearrow \underline{k}} G_\theta(l)$ denote the limit when l approaches \underline{k} from the left. If $G_\theta(l)$ is left-continuous at \underline{k} , then again (9) holds. Recall that $l(s) = \frac{dF_1}{dF_0}(s)$. For $l \in (0, \infty)$, if $G_\theta(l)$ is *not* left-continuous at \underline{k} , then $\int_{l(s)=\underline{k}} dF_\theta(s) > 0$ for both $\theta \in \{0, 1\}$. Then

$$\begin{aligned}
& \lim_{l \nearrow \underline{k}} G_0(l) - l^{-1} G_1(l) \\
& = \int_{l(s) < \underline{k}} dF_0(s) - \underline{k}^{-1} \int_{l(s) < \underline{k}} dF_1(s) \\
& = \int_{l(s) \leq \underline{k}} dF_0(s) - \int_{l(s) = \underline{k}} dF_0(s) \\
& \quad - \underline{k}^{-1} \left(\int_{l(s) \leq \underline{k}} dF_1(s) - \int_{l(s) = \underline{k}} dF_1(s) \right) \\
& = G_0(\underline{k}) - (\underline{k})^{-1} G_1(\underline{k}) - \int_{l(s) = \underline{k}} dF_0(s) + \underline{k}^{-1} \int_{l(s) = \underline{k}} dF_1(s) \\
& = G_0(\underline{k}) - (\underline{k})^{-1} G_1(\underline{k}) - \int_{l(s) = \underline{k}} dF_0(s) + \underline{k}^{-1} \int_{l(s) = \underline{k}} \frac{dF_1(s)}{dF_0(s)} dF_0(s) \\
& = G_0(\underline{k}) - (\underline{k})^{-1} G_1(\underline{k}) - \int_{l(s) = \underline{k}} dF_0(s) + \underline{k}^{-1} \int_{l(s) = \underline{k}} \underline{k} dF_0(s) \\
& = G_0(\underline{k}) - (\underline{k})^{-1} G_1(\underline{k}).
\end{aligned}$$

The same argument guarantees that

$$\lim_{T \rightarrow \infty} [1 - G_1(\bar{k}^T)] - \bar{k}^T [1 - G_0(\bar{k}^T)] = [1 - G_1(\bar{k})] - \bar{k} [1 - G_0(\bar{k})]. \quad (10)$$

Given (9) and (10),

$$\begin{aligned}
\lim_{T \rightarrow \infty} \Delta^T & = \frac{-v_0(x_0)}{2} [(1-2\varepsilon)x_0 [G_0(\underline{k}) - (\underline{k})^{-1} G_1(\underline{k})] - \varepsilon(1-2x_0)] \\
& \quad + \frac{v_1(x_1)}{2} [(1-2\varepsilon)(1-x_1) [[1-G_1(\bar{k})] - \bar{k} [1-G_0(\bar{k})]] - \varepsilon(2x_1-1)].
\end{aligned}$$

So Corollary 2 leads directly to (3).

A.8 Proof of Proposition 2

I present first the following proposition.

PROPOSITION 3 (Proposition 11 in Monzón and Rapp (2014)). *For all $l \in (\underline{l}, \bar{l})$, $G_\theta(l)$ satisfies*

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

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

$$\begin{aligned} [1 - G_1(k)] - k[1 - G_0(k)] &\geq [1 - G_1(k')] - k'[1 - G_0(k')] \\ G_0(k') - G_1(k')(k')^{-1} &\geq G_0(k) - G_1(k)(k)^{-1}. \end{aligned}$$

See Monzón and Rapp (2014) for the proof.

Let $NE_\delta = \{x \in [0, 1]^2 : d(x, NE) \leq \delta\}$ be the set of all points that are δ -close to elements of NE and let L^ε denote the set of limit points in a game with mistake probability $\varepsilon > 0$. I show first the following lemma.

LEMMA 11 (Limit set approaches NE). *For any $\delta > 0$, there exists $\tilde{\varepsilon} > 0$ such that $L^\varepsilon \subseteq NE_\delta$ for all $\varepsilon < \tilde{\varepsilon}$.*

PROOF. The proof is by contradiction. Assume that there exists (i) a sequence of mistake probabilities $\{\varepsilon^n\}_{n=1}^\infty$ with $\lim_{n \rightarrow \infty} \varepsilon^n = 0$ and (ii) an associated sequence $\{x^n\}_{n=1}^\infty$ with $x^n \in L^{\varepsilon^n}$ for all n , but (iii) $x^n \notin NE_\delta$ for all n . Since $x^n \in [0, 1]^2$ for all n , this sequence has a convergent subsequence $\{x^{n_m}\}_{m=1}^\infty$ with $\lim_{m \rightarrow \infty} x^{n_m} = \bar{x}$. If $v_0(\bar{x}_0) = v_1(\bar{x}_1) = 0$, then $\bar{x} \in NE$, so for m large enough, $x^{n_m} \in NE_\delta$. Then it must be the case that $v_\theta(\bar{x}_\theta) \neq 0$ for some θ .

Assume that $v_1(\bar{x}_1) > 0$. I show next that this requires $\bar{x}_1 = 1$. Pick \tilde{m} large enough so that $v_1(x_1^{n_m}) > 0$ for all $m > \tilde{m}$. For all m with $v_0(x_0^{n_m}) \geq 0$, Lemma 4 implies that $x^{n_m} = (1 - \varepsilon^{n_m}, 1 - \varepsilon^{n_m})$. So if $v_0(x_0^{n_m}) \geq 0$ infinitely often, then $\bar{x}_1 = 1$. Similarly, for all m with $v_0(x_0^{n_m}) < 0$, by Lemma 5, (3) must hold:

$$\begin{aligned} & \frac{-v_0(x_0^{n_m})}{2} \left[\overbrace{(1 - 2\varepsilon^{n_m})}^{\rightarrow 1} \overbrace{x_0^{n_m} [G_0(\underline{k}^{n_m}) - (\underline{k}^{n_m})^{-1} G_1(\underline{k}^{n_m})]}^{\geq 0} - \overbrace{\varepsilon(1 - 2x_0)}^{\rightarrow 0} \right] \\ & + \frac{v_1(x_1^{n_m})}{2} \left[\underbrace{(1 - 2\varepsilon^{n_m})}_{\rightarrow 1} \underbrace{(1 - x_1^{n_m})}_{\geq 0} \left[\underbrace{[1 - G_1(\bar{k}^{n_m})]}_{\geq 0} - \bar{k}^{n_m} \underbrace{[1 - G_0(\bar{k}^{n_m})]}_{\geq 0} \right] \right] \\ & - \underbrace{\varepsilon^{n_m} (2x_1^{n_m} - 1)}_{\rightarrow 0} \leq 0. \end{aligned} \tag{11}$$

Proposition 3 guarantees both that $[1 - G_1(\bar{k}^{n_m})] - \bar{k}^{n_m} [1 - G_0(\bar{k}^{n_m})] \geq 0$ and that $[G_0(\underline{k}^{n_m}) - (\underline{k}^{n_m})^{-1} G_1(\underline{k}^{n_m})] \geq 0$. Then, as (11) shows, when $\varepsilon^{n_m} \rightarrow 0$, only nonnegative terms may remain. Assume next that $\bar{x}_1 < 1$. Then $\lim_{m \rightarrow \infty} v_1(x_1^{n_m})(1 - x_1^{n_m}) = v_1(\bar{x}_1)(1 -$

$\bar{x}_1) > 0$. As $\bar{k} = -[v_0(x_0)(1 - x_0)]/[v_1(x_1)(1 - x_1)]$, this implies that $\lim_{m \rightarrow \infty} \bar{k}^{n_m} < \infty$. Since signals are of unbounded strength, then

$$\lim_{m \rightarrow \infty} [[1 - G_1(\bar{k}^{n_m})] - \bar{k}^{n_m} [1 - G_0(\bar{k}^{n_m})]] > 0.$$

To summarize, whenever $\bar{x}_1 < 1$, (11) is not satisfied for small enough ε^{n_m} . This proves that $\bar{x}_1 = 1$.

Analogous arguments (using also Lemma 6) imply that if $v_\theta(\bar{x}_\theta) > 0$, then $\bar{x}_\theta = 1$ and that if $v_\theta(\bar{x}_\theta) < 0$, then $\bar{x}_\theta = 0$. So $\bar{x} \in \text{NE}$ and, thus, I have reached a contradiction. \square

With Lemma 11, the proof of Proposition 2 is straightforward. Fix $\delta/2 > 0$ and let $\tilde{\varepsilon}$ be as given by Lemma 11. Write

$$\begin{aligned} d(X, \text{NE}) &= \min_{y \in \text{NE}} |X - y| = \min_{y \in \text{NE}} |X - l + l - y| \leq |X - l| + \min_{y \in \text{NE}} |l - y| \quad \text{for any } l \\ &\leq d(X, L^\varepsilon) + \min_{y \in \text{NE}} |l - y| \quad \text{for } l \in \arg \min_{l \in L^\varepsilon} |X - l| \\ &\leq d(X, L^\varepsilon) + \delta/2 \quad \forall \varepsilon < \tilde{\varepsilon}, \quad \text{by Lemma 11.} \end{aligned}$$

Then, for any σ , $\mathbf{P}_\sigma(d(X, \text{NE}) < \delta) \geq \mathbf{P}_\sigma(d(X, L^\varepsilon) < \delta/2)$. By Corollary 1, for all $\delta/2 > 0$ and all sequences of symmetric equilibria,

$$\lim_{T \rightarrow \infty} \mathbf{P}_{\sigma^{T,*}}(d(X, \text{NE}) < \delta) \geq \lim_{T \rightarrow \infty} \mathbf{P}_{\sigma^{T,*}}(d(X, L^\varepsilon) < \delta/2) = 1.$$

A.9 Alternative sampling rules

I consider first geometric sampling from all predecessors. Each individual observes the action of exactly one predecessor. Agent i in position $q(i) = t$ observes the action of a predecessor in position O_t . He receives sample $\xi_t = a_{O(t)}$, but does not learn the position O_t of the observed predecessor. The position O_t follows

$$\Pr(O_t = \tau) = \begin{cases} \frac{\gamma - 1}{\gamma} \frac{\gamma^\tau}{\gamma^{t-1} - 1} & \text{if } \gamma \neq 1 \\ \frac{1}{t-1} & \text{if } \gamma = 1 \end{cases}$$

for all $t \geq 2$. The first agent in the sequence receives an empty sample: $\xi_1 = \emptyset$.

The family of geometric sampling rules is characterized by a parameter $\gamma > 0$. This allows for two interesting cases: either those closer in the sequence are more likely to be observed ($\gamma > 1$), or those early in the sequence (and, thus, further from the observer) are more likely to be observed ($\gamma < 1$). The knife-edge case $\gamma = 1$ corresponds to uniform random sampling. Finally, as $\gamma \rightarrow \infty$, sampling approaches the observation of the immediate predecessor.

I present next an intermediate result that extends Proposition 1 to environments with geometric sampling whenever $\gamma \geq 1$. Lemma 12 provides a result analogous to Lemma 8. With Lemma 12 in hand, the proof of Proposition 1 also applies whenever $\gamma \geq 1$.

LEMMA 12. Assume $\gamma \geq 1$. There exists $c < \infty$ and $\delta < 1$ such that for any strategy σ ,

$$|\mathbf{P}_\sigma(a_{t+n} = 1 \mid a_t = 1, \theta) - \mathbf{P}_\sigma(a_{t+n} = 1, \theta)| \leq c\delta^n.$$

See the [Supplemental Material](#) for details.

Finally, [Lemma 10](#) also holds with geometric sampling when $\gamma > 1$. Since weights are geometric, $\lim_{T \rightarrow \infty} \pi_\theta^T - E_{\sigma^T}[X_\theta] = 0$. See [Proposition 3](#) in [Monzón and Rapp \(2014\)](#) for a formal argument. Then [Proposition 2](#) extends to geometric sampling rules. Strategic learning occurs if those closer are more likely to be observed.

The knife-edge case of $\gamma = 1$ (uniform random sampling) requires some further explanation. Although [Proposition 1](#) extends to it, the current proof of [Proposition 2](#) does not. [Lemma 10](#) does not hold when $\gamma = 1$. The reason for this is simple: with uniform random sampling, those early in the sequence are observed more often. Then the likelihood π_θ^T that an observed agent chooses action 1 and the (expected) proportion $E_{\sigma^T}[X_\theta]$ in the population choosing action 1 do not necessarily coincide. If instead those closer are more likely to be observed (even slightly), then the likelihood of sampling an agent who chooses action 1 approaches $E_{\sigma^T}[X_\theta]$ as the number of agents grows large.

Next, I study the case when the sample size is an increasing function $M(T) : \mathbb{N} \rightarrow \mathbb{N}$ of the number of players T . The following lemma shows that when the growth rate of the sample size is slow enough, the main results of this paper hold.

LEMMA 13. Let $M(T)$ be $o(\log(T))$. Then [Propositions 1 and 2](#) hold.

See the [Supplemental Material](#) for details.

REFERENCES

- Ali, S. Nageeb and Navin Kartik (2012), “Herding with collective preferences.” *Economic Theory*, 51, 601–626. [406]
- Arieli, Itai (2017), “Payoff externalities and social learning.” *Games and Economic Behavior*, 104, 392–410. [407]
- Banerjee, Abhijit V. (1992), “A simple model of herd behavior.” *Quarterly Journal of Economics*, 107, 797–817. [406]
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), “A theory of fads, fashion, custom, and cultural change as informational cascades.” *Journal of Political Economy*, 100, 992–1026. [406]
- Callander, Steven (2007), “Bandwagons and momentum in sequential voting.” *Review of Economic Studies*, 74, 653–684. [406, 419, 421]
- Cheng, Shih-Fen, Daniel M. Reeves, Yevgeniy Vorobeychik, and Michael P. Wellman (2004), “Notes on equilibria in symmetric games.” In *Proceedings of the 6th International Workshop on Game Theoretic and Decision Theoretic Agents GTDT 2004*, 71–78. [423, 424]
- Cripps, Martin W. and Caroline D. Thomas (forthcoming), “Strategic experimentation in queues.” *Theoretical Economics*. [407]

Dekel, Eddie and Michele Piccione (2000), “Sequential voting procedures in symmetric binary elections.” *Journal of Political Economy*, 108, 34–55. [406]

Eyster, Erik, Andrea Galeotti, Navin Kartik, and Matthew Rabin (2014), “Congested observational learning.” *Games and Economic Behavior*, 87, 519–538. [407, 419, 421]

Gallice, Andrea and Ignacio Monzón (forthcoming), “Cooperation in social dilemmas through position uncertainty.” *The Economic Journal*. [404, 419]

Kemeny, John G. and J. Laurie Snell (1983), *Finite Markov Chains*. Undergraduate Texts in Mathematics. Springer-Verlag. [424]

Monzón, Ignacio and Michael Rapp (2014), “Observational learning with position uncertainty.” *Journal of Economic Theory*, 154, 375–402. [406, 407, 423, 432, 434]

Smith, Lones and Peter Sørensen (2000), “Pathological outcomes of observational learning.” *Econometrica*, 68, 371–398. [406]

Co-editor Thomas Mariotti handled this manuscript.

Manuscript received 29 September, 2017; final version accepted 29 November, 2018; available online 30 November, 2018.