

# Social Learning Equilibria\*

Elchanan Mossel<sup>†</sup>, Manuel Mueller-Frank<sup>‡</sup>, Allan Sly<sup>§</sup> and Omer Tamuz<sup>¶</sup>

June 25, 2018

## Abstract

We consider social learning settings in which a group of agents face uncertainty regarding a state of the world, observe private signals, share the same utility function, and act in a general dynamic setting. We introduce Social Learning Equilibria, a static equilibrium concept that abstracts away from the details of the given extensive form, but nevertheless captures the corresponding asymptotic equilibrium behavior. We establish conditions for agreement, herding, and information aggregation in equilibrium, highlighting a connection between agreement and information aggregation.

**Keywords:** Consensus; Learning; Herding.

**JEL:** D83, D85.

## 1 Introduction

Social learning refers to the inference individuals draw from observing the behavior of others. This inference then in turn impacts their own behavior. Social learning has served

---

\*We thank seminar audiences in BI Norway, Caltech, Carnegie Mellon, Columbia, Irvine, Madison, Moscow, Paris School of Economics, Penn State and Stanford. We would also like to thank Itai Arieli, Laura Doval, Luciano Pomatto, Pablo Schenone, Nicolas Vieille and Xavier Vives for illuminating discussions. Manuel Mueller-Frank gratefully acknowledges the financial support of the Spanish Ministry of Economy and Competitiveness (Ref. ECO2015-63711-P). Elchanan Mossel was supported by NSF DMS-1737944, ONR N00014-17-1-2598, and CCF 1665252. Omer Tamuz was supported by a grant from the Simons Foundation (#419427).

<sup>†</sup>MIT.

<sup>‡</sup>IESE Business School.

<sup>§</sup>Princeton University.

<sup>¶</sup>Caltech.

as an explanation for economic phenomena such as herding<sup>1</sup>, bubbles and crashes in financial markets<sup>2</sup>, optimal contracting<sup>3</sup>, technology adoption<sup>4</sup> among many others.<sup>5</sup>

Many theoretical models of rational social learning have been based on a given dynamic social learning game, specified by a *social learning setting* and an *extensive form*. The setting consists of the players, their actions and common utility function, the state and signal spaces, and a commonly known probability distribution thereover. The extensive form specifies the order and frequency of decisions among players, and what each player knows at every given decision instant. These include sequential models<sup>6</sup>, models of learning on social networks<sup>7</sup> and more.

This approach has two inherent weaknesses. First, the analysis of asymptotic equilibrium behavior in dynamic games is not straightforward<sup>8</sup>, resulting in a limited range of tractable models and a focus on extremely stylized settings. Second, when trying to understand or predict behavior in real world social learning settings, the modeler might not know the exact nature of interaction among individuals, and the importance of each of the modeling assumptions is often unclear.

To address these issues we introduce a static equilibrium approach which we call *social learning equilibria* (SLE). We abstract away from the extensive form dynamics, and focus directly on the asymptotic steady state to which the dynamic converges. For a given social learning setting, an SLE includes a description of the information available to each agent, and an action chosen by each player, as a function of what she knows. An SLE does not include any details of the extensive form. However, the information available to the agents was presumably learned through participation in some social learning game. The equilibrium condition simply requires each agent's action to be optimal, given her information.

In this very general setting, our equilibrium approach allows us to tie together many of the classical ideas from social learning: most importantly work on agreeing to disagree and no trade (Aumann, 1976; Milgrom and Stokey, 1982), and herding and its relation to information aggregation (Bikhchandani et al., 1992; Smith and Sørensen, 2000). We gen-

---

<sup>1</sup>See Banerjee (1992); Bikhchandani et al. (1992).

<sup>2</sup>E.g., Scharfstein and Stein (1990); Welch (1992); Chari and Kehoe (2003).

<sup>3</sup>E.g., Khanna (1998); Arya et al. (2006).

<sup>4</sup>E.g., Walden and Browne (2002); Duan et al. (2009).

<sup>5</sup>Further references can be found in Bikhchandani et al. (1998), Chamley (2004), Vives (2010) and Jackson (2011).

<sup>6</sup>E.g., Banerjee (1992); Bikhchandani et al. (1992); Smith and Sørensen (2000).

<sup>7</sup>E.g., Parikh and Krasucki (1990); Gale and Kariv (2003); Acemoglu et al. (2011); Rosenberg et al. (2009); Mossel et al. (2015); Mueller-Frank (2013); Lobel and Sadler (2015).

<sup>8</sup>Gale and Kariv (2003): "The computational difficulty of solving the model is massive even in the case of three persons [...] This is an important subject for future research."

eralize results that were previously known for specific and stylized social learning games, showing that strong conclusions can be drawn from relatively weak assumptions on the extensive form. Importantly, we point out a deep connection between agreement and learning, which we call “we must learn to agree”: when private signals are unbounded, agreement can only be reached when large amounts of information are exchanged, and in particular enough information must be exchanged for agents to learn the correct action. Perhaps surprisingly, this holds regardless of the extensive form, generalizing a result of [Smith and Sørensen \(2000\)](#) for the classical sequential model.

For most of our results we focus attention on the *canonical setting* of social learning with countably many agents (i.e., a large group of agents) binary states, a common prior, conditionally i.i.d. signals and binary actions.

## A motivating example

Consider a large group of agents who each have to repeatedly choose between two actions. Initially, each agent receives a private signal. Then, in each discrete time period, each agent makes a choice and observes the others’ choices. The natural questions to ask are: Do they all eventually agree? And when they agree, do they agree on the correct action? Given answers to these questions, one might wonder how particular they are to this very stylized extensive form. For example, do they still hold if some or all agents also exchange information by talking to each other? Or if each agent acts in only some of the periods? A priori, such games are difficult to solve and analyze directly.

## SLE: A solution concept for social learning games

Instead of analyzing the extensive form dynamics, the fine details of which might be unknown to the modeler, we focus directly on the asymptotic state reached at the end, which is an SLE: As we show in [Theorem 5](#) (which is essentially a reformulation of a result due to [Rosenberg et al. \(2009\)](#)) the asymptotic behavior in any Nash equilibrium<sup>9</sup> of any social learning game is captured by an SLE.

A *complete social learning equilibrium* (CSLE) is an SLE in which each agent knows the other agents’ actions. In the above mentioned game everyone observes everyone else’s actions, and so this game converges to a CSLE. Perhaps surprisingly, this property of the SLE suffices to derive strong conclusions regarding the agents’ actions, as we now explain.

---

<sup>9</sup>The set of Nash equilibria include the perfect Bayesian equilibria—whatever their definition might be in this case.

## Agreement and learning in CSLE

Theorem 1 shows that every CSLE in a canonical setting satisfies *agreement*, i.e., all agents select the same action almost surely. Previous work<sup>10</sup> implies that agreement must hold unless agents are indifferent. Our contribution is to show that in large groups indifference is impossible, and so agreement always holds.

The probability of the agreement action being correct can be linked to the structure of private signals. As defined by Smith and Sørensen (2000), private signals are unbounded if the support of the probability of either state conditional on one signal contains both zero and one. Theorem 2 shows that every CSLE in a canonical setting with unbounded signals aggregates all private information, i.e., the agreement action is optimal conditional on the realized state.

These results imply that in the game described above all agents converge to the same action, and that, when signals are unbounded, they furthermore converge to the correct action. This implication does not require any analysis of the extensive form, but merely that agents know which action their peers converge to. Thus for any extensive form where this holds the conclusion does as well. This includes countless possible variations on the simple, stylized extensive form described above, some of which may potentially be intractable to detailed analysis.

The combination of Theorems 1 and 2 implies that in the case of CSLEs agreement implies information aggregation. We call this relation between agreement and information aggregation “we must learn to agree”. As we discuss next, this relation holds more generally.

## We must learn to agree

Our main result establishes a fundamental connection between herding, one of the most prominent concepts in the social learning literature, and information aggregation. We say that *herding* occurs in an SLE if almost surely all agents but a finite subset select the same action. We view herding as a weak form of agreement: it simply means that almost everyone agrees. We say that the herding action satisfies information aggregation if it is optimal, conditioned on the state.

Theorem 3 shows that in a canonical setting with unbounded signals, in every SLE that satisfies herding the herding action is optimal, conditioned on the state. This highlights a deep connection between the phenomena of agreement and of learning: when agents

---

<sup>10</sup>Aumann (1976); Milgrom and Stokey (1982); Sebenius and Geanakoplos (1983); Mueller-Frank (2013); Rosenberg et al. (2009).

exchange enough information to agree on actions, they must in fact exchange a very large amount of information, enough to learn the state. That is, the agents “must learn to agree.”

The combination of Theorem 3 and Theorem 5 implies that if a social learning game with unbounded signals satisfies herding, then it also satisfies information aggregation. Smith and Sørensen (2000) showed that unbounded signals imply information aggregation in the sequential model. Our result shows that—with unbounded signals and in large groups—information aggregation is independent of the exact extensive form, and even independent of the information agents have beyond their private signals: it holds whenever a herd forms.

## A sufficient condition for herd behavior

An important example of a social learning game is the sequential learning model of Bikhchandani et al. (1992), in which agents have to decide whether to (say) adopt or not adopt a new technology. They receive conditionally i.i.d. private signals, and each makes a decision in an exogenously determined order, after observing the choices of their predecessors. In this highly stylized setting many interesting results were proved regarding herding and learning, with a particularly important contribution by Smith and Sørensen (2000). However, it is natural to wonder what happens in more realistic settings. What if the agents come in groups that act together? Perhaps they exchange information with the people standing behind them or in front of them in line? Perhaps they are allowed to change their decision in the five periods following the first one in which they acted?

We say that an SLE is *weakly ordered* if there is a way to identify each agent with a natural number in such a way that  $i \leq j$  implies that  $j$  knows  $i$ 's action. This holds for—and is inspired by—the sequential learning model, but holds more generally, allowing agents to know much more than the actions of their predecessors.

Theorem 4 shows that every weakly ordered SLE satisfies herding. Thus the herding result of Bikhchandani et al. (1992) is, in this sense, extremely robust: regardless of the extensive form, any game (such as the original sequential learning game and its variants) that is weakly ordered will lead to a herd.

## Concentration of Dependence

A driving force behind our results is what we call the *Concentration of Dependence Principle*. Informally, this principle refers to the fact that when an event  $E$  is a function of i.i.d. random variables, then  $E$  is approximately independent of almost all the random

variables. While well known in probability theory, we believe that the value of applying the principle to economics might go beyond social learning applications.

In our social learning setting the concentration of dependence principle implies that social learning outcomes in large groups depend on the state, and beyond that only on a small number of signals. We use this observation to prove almost all of our results: to preclude indifference (and hence disagreement) in CSLEs, and to show that agreement implies information aggregation.

## Extensions

We consider several extensions of our results and the model. We first allow for heterogeneous priors and utility functions across agents and show that if there are finitely many types which are commonly known and if signals are unbounded then all agents reach agreement on the optimal action.

We also analyze the case of bounded signals, where the support of the belief conditional on one signal contains neither zero nor one. Here we borrow the concept of *information diffusion* introduced by [Lobel and Sadler \(2015\)](#) in the context of the sequential social learning model. We show that for bounded signals our theorems hold when one replaces information aggregation with the weaker notion of information diffusion.

We next show that all our results carry forward to settings in which we relax the assumption of conditional i.i.d. signals. Instead we assume that signals satisfy a *mixing* property. Informally, this means that conditional on the state, each agent's signal is almost independent of almost all the other agents' signals. Mixing signals capture settings in which agents who are close to each other—either geographically or temporally—observe the same or similar signals, but the signals of agents who are far away from each other are nearly independent.

Finally, we argue that our results hold for large finite groups, with a probability that goes to one as the group size goes to infinity.

## Related literature

The social learning literature is too large to comprehensively cite here. We limit the discussion to those papers whose results are most closely related. Our equilibrium approach is more in line with Aumann's approach ([1976](#)) of studying a static environment with common knowledge, as compared to later social learning papers (e.g., [Geanakoplos and Polemarchakis, 1982](#)), which analyze the process by which common knowledge is reached. Similarly to Aumann, we directly study the equilibrium, rather than specifying

the exact interaction structure and procedure by which the equilibrium is obtained. Indeed, in many other fields of economics the tendency is to study static equilibria directly rather than extensive forms.

Our equilibrium notion is conceptually very closely related to that of a rational expectations equilibrium (Grossman, 1981)<sup>11</sup>. In its original formulation (for example Radner, 1979) the concept of rational expectation equilibrium (henceforth REE) is applied to market environments where participants have private information. A forecast function maps signal vectors into a pricing vector which is commonly observed by all agents. A REE is then a forecast function such that markets clear and for (almost) all signal realizations the portfolios of agents maximize their expected utility conditional on the forecast function and their private signal.

There are three main differences between our notion of social learning equilibrium and REE. First, we differ from their particular form of forecast function, which imposes a single summary statistic that is commonly observed. Second, under REE not only do actions have to be individually optimal as in our setting, but they additionally have to satisfy a market clearing condition. This difference arises from the fact that in our social learning setting payoff externalities are absent, contrary to a market environment. Third and most importantly, we show (Theorem 5) how this static equilibrium notion can serve to understand asymptotic equilibrium behavior in dynamic social learning games, which, to the best of our knowledge, is not a general feature of REEs.

Minehart and Scotchmer (1999) introduce a concept of REE in a particular social learning setting. Despite some superficial similarities, their approach is essentially different from ours. For example, an equilibrium—as they define it—does not usually exist, and so they revert to an approximate equilibrium notion, in which they prove their main results.

The seminal papers of Banerjee (1992) and Bikhchandani et al. (1992) introduced the sequential social learning model where, in the canonical setting, agents make a one time choice observing the action chosen by all predecessors. They showed that a herd on the suboptimal action might emerge. Our Theorem 4 shows that herding is a more general feature of interaction among rational agents. Smith and Sørensen (2000) showed that the herding action in the sequential social learning model is almost surely optimal if signals are unbounded, and suboptimal with positive probability if signals are bounded. Our Theorem 3 shows that any herd by rational agents satisfies information aggregation, if signals are unbounded. Thus the relation between herding and information aggregation

---

<sup>11</sup>Similar ideas are also used in some versions of self-confirming equilibria (Rubinstein and Wolinsky, 1994; Dekel et al., 1999, 2004).



established by [Smith and Sørensen \(2000\)](#) is an extremely robust outcome of interaction among rational individuals. This result is similar to some results of optimality of REE ([Dutta and Morris, 1997](#); [DeMarzo and Skiadas, 1998, 1999](#)), which, unlike our results, already hold for a small number of players. A likewise similar approach is taken by [Ostrovsky \(2012\)](#) in studying information aggregation in dynamic markets.

Our [Theorem 7](#) is closely related to [Lobel and Sadler \(2015\)](#) who consider the sequential social learning model where each agent observes a random (possibly correlated) subset of her predecessors. For bounded signals, they introduce the notion of information diffusion, which is a weakening of information aggregation. Assume that the support of private beliefs is  $[1 - \beta, \beta]$ . An action satisfies information diffusion if it is optimal given the state with a probability of at least  $1 - \beta$ . They provide two sufficient conditions on the random observation structure such that information diffuses (respectively, fails to diffuse) in any equilibrium. Applying our results to this concept, we shed additional insight by connecting information diffusion to herding. [Theorem 3](#) together with [Theorem 5](#) implies that in any social learning game where herding occurs, the herding action satisfies information diffusion.

Finally, our [Theorem 1](#) extends the agreement results for settings of repeated interaction of [Gale and Kariv \(2003\)](#), [Mueller-Frank \(2013\)](#), and of [Rosenberg et al. \(2009\)](#) which all show that agreement occurs but in case of indifference among actions. Our [Theorem 1](#) shows that in the canonical settings in a CSLE indifference occurs with probability zero.

The rest of the paper is organized as follows. [Section §2](#) introduces the model and our equilibrium notion. [Section §3](#) presents our results on agreement and information aggregation in CSLEs. [Section §4](#) establishes our results on herding and information diffusion in SLEs. [Section §5](#) establishes the formal relation between social learning equilibria and asymptotic equilibrium behavior in social learning games. [Section §6](#) presents some extensions. [Section §7](#) concludes.

## 2 The Model

We consider a group of agents who must each choose an action under uncertainty about a state of nature. Each agent's utility depends only on her own action and the state, and agents are homogeneous in the sense of sharing the same utility function. Each agent observes a private signal, and additionally some information about the others' signals. A social learning equilibrium is simply a choice of action for each agent that maximizes her expected utility, given the information available to her; note that this information may



include the choices of others. We now define this formally.

### Social learning settings

A *social learning setting*  $(N, A, \Theta, u, S, \mu)$  is defined by a set of players  $N$ , a compact metrizable (common) action space  $A$ , a compact metrizable state space  $\Theta$ , a continuous utility function  $u : A \times \Theta \rightarrow \mathbb{R}$ , a measurable private signal space  $S$ , and finally a commonly known joint probability distribution  $\mu$  over  $\Theta \times S^N$ .

We will denote by  $\theta$  the random state of nature and by  $\bar{s} = (s_i)_{i \in N}$  the agents' private signals. When no ambiguity arises we will denote probabilities and expectations with respect to  $\mu$  by  $\mathbb{P}[\cdot]$  and  $\mathbb{E}[\cdot]$ , respectively. For some modeling applications it will furthermore be useful to add to this probability space a non-atomic random variable  $r$  that is independent of the rest.

### Social Learning Equilibria (SLE)

Each agent  $i$ , in addition to her private signal  $s_i$ , learns  $\ell_i$ , which is some function of  $\bar{s}$  (and possibly  $r$ ). Agent  $i$ 's (random) action is  $a_i$ . It takes values in  $A$ , and is some function of  $\ell_i$  and  $s_i$ . Equivalently,  $\ell_i$  and  $a_i$  are random variables that are, respectively,  $\sigma(\bar{s}, r)$ - and  $\sigma(\ell_i, s_i)$ -measurable.

Let  $\bar{\ell}$  and  $\bar{a}$  denote  $(\ell_i)_{i \in N}$  and  $(a_i)_{i \in N}$ , respectively. In a given social learning setting, a *social learning equilibrium* (or SLE) is a pair  $(\bar{\ell}, \bar{a})$  such that almost surely each agent's action  $a_i$  is a best response, given her information  $\ell_i$  and  $s_i$ :

$$a_i \in \operatorname{argmax}_{a \in A} \mathbb{E}[u(a, \theta) \mid \ell_i, s_i]. \quad (1)$$

It is useful to think of  $\ell_i$  as the information that player  $i$  has learned through participation in the equilibrium of some extensive form game. Likewise, one should think of  $a_i$  as the action that player  $i$  converged to in the same game, so that the pair  $(\bar{\ell}, \bar{a})$  captures the asymptotic state of the game in question. We elaborate on this in Section §5, where we show that this asymptotic state indeed satisfies the SLE condition (1).

So far we have put no restrictions on  $\bar{\ell}$ , and so, in this generality, one would not expect to prove interesting results. In the subsequent sections we will see how some relatively weak conditions on  $\bar{\ell}$  yield interesting properties of  $\bar{a}$ .

## Complete Social Learning Equilibria (CSLE)

The first class of social learning equilibria which we consider are *complete social learning equilibria* (or CSLE). In a CSLE  $\ell_i = \bar{a}$ . That is, each agent, in addition to her private signal, learns the actions of all other agents. Thus, in a given social learning setting, the actions  $\bar{a}$  are a CSLE if for each agent  $i$  it almost surely holds that

$$a_i \in \operatorname{argmax}_{a \in A} \mathbb{E}[u(a, \theta) \mid \bar{a}, s_i]. \quad (2)$$

To specify a CSLE it suffices to specify the actions  $\bar{a}$ , since  $\ell_i = \bar{a}$  for all  $i$ .

Note that a related, natural and more general class of SLEs are those in which  $\bar{a}$  is  $\sigma(\ell_i, s_i)$ -measurable. That is, those SLEs in which the agents all know each other's actions, and perhaps more information additionally. We will prove our results on CSLEs in this greater generality, but prefer to adhere to the definition above because of its simplicity and proximity to Nash equilibria.

### Existence

In every social learning setting there exists an SLE, and moreover a CSLE. This follows directly from the existence of an optimal action, given knowledge of all the private signals. For a CSLE that always exists, let  $a^* = a^*(\bar{s})$  be an action that maximizes expected utility conditional on  $\bar{s}$ , the entire collection of private signals, and set  $a_i = a^*$  for all  $i \in N$ . As  $a^*$  aggregates all private information we call such an equilibrium *information aggregating*.

## 3 Agreement and Information Aggregation in Complete Social Learning Equilibria

In this section we study complete social learning equilibria (CSLEs). We focus on a class of social learning settings which appears frequently in the literature: in *canonical settings*  $N$  is countably infinite,  $A = \Theta = \{0, 1\}$ , signals are informative and conditionally i.i.d., and  $u(a, \theta) = 1_{a=\theta}$ , so that the utility is 1 when the action matches the state, and 0 otherwise.

### Agreement

An SLE satisfies *agreement* if almost surely we have  $a_i = a_j$  for all pairs of agents  $i, j$ . Our first result establishes agreement as a property of any CSLE.

**Theorem 1.** *In a canonical setting every CSLE satisfies agreement.*

This result shows that Aumann’s seminal agreement result carries over to canonical social learning settings as a property of every CSLE. The conceptual reasoning behind the result, however, differs as no epistemic conditions beyond knowledge of the information structure and the social learning equilibrium are required. Previous results in the literature have established that agreement is achieved, except in cases of indifference (Milgrom and Stokey, 1982; Rosenberg et al., 2009; Mueller-Frank, 2013). Our contribution is to show that, for the case of CSLEs in canonical settings, indifference almost surely does not occur and hence agreement holds. This essentially follows from what we call the *Concentration of Dependence* principle, which we introduce now. This principle underlies almost all of our results.

Informally, concentration of dependence refers to the fact that when a decision or event is a function of i.i.d. signals then it significantly depends on only very few of them. The underlying mathematical fact is a well known phenomenon known as *mixing*, which we observe to have interesting implications in our settings.

Formally, we say that a random variable  $X$  is  $\varepsilon$ -independent of an event  $E$  if for every event  $F$  that depends only on  $X$  (i.e., if for every  $F \in \sigma(X)$ ) it holds that

$$\left| \mathbb{P}[E \cap F] - \mathbb{P}[E] \cdot \mathbb{P}[F] \right| < \varepsilon.$$

Note that  $X$  and  $E$  are independent if and only if the left hand side is zero.

**Lemma 1** (Concentration of Dependence Principle). *Let  $X_1, X_2, \dots$  be independent random variables, and let  $E$  be any event defined on the same probability space. Then except for at most  $1/\varepsilon^2$  many  $i$ ’s, each  $X_i$  is  $\varepsilon$ -independent of  $E$ .*

For the convenience of the reader we provide a proof of this fact in §. <sup>12</sup> In our canonical setting the private signals are i.i.d., conditional on the state. It thus follows from this principle that every event that depends on the private signals is approximately conditionally independent of almost all of them.

Let us now briefly outline the proof of Theorem 1. First note that whenever both actions are taken in equilibrium, it must be that all agents are indifferent between the actions. This follows from the same intuition that underlies the no trade theorem of Milgrom and Stokey (1982), as well as and similar results in social learning (e.g., Sebenius

---

<sup>12</sup>Readers who are unfamiliar with this idea may wish to engage with some examples. E.g., let  $X_1, \dots, X_n$  be i.i.d. fair coin tosses, and consider two possible events. The first is the event that the majority of  $X_i$ ’s equal  $H$ . It is easy to calculate and see that all the  $X_i$ ’s are very weakly correlated with  $E$ , and indeed intuitively this is clear. A less obvious example is when  $E$  is the event that an even number of  $X_i$ ’s are equal to  $H$ . Here, changing any  $X_i$  (while keeping the rest fixed) alters the indicator function of this event, and so it may seem that  $E$  strongly depends on each  $X_i$ . However,  $E$  is in fact independent of each  $X_i$ .

and Geanakoplos, 1983; Rosenberg et al., 2009; Mueller-Frank, 2013). This in turn implies that the disagreeing actions are equal to the true state with probability  $\frac{1}{2}$ . Denote the disagreement event by  $F$  and assume towards a contradiction that it has positive probability. For agent  $i$ , let  $b_i$  denote the optimal action of  $i$  conditional only on his private signal  $s_i$ . Note that since signals are informative,  $b_i$  is equal to the state of the world with a probability strictly larger than  $\frac{1}{2}$ . Consider a deviation strategy of each agent  $i$  such that he selects the equilibrium action whenever  $F$  does not occur and  $b_i$  otherwise. The Concentration of Dependence Principle implies that  $F$  is  $\varepsilon$ -independent of  $s_i$  for all but finitely many agents  $i$ . Hence for some agent  $i$  (in fact, all but finitely many) the probability of  $b_i$  being equal to the state conditional on  $F$  is strictly larger than  $\frac{1}{2}$ , establishing a contradiction.

### Information Aggregation

We next turn to the question of the learning properties of CSLEs. We have shown in the previous section that the agents agree on the same (random) action. Under which conditions is this agreement action optimal? In terms of our definitions from the previous section, we ask: under which conditions is an SLE information aggregating? Note that in a canonical setting an SLE is information aggregating if almost surely  $a_i = \theta$  for all  $i$ .

The *private belief*  $p_i$  of an agent is equal to her posterior probability conditional on her private signal only:

$$p_i = \mathbb{P}[\theta = 1 \mid s_i]. \quad (3)$$

As defined by Smith and Sørensen (2000), private signals are *unbounded* if the support of the private belief contains both 0 and 1. Similarly, private signals are *bounded* if the support of private beliefs contains neither 0 nor 1. Smith and Sørensen (2000) showed that in the sequential social learning model unbounded private signals are sufficient for agents eventually to select the action that corresponds to the true state. The following result relates the unbounded signal property to information aggregation in social learning equilibria.

**Theorem 2.** *In a canonical setting with unbounded signals every CSLE is information aggregating.*

Combining Theorem 1 and Theorem 2 we learn that in the case of CSLEs agreement implies information aggregation. This relation between agreement and information aggregation holds much more generally, as we will establish in the next section.

The proof of Theorem 2 is also driven by the Concentration of Dependence Principle. By Theorem 1, there is some (random) agreement action  $a_0$  that all players take. Consider

(towards a contradiction) the case in which the probability that  $a_0 = \theta$  was not 1, but some  $q < 1$ . A player  $i$  can consider the deviation in which, instead of always choosing  $a_0$ , she chooses  $a_0$  when her private signal is weak, but follows her private signal whenever her private belief  $p_i$  is strong. By strong we mean either greater than  $q$  (in which case she would take action 1) or less than  $1 - q$  (in which case she would take action 0). Because signals are unbounded, this occurs with positive probability.

By the Concentration of Dependence Principle,  $a_0$  is essentially a function of some finite number of private signals, and almost all players  $i$  have a private signal that is almost independent of the agreement action  $a_0$ . Therefore this deviation is profitable for some (in fact, almost all) players. Thus it is impossible that in equilibrium  $q < 1$ , and so in equilibrium  $a_0 = \theta$  almost surely.

## 4 Herding, Agreement and Information Aggregation

Arguably the most prominent result in the social learning literature is the herding result established by [Bikhchandani et al. \(1992\)](#) in the canonical sequential social learning model. They show that if agents make an irreversible binary decision in strict sequential order, observing all the actions taken before them, then eventually all agents take the same action, i.e., a herd occurs. Furthermore, the action chosen by this herd is not necessarily optimal, even though the information contained in the pooled private signals suffices to choose the optimal action. Later, [Smith and Sørensen \(2000\)](#) showed that when signals are unbounded a herd still occurs, but the action chosen by the herd is almost surely optimal.

In analogy to this herding phenomenon, we say that an SLE satisfies *herding* if there is almost surely a cofinite set of agents who choose the same action. We call this (random) action the *herding action*. In a canonical setting, we say that the herding action is aggregating if it is equal to  $\theta$  with probability one.

We would like to emphasize that despite the image that the term “herd” evokes, herding does not imply that the agents take a mindless, suboptimal action; indeed, the action chosen by the herd can be correct with probability one, as in the case unbounded signals in the model of [Smith and Sørensen \(2000\)](#). Accordingly, we think of a herding as a form of weak *agreement*: a herding SLE is one in which almost all the agents agree.

## We must learn to agree

Our first result of this section highlights a deep connection between agreement and learning: one cannot agree without learning. In other words, in order for agents to agree they must exchange a large amount of information, and in particular such a large amount that they learn much about the state in the process. We thus offer the phrase “we must learn to agree” not as an imperative but as an observation.

**Theorem 3.** *In a canonical setting, and when signals are unbounded, in every SLE that satisfies herding, the herding action is information aggregating.*

The proof of this theorem again relies on the Concentration of Dependence Principle, and is similar to the proof of Theorem 2. The herding action is approximately independent of almost all the private signals. Yet, it is taken by almost all the agents. Therefore, there will be an agent that takes the herding action with very high probability, and whose signal is almost completely independent from it. Hence such an agent would prefer to follow her own private signal whenever doing so is more likely to be correct than following the herd. But in equilibrium this agent does follow the herd, and so it must be that her private signals never give an indication that is stronger than the information contained in the herding action. But this is impossible when signals are unbounded.

This result is related to similar results for rational expectations equilibria (e.g., [Dutta and Morris, 1997](#); [DeMarzo and Skiadas, 1998, 1999](#); [Ostrovsky, 2012](#)). There, however, agreement implies efficient aggregation of information even for a small number of players and bounded signals, whereas in our setting this holds less generally and crucially depends on both the large size of the group and the unboundedness of the signals.

## Herding in weakly ordered SLEs

Since herding implies that almost all agents agree, it can be seen as a weak form of agreement. Above we have shown that CSLEs satisfy agreement: if all agents observe each other’s actions then they all agree. In this section we relax the complete observation assumption of CSLEs to a weaker condition, and show that it implies herding rather than agreement.

We say that an SLE is *weakly ordered* if each agent  $i$  can be assigned a natural number  $n(i)$  in such a way that if  $n(i) \leq n(j)$  then agent  $j$  observes  $i$ ’s action:  $a_i$  is  $\sigma(\ell_j)$ -measurable. Abstracting away from the extensive form, the classical sequential models of [Bikhchandani et al. \(1992\)](#) and [Smith and Sørensen \(2000\)](#) are a particular example in which the set of agents is identified with the natural numbers and where each agent learns only the

actions of her predecessors. Weakly ordered SLEs are a larger class that allows agents to furthermore have additional information, beyond the actions of their predecessors. They can thus be used to model the asymptotic state of games with richer extensive forms: perhaps the agents come in groups that act together; perhaps they exchange information with the people standing behind them or in front of them in line; or maybe they are not restricted to only acting once.

We show that every weakly ordered SLE satisfies *herding*

**Theorem 4.** *In a canonical setting every weakly ordered SLE satisfies herding.*

The proof of Theorem 4 uses similar ideas to that of Theorem 3, but involves a number of additional steps. Here, one must first observe that if both actions are taken infinitely often then agents must asymptotically be indifferent. If this occurs with positive probability, then eventually agents will be able to guess (correctly with high probability) that this will happen. Since—again asymptotically—almost all agents have signals that are independent of this event, they would choose to ignore it and follow their own private signals. But then they would not be indifferent, and thus this cannot happen with positive probability.

A direct corollary of Theorems 3 and 4 is the following corollary, which is a generalization of the results of [Smith and Sørensen \(2000\)](#).

**Corollary 1.** *In a canonical setting, and when signals are unbounded, then every weakly ordered SLE satisfies herding, and the herding action is information aggregating.*

## 5 Social Learning Equilibria and Social Learning Games

In this section we consider a large class of *social learning games*. A social learning game is a dynamic game with incomplete information in which agents choose actions and observe information about other agents' actions and signals. Its definition includes a *social learning settings*—as in the definition of an SLE—and a description of the extensive form. This class comprises many models studied in the literature, including sequential learning models and models of repeated interaction on social networks.

The main result of this section relates social learning games to SLEs. We show that the asymptotic equilibrium behavior of agents in any social learning game is captured by an SLE: for any distribution over asymptotic equilibrium action profiles of a social learning game there exists an SLE with a matching distribution over action profiles.

This correspondence provides motivation for studying SLEs, and also allows to understand the long-run behavior of agents in many dynamic settings.



## Social learning games

A *social learning game* includes a social learning setting  $(N, A, \Theta, u, S, \mu)$ , together with a description of the extensive form by which agents interact and learn. The extensive form consists of the tuple  $(T, \mathbf{k}, \delta)$ . For each agent  $i$  the set  $T_i \subseteq \{1, 2, \dots\}$  denotes the set of *action times*  $i$ , i.e., the set of time periods in which agent  $i$  exogenously “wakes up”, receives information and takes an action. The set  $T = (T_i)_{i \in N}$  denotes the tuple of action times. For each agent  $i$  and time  $t \in T_i$ , let  $\mathbf{k}_{i,t}$  be the information learned by agent  $i$  at time  $t$ , and let  $a_{i,t}$  be the action taken by agent  $i$  at time  $t$ . We denote by

$$\mathbf{k}_i^t = \{\mathbf{k}_{i,\tau} : \tau \leq t, \tau \in T_i\}$$

the information observed by agent  $i$  by time  $t$ , and by

$$\mathbf{k}_i = \{\mathbf{k}_{i,t} : t \in T_i\}$$

all the information observed by her, excluding her signal. We denote by

$$h_i^t = \{a_{i,\tau} : \tau < t, \tau \in T_i\}$$

the actions taken by agent  $i$  *before* time  $t$ , and by

$$h^t = (h_i^t)_{i \in N}$$

all the actions taken by all the agents before time  $t$ .

The information  $\mathbf{k}_{i,t}$  is some function of the agents’ actions before time  $t$ , the private signals themselves, as well as the additional independent random variable  $r$ , and takes values in some measurable space:

$$\mathbf{k}_{i,t} = \mathbf{k}_{i,t}(h^t, \bar{s}, r).$$

The (possible) dependence on  $r$  allows this framework to include mixed strategies and random observation sets, such as observing a random subset of the previously chosen actions.

The strategy of agent  $i$  at time  $t \in T_i$  is denoted by  $\sigma_{i,t}$ , takes values in  $A$ , and is some function of the information known to agent  $i$  at time  $t$ , which consists of  $\mathbf{k}_i^t$  and her private signal  $s_i$ :

$$\sigma_{i,t} = \sigma_{i,t}(\mathbf{k}_i^t, s_i).$$

The collection of maps  $\sigma_i = (\sigma_{i,t})_{t \in T_i}$  is player  $i$ 's strategy, and the tuple of strategies across all agents,  $(\sigma_i)_{i \in N}$ , is the strategy profile. The history  $(h^t)_{t \in \mathbb{N}}$  is generated according to  $(\sigma_i)_{i \in N}$ .

Finally,  $\delta$  is the common discount factor, and agent  $i$ 's discounted expected utility is

$$\sum_{t \in T_i} \delta^t \cdot \mathbb{E}_\sigma[u(a_{i,t}, \theta)].$$

A strategy profile  $\sigma$  is a *Nash equilibrium* if for each agent  $i$  her strategy  $\sigma_i$  maximizes her discounted expected utility given  $\sigma_{-i}$ , among all possible strategies for player  $i$ .

If agents are myopic, i.e.  $\delta = 0$ , a strategy profile  $\sigma$  is a Nash equilibrium if for each agent  $i$ , given  $\sigma_{-i}$ , in each period  $t$  her strategy  $\sigma_{i,t}$  is such that her action  $a_{i,t}$  maximizes her expected utility in period  $t \in T_i$  conditional on  $\mathbf{k}_i^t$  and  $s_i$ .

This definition of a social learning game is rather general and captures a variety of different models. Most prominently it captures the sequential social learning model of [Bikhchandani et al. \(1992\)](#). To see this simply set  $T_i = \{i\}$  for every agent  $i$  and

$$\mathbf{k}_{i,i} = \mathbf{k}_i = \{a_{j,j} : j < i\}.$$

The sequential social learning models of [Acemoglu et al. \(2011\)](#), [Lobel and Sadler \(2015\)](#) and others are likewise included in this framework. Here we have  $T_i = \{i\}$  again, but  $\mathbf{k}_{i,t}$  does not include all the actions of the predecessors, but rather only those of a random subset of the predecessors. The models of repeated interaction on social networks of [Gale and Kariv \(2003\)](#), [Mossel et al. \(2014\)](#) and [Mossel et al. \(2015\)](#) can be captured by setting  $T_i = \mathbb{N}$  for all agents  $i$  and where  $\mathbf{k}_{i,t}$  contains the last period actions of all the neighbors of agent  $i$ . [Rosenberg et al. \(2009\)](#) study a more general model that is not subsumed by the framework, but still shares many similarities. In fact, the proof of our result for this section, [Theorem 5](#), follows exactly the proof of their [Proposition 2.1](#).

Finally, the models of repeated communication of beliefs in a social network analyzed in [Geanakoplos and Polemarchakis \(1982\)](#), [Parikh and Krasucki \(1990\)](#), and [Mueller-Frank \(2013\)](#) can be captured by a squared loss utility function and a discount factor equal to zero, hence inducing myopic behavior.

For a given strategy  $(\sigma_{i,t})_{t \in T_i}$ , let  $\bar{A}_i$  denote the (random) set of *accumulation points* of agent  $i$ 's realized actions; if  $T_i$  is finite, then let  $\bar{A}_i$  be the singleton that contains only the last period action of agent  $i$ . If  $T_i$  is infinite and  $A$  finite, then  $\bar{A}_i$  consists of the actions chosen infinitely often.

Given these definitions, we are ready to establish the relation between Nash equilibria

of the social learning game and SLEs. As we mention above, this theorem is essentially due to [Rosenberg et al. \(2009\)](#).

**Theorem 5.** *Consider a social learning game, i.e., a social learning setting and extensive form  $(T, \mathbf{k}, \delta)$ , and a corresponding Nash equilibrium  $\sigma$ . Let  $\bar{a}$  be any (random) action such that  $a_i \in \bar{A}_i$ , and let  $\ell_i = \mathbf{k}_i$ , where  $\mathbf{k}_i$  is generated according to  $\sigma$ . Then  $(\bar{\ell}, \bar{a})$  is an SLE.*

Thus, the asymptotic behavior in any Nash equilibrium<sup>13</sup> of a social learning game is captured by an SLE. Given the formulation of the SLE, the result follows if any limit action by any agent  $i$  is optimal conditional on his limit information  $\mathbf{k}_i$ . [Rosenberg et al. \(2009\)](#) have shown this to be true. We provide a version of their proof adjusted to our language and notation in the appendix, establishing a link between their result and our concept of SLE.

### Learning and agreeing in social learning games

In [Theorem 5](#) we showed that the asymptotic outcomes of social learning games correspond to social learning equilibria. That is, if for each agent  $i$  we let  $\ell_i$  denote the information  $i$  has learned by participating in the game, and if we let  $a_i$  be the action that it converged to (or some limit action in lieu of convergence), then  $(\bar{\ell}, \bar{a})$  is an SLE. We can therefore apply our results on SLEs to social learning games, with wide-reaching implications.

Consider any social learning game (in a canonical setting) in which all the agents act in each period and observe each other’s actions; they could, additionally, exchange information in other ways. Then each agent knows the limit actions of each of the others, and hence any such game converges to a CSLE. It therefore follows from [Theorem 1](#) that all agents must always converge to the same action. And if private signals are unbounded, then by [Theorem 2](#) they must all converge to the correct action.

A straightforward application of [Theorem 3](#) to social learning games implies that in every social learning game which is played in a canonical setting with unbounded signals, it holds that if herding occurs (i.e., if a cofinite set of agents *converges* to the same action) then this cofinite set of agents *converges* to the correct action.

[Theorem 4](#) implies that herding is indeed the outcome across a large spectrum of social learning games: it suffices that one can assign a natural number  $n(i)$  to each player  $i$  such that if  $n(i) \leq n(j)$  then  $j$  observes which actions  $i$  converges to. This generalizes the results of [Bikhchandani et al. \(1992\)](#), highlighting the deeper forces that drive them:

---

<sup>13</sup>The set of Nash equilibria include the perfect Bayesian equilibria—whatever their definition might be in this case.

herding (e.g., in the classical sequential model) is not a feature of the sequential timing of actions, but rather of the observation structure of agents. In particular, any social learning game with a weakly ordered observation structure satisfies herding as an asymptotic equilibrium outcome.

## 6 Extensions

### 6.1 Heterogeneous preferences and priors

We now relax the homogeneity assumption and consider agents who have different utility functions or prior beliefs. For the former, assume that agents share a common prior over the binary state space, have the same belief regarding the conditional signal distributions, but whose utility for matching the action to the state varies with the state (but for whom it is still strictly preferable to match the state than to mismatch). For the latter, let agents share the same utility function and the same belief regarding the conditional signal distributions but assume that their prior belief on the state space differs.

Assume that there are a finite number of different types of agents (here the type refers to the agent's utility function or prior) and assume that the agents' types are common knowledge. The following can be established by making slight adjustments to our proofs.

**Theorem 6.** *In a canonical setting with finitely many commonly known types and where signals are unbounded, every CSLE satisfies information aggregation.*

If signals are unbounded then in a CSLE all agents agree on the same action, and additionally this action is optimal. Thus unbounded signals overcome heterogeneity in priors and payoffs. Theorem 6 is interesting to view in light of Aumann's "agreeing to disagree" result (1976). He showed that if agents share a common prior then common knowledge of posteriors implies agreement. Theorem 6 shows that if signals are unbounded then (common) knowledge of actions implies that agreement holds among an infinite group of agents, even if priors and utility functions differ.

### 6.2 Bounded signals

Recall that Theorem 2 shows that in a canonical setting with unbounded signals every CSLE is information aggregating.

What can be said about information aggregation when signals are bounded? Since independent of the signal structure there always exists an information aggregating equi-

librium<sup>14</sup>, the question is what is the worst possible equilibrium outcome in terms of learning. To answer this, we borrow the notion of information diffusion introduced by Lobel and Sadler (2015) in context of the sequential social learning model. Consider the support of the private belief and let its convex hull be  $[\beta_L, \beta_H]$ . For simplicity assume that the support is symmetric, i.e.,  $[\beta, 1 - \beta]$ .

An SLE in a canonical setting satisfies *information diffusion* if the probability of any agent's action being optimal is at least  $1 - \beta$ :

$$\mathbb{P}[a_i \in \operatorname{argmax}_{a \in A} u(a, \theta)] \geq 1 - \beta.$$

As Lobel and Sadler highlight, the notion of information diffusion is particularly insightful if strong signals, i.e., those that induce a posterior belief close to  $\beta$  or  $1 - \beta$ , are rare.

The next theorem, is a generalization of Theorem 2 to the bounded signal setting.

**Theorem 7.** *In a canonical setting every CSLE satisfies information diffusion.*

Likewise, the next theorem is a generalization of Theorem 3 to the bounded signal setting.

**Theorem 8.** *In a canonical setting every SLE that satisfies herding also satisfies information diffusion.*

### 6.3 Mixing signals and canonical\* settings

All of our results hold when we relax the requirement that signals are i.i.d., and require only a form of *mixing*.

We say that a sequence  $X_1, X_2, \dots$  of random variables is mixing if (1) the marginal distributions of the  $X_i$  are all identical, and (2) for each  $\varepsilon > 0$  there is a  $n(\varepsilon)$  such that for each agent  $i$  there are at most  $n(\varepsilon)$  agents  $j$  such that  $X_i$  and  $X_j$  are not  $\varepsilon$ -independent; that is, for every events  $E \in \sigma(X_i)$  and  $F \in \sigma(X_j)$  it holds that

$$\left| \mathbb{P}[E \cap F] - \mathbb{P}[E] \cdot \mathbb{P}[F] \right| < \varepsilon.$$

We say that private signals are mixing when they are mixing random variables, conditional on the state.

---

<sup>14</sup>To see this, consider the SLE where the assigned consensus action is optimal conditional on the realized signal vector.

Intuitively, private signals are mixing when for each agent  $i$  there are only finitely many other agents with whom  $i$  has a significantly correlated signal. One obvious example of mixing signals are i.i.d. signals. Other examples arise naturally when agents who are close to each other—either geographically or temporally—observe same or similar signals, but agents who are far away from each other observe only very weakly related signals.

The following lemma captures the Concentration of Dependence Principle for mixing random variables.

**Lemma 2** (Concentration of Dependence for Mixing Random Variables). *Let  $X_1, X_2, \dots$  be mixing random variables, and let  $E$  be any event defined on the same probability space. Then for every  $\varepsilon > 0$  there is an  $m(\varepsilon)$  such that, except for at most  $m(\varepsilon)$  many  $i$ 's, each  $X_i$  is  $\varepsilon$ -independent of  $E$ .*

In *canonical\** settings  $N$  is countably infinite,  $A = \Theta = \{0, 1\}$ , the utility function assigns 1 if the action of an agent matches the state and 0 otherwise (as in canonical settings), but signals are informative and mixing, rather than just i.i.d. As mentioned above, all of our results hold when we relax the requirement that signals are i.i.d., and require only mixing; in the Appendix we prove our theorems in this generality.

## 6.4 Large finite groups

All of our theorems are proved in settings with infinitely many agents. Analogous qualified statements for large finite groups follow from our proofs.

For example, Theorem 2 states that in a canonical setting, and when signals are unbounded, every CSLE is aggregating, so that

$$\mathbb{P}[a_i = \theta \text{ for all } i] = 1.$$

A careful reading of the proof shows that if we fix the marginal distributions of the (unbounded) private signals, then for every  $\varepsilon > 0$  there is a  $k(\varepsilon)$  such that in every CSLE with more than  $k(\varepsilon)$  players it holds that

$$\mathbb{P}[a_i = \theta \text{ for all } i] > 1 - \varepsilon.$$

## 7 Conclusion

We introduced the concept of social learning equilibrium as a useful tool to analyze social learning. The advantage over the conventional approach is that results can be derived and predictions can be made without knowing the exact dynamic of the interaction structure. We provide agreement, herding and information aggregation results of social learning equilibria that unify and shed additional insight on the social learning literature. In particular, we show that the relation between unbounded signals and the optimality of the herding action established by [Smith and Sørensen \(2000\)](#) holds much more generally. In fact, in any social learning environment with unbounded signals the action selected in a Bayesian herd is optimal.

The main value deriving from our analysis, however, is to show that the asymptotic equilibrium behavior of any social learning game might be analyzed via the static concept of social learning equilibrium, greatly simplifying the analysis.

There are several natural avenues for future research. One concerns the extension of our analysis to capture payoff externalities. Another concerns a weaker notion of herding. In this paper we focused on a strong notion of herding where cofinite agreement occurs almost surely, as in the sequential social learning models of [Bikhchandani et al. \(1992\)](#) and [Smith and Sørensen \(2000\)](#). Instead, the sequential models with random observation structures by [Acemoglu et al. \(2011\)](#) and [Lobel and Sadler \(2015\)](#) only satisfy herding in probability, where agents can be ordered such that the probability of selecting a herding action converges to one along the order. For random observation structures clearly our weak order condition fails. Thus an interesting question concerns the structural conditions on the SLEs that guarantee herding in probability. This is left for future work.

## References

- Daron Acemoglu, Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar. Bayesian learning in social networks. *The Review of Economic Studies*, 78(4):1201–1236, 2011.
- Anil Arya, Jonathan Glover, and Brian Mittendorf. Hierarchical reporting, aggregation, and information cascades. *Managerial and Decision Economics*, 27(5):355–362, 2006.
- Robert J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976. ISSN 0090-5364.



- Abhijit V Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.
- Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, 100(5):992–1026, 1992.
- Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of economic perspectives*, 12(3):151–170, 1998.
- Christophe Chamley. *Rational herds: Economic models of social learning*. Cambridge University Press, 2004.
- Varadarajan V Chari and Patrick J Kehoe. Hot money. *Journal of Political Economy*, 111(6):1262–1292, 2003.
- Eddie Dekel, Drew Fudenberg, and David K Levine. Payoff information and self-confirming equilibrium. *Journal of Economic Theory*, 89(2):165–185, 1999.
- Eddie Dekel, Drew Fudenberg, and David K Levine. Learning to play bayesian games. *Games and Economic Behavior*, 46(2):282–303, 2004.
- Peter DeMarzo and Costis Skiadas. Aggregation, determinacy, and informational efficiency for a class of economies with asymmetric information. *journal of economic theory*, 80(1):123–152, 1998.
- Peter DeMarzo and Costis Skiadas. On the uniqueness of fully informative rational expectations equilibria. *Economic Theory*, 13(1):1–24, 1999.
- Wenjing Duan, Bin Gu, and Andrew B Whinston. Informational cascades and software adoption on the internet: an empirical investigation. *MIS quarterly*, 33(1):23–48, 2009.
- Jayasri Dutta and Stephen Morris. The revelation of information and self-fulfilling beliefs. *journal of economic theory*, 73(1):231–244, 1997.
- Douglas Gale and Shachar Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 45(2):329–346, 2003.
- John D Geanakoplos and Heraklis M Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 28(1):192–200, 1982.

- Sanford J Grossman. An introduction to the theory of rational expectations under asymmetric information. *The Review of Economic Studies*, 48(4):541–559, 1981.
- Matthew O Jackson. An overview of social networks and economic applications. In *Handbook of social economics*, volume 1, pages 511–585. Elsevier, 2011.
- Naveen Khanna. Optimal contracting with moral hazard and cascading. *The Review of Financial Studies*, 11(3):559–596, 1998.
- Ilan Lobel and Evan Sadler. Information diffusion in networks through social learning. *Theoretical Economics*, 10(3):807–851, 2015.
- Paul Milgrom and Nancy Stokey. Information, trade and common knowledge. *Journal of economic theory*, 26(1):17–27, 1982.
- Deborah Minehart and Suzanne Scotchmer. Ex post regret and the decentralized sharing of information. *Games and Economic Behavior*, 27(1):114–131, 1999.
- Elchanan Mossel, Allan Sly, and Omer Tamuz. Asymptotic learning on bayesian social networks. *Probability Theory and Related Fields*, 158(1-2):127–157, 2014.
- Elchanan Mossel, Allan Sly, and Omer Tamuz. Strategic learning and the topology of social networks. *Econometrica*, 83(5):1755–1794, 2015.
- Manuel Mueller-Frank. A general framework for rational learning in social networks. *Theoretical Economics*, 8(1):1–40, 2013.
- Michael Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–2647, 2012.
- Rohit Parikh and Paul Krasucki. Communication, consensus, and knowledge. *Journal of Economic Theory*, 52(1):178–189, 1990. ISSN 0022-0531.
- Roy Radner. Rational expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica: Journal of the Econometric Society*, 47(3):655–678, 1979.
- Dinah Rosenberg, Eilon Solan, and Nicolas Vieille. Informational externalities and emergence of consensus. *Games and Economic Behavior*, 66(2):979–994, 2009. ISSN 0899-8256.
- Ariel Rubinstein and Asher Wolinsky. Rationalizable conjectural equilibrium: between nash and rationalizability. *Games and Economic Behavior*, 6(2):299–311, 1994.

David S Scharfstein and Jeremy C Stein. Herd behavior and investment. *The American Economic Review*, 80(3):465–479, 1990.

James K Sebenius and John Geanakoplos. Don't bet on it: Contingent agreements with asymmetric information. *Journal of the American Statistical Association*, 78(382):424–426, 1983.

Lones Smith and Peter Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.

Xavier Vives. *Information and learning in markets: the impact of market microstructure*. Princeton University Press, 2010.

Eric A. Walden and Glenn J. Browne. Information cascades in the adoption of new technology. In *Proceedings of the Twenty-Third International Conference on Information Systems*, pages 435–443, 2002.

Ivo Welch. Sequential sales, learning, and cascades. *The Journal of finance*, 47(2):695–732, 1992.

## A Agreement and Information Aggregation in Complete Social Learning Equilibria

In this section we prove Theorem 1. We start with the following lemma, which is essentially a formulation of the No Trade Theorem of [Milgrom and Stokey \(1982\)](#). This lemma states that when there is disagreement then players must be indifferent.

In the context of an SLE, denote agent  $i$ 's equilibrium belief by

$$q_i = \mathbb{P}[\theta = 1 \mid \ell_i, s_i],$$

and let  $D$  be the event that  $a_i \neq a_j$  for some  $i, j \in N$ .

**Lemma 3.** *Let  $A = \Theta = \{0, 1\}$  and  $u(a, \theta) = 1_{a=\theta}$  (as in a canonical setting, but with no restrictions on the signals). In any SLE, if the event  $D$  has positive probability, then conditioned on  $D$  it almost surely holds that  $q_i = 1/2$  for all  $i$ .*

**Proof.** Consider an outside observer who observes all the agents' actions  $\bar{a}$ . Her belief is

$$q_* = \mathbb{P}[\theta = 1 \mid \bar{a}].$$

Since  $\bar{a}$  is  $\sigma(\ell_i, s_i)$ -measurable, it follows from the law of total expectations that for every  $i$

$$q_* = \mathbb{E}[q_i | \bar{a}]. \quad (4)$$

Since 1 is the action that is optimal for beliefs above  $1/2$ , we have that  $a_i = 1$  implies that  $q_i \geq 1/2$ . Likewise,  $a_i = 0$  implies  $q_i \leq 1/2$ . Hence the claim follows by (4). ■

Thus, to prove Theorem 1, it suffices to show that indifference is impossible, which, as we now show, follows from mixing. In fact, we prove it more generally, for canonical\* settings (see §6.3).

*Proof of Theorem 1.* It follows from Lemma 3 that to prove our claim it suffices to show that the *indifference event*  $F = \{q_i = 1/2 \text{ for all } i\}$  occurs with probability 0.

Assume towards a contradiction that  $F$  has positive probability, and consider, for each player  $i$ , the deviation of following her private signal only whenever  $F$  occurs, and choosing action  $a_i$  otherwise. That is, choosing

$$b_i \in \operatorname{argmax}_{a \in A} \mathbb{P}[\theta = a | s_i],$$

whenever  $F$  occurs, and  $a_i$  otherwise. The profit player  $i$  stands to gain from this deviation is

$$P = \mathbb{P}[b_i = \theta, F] - \mathbb{P}[a_i = \theta, F].$$

The second term is simply equal to  $\frac{1}{2}\mathbb{P}[F]$ , since conditioned on  $F$  player  $i$  is indifferent, and so her expected utility from any action is  $1/2$ .

Let

$$Q = \mathbb{P}[b_i = \theta]$$

be the expected utility of a player who always follows her own signal. Then  $Q > 1/2$ , since signals are informative. By the Concentration of Dependence Principle,

$$\lim_{i \rightarrow \infty} \mathbb{P}[b_i = \theta, F | \theta] = \mathbb{P}[b_i = \theta | \theta] \cdot \mathbb{P}[F | \theta],$$

where the right-hand side does not depend on  $i$ , since the  $b_i$ 's are identically distributed.

Hence also unconditionally

$$\lim_{i \rightarrow \infty} \mathbb{P}[b_i = \theta, F] = \mathbb{P}[b_i = \theta] \cdot \mathbb{P}[F],$$

and we have that

$$\lim_{i \rightarrow \infty} \mathbb{P}[b_i = \theta, F] = Q \cdot \mathbb{P}[F] > \frac{1}{2} \mathbb{P}[F] = \mathbb{P}[a_i = \theta, F].$$

Hence for some agent  $i$  (in fact, all but finitely many),  $b_i$  is a profitable deviation, and we have reached a contradiction with our equilibrium assumption. ■

## B Herding

In this section we prove Theorem 4, in the more general framework of canonical\* settings (see §6.3).

Let

$$x_i = \mathbb{P}[\theta = 1 \mid a_1, \dots, a_i]$$

be the sequence of *public beliefs*, and let

$$q_i = \mathbb{P}[\theta = 1 \mid \ell_i, s_i]$$

be agent  $i$ 's equilibrium belief. Note that, since each agent  $i$  knows  $\{a_1, \dots, a_i\}$ ,

$$x_i = \mathbb{E}[q_i \mid a_1, \dots, a_i], \tag{5}$$

by the law of total expectations.

Note that the action 1 is optimal for beliefs  $1/2$  and higher, and the action 0 is optimal for beliefs  $1/2$  and lower. Therefore, and since  $\bar{a}$  is an equilibrium,

$$a_i = 1 \Rightarrow q_i \geq 1/2 \quad \text{and} \quad a_i = 0 \Rightarrow q_i \leq 1/2 \tag{6}$$

and

$$\mathbb{E}[u(a_i, \theta) \mid q_i] = \mathbb{P}[a_i = \theta \mid q_i] = \max\{q_i, 1 - q_i\}. \tag{7}$$

We start with two simple claims regarding  $a_i$  and  $x_i$ .

**Claim 1.** *If  $a_i = 1$  then  $x_i \geq 1/2$ . If  $a_i = 0$  then  $x_i \leq 1/2$ .*

**Proof.** By (6) we have that  $q_i \geq 1/2$  conditioned on  $a_i = 1$ . Hence, by (5),  $x_i \geq 1/2$  conditioned on  $a_i = 1$ . An analogous argument holds for the case  $a_i = 0$ . ■

**Claim 2.**  $\mathbb{P}[a_i = \theta \mid x_i] = \max\{x_i, 1 - x_i\}$ .

**Proof.** By Claim 1

$$\mathbb{P}[\theta = a_i \mid x_i] = \begin{cases} \mathbb{P}[\theta = 1 \mid x_i] & \text{if } x_i > 1/2 \\ \mathbb{P}[\theta = 0 \mid x_i] & \text{if } x_i < 1/2 \\ \mathbb{P}[\theta = a_i \mid x_i] & \text{if } x_i = 1/2. \end{cases}$$

By (5) and (6), if  $x_i = 1/2$  then  $x_i = q_i$ . Therefore, and since  $\mathbb{P}[\theta = 1 \mid x_i] = x_i$ , and  $\mathbb{P}[a_i = \theta \mid q_i = 1/2] = 1/2$  by (7),

$$\mathbb{P}[\theta = a_i \mid x_i] = \begin{cases} x_i & \text{if } x_i > 1/2 \\ 1 - x_i & \text{if } x_i < 1/2 \\ 1/2 & \text{if } x_i = 1/2. \end{cases}$$

Thus  $\mathbb{P}[\theta = a_i \mid x_i] = \max\{x_i, 1 - x_i\}$ . ■

Let

$$x = \mathbb{P}[\theta = 1 \mid \bar{a}],$$

and note that  $x_i$  is a bounded martingale that converges a.s. to  $x$ . It thus follows from Claim 1 that conditioned on  $a_i$  taking both values infinitely often it holds that  $x = 1/2$ . Thus, to prove our theorem, we will show that the probability of  $x = 1/2$  is zero. Accordingly, define the event

$$F^0 = \{x = 1/2\},$$

and for  $\varepsilon > 0$  define the events

$$F_i^\varepsilon = \{x_i \in (1/2 - \varepsilon, 1/2 + \varepsilon)\}.$$

The event  $F_i^\varepsilon$  is the event that the public belief  $x_i$  is close to  $1/2$ . Since the sequence  $(x_i)_i$

converges a.s. to  $x$ , we have that

$$\lim_{i \rightarrow \infty} \mathbb{P}[F_0 \setminus F_\varepsilon^i] = 0 \quad (8)$$

for every  $\varepsilon > 0$ , and that

$$\lim_{\varepsilon \rightarrow 0} \limsup_{i \rightarrow \infty} \mathbb{P}[F_i^\varepsilon] = \mathbb{P}[F^0]. \quad (9)$$

Thus, to prove that  $\mathbb{P}[F^0] = 0$ —which, as we explained above, proves the claim—it suffices to show that the left hand side of the above expression vanishes.

To this end, let

$$b_i = b(s_i) \in \operatorname{argmax}_{a \in A} \mathbb{P}[\theta = a \mid s_i]$$

be an optimal action chosen given agent  $i$ 's private signal only. These are all chosen using the same function  $b$ , and so, since the private signals  $s_i$  are identically distributed (but not necessarily independently), the random variables  $b_i$  are also identically distributed.

Imagine that player  $i$  chooses  $b_i$  instead of  $a_i$ , whenever  $F_i^\varepsilon$  occurs. Then player  $i$ 's gain in expected utility from this deviation is

$$\mathbb{P}[b_i = \theta, F_i^\varepsilon] - \mathbb{P}[a_i = \theta, F_i^\varepsilon].$$

We prove that the left-hand side of (9) vanishes by showing that

$$\lim_{\varepsilon \rightarrow 0} \limsup_{i \rightarrow \infty} \mathbb{P}[b_i = \theta, F_i^\varepsilon] - \mathbb{P}[a_i = \theta, F_i^\varepsilon] > 0,$$

and thus this is a profitable deviation for some  $\varepsilon$  small enough and  $i$  large enough, contradicting the assumption that  $\bar{a}$  is an SLE.

To this end, we note that

$$\mathbb{P}[b_i = \theta, F_i^\varepsilon] \geq \mathbb{P}[b_i = \theta, F^0] - \mathbb{P}[b_i = \theta, F^0 \setminus F_i^\varepsilon],$$

since

$$F^0 \setminus (F^0 \setminus F_i^\varepsilon) = F^0 \cap F_i^\varepsilon \subseteq F_i^\varepsilon.$$



It thus follows by (8) that

$$\liminf_{i \rightarrow \infty} \mathbb{P}[b_i = \theta, F_i^\varepsilon] \geq \liminf_{i \rightarrow \infty} \mathbb{P}[b_i = \theta, F^0].$$

Now, conditioned on  $\theta$ , the  $s_i$ 's are i.i.d., and in particular mixing. Hence it follows from the Concentration of Dependence Principle that

$$\lim_{i \rightarrow \infty} \mathbb{P}[b_i = \theta, F^0 \mid \theta] = \mathbb{P}[b_i = \theta \mid \theta] \cdot \mathbb{P}[F^0 \mid \theta],$$

where the right-hand side does not depend on  $i$ , since the  $b_i$ 's are identically distributed. Hence

$$\lim_{i \rightarrow \infty} \mathbb{P}[b_i = \theta, F^0] = \mathbb{P}[b_i = \theta] \cdot \mathbb{P}[F^0].$$

Since private signals are informative, it follows that  $\mathbb{P}[b_i = \theta] > 1/2$ , and so we have that

$$\liminf_{\varepsilon \rightarrow 0} \liminf_{i \rightarrow \infty} \mathbb{P}[b_i = \theta, F_i^\varepsilon] > \frac{1}{2} \mathbb{P}[F^0]. \quad (10)$$

Now,

$$\begin{aligned} \mathbb{P}[a_i = \theta \mid F_i^\varepsilon] &= \mathbb{E}[\mathbb{P}[a_i = \theta \mid x_i] \mid F_i^\varepsilon] \\ &= \mathbb{E}[\max\{x_i, 1 - x_i\} \mid F_i^\varepsilon], \end{aligned}$$

where the second equality is an application of Claim 2. Since  $x_i \in (1/2 - \varepsilon, 1/2 + \varepsilon)$  conditioned on  $F_i^\varepsilon$ , we get that

$$\mathbb{P}[a_i = \theta, F_i^\varepsilon] < \left(\frac{1}{2} + \varepsilon\right) \cdot \mathbb{P}[F_i^\varepsilon].$$

Therefore, by (9),

$$\lim_{\varepsilon \rightarrow 0} \limsup_{i \rightarrow \infty} \mathbb{P}[a_i = \theta, F_i^\varepsilon] \leq \frac{1}{2} \cdot \mathbb{P}[F^0].$$

Therefore, in combination with (10), the expected profit from deviating from  $a_i$  to  $b_i$  on  $F_i^\varepsilon$  satisfies

$$\lim_{\varepsilon \rightarrow 0} \limsup_{i \rightarrow \infty} \mathbb{P}[b_i = \theta, F_i^\varepsilon] - \mathbb{P}[a_i = \theta, F_i^\varepsilon] > 0,$$

and thus this is a profitable deviation for some  $\varepsilon$  small enough and  $i$  large enough. Hence

it follows that  $F^0$  has probability zero, concluding the proof of Theorem 4.

## C Concentration of Dependence

*Proof of Lemmas 1 and 2.* Choose any  $\varepsilon > 0$  and let  $X_{i_1}, X_{i_2}, \dots, X_{i_k}$  be  $k$  random variables that are not  $\varepsilon$ -independent of  $E$ . Without loss of generality we may assume that  $(i_1, i_2, \dots, i_k) = (1, \dots, k)$ . Let  $F_1, \dots, F_k$  be events that witness the violation of  $\varepsilon$ -independence, so that for  $i = 1, \dots, k$

$$\left| \mathbb{P}[E \cap F_i] - \mathbb{P}[E] \cdot \mathbb{P}[F_i] \right| \geq \varepsilon, \quad (11)$$

and each  $F_i$  is in  $\sigma(X_i)$ .

Let  $Y$  be the indicator of the event  $E$ . Then we can write (11) as

$$\left| \text{Cov}(F_i, Y) \right| \geq \varepsilon.$$

Let  $\eta_i \in \{-1, +1\}$  equal the sign of  $\text{Cov}(F_i, Y)$ . Then (11) is equivalent to

$$\text{Cov}(\eta_i F_i, Y) \geq \varepsilon.$$

Summing over  $i$  we get

$$\sum_{i=1}^k \text{Cov}(\eta_i F_i, Y) \geq k\varepsilon.$$

By additivity of covariance, it follows that

$$\text{Cov}\left(\sum_{i=1}^k \eta_i F_i, Y\right) \geq k\varepsilon.$$

By the Cauchy-Schwarz inequality it follows that

$$\sqrt{\text{Var}\left(\sum_{i=1}^k \eta_i F_i\right) \cdot \text{Var}(Y)} \geq k\varepsilon.$$

Denote  $I = \sum_{i=1}^k \eta_i F_i$ , and note that  $\text{Var}(Y) \leq 1$ , since  $Y \in \{0, 1\}$ . So, squaring both sides

yields

$$\text{Var}(I) \geq k^2 \varepsilon^2. \tag{12}$$

Consider first the i.i.d. case. Then  $\text{Var}(I)$  is at most  $k$ , since  $I$  is the variance of  $k$  independent random variables, each with variance at most 1. Hence we have that

$$k \geq k^2 \varepsilon^2$$

or  $k \leq 1/\varepsilon^2$ . This completes the proof of Lemma 1.

Consider now the mixing case. Write

$$\text{Var}(I) = \sum_{i=1}^k \sum_{j=1}^k \text{Cov}(\eta_i F_i, \eta_j F_j)$$

Now, by the mixing property, for each  $i$  there are at most  $n = n(\varepsilon^2/2)$  possible  $j$ 's such that  $\text{Cov}(\eta_i F_i, \eta_j F_j) > \varepsilon^2/2$ , in which case it is at most 1. Hence

$$\text{Var}(I) \leq k(n \cdot 1 + (k - n) \cdot \varepsilon^2/2) \leq kn + k^2 \varepsilon^2/2.$$

Applying this back into (12) yields

$$kn + k^2 \varepsilon^2/2 \geq k^2 \varepsilon^2.$$

Rearranging yields

$$k \leq 2n\varepsilon^2,$$

and so we can have proved Lemma 2, with  $m(\varepsilon) = 2n(\varepsilon)/\varepsilon^2$ . ■

## D We must learn to agree

We say that an SLE satisfies *herding in probability* if there is a random variable  $a^*$  taking values in  $A$  such that

$$\lim_{i \rightarrow \infty} \mathbb{P}[a_i = a^*] = 1.$$

Here the limit is taken by arbitrarily identifying the agents with the set of natural numbers.

Theorem 3 is a consequence of the following, stronger statement that applies to herding in probability, rather than (almost sure) herding in which a cofinite set of agents chooses the same action.

**Theorem 9.** *In canonical\* setting, every SLE that satisfies herding in probability also satisfies information diffusion.*

Let

$$p = \mathbb{P}[a^* = \theta].$$

It follows from herding in probability that  $\lim_i \mathbb{P}[a_i = a^*] = 1$ , and so

$$\begin{aligned} \lim_i \mathbb{P}[a_i = \theta] &= \lim_i \mathbb{P}[a_i = \theta, a_i = a^*] + \mathbb{P}[a_i = \theta, a_i \neq a^*] \\ &= \lim_i \mathbb{P}[a^* = \theta, a_i = a^*] + \mathbb{P}[a_i = \theta, a_i \neq a^*] \\ &= p. \end{aligned}$$

Assume by contradiction that

$$p \leq 1 - \beta - 2\varepsilon$$

for some  $\varepsilon > 0$ .

As in the proof of Theorem 4, let

$$b_i = b(s_i) \in \operatorname{argmax}_{a \in A} \mathbb{P}[\theta = a \mid s_i]$$

be an optimal action chosen given agent  $i$ 's private signal only. These are all chosen using the same function  $b$ , and so, since the private signals  $s_i$  are identically distributed (but not necessarily independently), the random variables  $b_i$  are also identically distributed. Let  $B_i$  be the event that  $\mathbb{P}[b_i = \theta \mid s_i] > 1 - \beta - \varepsilon$ . Since the  $b_i$ 's are identically distributed, all of the events  $B_i$  have the same probability. Furthermore, this probability is positive, by our assumption on the support of the private signals.

Imagine that agent  $i$  deviates and chooses  $b_i$  whenever  $B_i$  occurs, and otherwise follows  $a_i$ . Then her expected gain in utility is

$$\mathbb{P}[b_i = \theta, B_i] - \mathbb{P}[a_i = \theta, B_i].$$

To bound the first term, we note that, by the definition of  $B_i$ ,

$$\mathbb{P}[b_i = \theta, B_i] \geq (1 - \beta - \varepsilon)\mathbb{P}[B_i].$$

To bound the second term, we write

$$\begin{aligned} \mathbb{P}[a_i = \theta, B_i] &= \mathbb{P}[a_i = \theta, a_i = a^*, B_i] + \mathbb{P}[a_i = \theta, a_i \neq a^*, B_i] \\ &= \mathbb{P}[a^* = \theta, a_i = a^*, B_i] + \mathbb{P}[a_i = \theta, a_i \neq a^*, B_i] \end{aligned}$$

Since  $\bar{a}$  satisfies herding in probability,  $\lim_i \mathbb{P}[a_i = a^*] = 1$ , and so it follows that

$$\limsup_i \mathbb{P}[a_i = \theta, B_i] = \limsup_i \mathbb{P}[a^* = \theta, B_i].$$

Since private signals are conditionally mixing, it follows from Lemma 2 that

$$\limsup_i \mathbb{P}[a^* = \theta, B_i] = \mathbb{P}[a^* = \theta]\mathbb{P}[B_i] = p \cdot \mathbb{P}[B_i],$$

where the right-hand side does not depend on  $i$ , since the events  $B_i$  all have the same probability. We have thus shown that

$$\limsup_i \mathbb{P}[a_i = \theta, B_i] = p \cdot \mathbb{P}[B_i],$$

Combining the bounds on the two terms we get that the expected gain in utility is

$$\liminf_i \mathbb{P}[b_i = \theta, B_i] - \mathbb{P}[a_i = \theta, B_i] \geq (1 - \beta - \varepsilon - p)\mathbb{P}[B_i].$$

Since we assumed that  $p \leq 1 - \beta - 2\varepsilon$  we have that this is at least  $\varepsilon\mathbb{P}[B_i]$ , and in particular positive. Thus  $\bar{a}$  is not an equilibrium, as for some  $i$  large enough player  $i$  would have a profitable deviation. This completes the proof of Theorem 9.

## E Proof of Theorem 5

This proof is essentially a recasting of the proof of Proposition 2.1 in [Rosenberg et al. \(2009\)](#) to our language and notation.

Fix an agent  $i$ . The case that  $\delta = 0$  or  $T_i$  is finite is immediate. We thus assume henceforth that  $\delta > 0$  and  $|T_i| = \infty$ .

Let

$$v_i = \max_{a \in A} \mathbb{E}[u(a, \theta) \mid \mathbf{k}_i, s_i]$$

be the maximum expected utility agent  $i$  can guarantee given what she (asymptotically) knows at the end of the game.

Fix  $(\mathbf{k}_i, s_i)$  and  $\varepsilon > 0$ , and let  $U, W \subseteq A$  be the sets of actions given by

$$U = \{a \in A : \mathbb{E}[u(a, \theta) \mid \mathbf{k}_i, s_i] > v_i - \varepsilon\}$$

and

$$W = \{b \in A : \mathbb{E}[u(b, \theta) \mid \mathbf{k}_i, s_i] < v_i - 3\varepsilon\}.$$

That is,  $U$  is the set of actions that is  $\varepsilon$ -optimal, and  $W$  is the set of actions that is  $3\varepsilon$ -suboptimal—conditioned on the information available to the player at the end of the game.

Note that the sets  $U$  and  $W$  are open and disjoint and that utilities are continuous. It therefore follows from the martingale convergence theorem that, for any  $\eta > 0$  and  $t \in T_i$  large enough, it holds for every  $a \in U$  and  $b \in W$  that

$$\mathbb{P}[\mathbb{P}[u(a, \theta) > u(b, \theta) + \varepsilon \mid \mathbf{k}_{i,t}, s_i] > 1 - \eta \mid \mathbf{k}_i, s_i] = 1.$$

That is, for large enough  $t$ , the agent will almost surely assign high probability to the event that any action  $a \in U$  yields at least  $\varepsilon$  more utility than any  $b \in W$ .

It follows that choosing any  $b \in W$  will, for large enough  $t$ , result in an expected utility loss of at least  $\varepsilon \cdot (1 - \eta) \cdot (1 - \delta)$  in the subgame starting at  $t$ , which, for  $\eta$  small enough, is greater than  $\delta \cdot \max_{a, \theta} u(a, \theta)$ , and thus greater than the continuation utility of any strategy. It follows that, in equilibrium and conditioned on almost every  $(\mathbf{k}_i, s_i)$ , the agent eventually stops choosing actions in  $W$ . Since this holds for every  $\varepsilon$ , it follows that every limit point of actions taken by the agents almost surely maximizes her expected utility, conditioned on  $(\mathbf{k}_i, s_i)$ .