MULTI-DIMENSIONAL SOCIAL LEARNING

Itai Arieliy

Manuel Mueller-Frank

The Public-Private Sector Research Center is a Research Center based at IESE Business School. Its mission is to develop research that analyzes the relationships between the private and public sectors primarily in the following areas: regulation and competition, innovation, regional economy and industrial politics and health economics.

Research results are disseminated through publications, conferences and colloquia. These activities are aimed to foster cooperation between the private sector and public administrations, as well as the exchange of ideas and initiatives.

The sponsors of the Public-Private Sector Research Center are the following:

- Ajuntament de Barcelona
- Departament d' Economia i Coneixement de la Generalitat de Catalunya
- Departament d' Empresa i Ocupació de la Generalitat de Catalunya
- Diputació de Barcelona
- EVERIS
- Fundació AGBAR
- Institut Català de les Empreses Culturals (ICEC)
- PricewaterhouseCoopers
- Sanofi

The contents of this publication reflect the conclusions and findings of the individual authors and not the opinions of the Center's sponsors.

Multi–Dimensional Social Learning[*]

Itai Arieli[†], Manuel Mueller-Frank[‡]

October 2, 2014

Abstract

This paper provides a model of social learning where the order in which actions are taken is determined by an m-dimensional integer lattice rather than along a line as in the sequential social learning model. The observation structure is determined by a random network. Every agent links to each of his preceding lattice neighbors independently with probability p, and observes the actions of all agents that are reachable via a directed path in the realized social network. We establish a strong discontinuity of learning with respect to the linkage probability. If p is close to but different from one an arbitrary high proportion of agents select the optimal action in the limit, for any informative signal structure. For bounded signals and a linkage probability equal to one, however, there exists a positive probability that all agents select the suboptimal action. We also show that for every $p < 1$ and every unbounded signal structure there exists a bounded signal structure with strictly better learning properties. Finally, we address robustness by providing a general characterization of multi-dimensional structures for which our main results carry forward.

1

1 Introduction

Social learning forms an important part of daily life. We observe the choices of others and take this into account when making decisions. Typically, our choices are influenced by our own private information as well as the information we infer from the actions of those we observe. One might then conclude that as the society grows large, the majority of individuals take the optimal action with a probability that goes to one. But the validity of this conclusion depends on the specific interaction structure among the individuals.

The standard sequential social learning model (henceforth SSLM) as introduced by Bikhchandani, Hirshleifer and Welch [8] and Banerjee [6] feature a countable set of fully rational agents making a one-time, binary choice in a predetermined sequence. Each agent receives a conditional independent and identically distributed private signal concerning the realized state of the world and observes the choices of all his predecessors. These seminal papers show that the probability of the majority taking the correct action might fail to converge to one as the society grows. More precisely, from a finite time onward all subsequent agents may take the suboptimal action.

The SSLM relies on two main features. The first feature is a strict sequential order of choices, where one agent per period makes a choice. The second feature is that all agents commonly know the history of choices of all their predecessors. These main features of the SSLM fail to adequately represent many real world environments where groups of individuals act simultaneously and the size of the groups is growing over time. Consider for example the adoption of a new product. As time goes by an increasing number of individuals become aware of the existence of the product and then face the decision of whether to adopt it or not. The size of the group that decides whether or not to adopt in any given period hence grows over time. The active user base of online networking sites such as Facebook and Twitter might serve as a real world example as it was subject to spectacular growth over a number of years from their inception.

This paper introduces a model which features (i) a multi-dimension choice order, where in each discrete point in time a group of agents simultaneously take an irreversible action, (ii) a growth in the group size over time, and (iii) limited observability of the history for each agent.

The central phenomena in the social learning literature are information cascades, where

2

from a finite time period onward all agents ignore their private signal and base their choice only on the social information that is accumulated in the history of previous choices, and asymptotic learning, where the probability of selecting the correct action goes to one over time. Bikhchandani, Hirshleifer and Welch [8] show that in environments with finitely many signals information cascades always occur. Smith and Sorensen [22] establish that asymptotic learning holds always, if the signal structure is unbounded, and fails always, if the signal structure is bounded.

In the context of groups of agents acting in each period, it might be too strong to require all of them to act optimally or independent of their signal. Instead, it is more natural to focus on the proportion of agents. Our novel notion of learning, α-proportional asymptotic learning is satisfied if the probability, that a proportion of at least α of the agents deciding at time t select the correct action, converges to one as t grows to infinity. Similarly, a γ-proportional information cascade occurs if the probability, that a proportion of at least γ of agents in round t select their action independent of their private signal, goes to one as t grows.

The general question we are interested in is whether the properties of social learning that have been established in the literature carry forward to the more general multi-dimensional structure. In particular, we are concerned with sufficient conditions for the occurrence of proportional information cascades and proportional asymptotic learning.

A standard mathematical structure that features a multi-dimensional choice order and generational growth is the lattice which is an often used structure in the economics literature. One example herefore is the contagion literature that uses lattices to describe the local interactions among players, see for example Morris [19], and Lee and Valentinyi [16].

In the central part of the paper we consider the following model. The agents are organized in the non-negative orthant of an m-dimensional integer lattice. The round in which a given agent acts, that is selects one action out of a binary set, is determined by the lattice distance to the origin. In round 1 the agent placed in the origin (at $\mathbf{0}$) acts. All agents whose lattice distance from the origin is exactly $t - 1$ act in round t. The number of agents selecting an action in a given round is increasing from round to round, whenever the lattice dimension m is greater than one. The observation structure is determined by a social network that builds upon the lattice structure. Each agent links with each of his (lattice) neighboring predecessors

independently with probability p and observes the actions chosen by all previous agents that form part of his extended social network. That is, an agent observes the actions of those that can be reached by a directed path initiating from the agent. The random graph underlying our observation structure is known as lattice percolation and is well studied in probability theory.[1] Note that the SSLM is captured by a one-dimensional lattice with linkage probability $p = 1$ as actions are taken along a line and hence in a strict sequential order. One may therefore view the approach of representing the order of choices and the observational structure via a lattice as a natural generalization of the SSLM.

We consider both the case of a deterministic social network, $p = 1$, and the random social network case with a linkage probability $p < 1$. It is now apparent why the traditional definitions of learning and cascades are too strict. For a linkage probability $p < 1$ both asymptotic learning as well as information cascades are impossible due to the fact that each given agent observes no history with positive (bounded away from zero) probability in which case he bases his action only on his private signal. The proportion of agents that act optimally, or independent of their signal, is hence a natural object to consider.

We focus primarily on a lattice structure since it has some elegant features that ease the analysis. From an economic perspective, however, there is no particular reason to focus on a lattice structure as opposed to a different multi-dimensional structure that features a growing number of decisions per period over time. A natural question then is whether our analysis is robust or depends on the properties of the lattice structure. As we show in section 5, our main results carry forward to a more general multi-dimensional observation structure.

1.1   Summary of Results: Deterministic Social Network

We first characterize asymptotic learning in the deterministic social network with a linkage probability $p = 1$. This characterization depends upon the signal structure. First, we establish that α-proportional asymptotic learning is satisfied for any $\alpha < 1$ if private signals are unbounded using a result from Acemoglu et al. [1]. Second, we show that α-proportional asymptotic learning fails for every $\alpha > 0$ if private signals are bounded. As defined in Smith and Sorensen [22], private signals are unbounded if the support of the conditional private proba-

---

[1]See Durrett [12] and Bollobas and Riordan [9]

bility (of either state) contains 0 and 1, and bounded if the support includes neither 0 nor 1. For $p = 1$, our result on asymptotic learning is qualitatively very close to Smith and Sorensen's [22] characterization in the SSLM which establishes that asymptotic learning fails for bounded signals and is satisfied for unbounded signals, in any equilibrium.

We find a crucial distinction to the SSLM in regards to the occurrence of information cascades. Herrera and Hoerner [15] consider the bounded signal case and characterize the absence of information cascades in terms of properties of the signal structure.[2] In contrast, we establish that there exists a positive probability such that from a finite time on all agents, that do not lie on the boundary of the lattice, cascade. This implies that, in the limit, an arbitrary high proportion of agents cascade for any bounded signal structure, with positive probability.

1.2 Summary of Results: Random Social Network

Next we consider the case of a random social network. Our main result establishes that for every $\alpha > 0$ there exists a corresponding linkage probability $p(\alpha)$ such that α-proportional asymptotic learning is satisfied for all $p \in (p(\alpha), 1)$ in any Perfect Bayesian equilibrium of the game, *independent of the signal structure.* Hence as p goes to 1, an arbitrarily large proportion of agents in late enough generations act optimally.

The intuition for the result is the following. We establish that under the lattice percolation process underlying the observation structure an arbitrary high proportion of late movers observe unboundedly many agents that are isolated, i.e. who do not observe anybody. Since the actions of the isolated agents are conditionally independent, observing unboundedly many such agents is sufficient to learn the true state of the world.

Our observation structure features a strong discontinuity at the linkage probability $p = 1$ in the bounded signals case. While α-proportional asymptotic learning fails for every $\alpha > 0$ at $p = 1$, choosing a linkage probability p sufficiently close to one, induces the proportion of agents selecting the correct action to approximate one, with a probability that goes to one. Therefore, introducing infinitesimal observational noise via a random social network can

---

[2]By definition information cascades can only occur for bounded signals as under unbounded signals there exists a positive probability that an agent overturns the decision of his predecessor, for any prior probability measure.

significantly improve learning outcomes in a setting with bounded private signals.

Finally, we show that for every $\gamma < 1$ there exists a linkage probability $p(\gamma)$ such that for every $p \in (p(\gamma), 1)$ a γ-proportional information cascade occurs in any Perfect Bayesian equilibrium of the game. In other words, adjusting for the linkage probability an arbitrary high proportion of agents ignore their signal in the long run with probability one.

1.3   Discussion of Results

There are three conclusions from our analysis worth emphasizing. First, allowing for multi-dimensional decision making leads to a clean relationship between observational noise and the proportion of asymptotic learners. In a social learning environment where the order of agents' choices is represented by a multi-dimensional model, the linkage probability p becomes the key determinant of proportional asymptotic learning rather than properties of the signal space as in the standard strict sequential model. To be more precise, a specific p provides a lower bound for the proportion of agents acting optimally in any informative signal structure. As p goes to one, this lower bound approaches one.

Second, our model provides a new insight on the learning properties of bounded versus unbounded signals. Throughout the literature, asymptotic learning occurs in any equilibrium regardless of the signal distribution as long as it is unbounded and typically fails whenever signals are bounded. In contrast, we show that for any linkage probability $p < 1$ in some cases bounded signal distributions achieve superior learning in equilibrium than unbounded signal distributions. More precisely, for every unbounded signal distribution there exists a bounded signal distribution for which the probability that the proportion of agents taking the correct action is asymptotically larger than the corresponding proportion under the unbounded distribution.

Third, within the multi-dimensional framework we show an interesting relationship between the occurrence of information cascades and asymptotic learning. In the strict sequential setting information cascades preclude asymptotic learning: if an agent n selects an action independent of his signal, then so will all later agents. As this action is only based upon finitely many signals, those of agent n's predecessors, the probability of the correct action being chosen is bounded away from one. In contrast, our results establish a coexistence of proportional information

cascades and proportional asymptotic learning under bounded signals and lattice dimensions greater than one. Controlling for the linkage probability, both the proportion of agents that select an action independent of their signal and the proportion of agents that select the correct action may exceed any threshold in the long run with probability one.

1.4   Relation to Literature

The most closely related papers are Acemoglu, Dahleh, Lobel and Ozdaglar [1], and Lobel and Sadler [17]. They study strict sequential decision making in a framework that relaxes the assumption of common observability of the history. Both papers consider a framework where the set of agents that a given agent observes is randomly drawn among all his predecessors. Acemoglu et al. [1] assume that these neighborhood draws are independent across agents and independent of the signals. Their main result demonstrates the power of unbounded signals, by showing that asymptotic learning holds for a large variety of such random network topologies.

Lobel and Sadler [17] generalize the approach of Acemoglu et al. [1] by allowing dependence of the neighborhood draws across agents and by introducing a weaker concept of learning, where agents are not required to select the correct action with probability one but only with a probability that an expert selects the correct action.

Interestingly, for $p = 1$ the m-dimensional lattice is included in the class of observation structures for which Lobel and Sadler [17] show that their weaker notion of learning is satisfied. For $p < 1$, however, the random observation structure generated by the lattice is included in the class of observation structure for which their concept of learning fails. This follows from the fact that for $p < 1$ any given agent observes no history with a probability that is bounded away from zero across agents. We note that the general multi-dimensional structure for $p = 1$ considered in Section 5 is incomparable with the two classes characterized by Lobel and Sadler [17].

Our main results (Theorem 2 and Theorem 4) concerning the learning properties in the multi-dimensional framework are driven by the feature that most agents have access to an over time growing subset of isolated decision makers. That is, decision makers that take their action in isolation and as a result reveal more about their private information. This is the main characteristic that enables learning in our model. The idea of *sacrificial lambs* is not

new to the literature. In fact, Smith [21] demonstrates that an infinite but proportionally vanishing set of isolated agents (or "sacrificial lambs") enables learning. Acemoglu et al. [1] generalize this idea and characterize sufficient conditions on the random observation structure such that learning holds under bounded signals. This is done via a carefully designed random structure such that every network realization has the property considered by Smith. In the experimental literature, Anderson and Holt [2], and Goeree, Palfrey, Rogers and McKelvey [13] have shown that individuals overturn cascades, that is act based on their signal and as such play the role of isolated agents. Our contribution in this context is to consider a mathematically clear and natural structure and study the relation between observational noise, captured by p, and asymptotic learning.

The rest of the paper is organized as follows. Section 2 introduces the model and our concepts of information cascades and learning. Section 3 considers the deterministic social network case. Section 4 analyzes random social networks and provides our main results. Section 5 considers a more general multi-dimensional structure for which our main result holds. Section 6 concludes. All proofs are relegated to the appendix.

2   A Multi-Dimensional Model of Sequential Choice

A countable infinite set of agents N are organized in the non-negative orthant of an m-dimensional integer lattice $Z_+^m \subset \mathbb{R}^m$ whose vertices are m-tuples of integers (see Figure 1 below for an illustration of the two-dimensional case). Each agent $\mathbf{x}$ makes a single, irreversible decision, $a_{\mathbf{x}} \in \{0, 1\}$ under uncertainty which is represented by a binary state space $\Omega = \{0, 1\}$. The true state of the world is drawn according to the uniform prior in round $t = 0$. The payoff of each agent $\mathbf{x}$ is given by

$$u_{\mathbf{x}}(a_{\mathbf{x}}, \omega) = \begin{cases} 1 \text{ if } a_{\mathbf{x}} = \omega \\ 0 \text{ otherwise} \end{cases}.$$

Agents do not know the realized state but each agent i observes a private signal $s_{\mathbf{x}}$ belonging to a measurable signal space S. The distribution F_ω according to which the signal $s_{\mathbf{x}}$ is drawn depends on the state of the world, $F_0, F_1 \in \Delta(S)$. Signals are conditionally independent and identically distributed across agents. We assume that the probability measures F_0 and F_1 are

absolutely continuous with respect to each other but not identical. This implies that signals have some information value but are not perfectly informative regarding the realized state of the world.
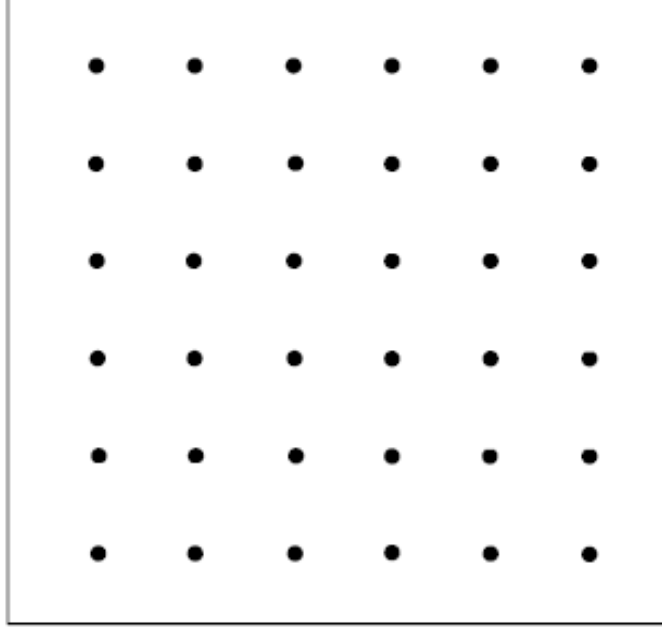

Figure 1: The 2-dimensional integer lattice.

The lattice determines the order of choices of the agents. In what follows we identify each agent i with a corresponding point $\mathbf{x}$ in the lattice. Define the lattice distance function $d_{Z_+^m} : Z_+^m \times Z_+^m \to \mathbb{R}$ as follows: for two points $\mathbf{x}, \mathbf{y} \in Z_+^m$ the lattice distance $d_{Z_+^m}(\mathbf{x}, \mathbf{y})$ is given by

$$d_{Z_+^m}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^{m} |\mathbf{x}_i - \mathbf{y}_i|.$$

There are countable decision rounds. In round $t \in 1, 2, ...$ all agents with lattice distance of $t - 1$ from agent $\mathbf{0}$ make their decision. In particular, this means that the number of agents acting in a given round is increasing with the number of rounds.

2.1 The Observation Structure

In the SSLM agents observe the history of choices of their predecessors. By organizing agents in a lattice, one can distinguish between the order in which actions are taken and the observed actions of a given agent. Our observation structure is determined by the following random

directed social network. For any pair of agents $\mathbf{x}, \mathbf{y}$ such that $\mathbf{x} = \mathbf{y} + \mathbf{e}^j$ for some $j = 1, ..., m$, where $\mathbf{e}^j$ is the unit vector in $\mathbb{R}^m$, the edge from $\mathbf{x}$ to $\mathbf{y}$ is realized independently with linkage probability p. Any other edge $\mathbf{xz}$ with $\mathbf{x} \neq \mathbf{z} + \mathbf{e}^j$ is realized with probability zero. The deterministic social network, $p = 1$, is denoted by $G_1 = (Z_+^m, E^m)$. The nodes of the network are given by the agents and the set of edges of the graph is given by

$$E^m = \left\{ \mathbf{xy} : \mathbf{x} = \mathbf{y} + \mathbf{e}^j \text{ for some } j = 1, ..., m \right\}.$$

In the random network G_p with linkage parameter $p \in (0, 1)$ each edge in E^m is formed independently with probability p. A realization of the network is denoted by G, while G_1 denotes the deterministic network, and G_p the random variable generating the network. We assume that the random network is drawn at time $t = 0$ independently of the state of the world ω.

The observation structure is linked to the social network as follows. A given agent $\mathbf{x}$ observes the previous actions of all agents that form part of his extended social network, i.e. of all those agents that can be reached from $\mathbf{x}$ through a directed path. The observation structure is commonly known among all agents while the realized actions are not. In the one-dimensional case, $m = 1$, the network G_1 equals an infinite line network with agent $\mathbf{0}$ as the origin. Therefore, the one-dimensional case corresponds to the standard sequential social learning setting: agents decide in strict sequential order and observe the actions of all their predecessors. By increasing the dimension m, we increase the number of agents that act simultaneously in each round. We restrict attention to dimensions larger than one.

Figure 2 displays the deterministic social network G_1 and the observation set of agent $\mathbf{x} = (2, 3)$ who lies at a lattice distance of 5 from the origin and therefore makes his choice in period $t = 6$. All agents denoted by a circle with empty interior lie on a directed path from $\mathbf{x}$ and their actions are observed by agent $\mathbf{x}$. In contrast to the SSLM, for $m \geq 2$ agents do not observe the actions chosen by all their predecessors even in the deterministic network. Consider for example the agent at $(0, 4)$ who acts prior to $\mathbf{x}$ but is not observed by him. Figure 3 displays a random realization of the social network G_p and the agents whose actions are observed by agent $\mathbf{x}$.
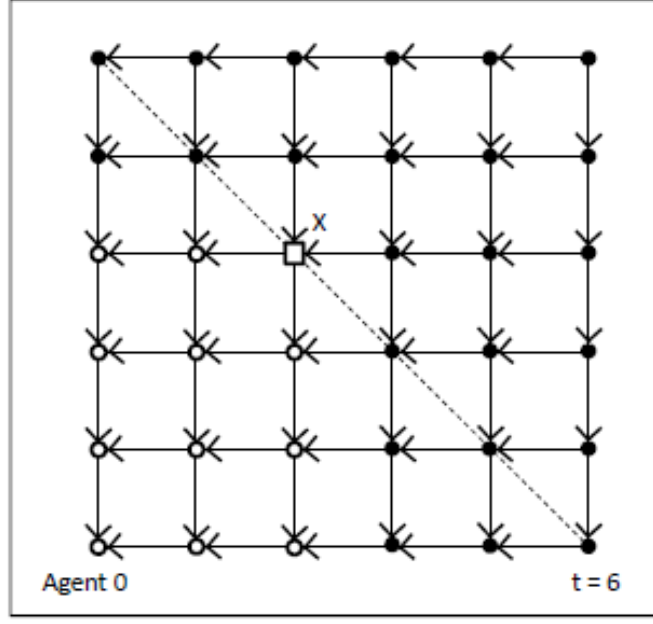
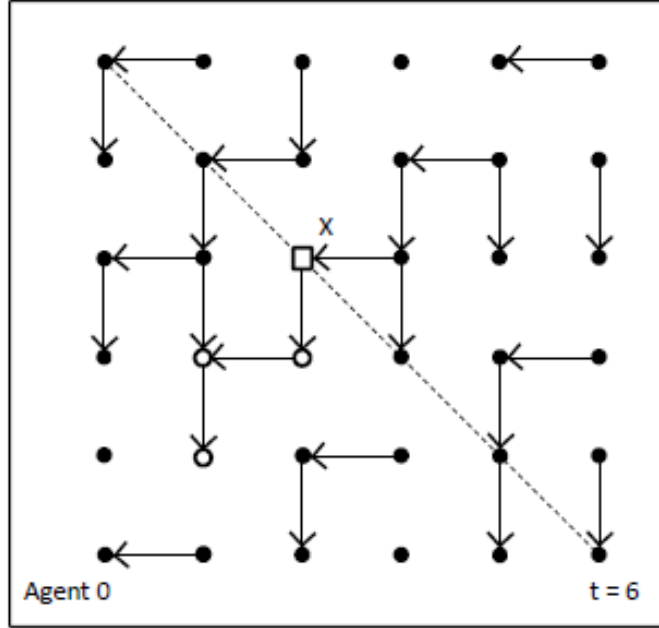Figure 2: The observation set of agent $(2, 3)$ in G_1.

Figure 3: The realized random observation set of agent $(2, 3)$.

For a graph G, the distance $d_G(\mathbf{x}, \mathbf{y})$ between nodes $\mathbf{x}$ and $\mathbf{y}$ is equal to the length of the directed path from $\mathbf{x}$ to $\mathbf{y}$. If there exists no connecting directed path the distance is set to infinity. Formally, each agent $\mathbf{x}$ observes the actions of all agents in $B_G(\mathbf{x})$ where

$$B_G(\mathbf{x}) = \{\mathbf{y} \in Z_+^m : d_G(\mathbf{y}, \mathbf{x}) \text{ is finite}\}.$$

11

The set $B_G(\mathbf{x})$ describes the extended social network of agent $\mathbf{x}$.

2.2 Information Sets, Strategies and Equilibrium

Denote the game with the above random observation structure, linkage probability p and lattice dimension m by Γ_p^m. The information set $I_{\mathbf{x}}$ of agent $\mathbf{x}$ is given by his signal realization $s_{\mathbf{x}}$, and the actions chosen by the agents he observes,

$$I_{\mathbf{x}} = \{(s_{\mathbf{x}}, \{a_{\mathbf{y}}\}) \text{ for } \mathbf{y} \in B_G(\mathbf{x})\}.$$

Denote by $\mathcal{I}_{\mathbf{x}}$ the set of all possible information sets of agent $\mathbf{x}$. A strategy $\sigma_{\mathbf{x}}$ of agent $\mathbf{x}$ is a measurable mapping that assigns a (mixed) action to each possible information set,

$$\sigma_{\mathbf{x}} : \mathcal{I}_{\mathbf{x}} \to \Delta(A).$$

A strategy profile is given by the set of strategies of all agents, $\sigma = \langle \sigma_{\mathbf{x}} \rangle_{\mathbf{x} \in Z_+^m}$. As common in the literature we solve the game for its Perfect Bayesian equilibria, the set of strategy profiles σ such that, for each agent $\mathbf{x}$, $\sigma_{\mathbf{x}}$ maximizes the expected utility of $\mathbf{x}$ given the strategies of all other agents.

2.3 Proportional Information Cascades and Proportional Asymptotic Learning

An information cascade, where from a finite time period onward all agents select the same action independent of their private signal realization, is a central phenomenon in the social learning literature. Since the action chosen in such an information cascade is based only upon the private signals of a finite set of agents there exists a positive probability that the inferior action is being chosen. Bikhchandani, Hirshleifer and Welch [8] establish that if the signal support is finite then information cascades occur with probability one. In particular, with positive probability the suboptimal action is chosen from a finite period on. The equilibrium property which contrasts information cascades is that of asymptotic learning. In the strict sequential setting asymptotic learning occurs if the probability with which agents select the optimal action converges to one over time.

We provide an alternative definition for both concepts which captures the multi-dimensional order in which actions are taken. Note that in a random network setting each agent observes an empty history with positive probability for any lattice dimension m and for every linkage probability $p < 1$. Hence, as each agent acts only based upon his signal with positive probability, neither asymptotic learning nor information cascades can occur. The question arises of how one should define alternative notions of asymptotic learning and information cascades that are conceptually close to the standard definitions.

Consider the set of agents that act in period t. Since not all agents can act independent of their signal nor select the correct action with probability 1, as argued above, we focus on the proportion of agents doing so. From the perspective of learning, the best possible outcome is that a large proportion of agents acting in round t select the correct action as t grows.

Let us first define the appropriate notion of proportional asymptotic learning. The game Γ_p^m together with given a strategy profile σ define a probability measure $\mathbf{P}_{\sigma,p}^m$ over the sequence of actions $\langle a_{\mathbf{x}} \rangle_{\mathbf{x} \in Z_+^m}$. Denote by B_t the set of agents with a lattice distance of precisely t from the origin. These are the agents that take their decision at time $t+1$. Let b_t denote the size of B_t and let r_t be the random variable that represents the number of agents within B_t whose action matches the realized state.[3]

Definition 1. Consider the game Γ_p^m. A strategy profile σ satisfies α-*proportional asymptotic learning* if

$$\lim_{t \to \infty} \mathbf{P}_{\sigma,p}^m \left(\frac{r_t}{b_t} \geq \alpha \right) = 1.$$

Under α-proportional asymptotic learning the probability of a proportion of at least α of the agents in B_t selecting the correct action converges to one with the number of rounds t growing to infinity.

Next we define the proportional information cascade concept. An agent cascades if his action is independent of his private signal. Formally, let $h_{\mathbf{x}}$ denote a realization of the observed history of agent $\mathbf{x}$. That is, the information set of agent $\mathbf{x}$ is given by $I_{\mathbf{x}} = \{s_{\mathbf{x}}, h_{\mathbf{x}}\}$.

Definition 2. Consider a strategy profile σ. Agent $\mathbf{x}$ *cascades* given history $h_{\mathbf{x}}$, if there exists an action a such that $\sigma_{\mathbf{x}}(s_{\mathbf{x}}, h_{\mathbf{x}}) = a$ for almost every signal realization $s_{\mathbf{x}} \in S$.

---

[3]Note that both b_t and r_t depend on the dimension of the lattice. To simplify the exposition we omit this dependence.

Let c_t be the random variable that represents the number of agents within B_t that cascade, i.e. select an action independent of their signal. The proportion of agents that ignore their signal in period t is represented by the random variable $\frac{c_t}{b_t}$. Proportional information cascades are now defined as follows.

Definition 3. Consider the game Γ_p^m. A γ-*proportional information cascade* occurs under strategy profile $\boldsymbol{\sigma}$ if

$$\lim_{t \to \infty} \mathbf{P}_{\sigma,p}^m \left(\frac{c_t}{b_t} \geq \gamma \right) = 1.$$

In words, a γ-proportional information cascade occurs if the probability that a proportion of at least γ of the agents in B_t ignoring their signal converges to one with the number of rounds.

3 The Deterministic Social Network Case

We first consider the case of a deterministic social network. Our objective is twofold. First, we aim at characterizing the conditions under which proportional asymptotic learning is satisfied. Second, we will study the phenomenon of proportional information cascades in our multi-dimensional observation structure.

Smith and Sorensen [22] establish that in the SSLM asymptotic learning critically hinges upon properties of the signal space. They show that asymptotic learning occurs in any equilibrium if private signals are unbounded and fails in any equilibrium if private signals are bounded. Boundedness and unboundedness are properties of the support of the private beliefs induced by the signals. Let q_1 be the private belief of agent 1, i.e. $q_1 = \Pr(\omega = 1 \,|\, s_1)$. Formally, bounded and unbounded signals are defined as follows.

Definition 4. Let $\underline{\beta}, \overline{\beta}$ be defined as follows

$$
\begin{aligned}
\underline{\beta} &= \inf \left\{ r \in [0,1] : \Pr(q_1 \leq r) > 0 \right\}, \text{ and} \\
\overline{\beta} &= \sup \left\{ r \in [0,1] : \Pr(q_1 \leq r) < 1 \right\}.
\end{aligned}
$$

The signal structure has *bounded private signals* if $0 < \underline{\beta} < \overline{\beta} < 1$ and *unbounded private signals* if $\underline{\beta} = 1 - \overline{\beta} = 0$.

The following theorem characterizes asymptotic learning in the deterministic case for any finite lattice dimension m as a function of the signal structure.

Theorem 1. *For any Perfect Bayesian equilibrium σ of the game Γ_1^m the following holds:*

- *If private signals are unbounded, then σ satisfies α-proportional asymptotic learning for any $\alpha < 1$.*

- *If private signals are bounded, then σ fails α-proportional asymptotic learning for any $\alpha > 0$.*

The theorem establishes that in the deterministic multi-dimensional social learning framework asymptotic learning still critically depends upon whether signals are unbounded. If signals are unbounded then an arbitrary high proportion of agents take the correct action in the limit with probability 1. However, if signals are bounded, then an arbitrary high proportion of agents select the incorrect action with positive probability. There are two ways to interpret the negative result. On the one hand, failure of learning under bounded signals is quite strong as one cannot assure any positive proportion of agents taking the correct action in the limit. On the other hand, depending on the characteristics of the bounded signal structure, the event that α-proportional asymptotic learning fails for any $\alpha > 0$ might have arbitrary low probability.

Next we address the possibility of information cascades in our general framework. As unbounded signals imply asymptotic learning we restrict attention to the case of bounded signals. The question of whether information cascades are a robust phenomenon, i.e. occur for more general information and signal structures than the finite signal case considered in Bikhchandani, Hirshleifer and Welch [8] has been discussed since. In a recent working paper, Herrera and Hoerner [15] provide a characterization for the absence of information cascades in terms of the signal structure. They consider bounded signal distributions that satisfy the strict monotone likelihood ratio property and show that information cascades occur with probability zero if and only if the signal distributions satisfy the Increasing Hazard Ratio Property (IHRP) and the Increasing Failure Ratio Property (IFRP). The minimal coordinate of $\mathbf{x}$ is denoted by $\mathbf{x}_-$.

Proposition 1. *Let ω be the true state of the world, the following event holds with positive probability $p_\omega > 0$: All agents select action $1 - \omega$, and in addition there exists a time t' such that for every $t > t'$ all agents $\mathbf{x} \in B_t$ cascade for whom $\mathbf{x}_- > 0$.*

Hence with positive probability from some finite time period on all agents, that lie on the boundary of the lattice, cascade. The proposition implies the following corollary.

Corollary 1. *Consider a game Γ_1^m with a bounded private signal structure. For any Perfect Bayesian equilibrium σ and any $\gamma < 1$ we have*

$$\lim_{t \to \infty} \mathbf{P}_{\sigma,1}^m \left(\frac{c_t}{b_t} > \gamma \right) \geq r > 0.$$

The proposition and corollary highlight a stark distinction between our multi-dimensional model and the strict sequential setting in regards to the occurrence of information cascades. Here, for any bounded signal structures an arbitrary high proportion of agents cascade in the long run with positive probability bounded away from zero. This differs from Herrera and Hoerner's [15] result under which cascades occur with probability zero for certain bounded signal structures.

The analysis of the deterministic observation structure uncovered that while the characterization of asymptotic learning is conceptually almost identical to the one in the SSLM, in terms of information cascades the strict sequential setting differs substantially from any higher dimensional observation structure.

4   The Random Social Network Case

In this section, we relax the assumption of perfect network connectivity by considering a linkage probability p smaller than one which makes the social network a random directed graph.

The main question of the paper concerns the possibility of proportional asymptotic learning in our multi-dimensional environment of choice order in random networks. The following theorem provides the central result of this paper.

Theorem 2. *For every $\alpha < 1$ there exists a linkage probability $0 < p(\alpha) < 1$ such that for every dimension $m \geq 2$ and every $p \in (p(\alpha), 1)$ any Perfect Bayesian equilibrium σ of Γ_p^m satisfies α-proportional asymptotic learning.*

Here the cutoff $p(\alpha)$ is the same *for all* informative signal structures. The theorem states that if the network connectivity is sufficiently high, that is the linkage probability p sufficiently close to one, then an arbitrary high proportion of agents select the optimal action in the long run

16

with probability one. This result sharply contrasts the equilibrium behavior in the deterministic social network under bounded signals established in Theorem 1 as well as the characterization of asymptotic learning in the standard sequential model. Note that for $p = 1$ there exists a positive probability that an arbitrary high proportion of agents select the suboptimal action in the limit. Introducing a minimal degree of observational noise, however, leads to an arbitrary high proportion of agents selecting the optimal action in the limit with probability one.[4] In the case of bounded signals, the multi-dimensional framework therefore features a strong discontinuity at $p = 1$ in terms of the learning properties of the equilibria.

The equilibrium behavior also differs significantly under bounded signals in comparison with the SSLM. While asymptotic learning fails under bounded signals in the strict sequential model, in our lattice structure asymptotic learning might occur in arbitrary high proportion in the limit even for bounded signals. The key parameter determining learning in the higher dimensional setting is the linkage probability. More precisely, the linkage probability provides a lower bound for the proportion of agents taking the correct action in the limit, for any informative signal structure.

A necessary condition for proportional asymptotic learning is that a proportion of agents in a given round have access to an over time growing history. Under unbounded signals this condition is also sufficient. In the case of bounded signals, however, observing an infinite history does not reveal the true state with probability one and hence is not sufficient for taking the correct action. However, observing an infinite set of actions taken in isolation is sufficient for taking the correct action with probability one. In this context, changing the linkage probability p has two opposing effects. First, increasing p increases the connectivity of the social network and hence the expected size of the observed history of a given agent. Second, increasing p decreases the probability that a given agent acts in isolation.[5] The crucial distinction between the standard framework where the observation structure is represented by a line and our lattice approach can now be easily understood. In the line, for any linkage probability $p < 1$ the probability of an infinite component (an observation path of infinite length) is equal to zero. In a higher dimensional lattice, however, this is not the case. A standard result in percolation

---

[4] To be clear, the theorem does not state that random structures are superior to the deterministic structure for any $p < 1$, but only for a close neighborhood of $p = 1$.

[5] The implication of these two competing effects is discussed in the subsection below.

17

theory establishes that the probability of an infinite component including the origin goes to one with the linkage probability going to one.

The next question we address concerns the occurrence of γ-proportional information cascades. Consider the case of unbounded signals. No agent ever takes an action independent of his signal as for every prior belief either action may be optimal conditional on the realized signal, with positive probability. Hence our question can be refined towards the occurrence of γ-proportional information cascades under bounded signals.

Intuitively an information cascade seems to prevent asymptotic learning as agents select actions disregarding their private signals. In the one-dimensional setting this intuition is in fact correct; the occurrence of an information cascade makes learning impossible. Theorem 2 establishes a relation between high linkage probabilities and an arbitrary high proportion of asymptotic learners. Therefore, γ-proportional information cascades may be impossible in an environment with a high linkage probability. Contrary to this intuition we establish the following result.

Theorem 3. *Let private signals be bounded. For every $\gamma < 1$ there exists a linkage probability $p(\gamma)$ such that for every dimension $m \geq 2$ and for every $p \in (p(\gamma), 1)$, a γ-proportional information cascade occurs in any Perfect Bayesian equilibrium σ of Γ_p^m.*

Theorem 3 establishes that if the linkage probability is sufficiently high and thus the observational noise sufficiently low, then an arbitrary high proportion of agents select their action independent of their signal in the long run. Our two results provide an interesting insight on the relationship between information cascades and asymptotic learning. While they exclude each other in the standard framework, the proportional concepts coexist under bounded signals. To be more precise, a high proportion of asymptotic learners does not only coexists with a high proportion of cascading agents but in fact requires it (under bounded signals). To see this, note that as time increases a high proportion of agents select the correct action with arbitrary high probability by Theorem 2. This implies that their beliefs are arbitrary close to the true belief. But as private signals are bounded most agents will eventually disregard their private signal and base their decision solely on the observed history.

4.1 Outline of the Proof of Theorem 2

The proof of our main theorem consists of several parts. Each part corresponds to a different challenge in establishing the result. The main challenge is to show that as the linkage probability p goes to one most agents observe a growing set of agents that take their action in isolation, i.e. based only upon their signal. From this we can deduce that as p goes to one most agents select the correct action.

A priori it is not clear that this occurs as the linkage probability p has two contrary effects as discussed above. On the one hand as p approaches one the number of agents that a given agent observes is growing which suggests that the agent might observe more isolated agents. On the other hand, however, when p grows the probability of a given agent to be isolated decreases.

We approach this problem by making use of the percolation structure of the random network. More concretely, we rely on a result from the theory of percolation showing that, when reverting the direction of the edges, the probability of the origin being part of an infinite connected component goes to one with p. Fix an integer n. We show that for every $\alpha < 1$ there exists a linkage probability $p(\alpha)$ such that for every $p \in (p(\alpha), 1)$ and for every large enough t, most agents in B_t have a path in G to some agents in B_{t-n} with a probability greater than α. That is, with a probability greater than α most agents in B_t can observe at least some agents in B_{t-n}. Moreover, $p(\alpha)$ is independent of n. Using Chebyshev's inequality we then show that the proportion of agents in B_t that have access to agents in B_{t-n} exceeds α with a probability that goes to one as t grows to infinity.

Based on the above we establish that for every linkage probability $p \in (p(\alpha), 1)$ and every $k \in \mathbb{N}$ a proportion greater than α of agents in B_t observe at least k isolated agents with a probability that goes to one when t goes to infinity. Therefore, a proportion greater than α observe an arbitrary high number of isolated agents with a probability approaching one as t goes to infinity. Since the action of an isolated agent is solely based on his signal, observing a growing set of isolated agents reveals the true state of the world with a probability that approaches one. This implies that the action of those agents converge to the correct action with a probability converging to one.

4.2 Bounded Versus Unbounded Signals

Throughout the existing literature on social learning unbounded signals induce superior asymptotic learning properties than bounded signals (see for example Smith and Sorensen [22], Acemoglu et al. [1], and Arieli and Mueller-Frank [5]).[6] This holds true even when the unbounded signal distribution is non informative with high probability.[7]

In our multi-dimensional model, however, bounded signals can outperform unbounded signals and achieve better asymptotic learning result whenever the linkage probability p is smaller than 1. To be more precise, for any $p < 1$ and any unbounded signal structure there exists a bounded signal structure for which a higher proportion of late movers select the optimal action than in the unbounded case, with a probability that goes to one over time.

Let us informally explain the logic behind this claim. Fix a value of $p < 1$, and recall that B_t is the set of agents that act in period t. It follows from the same logic as applied in the proof of Theorem 2, that, when t grows large, agents in B_t can be divided into two groups. The first group comprises those agents that observe unboundedly many agents from previous periods. The second group comprises those agents that observe boundedly many agents from previous periods. For a fixed $0 < p < 1$ the second group has a proportion of $\delta = \delta(p) > 0$ and the first group has proportion of $1 - \delta$.[8]

We note that agents belonging to the first group have access to unboundedly many agents that take their decision in isolation. It therefore holds that the probability with which a member of the first group takes the optimal action approaches one. Given any signal distribution we can write the limit proportion of agents that play optimally as follows:

$$(1 - \delta) + \delta q,$$

where q is the proportion of agents that act optimally within the second group. We note that as signals do not reveal the true state, it must be the case that $q < 1$ even for unbounded signal distributions. For any given unbounded signal distribution one can clearly find another bounded

---

[6]Unbounded signals also play an important role in models where agents interact repeatedly. See Mossel, Sly and Tamuz [20].

[7]Consider for example the case where the unbounded signal provides no valuable information to the agents with probability 0.99.

[8]$\delta(p)$ goes to zero as p approaches one.

distribution for which the corresponding q is arbitrarily close to 1, and hence outperform any given unbounded signal distribution.

4.3   Simulation of Equilibrium Behavior

In the following we use simulations to visualize the properties of equilibrium behavior in our framework. We consider the two-dimensional lattice and binary signals with a precision of .6, i.e. $\Pr(s = \omega | \omega) = .6$ for $\omega = 0, 1$ and the choices of agents within the first 200 rounds. We consider three different linkage probabilities, $p \in \{.8, .95, 1\}$. In the underlying equilibrium strategies each agent selects the action that equals his signal if he is indifferent among both actions. The figures report the proportion of agents selecting the correct action per round and the proportion of agents that cascade per round. The proportion is the average computed from 1000 simulations.
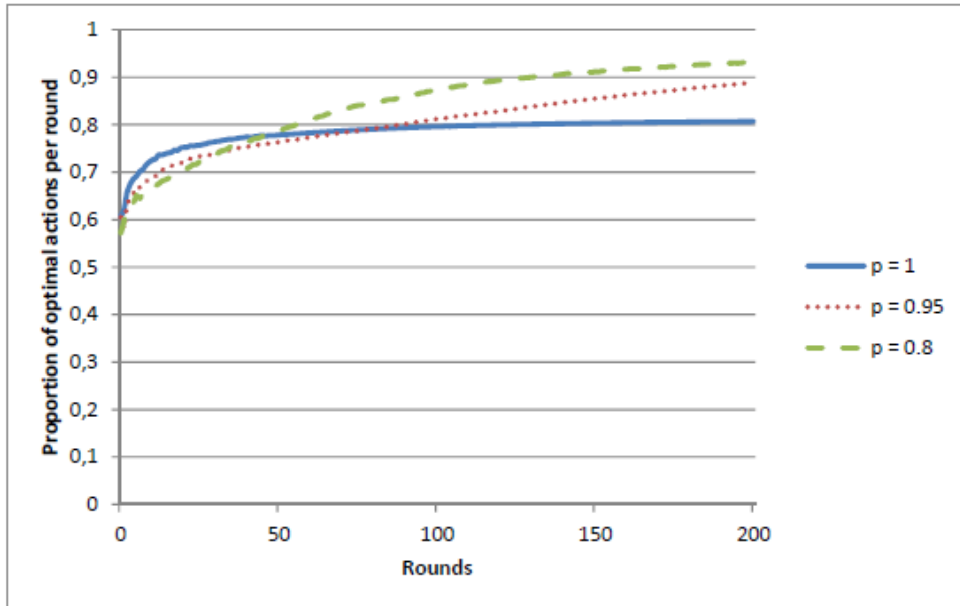
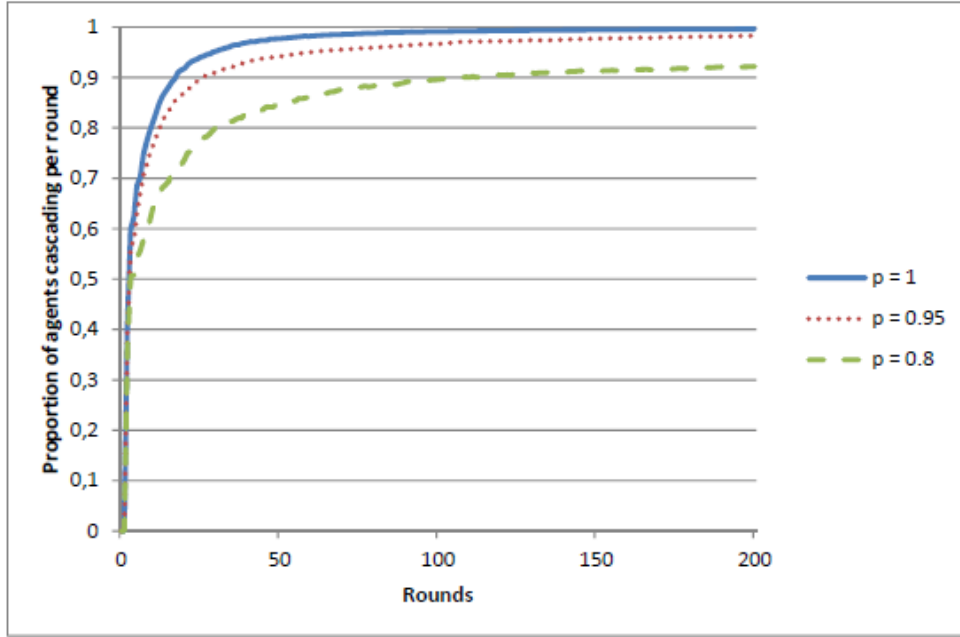

Figure 4: Proportion of optimal actions.

Figure 5: Proportion of cascading agents.

The graphs highlight two of the main results of this paper. First, the introduction of observational noise improves long run learning as can be seen in figure 4. Second, comparing the graphs for the proportion of optimal actions and the proportion of agents cascading we can clearly see the coexistence of proportional cascades and "asymptotic" learning. Finally, note that from periods 50-200 the linkage probability $p = .8$ outperforms $p = .95$ in the sense of a higher proportion of agents select the correct action (see figure 4). Nevertheless, the slope in round 200 of the graph corresponding to $p = .95$ is steeper than that of the graph corresponding to $p = .8$. This highlights a trade-off between speed of learning and accuracy of learning. While values of p close to 1 achieve a higher proportion of asymptotic learners, in early rounds a higher proportion of agents cascade, as can be seen in figure 5, partly on the incorrect action.

5   Extension: A General Multi-Dimensional Model

In this section we address the robustness of our results. The multi-dimensional model we considered so far is based upon a lattice percolation process. While being a standard type of random graph, there are no compelling economic reasons to emphasize this random graph structure over others which feature a growing number of decision makers over time. In this

section we provide a more general observation structure for which Theorem 2 and Theorem 3 hold.

Let $G = (N, E)$ be a an infinite directed deterministic graph where N is a countable set of vertices corresponding to the set of agents, and E is the set of directed edges. Assume that N is partitioned into disjoint groups B_t, the set of agents that take their action in period t,

$$N = \bigcup_{t \in \mathbb{N}} B_t.$$

Assume further that the set of agents in B_{t+1} is directly linked only to the agents in B_t. For any agent $\mathbf{x} \in N$, let $N_{\mathbf{x}}$ be the set of agents which agent $\mathbf{x}$ is directly linked to, i.e. those acting one period before him and whom he observes.[9] For any subset $V \subset B_t$, we let $D(V) \subset B_{t-1}$ be the subset of agents for whom there exists a directed edge from at least one agent in V, and let $d(V) = |D(V)|$ be the size of $D(V)$. Let $M \geq 2$ denote a natural number.

Definition 5. A subset $V_t^M \subset B_t$ satisfies *M-boundedness* if the following conditions hold:

1. For any subset $V \subset V_t^M$ it holds that $|V| + 1 \leq d(V) \leq M \times |V|$,

2. for any $\mathbf{x} \in V_t^M$ the subset $\{\mathbf{y} \in V_t^M : \mathbf{y} \neq \mathbf{x}, \ N_{\mathbf{y}} \cap N_{\mathbf{x}} \neq \emptyset\}$ is smaller than M.

The first condition says that any subset V of V_t^M collectively observes at least $|V| + 1$ agents from B_{t-1} and at most $M \times |V|$. The second condition says that for every $\mathbf{x} \in V_t^M$ the number of agents $\mathbf{y} \in V_t^M$, different from $\mathbf{x}$, that share a common observed agent with $\mathbf{x}$ is less than M.

Definition 6. A graph G has a *bounded interdependent* observation structure if there exists a natural number $M \geq 2$ such that the following conditions hold:

(i) $\lim_{t \to \infty} |B_t| = \infty$,

(ii) there exists a sequence of M-bounded subsets $\{V_t^M\}_{t \geq 1}$ such that $V_t^M \subset B_t$ and

$$\lim_{t \to \infty} \frac{|V_t^M|}{|B_t|} = 1.$$

---

[9] A direct link refers to an edge.

Note first that any m-dimensional lattice $\mathbb{Z}_+^m$ has a bounded interdependent observation structure with $M = m$. To see this, let V_t^M be the set of all interior points in B_t, i.e., the points that lie at a distance of t from the origin for whom all lattice coordinates are positive.

Next we introduce a random graph structure similarly as in the lattice case. For a deterministic graph G, consider the percolation graph G_p where each edge in G is realized with probability p, where $0 < p < 1$. The percolation graph G_p induces a game Γ_p where every agent $\mathbf{x}$ observes the set of agents that he is linked to by a directed path, and where all agents at a distance of t (in graph G) from the initial agents act in period $t + 1$.

Theorem 4. *Consider the percolation graph induced by a graph with bounded interdependent observation structure. For every $\alpha < 1$ there exists a linkage probability $0 < p(\alpha) < 1$ such that for every $p \in (p(\alpha), 1)$ any Perfect Bayesian equilibrium σ of Γ_p satisfies α-proportional asymptotic learning.*

Theorem 3 is generalized in the same manner. The proof of Theorem 4 goes along similar lines to the proof of Theorem 2 and is therefore omitted. To give some intuition, note that for every n the observation structure of agent (n, n) in the two-dimensional lattice can be embedded in the observation structure of almost every x in V_t^M for every large enough t. Percolation based reasoning similar to that employed in the proof of Theorem 2, establishes the existence of a subset of agents in V_t^M, with a proportion that can be made arbitrarily close to 1 as a function of p, such that all of its members observe unboundedly many isolated agents for growing t.

6   Conclusion

This paper contributes to the literature on social learning by providing a natural generalization of the strict sequential setting. We represent the order of choices of agents via a general integer lattice structure rather than via a line. When the lattice dimension is higher than one, the number of agents acting in a given round is increasing over time. Based on this feature of our model we formulate the concepts of proportional information cascades and proportional asymptotic learning which require only a proportion of agents in a given round to cascade or select the correct action.

The introduction of observational noise leads to a neat characterization of proportional

asymptotic learning. The occurrence of proportional asymptotic learning depends uniquely on the linkage probability of the random network rather than on properties of the signal space, as in Smith and Sorensen [22] and Acemoglu, Dahleh, Lobel and Ozdaglar [1]. More precisely, the linkage probability provides a lower bound for the proportion of agents taking the correct action in the limit, for any informative signal structure. Therefore, introducing a minimal amount of noise via random social connections in an environment with bounded private signals results in a substantial improvement of information aggregation in equilibrium. That is, an infinitesimal amount of noise in the observation structure allows for an arbitrary high proportion of agents selecting the correct action in the limit.

Our different modelling approach leads to phenomena that fail to occur in the strict sequential model. In our deterministic multi-dimensional model with bounded signals there exists a positive probability for which an arbitrary high proportion of agents act independent of their signal in the long run, in contrast to the recent characterization of information cascades in the strict sequential framework provided by Herrera and Hoerner [15]. Finally, we show that in our multi-dimensional model with random observations, bounded signal structures can outperform unbounded signals in terms of their asymptotic learning properties.

A   Appendix

A.1   Proof of Proposition 1

Proof of Proposition 1. Assume for simplicity that the two agents $(0,1)$ and $(1,0)$ are playing according to the same strategy, and that for every action a played by agent $(0,0)$ there exists a positive probability that the action $1 - a$ is played by either of the agents. That is, we assume that neither agent $(0,1)$ nor agent $(1,0)$ are cascading.[10] This assumption implies that,

$$\mathbf{P}_{\sigma,1}(a_{\mathbf{x}} = 1 | \omega = 1, a_{(0,0)} = 1) > \mathbf{P}_{\sigma,1}(a_{\mathbf{x}} = 1 | \omega = 0, a_{(0,0)} = 1) \tag{1}$$

for $\mathbf{x} \in \{(0,1), (1,0)\}$.

For every agent $\mathbf{x}$ we let $p_{\mathbf{x}}$ be the public belief of agent $\mathbf{x}$ prior to receiving his signal given that all the players in his observation set play action 1. Clearly, the sequence $\{p_{(0,t)}\}_t$ is

---

[10]The assumption makes the analysis more transparent and saves notational complications.

increasing sequence bounded by 1 and therefore converges to a limit q'. By Lemma 2 in Smith and Sorensen [22],

$$q' \geq 1 - \underline{\beta}. \tag{2}$$

A similar argument holds for the sequence $\{p_{(t,0)}\}_t$.

For every t let $p_t = \min\{p_{(0,t)}, p_{(t,0)}\}$ and $q = \lim_{t \to \infty} p_t$. Define *the likelihood* ratio for every player $\mathbf{x}$ by,

$$L_{\mathbf{x}} = \frac{p_{\mathbf{x}}}{1 - p_{\mathbf{x}}}.$$

Let $L_t = \frac{p_t}{1-p_t}$, and $\overline{L} = \frac{1-\overline{\beta}}{\overline{\beta}}$. By definition if $L_{\mathbf{x}} > \overline{L}$ then agent $\mathbf{x}$ chooses action 1 with probability 1 regardless of his signal.

Since $p_t \to q \geq 1 - \underline{\beta}$, we must have $\lim_t L_t \geq \overline{L}$. Hence by equation (1) there exists t_0 such that

$$L_* := L_{t_0+1} \cdot \frac{\mathbf{P}_{\sigma,1}(a_{\mathbf{x}} = 1 | \omega = 1, a_{(0,0)} = 1)}{\mathbf{P}_{\sigma,1}(a_{\mathbf{x}} = 1 | \omega = 0, a_{(0,0)} = 1)} > \overline{L}.$$

Let R be the event that all agents $\mathbf{x}$ such that $\mathbf{x}_- = \min \mathbf{x} = 0$, and all agents prior to and at time $2t_0$ chose action 1. Note that agent $\mathbf{x} = (t, 0) \in Z_+^2$ observes all agents $(l, 0)$ for $l < t$. Similarly, agent $\mathbf{x} = (0, t)$ observes $(0, l)$ for $l < t$. Hence, the observation structure of every agent $\mathbf{x}$ with $\mathbf{x}_- = 0$ is identical to the strict sequential model. Since signals are bounded, Theorem 3 in Smith and Sorensen [22] implies that with positive probability all agents $\mathbf{x}$ such that $\mathbf{x}_- = 0$ take action 1. Therefore the event R occurs with positive probability under σ conditional on the true state being 0.

We shall show that given the event R all agents choose action 1 under σ. Moreover we shall further show that all agents $\mathbf{x} \in B_t$ such that $t > 2t_0$ and $\mathbf{x}_- > 0$ cascade.

Let $\mathbf{x} \in B_{2t_0+1}$. If $\mathbf{x}_- = 0$ then clearly given the event R we must have that $a_{\mathbf{x}} = 1$. Assume that $\mathbf{x}_- > 0$. Let,

$$K_1 = \{(0,0), (0,1), \ldots, (0, t_0)\} \cup \{(1, 0)\},$$

and

$$K_2 = \{(0,0), (1,0), \ldots, (t_0, 0)\} \cup \{(0, 1)\}.$$

By definition the observation set $B_G(\mathbf{x})$ of agent $\mathbf{x}$ contains one of these two sets. Assume for simplicity that it contains the agents in K_1. Let L_1 be the likelihood ratio obtained from

observing K_1. By definition the set $K_1 \setminus \{(1,0)\} = B_G((0, t_0 + 1))$. Since given the event R all agents in K_1 take action 1 we can deduce from Bayes rule that,

$$L_1 \geq L_{t_0+1} \cdot \frac{\mathbf{P}_{\sigma,1}(a_{\mathbf{x}} = 1 | \omega = 1, a_{(0,0)} = 1)}{\mathbf{P}_{\sigma,1}(a_{\mathbf{x}} = 1 | \omega = 0, a_{(0,0)} = 1)} > \overline{L}.$$

Since $K_1 \subset B_G(\mathbf{x})$ and all agents in $B_G(\mathbf{x})$ take action 1 we must have that,

$$L_{\mathbf{x}} \geq L_1. \tag{3}$$

Hence agent $\mathbf{x}$ cascades. A simple inductive argument shows that given R it holds for all agent $\mathbf{x} \in B_t$ such that $t > 2t_0$, and $\mathbf{x}_- > 0$ that,

$$L_{\mathbf{x}} \geq L_1 > \overline{L}.$$

This concludes the proof of the lemma and the proof of Theorem 1 $\qquad\qquad\square$

A.2 Proof of Theorem 1

Proof of the first part of Theorem 1. We first show that if private signals are unbounded, then the equilibrium σ satisfies α-proportional asymptotic learning for any $\alpha < 1$. Let $f : Z_+^m \to \mathbb{N}$ be a bijection that assigns a unique natural number to every agent $\mathbf{x}$. We shall identify every agent $\mathbf{x}$ with $f(\mathbf{x})$. For every n, let Y_n be a Bernoulli random variable that corresponds to the payoff of agent n.

By Theorem 2 in Acemoglu et al. [1], $\lim_{n \to \infty} \mathbf{P}_{\sigma,1}(Y_n = 1) = 1$. Let $\epsilon > 0$.[11] There exists a large enough n_0 such that for every $n > n_0$,

$$E(Y_n) > 1 - \epsilon. \tag{4}$$

Let t_0 be sufficiently large such that for every $t > t_0$ and $\mathbf{x} \in B_t$ it holds that $f(\mathbf{x}) > n_0$. Let $b_t = |B_t|$ and let $S_t = \frac{\sum_{\mathbf{x} \in B_t} Y_{\mathbf{x}}}{b_t}$. For every $t > t_0$ we get from (4) and the linearity of the expectation operator that,

$$E(S_t) > 1 - \epsilon.$$

Since S_t is bounded above by 1 and since ϵ is arbitrary we get that,

$$\lim_{t \to \infty} E(S_t) = 1. \tag{5}$$

---

[11] For expositional clarity, we omit the dependence of the probability expression of m.

Again since S_t is bounded above by 1, equation (5) implies that for every $\alpha < 1$,

$$\lim_{t \to \infty} \mathbf{P}_{\sigma,1}(S_t < \alpha) = 0.$$

Hence for every $\alpha < 1$

$$\lim_{t \to \infty} \mathbf{P}_{\sigma,1}(S_t \geq \alpha) = 1,$$

as claimed.

We now turn to the second part of Theorem 1. Assume for simplicity that the lattice dimension m equals two, a similar argument holds for any $m > 2$. Fix a Bayesian equilibrium σ. By Proposition 1, if the true state of the world is ω then there exists a positive probability $p_\omega > 0$ such that all agents take action $1 - \omega$. It then follows that α proportional learning fails for every $\alpha > 0$. More generally we establish the following lemma which implies the second part of the theorem. $\square$

A.3 Proof of Corollary 1

Proof of Corollary 1. By Proposition 1 there exists a positive probability $p > 0$ and a time t' such that for every $t > t'$, all agents $\mathbf{x}$ cascade for whom $\mathbf{x}_- > 0$. Hence with probability p for every time $t > t'$ the proportion of agents that cascade is $\frac{t-1}{t+1}$. In general, for lattice dimension m this proportion equals $\frac{b_t - m}{b_t}$, and goes to 1 as t goes to infinity. Recall that c_t is the number of cascading agents at time t. We have that for every $\gamma < 1$, there exists a $\hat{t}$ such that for every $t > \hat{t}$

$$\mathbf{P}_{\sigma,1}(\frac{c_t}{b_t} > \gamma) > 0.$$

$\square$

A.4 Proof of Theorem 2

For simplicity we assume that the dimension of the lattice is 2.[12] Consider the following infinite random directed graph. Call each $\mathbf{x} \in Z_2^+$ a node. Assume that for each node $\mathbf{x}$ each of the edges $(\mathbf{x}, \mathbf{x} + (0,1))$, $(\mathbf{x}, \mathbf{x} + (1,0))$ is formed with independent probability p. Assume that the edges are formed independently across all nodes. This random graph structure is known as

---

[12] See the end of the appendix for a discussion of the general case.

oriented percolation and differs from our random network only by reverting the direction of the edges. For a certain realization of the random graph, say that a node $\mathbf{x} \in Z^2$ is reachable from the origin if there exists a directed path leading from $\mathbf{0}$ to $\mathbf{x}$. We let ξ_t^0 be the random set of all nodes that are reachable and lie in a distance of t from the origin.

The following proposition is a straightforward corollary from Durrett [12].[13]

Proposition 2. *Let Υ be the event that there exists an infinite path starting at the origin. Then $\mathbf{P}_p(\Upsilon) \uparrow 1$ as $p \uparrow 1$.*

As a corollary from the proposition one may deduce that the size of the set ξ_t^0 is growing to infinity with probability that approaches 1 when p tends to 1. To make the proof self contained we shall formally introduce this argument in the following lemma.

Lemma 1. *For every $\alpha < 1$ there exists a positive $p(\alpha) < 1$ such that for all $p \in (p(\alpha), 1)$ the following holds. For every integer l there exists $t(l)$ such that*

$$\mathbf{P}_p(\forall t \geq t(l), \; |\xi_t^0| \geq l) > \alpha.$$

Proof of Lemma 1. Assume by contradiction that the lemma does not hold. As a result, the following must be true:

- There exists $\epsilon' > 0$ such that for every $\delta > 0$ there exists a $p \in (1 - \delta, 1)$, and $M_p > 0$ such that the following holds:

$$\mathbf{P}_p(|\xi_t^0| \leq M_p \; i.o. \;) \geq \epsilon'.$$

By Proposition 2 $\mathbf{P}_p(\Upsilon) \to_{p \to 1} 1$. We can therefore find $\epsilon_0 > 0$, and large enough p such that

$$\mathbf{P}_p(1 \leq |\xi_t^0| \leq M_p \; i.o.) \geq \epsilon_0.$$

For each k let A_k be the event that the inequality $1 \leq |\xi_t^0| \leq M_p$ holds for at least k distinct times. For $l \geq k$, let $A_{k,l}$ be the event that at time l the inequality $1 \leq |\xi_t^0| \leq M_p$ holds for the k-th time between times 0 and l. Note that $A_k = \cup_{l \geq k} A_{k,l}$ and that $A_{k,l} \cap A_{k,m} = \varnothing$ for $m \neq l$. Let $\delta = 1 - (1 - p)^{2M_p}$. We shall show by induction that $\mathbf{P}_p(A_k) \leq \delta^{k-1}$. For $k = 1$ we

---
[13]This follows, for example, from the proof on page 1005.

obviously have $\mathbf{P}_p(A_1) \leq 1$. Assume that $\mathbf{P}_p(A_k) \leq \delta^{k-1}$. We shall show the following holds for $\mathbf{P}_p(A_{k+1})$:

$$\mathbf{P}_p(A_{k+1})$$

$$= \sum_{l=k}^{\infty} \mathbf{P}_p(A_{k,l})\mathbf{P}_p(A_{k+1}|A_{k,l}) \tag{6}$$

$$\leq \sum_{l=k}^{\infty} \mathbf{P}_p(A_{k,l})\delta \tag{7}$$

$$= \mathbf{P}_p(A_k)\delta \leq \delta^k \tag{8}$$

Equation (6) is clear from the law of total probability and the fact that $A_{k+1} \subset A_k$. $\mathbf{P}_p(|\xi_{l+1}^0| = 0| \ |\xi_l^0| \leq M_p) \geq 1 - \delta$ together with

$$A_{k+1} \cap A_{k,l} \subset \{|\xi_{l+1}^0| = 0\}^c \cap A_{k,l}$$

imply that

$$\mathbf{P}_p(A_{k+1}|A_{k,l}) \leq \delta.$$

Hence, inequality (7) follows. Inequality (8) follows from the induction hypothesis. Hence since $\{A_k\}_k$ is a decreasing family of events we have,

$$\mathbf{P}_p(1 \leq |\xi_t^0| \leq M_p \ i.o.) = \mathbf{P}_p(\bigcap_k A_k) = \lim_k \mathbf{P}_p(A_k) = 0.$$

This stands in contradiction to the assumption.

$\square$

Fix $\alpha > 0$, and $p \in (p(\alpha), 1)$. Let $\mathbf{x} = (i, j)$. Assume that $i, j \geq n$ and let $t = i + j + 1$. The set $\xi_n^{\mathbf{x}} = B(\mathbf{x}) \cap B_{t-n}$ represents the set of agents that can be observed by $\mathbf{x}$ and make their decision n periods prior to $\mathbf{x}$, at time $t - n$. A critical point of our analysis is the fact that $|\xi_n^{\mathbf{x}}|$ and $|\xi_n^0|$ have an identical distribution. That is the number agents in B_n that can be reached from the origin has the same distribution as the number of agents in B_{t-n} that can be observed by $\mathbf{x}$.

Let $T^{n,l}$ be the subset of agents $\mathbf{x}$ for whom $|\xi_n^{\mathbf{x}}| \geq l$. Given $\mathbf{x} = (i, j)$, using the above observation and Lemma 1, one can deduce that if $n \geq t(l)$ and $i, j \geq n$ then the probability that $\mathbf{x} \in T^{n,l}$ is greater than α. Formally we have the following corollary.

Corollary 2. *Let $p(\alpha)$ be defined as in Lemma 1. For every $p \in (p(\alpha), 1)$ and $l \in \mathbb{N}$ there exists $t(l)$ such that for every $n \geq t(l)$, and an agent $\mathbf{x} = (i, j) \in B_t$ such that $i, j \geq n$*

$$\mathbf{P}_p(\mathbf{x} \in T^{n,l}) > \alpha.$$

Proof of Corollary 2. Let $p \in (p(\alpha), 1)$. By Lemma 1 there exists $t(l)$ such that for every $n \geq t(l)$ it holds that $\mathbf{P}_{\sigma,p}(|\xi_n^0| \geq l) > \alpha$. Since the set $|\xi_n^{\mathbf{x}}|$ has identical distribution to $|\xi_n^0|$ we have that,

$$\mathbf{P}_p(\mathbf{x} \in T^{n,l}) > \alpha.$$

$\square$

Call agent $\mathbf{x}$ isolated if he does not observe any other agents. Let R_t^k be the set of agents in B_t that can observe at least k isolated agents, and let r_t^k be the size of R_t^k.

The following lemma shows that if p is large enough, then for every k the proportion of agents that can observe at least k isolated agents exceeds α with probability that goes to 1 when t increases.

Lemma 2. *For every $p \in (p(\alpha), 1)$, and k*

$$\lim_{t \to \infty} \mathbf{P}_p\left(\frac{r_t^k}{b_t} > \alpha\right) = 1.$$

It is important to note that $p(\alpha)$ depends only on α and not on k. Before proving the lemma we establish the following auxiliary result.

Lemma 3. *For every $t \in \mathbb{N}$ let $\{X_i^t\}_{1 \leq i \leq m_t}$ be a sequence of Bernoulli random variables for which there exists $\epsilon > 0$, and $\alpha > 0$ such that $E(X_i^t) \geq \alpha + \epsilon$ for every i. Assume that there exists an integer n such that for every i the random variable X_i^t depends on at most n other random variables from $\{X_i^t\}_{1 \leq i \leq m_t}$, and that $m_t \to_{t \to \infty} \infty$. Then,*

$$\lim_{t \to \infty} \mathbf{P}\left(\frac{1}{m_t} \sum_{i=1}^{m_t} X_i^t > \alpha\right) = 1.$$

Proof. The proof follows from Theorem 2 in Andrews [3]. It also follows directly from Chebyshev inequality. $\square$

Proof of Lemma 2. Note first that the probability that a given agent $\mathbf{x} \in B_t$ observes at least k isolated agents from B_{t-n}, is larger or equal to the following expression

$$\mathbf{P}_p(\mathbf{x} \in T^{n,l})\mathbf{P}_p(\text{at least } k \text{ agents out of } l \text{ are isolates}). \tag{9}$$

The probability of each $\mathbf{y} \in \xi_n^{\mathbf{x}}$ to be isolated is $(1-p)^2$. This happens independently for all agents in $\xi_n^{\mathbf{x}}$. Hence, when l grows and $n \geq t(l)$, the probability that a given agent $\mathbf{x} \in T^{n,l}$ observes at least k isolated agents approaches 1. By Corollary 2 for every $n \geq t(l)$ and agent $\mathbf{x} = (i,j)$ such that $i,j \geq n$ it holds that $\mathbf{P}_p(\mathbf{x} \in T^{n,l}) > \alpha$. Hence for every k we can find large enough $l = l(k)$ such that for every for every $n \geq t(l)$ and agent $\mathbf{x} = (i,j)$ such that $i,j \geq n$,

$$\mathbf{P}_p(\mathbf{x} \in T^{n,l})\mathbf{P}_p(\text{at least } k \text{ agents out of } l \text{ are isolates}) > \alpha. \tag{10}$$

We shall identify B_t with the set $\{0,\ldots,t\}$ where each number $0 \leq i \leq t$ corresponds to the agent $\mathbf{x} = (i, t-i) \in B_t$. Fix k, l such that $l \geq l(k)$, and $n \geq t(l)$. For every agent $i \in B_t$ define a random variable h_i^t to be equal 1 if the agent $i \in B_t$ observes at least k isolated nodes in ξ_n^i. By equation (10), for every large t and $n \leq i \leq t-n$ it holds that $\mathbf{P}_{\sigma,p}(h_i^t = 1) > \alpha$.

In addition note that if $0 \leq i,j \leq t$ are such that $|i - j| > n$ then $\xi_n^i \cap \xi_n^j = \varnothing$ with probability 1 hence h_i^t and h_j^t are independent random variables. Based on Lemma 3 it follows that:

$$\lim_{t \to \infty} \mathbf{P}_p\Big(\sum_{n \leq i \leq t-n} \frac{h_i^t}{t - 2n + 1} > \alpha \Big) = 1. \tag{11}$$

Hence in particular since n is finite,[14]

$$\lim_{t \to \infty} \mathbf{P}_p\Big(\sum_{0 \leq i \leq t} \frac{h_i^t}{t + 1} > \alpha \Big) = 1. \tag{12}$$

This concludes the proof of the lemma as

$$\mathbf{P}_p\Big(\frac{r_t^k}{b_t} > \alpha\Big) \geq \mathbf{P}_p\Big(\sum_{0 \leq i \leq t} \frac{h_i^t}{t + 1} > \alpha \Big) = 1.$$

$\square$

By Lemma 1 one can deduce that for $p > p(\alpha)$ a proportion α of agents in B_t has an access to a growing number of isolated points. This in turn means that this set of agents can infer the

---

[14]Here we rely on the fact that $b_t = t + 1$.

true state, and hence take the correct action, with arbitrarily high precision which is exactly what we wanted to prove. Establishing this formally requires a little more work.

Assume that a given agent observes only isolated agents. If the agent knows the strategy of the isolated agents then, as k grows to infinity, the agent learns the true state of the world with arbitrarily high probability. Therefore, there exists a sequence $\{q_k\}_k$ converging to 1 such that if a given agent observes at least k isolated agents then he assigns the true state of the world a probability that is greater than $\frac{1}{2}$ with probability of at least q_k.[15]

We fix an equilibrium strategy σ of Γ_p^2. For every agent $i \in B_t$ we let A_i be the set of isolated agents that i can observe. For every i we assign a random variable X_i^t that equals 1 if conditional on observing only the actions of the agents in A_i the true true state has a probability that is greater than $\frac{1}{2}$. We then have the following corollary.

Corollary 3. *For every $p \in (p(\alpha), 1)$,*[16]

$$\lim_{k \to \infty} \lim_{t \to \infty} E_{\sigma,p}\left(\frac{\sum_{i \in R_t^k} X_i^t}{r_t^k}\right) = 1.$$

Proof of Corollary 3. Note that $\mathbf{P}_{\sigma,p}(X_i^t = 1 | i \in R_t^k) \geq q_k$. For every k we have by Lemma 2 that

$$\lim_{t \to \infty} \mathbf{P}_{\sigma,p}(r_t^k = 0) = 0.$$

Hence for every k,

$$\lim_{t \to \infty} E_{\sigma,p}\left(\frac{\sum_{i \in R_t^k} X_i^t}{r_t^k}\right) \geq q_k.$$

Corollary 3 follows since $\lim_{k \to \infty} q_k = 1$. $\qquad\square$

We now let Y_i^t be the random variable that corresponds to the payoff of agent $i \in B_t$. It must hold that $E_{\sigma,p}(X_i^t) \leq E_{\sigma,p}(Y_i^t)$. This follows since observing more agents only improves the expected payoff. Lemma 2 implies the following corollary.

Corollary 4. *For every $p \in (p(\alpha), 1)$,*

$$\lim_{k \to \infty} \lim_{t \to \infty} E_{\sigma,p}\left(\frac{\sum_{i \in R_t^k} Y_i^t}{r_t^k}\right) = 1.$$

---

[15]q_k may be taken independently of the precise strategy of the isolated agents, as long they are best replying with respect to their private belief.

[16]Set $\frac{\sum_{i \in R_t^k} X_i^t}{r_t^k}$ to be 0 if $r_n^k = 0$.

Proof of Theorem 2. Lemma 2 shows that

$$\lim_{k\to\infty}\lim_{t\to\infty}\mathbf{P}_{\sigma,p}\left(\frac{r_t^k}{t+1} > \alpha\right) = 1.$$

Since for every k and t, the random variable $\frac{\sum_{i\in R_t^k} Y_i^t}{r_t^k}$ is bounded above by 1 it follows from Corollary 4 that for every $\varepsilon > 0$,

$$\liminf_{k\to\infty}\liminf_{t\to\infty}\mathbf{P}_{\sigma,p}\left(\frac{\sum_{i\in R_t^k} Y_i^t}{r_t^k} < 1 - \varepsilon\right) = 0.$$

Therefore it must hold that,

$$\lim_{k\to\infty}\lim_{t\to\infty}\mathbf{P}_{\sigma,p}\left(\frac{\sum_{i\in R_t^k} Y_i^t}{r_t^k} = 1\right) = 1.$$

As a result of Lemma 2 we can deduce that,

$$\lim_{k\to\infty}\lim_{t\to\infty}\mathbf{P}_{\sigma,p}\left(\frac{\sum_{i\in R_t^k} Y_i^t}{t+1} > \alpha\right) = \lim_{k\to\infty}\lim_{t\to\infty}\mathbf{P}_{\sigma,p}\left(\frac{\sum_{i\in R_t^k} Y_i^t}{r_t^k}\frac{r_t^k}{t+1} > \alpha\right) = 1.$$

Hence it follows that for every $\epsilon > 0$ there exists a k such that,

$$\lim_{t\to\infty}\mathbf{P}_{\sigma,p}\left(\frac{\sum_{i\in R_t^k} Y_i^t}{t+1} > \alpha\right) \geq 1 - \epsilon,$$

Note that for every t,

$$\frac{\sum_{i\in R_t^k} Y_i^t}{t+1} \leq \frac{\sum_{i\in B_t} Y_i^t}{t+1}.$$

Therefore,

$$\lim_{t\to\infty}\mathbf{P}_{\sigma,p}\left(\frac{\sum_{i\in B_t} Y_i^t}{t+1} > \alpha\right) \geq 1 - \epsilon,$$

Since ϵ was arbitrary it follows that,

$$\lim_{t\to\infty}\mathbf{P}_{\sigma,p}\left(\frac{\sum_{i\in B_t} Y_i^t}{t+1} > \alpha\right) = 1.$$

which concludes the proof of our main theorem. $\qquad\square$

Proof of Theorem 3. Theorem 3 readily follows from the proof of Theorem 2. Lemma 3 demonstrates that for every $\alpha < 1$ and $p > p(\alpha)$ it holds for every k that in the limit an α proportion of agents observe at least k isolated agents with probability one. When k grows the prior probability of every agent in R_t^k converges to the true state. Therefore in particular, as signals are bounded, a proportion of at least α ignores their private signals as t grows. $\qquad\square$

We shall briefly discuss the generalization of the proof to the case where the dimension of the lattice is $m > 2$. Proposition 2 holds for any $m > 2$ since a 2-dimensional lattice containing the origin is embedded in any lattice of dimension $m > 2$. Therefore, Lemma 1 and Corollary 2 are also valid for any $m > 2$. For Lemma 2 note that in this case the probability that a given node is isolated decreases to $(1-p)^m$. But the proof of the lemma only relies on the fact that this probability is bounded away from 0 independently for every node. In addition note that the sum in equation (11) should be taken over all agents $\mathbf{x} \in B_t$ for which $\mathbf{x}_- \geq n$. Since for any fixed n, the proportion of such agents in B_t tends to 1 as t grow, equation (12) is also valid for the general case when we replace $t+1$ with b_t. The rest of the proof follows without change.

References

[1] D. Acemoglu, M.A. Dahleh, I. Lobel, A. Ozdaglar (2010), "Bayesian Learning in Social Networks", *Review of Economic Studies* 78, 1-34.

[2] L.R. Anderson, C. Holt (1997), "Information Cascades in the Laboratory", *American Economic Review* 87 ,847-862

[3] D.W.K. Andrews (1988) "Law of Large Numbers for Dependent Non-Identically Distributed Random Variables", *Econometric Theory* 4, 458-467.

[4] I. Arieli, M. Mueller-Frank (2013), "Inferring Beliefs from Actions", working paper, IESE Business School.

[5] I. Arieli, M. Mueller-Frank (2014), "Learning what matters", working paper, IESE Business School.

[6] A.V. Banerjee (1992), "A Simple Model of Herd Behavior", *Quarterly Journal of Economics* 107, 797-817.

[7] A.V. Banerjee, Drew Fudenberg (2004), "Word-of-Mouth Learning", *Games and Economic Behavior* 46, 797-817.

[8] S. Bikhchandani, D. Hirshleifer, I. Welch (1992), "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades", *The Journal of Political Economy* 100, 151-170.

[9] B. Bollobas, O. Riordan (2006), *Percolation*, Cambridge University Press.

[10] B. Celen, S. Kariv (2004), "Observational Learning under Imperfect Information", *Games and Economic Behavior* 47, 72-86.

[11] T.G. Conley, C.R. Udry (2010), "Learning about a New Technology: Pineapple in Ghana", *American Economic Review* 100, 35-69.

[12] R. Durrett (1984), "Oriented Percolation in Two Dimensions", *The Annals of Probability* 12, 999-1040.

[13] J.K. Goeree, T.R. Palfrey, B.W. Rogers, R.D. McKelvey (2007), "Self-Correcting Information Cascades", *Review of Economic Studies* 74, 733-762.

[14] Z. Griliches (1957), "Hybrid Corn: An Exploration in the Economics of Technological Change", *Econometrica* 25, 501-522.

[15] H. Herrera, J. Hoerner (2013), "A Necessary and Sufficient Condition for Information Cascades", working paper, Yale University.

[16] I.H. Lee, A. Valentinyi (2000), "Noisy Contagion Without Mutation", *Review of Economic Studies* 67, 47-56.

[17] I. Lobel, E. Sadler (2013), "Social Learning and Aggregate Network Uncertainty", working paper, New York University.

[18] E. Mansfield (1961), "Technical Change and the Rate of Imitation", *Econometrica* 29, 741-766.

[19] S. Morris (2000), "Contagion", *Review of Economic Studies* 67, 57-78.

[20] E. Mossel, A. Sly, O. Tamuz (2012), "On Agreement and Learning", working paper U.C. Berkeley.

[21] L. Smith (1991), "Essays on Dynamic Models of Equilibrium and Learning", *Ph.D. thesis*, University of Chicago

[22] L. Smith, P. Sorensen (2000), "Pathological Outcomes of Observational Learning", *Econometrica* 68, 371-398.