

Opinion exchange dynamics

Elchanan Mossel

Massachusetts Institute of Technology
e-mail: elmos@mit.edu

and

Omer Tamuz

California Institute of Technology
e-mail: tamuz@caltech.edu

Received January 2014.

Contents

1	Introduction	156
1.1	Modeling opinion exchange	156
1.1.1	Modeling approaches	157
1.2	Mathematical connections	158
1.3	Related literature	158
1.4	Framework	158
1.5	General definitions	159
1.5.1	Agents, state of the world and private signals	159
1.5.2	The social network	159
1.5.3	Time periods and actions	159
1.5.4	Extensions, generalizations, variations and special cases	160
1.6	Questions	160
1.7	Acknowledgments	161
2	Heuristic models	161
2.1	The DeGroot model	161
2.1.1	Definition	161
2.1.2	Questions and answers	162
2.1.3	Results	162
2.1.4	Degroot with cheaters and bribes	164
2.1.5	The case of infinite graphs	164
2.2	The voter model	165
2.2.1	Definition	165
2.2.2	Questions and answers	165
2.2.3	Results	166
2.2.4	A variant of the voter model	167
2.3	Deterministic iterated dynamics	168
2.3.1	Definition	168

2.3.2	Questions and answers	169
2.3.3	Convergence	170
2.3.4	Retention of information	171
3	Bayesian models	174
3.0.5	Toy model: Continuous actions	174
3.0.6	Definitions and some observations	175
3.1	Agreement	178
3.2	Continuous utility models	179
3.3	Bounds on number of rounds in finite probability spaces	179
3.4	From agreement to learning	180
3.4.1	Agreement on beliefs	181
3.4.2	Agreement on actions	182
3.5	Sequential models	185
3.5.1	The external observer at infinity	186
3.5.2	The agents' calculation	187
3.5.3	The Markov chain and the martingale	187
3.5.4	Information cascades, convergence and learning	188
3.6	Learning from discrete actions on networks	189
3.6.1	Additional general notation	190
3.6.2	Sequences of rooted graphs and their limits	191
3.6.3	Coupling isomorphic balls	194
3.6.4	δ -independence	194
3.6.5	Asymptotic learning	195
3.6.6	Example of non-atomic private beliefs leading to non-learning	201
	References	202

1. Introduction

1.1. Modeling opinion exchange

The exchange of opinions between individuals is a fundamental social interaction that plays a role in nearly any social, political and economic process. While it is unlikely that a simple mathematical model can accurately describe the exchange of opinions between two people, one could hope to gain some insights on emergent phenomena that affect large groups of people.

Moreover, many models in this field are an excellent playground for mathematicians, especially those working in probability, algorithms and combinatorics. The goal of this survey is to introduce such models to mathematicians, and especially to those working in discrete mathematics, information theory, optimization, probability and statistics.

1.1.1. Modeling approaches

Many of the models we discuss in the survey comes from the literature in theoretical economics. In microeconomic theory, the main paradigm of modeling human interaction is by a *game*, in which participants are rational agents, choosing their moves optimally and responding to the strategies of their peers. A particularly interesting class of games is that of probabilistic *Bayesian games*, in which players also take into account the uncertainty and randomness of the world. We study Bayesian models in Section 3.

Another class of models, which have a more explicit combinatorial description, are what we refer to as *heuristic* models. These consider the dynamics that emerge when agents are assumed to utilize some (usually simple) update rule or algorithm when interacting with each other. Economists often justify such models as describing agents with *bounded rationality*. We study such models in Section 2.

It is interesting that both of these approaches are often justified by an *Occam's razor* argument. To justify the heuristic models, the argument is that assuming that people use a simple heuristic satisfies Occam's razor. Indeed, it is undeniable that the simpler the heuristic, the weaker the assumption. On the other hand, the Bayesian argument is that even by choosing a simple heuristic one has too much freedom to reverse engineer any desired result. Bayesians therefore opt to only assume that agents are rational. This, however, may result in extremely complicated behavior.

There exists several other natural dichotomies and sub-dichotomies. In rational models, one can assume that agents tell each other their opinions. A more common assumption in Economics is that agents learn by observing each other's *actions*; these are choices that an individual makes that not only reflect their belief, but also carry potential gain or penalty. For example, in financial markets one could assume that traders tell each other their value estimates, but a perhaps more natural setting is that they learn about these values by seeing which actual bids their peers place, since the latter are costly to manipulate. Hence the adage "actions speak louder than words."

Some actions can be more revealing than others. A bid by a trader could reveal the value the trader believes the asset carries, but in a different setting it could perhaps just reveal whether the trader thinks that the asset is currently overpriced or underpriced. In other models an action could perhaps reveal *all* that an agent knows. We shall see that widely disparate outcomes can result in models that differ only by how revealing the actions are.

Although the distinction between *opinions*, *beliefs* and *actions* is sometimes blurry, we shall follow the convention of having agents learn from each other's actions. While in some models this will only be a matter of nomenclature, in others this will prove to be a pivotal choice. The term **belief** will be reserved for a technical definition (see below), and we shall not use *opinion*, except informally.

1.2. *Mathematical connections*

Many of the models of information exchange on networks are intimately related to nice mathematical concepts, often coming from probability, discrete mathematics, optimization and information theory. We will see how the theories of Markov chains, martingale arguments, influences and graph limits all play a crucial role in analyzing the models we describe in these notes. Some of the arguments and models we present may fit well as classroom materials or exercises in a graduate course in probability.

1.3. *Related literature*

It is impossible to cover the huge body of work related to information exchange in networks. We will cite some relevant papers at each section. Mathematicians reading the economics literature may benefit from keeping the following two comments in mind:

- The focus in economics is not the mathematics, but the economics, and in particular the justification of the model and the interpretation of the results. Thus important papers may contain little or no new mathematics. Of course, many papers do contain interesting mathematics.
- For mathematicians who are used to models coming from natural sciences, the models in the economics literature will often look like very rough approximation and the conclusions drawn in terms of real life networks unjustified. Our view is that the models have very limited implication towards real life and can serve as most as allegories. We refer the readers who are interested in this point to Rubinstein's book "Economic Fables" [39].

1.4. *Framework*

The majority of models we consider share the a common underlying framework, which describes a set of agents, a state of the world, and the information the agents have regarding this state. We describe it formally in Section 1.5 below, and shall note explicitly whenever we depart from it.

We will take a probabilistic / statistical point of view in studying models. In particular we will assume that the model includes a random variable S which is the *true state of the world*. It is this S that all agents want to learn. For some of the models, and in particular the rational, economic models, this is a natural and even necessary modeling choice. For some other models - the voter model, for example (Section 2.2), this is a somewhat artificial choice. However, it helps us take a single perspective by asking, for each model, how well it performs as a statistical procedure aimed at estimating S . Somewhat surprisingly, we will reach similarly flavored conclusions in widely differing settings. In particular, a repeated phenomenon that we observe is that *egalitarianism*, or *decentralization* facilitates the flow of information in social networks, in both game-theoretical and heuristic models.

1.5. General definitions

1.5.1. Agents, state of the world and private signals

Let V be a countable set of agents, which we take to be $\{1, 2, \dots, n\}$ in the finite case and $\mathbb{N} = \{1, 2, \dots\}$ in the infinite case. Let $\{0, 1\}$ be the set of possible values of the *state of the world* S .

Let Ω be a compact metric space equipped with the Borel sigma-algebra. For example, and without much loss of generality, Ω could be taken to equal the closed interval $[0, 1]$. Let $W_i \in \Omega$ be agent i 's private signal, and denote $\bar{W} = (W_1, W_2, \dots)$.

Fix μ_0 and μ_1 , two mutually absolutely continuous measures on Ω . We assume that S is distributed uniformly, and that conditioned on S , the W_i 's are i.i.d. μ_S : when $S = 0$ then $\bar{W} \sim \mu_0^V$, and when $S = 1$ then $\bar{W} \sim \mu_1^V$.

More formally, let δ_0 and δ_1 be the distributions on $\{0, 1\}$ such that $\delta_0(0) = \delta_1(1) = 1$. We consider the probability space $\{0, 1\} \times \Omega^V$, with the measure \mathbb{P} defined by

$$\mathbb{P} = \mathbb{P}_{\mu_0, \mu_1, V} = \frac{1}{2} \delta_0 \times \mu_0^V + \frac{1}{2} \delta_1 \times \mu_1^V,$$

and let

$$(S, \bar{W}) \sim \mathbb{P}.$$

1.5.2. The social network

A **social network** $G = (V, E)$ is a directed graph, with V the set of agents. The set of neighbors of $i \in V$ is $\partial i = \{j : (i, j) \in E\} \cup \{i\}$ (i.e., ∂i includes i). The **out-degree** of i is given by $|\partial i|$. The **degree** of G is given by $\sup_{i \in V} |\partial i|$.

We make the following assumption on G .

Assumption 1.1. *We assume throughout that G is simple and strongly connected, and that each out-degree is finite.*

We recall that a graph is strongly connected if for every two nodes i, j there exists a directed path from i to j . Finite out-degrees mean that an agent observes the actions of a finite number of other agents. We do allow infinite **in-degrees**; this corresponds to agents whose actions are observed by infinitely many other agents. In the different models that we consider we impose various other constraints on the social network.

1.5.3. Time periods and actions

We consider the discrete time periods $t = 0, 1, 2, \dots$, where in each period each agent $i \in V$ has to choose an action $A_t^i \in \{0, 1\}$. This action is a function of agent i 's private signal, as well as the actions of its neighbors in previous time periods, and so can be thought of as a function from $\Omega \times \{0, 1\}^{|\partial i| \cdot t}$ to $\{0, 1\}$. The exact functional dependence varies among the models.

1.5.4. Extensions, generalizations, variations and special cases

The framework presented above admits some natural extensions, generalizations and variations. Conversely, some special cases deserve particular attention. Indeed, some of the results we describe apply more generally, while others do not apply more generally, or apply only to special cases. We discuss these matters when describing each model.

- **The state of the world** can take values from sets larger than $\{0, 1\}$, including larger finite sets, countably infinite sets or continuums.
- **The agents' private signals** may not be i.i.d. conditioned on S : they may be independent but not identical, they may be identical but not independent, or they may have a general joint distribution.

An interesting special case is when the space of private signals is equal to the space of the states of the world. In this case one can think of the private signals as each agent's initial guess of S .

- A number of models consider only undirected **social networks**, that is, symmetric social networks in which $(i, j) \in E \Leftrightarrow (j, i) \in E$.
- More general network model include weighted directed models where different directed edges have different weights.
- **Time** can be continuous. In this case we assume that each agent is equipped with an i.i.d. *Poisson clock* according to which it “wakes up” and acts. In the finite case this is equivalent to having a single, uniformly chosen random agent act in each discrete time period. It is also possible to define more general continuous time processes.
- **Actions** can take more values than $\{0, 1\}$. In particular we shall consider the case that actions take values in $[0, 1]$.

In order to model randomized behavior of the agents, we shall also consider actions that are not measurable in the private signal, but depend also on some additional randomness. This will require the appropriate extension of the measure \mathbb{P} to a larger probability space.

1.6. Questions

The main phenomena that we shall study are *convergence*, *agreement*, *unanimity*, *learning* and more.

- **Convergence.** We say that agent i *converges* when $\lim_t A_t^i$ exists. We say that the entire process converges when all agents converge.

The question of convergence will arise in all the models we study, and its answer in the positive will often be a requirement for subsequent treatment. When we do have convergence we define

$$A_\infty^i = \lim_{t \rightarrow \infty} A_t^i.$$

- **Agreement and unanimity.** We say that agents i and j *agree* when $\lim_t A_t^i = \lim_t A_t^j$. *Unanimity* is the event that i and j agree for all pairs of agents i and j . In this case we can define

$$A_\infty = A_\infty^i,$$

where the choice of i on the r.h.s. is immaterial.

- **Learning.** We say that agent i *learns* S when $A_\infty^i = S$, and that learning occurs in a model when all agents learn. In cases where we allow actions in $[0, 1]$, we will say that i learns whenever $\text{round}(A_\infty^i) = S$, where $\text{round}(\cdot)$ denotes rounding to the nearest integer, with $\text{round}(1/2) = 1/2$.

We will also explore the notion of *asymptotic learning*. This is said to occur for a sequence of graph $\{G_n\}_{n=1}^\infty$ if the agents on G_n learn with probability approaching one as n tends to infinity.

A recurring theme will be the relation between these questions and the *geometry* or *topology* of the social network. We shall see that indeed different networks may exhibit different behaviors in these regards, and that in particular, and across very different settings, *decentralized* or *egalitarian* networks tend to promote *learning*.

1.7. Acknowledgments

Allan Sly is our main collaborator in this field. We are grateful to him for allowing us to include some of our joint results, as well as for all that we learned from him. The manuscript was prepared for the 9th Probability Summer School in Cornell, which took place in July 2013. We are grateful to Laurent Saloff-Coste and Lionel Levine for organizing the school and for the participants for helpful comments and discussions. We would like to thank Shachar Kariv for introducing us to this field, and Eilon Solan for encouraging us to continue working in it. The research of Elchanan Mossel is partially supported by NSF grants DMS 1106999 and CCF 1320105, and by ONR grant N000141110140. Omer Tamuz was supported by a Google Europe Fellowship in Social Computing.

2. Heuristic models

2.1. The DeGroot model

The first model we describe was pioneered by Morris DeGroot in 1974 [13]. DeGroot's contribution was to take standard results in the theory of Markov Processes (See, e.g., Doob [15]) and apply them in the social setting. The basic idea for these models is that people repeatedly average their neighbors' actions. This model has been studied extensively in the economics literature. The question of *learning* in this model has been studied by Golub and Jackson [22].

2.1.1. Definition

Following our general framework (Section 1.5), we shall consider a state of the world $S \in \{0, 1\}$ with conditionally i.i.d. private signals. The distribution of private signals is what we shall henceforth refer to as *Bernoulli* private signals:

for some $\frac{1}{2} > \delta > 0$, $\mu_i(S) = \frac{1}{2} + \delta$ and $\mu_i(1-S) = \frac{1}{2} - \delta$, for $i = 0, 1$. Obviously this is equivalent to setting $\mathbb{P}[W_i = S | S] = \frac{1}{2} + \delta$.

In the DeGroot model, we let the actions take values in $[0, 1]$. In particular, we define the actions as follows:

$$A_0^i = W_i$$

and for $t > 0$

$$A_t^i = \sum_{j \in \partial i} w(i, j) A_{t-1}^j, \quad (2.1)$$

where we make the following three assumptions:

1. $\sum_{j \in \partial i} w(i, j) = 1$ for all $i \in V$.
2. $i \in \partial i$ for all $i \in V$.
3. $w(i, j) > 0$ for all $(i, j) \in E$.

The last two assumptions are non-standard, and, in fact, not strictly necessary. We make them to facilitate the presentation of the results for this model.

We assume that the social network G is finite. We consider both the general case of a directed strongly connected network, and the special case of an undirected network.

2.1.2. Questions and answers

We shall ask, with regards to the DeGroot model, the same three questions that appear in Section 1.6.

1. **Convergence.** Is it the case that agents' actions converge? That is, does, for each agent i , the limit $\lim_t A_t^i$ exist almost surely? We shall show that this is indeed the case.
2. **Agreement.** Do all agents eventually reach agreement? That is, does $A_\infty^i = A_\infty^j$ for all $(i, j) \in V$? Again, we answer this question in the positive.
3. **Learning.** Do all agents learn? In the case of continuous actions we say that agent i has learned S if $\text{round}(A_\infty^i) = S$. Since we have agreement in this model, it follows that either all agents learn or all do not learn. We will show that the answer to this question depends on the topology of the social network, and that, in particular, a certain form of *egalitarianism* is a sufficient condition for learning with high probability.

2.1.3. Results

The key to the analysis of the DeGroot model is the realization that (2.1) describes a transformation from the actions at time $t - 1$ to the actions at time t that is the *Markov operator* P_w of the a random walk on the graph G . However, while usually the analysis of random walks deals with action of P_w on

distributions from the right, here we act on functions from the left [16]. While this is an important difference, it is still easy to derive properties of the DeGroot process from the theory of Markov chains (see, e.g., Doob [15]).

Note first, that assumptions (2) and (3) on (2.1) make this Markov chain irreducible and a-periodic. Since, for a node j

$$A_t^j = \mathbb{E} \left[W_{X_t^j} \right],$$

where X_t^j is the Markov chain started at j and run for t steps, it follows that $A_\infty^j := \lim_t A_t^j$ is nothing but the expected value of the private signals, according to the stationary distribution of the chain. We thus obtain

Theorem 2.1 (Convergence and agreement in the DeGroot model). *For each $j \in V$,*

$$A_\infty := \lim_t A_t^j = \sum_{i \in V} \alpha_i W_i,$$

where $\alpha = (\alpha_1, \dots, \alpha_n)$ is the stationary distribution of the Markov chain described by P_w .

Recall that α is the left eigenvector of P_w corresponding to eigenvalue 1, normalized in ℓ^1 . In the internet age, the vector α is also known as the *PageRank* vector [37]. It is the asymptotic probability of finding a random walker at a given node after infinitely many steps of the random walk. Note that α is not random; it is fixed and depends only on the weights w . Note also that Theorem 2.1 holds for *any* real valued starting actions, and not just ones picked from the distribution described above. To gain some insight into the result, let us consider the case of undirected graphs and simple (lazy) random walks. For these, it can be shown that

$$\alpha_i = \frac{|\partial i|}{\sum_j |\partial j|}.$$

Recall that $\mathbb{P} [A_0^i = S] = \frac{1}{2} + \delta$. We observe the following.

Proposition 2.2 (Learning in the DeGroot model). *For a set of weights w , let $p_w(\delta) = \mathbb{P}[\text{round}(A_\infty) = S]$. Then:*

- p_w is a monotone function of δ with $p_w(0) = 1/2$ and $p_w(1/2) = 1$.
- For a fixed $0 < \delta < 1/2$, among all w 's on graphs of size n , $p_w(\delta)$ is maximized when the stationary distribution of G is uniform.

Proof. • The first part follows by coupling. Note that we can couple the processes with $\delta_1 < \delta_2$ such that the value is S is the same and moreover, whenever $W_i = S$ in the δ_1 process we also have $W_1 = S$ in the δ_2 process. Now, since the vector α is independent of δ and $A_\infty = \sum_i \alpha_i W_i$, the coupling above results in $|A_\infty - S|$ being smaller in the δ_2 process than it is in the δ_1 process.

- The second part follows from the Neyman-Pearson lemma in statistics. This lemma states that among *all* possible estimators, the one that maximizes the probability that S is reconstructed correctly is given by

$$\hat{S} = \text{round} \left(\frac{1}{n} \sum_i W_i \right) \quad \square$$

We note that an upper bound on $p_w(\delta)$ can be obtained using Hoeffding's inequality [23]. We leave this as an exercise to the reader.

Finally, the following proposition is again a consequence of well known results on Markov chains. See the books by Saloff-Coste [40] or Levin, Peres and Wilmer [26] for basic definitions.

Proposition 2.3 (Rate of Convergence in the Degroot Model). *Suppose that at time t , the total variation distance between the chain started at i and run for t steps and the stationary distribution is at most ϵ . Then a.s.:*

$$\max_i |A_t^i - A_\infty| \leq 2\epsilon\delta.$$

Proof. Note that

$$A_t^i - A_\infty = \mathbb{E} \left[W_{X_t^i} - W_{X_\infty} \right].$$

Since we can couple the distributions of X_t and X_∞ so that they disagree with probability at most ϵ and the maximal difference between any two private signals is at most δ , the proof follows. \square

2.1.4. Degroot with cheaters and bribes

A cheater is an agent who plays a fixed action.

- **Exercise.** Consider the DeGroot model with a single cheater who picks some fixed action. What does the process converge to?
- **Exercise.** Consider the DeGroot model with k cheaters, each with some (perhaps different) fixed action. What does the model converge to?
- **Research problem.** Consider the following zero sum game.¹ A and B are two companies. Each company's strategy is a choice of k cheaters (cheaters chosen by both play honestly), for whom the company can choose a fixed value in $[0, 1]$. The utility of company A is the sum of the players' limit actions, and the utility of company B is minus the utility of A . What are the equilibria of this game?

2.1.5. The case of infinite graphs

Consider the DeGroot model on an infinite graph, with a simple random walk.

¹We do not formally define games here. A good introduction is Osborne and Rubinstein's textbook [35].

- **Easy exercise.** Give an example of specific private signals for which the limit A_∞ doesn't exist.
- **Easy exercise.** Prove that A_∞ exists and is equal to S on non-amenable graphs a.s. A graph is non-amenable if the Markov operator $P_w : \ell^2(V) \rightarrow \ell^2(V)$ has norm strictly less than 1.
- **Harder exercise.** Prove that A_∞ exists and is equal to S on general infinite graphs.

2.2. The voter model

This model was described by P. Clifford and A. Sudbury [11] in the context of a spatial conflict where animals fight over territory (1973) and further analyzed by R.A. Holley and T.M. Liggett [24].

2.2.1. Definition

As in the DeGroot model above, we shall consider a state of the world $S \in \{0, 1\}$ with conditionally i.i.d. *Bernoulli* private signals, so that $\mathbb{P}[W_i = S] = \frac{1}{2} + \delta$.

We consider binary actions and define them in a way that resembles our definition of the DeGroot model. We let:

$$A_0^i = W_i$$

and for $t > 0$, all i and all $j \in \partial i$,

$$\mathbb{P}\left[A_t^i = A_{t-1}^j\right] = w(i, j), \quad (2.2)$$

so that in each round each agent chooses a neighboring agent to emulate. We make the following assumptions:

1. All choices are independent.
2. $\sum_{j \in \partial i} w(i, j) = 1$ for all $i \in V$.
3. $i \in \partial i$ for all $i \in V$.
4. $w(i, j) > 0$ for all $(i, j) \in E$.

As in the DeGroot model, the last two assumptions are non-standard, and are made to facilitate the presentation of the results for this model.

We assume that the social network G is finite. We consider both the general case of a directed strongly connected network, and the special case of an undirected network.

2.2.2. Questions and answers

We shall ask, with regards to the voter model, the same three questions that appear in Section 1.6.

1. **Convergence.** Does, for each agent i , the limit $\lim_t A_t^i$ exist almost surely? We shall show that this is indeed the case.

2. **Agreement.** Does $A_\infty^i = A_\infty^j$ for all $(i, j) \in V$? Again, we answer this question in the positive.
3. **Learning.** In the case of discrete actions we say that agent i has learned S if $A_\infty^i = S$. Since we have agreement in this model, it follows that either all agents learn or all do not learn. Unlike other models we have discussed, we will show that the answer here is no. Even for large egalitarian networks, learning doesn't necessarily hold. We will later discuss a variant of the voter model where learning holds.

2.2.3. Results

We first note that

Proposition 2.4. *In the voter model with assumptions (2.2) all agents converge to the same action.*

Proof. The voter model is a Markov chain. Clearly the states where $A_t^i = 0$ for all i and the state where $A_t^i = 1$ for all i are absorbing states of the chain. Moreover, it is easy to see that for any other state, there is a sequence of moves of the chain, each occurring with positive probability, that lead to the all 0 / all 1 state. From this it follows that the chain will always converge to either the all 0 or all 1 state. \square

We next wish to ask what is the probability that the agents learned S ? For the voter model this chance is never very high as the following proposition shows:

Theorem 2.5 ((Non) Learning in the Voter model). *Let A_∞ denote the limit action for all the agents in the voter model. Then:*

$$\mathbb{P}[A_\infty = 1|W] = \sum_{i \in V} \alpha_i W_i, \quad (2.3)$$

and

$$\mathbb{P}[A_\infty = S|W] = \sum_{i \in V} \alpha_i 1(W_i = S). \quad (2.4)$$

where $\alpha = (\alpha_1, \dots, \alpha_n)$ is the stationary distribution of the Markov chain described by P_w . Moreover,

$$\mathbb{P}[A_\infty = S] = \frac{1}{2} + \delta. \quad (2.5)$$

Proof. Note that (2.4) follows immediately from (2.3) and that (2.5) follows from (2.4) by taking expectation over W . To prove (2.3) we build upon a connection to the DeGroot model. Let D_t^i denote the action of agent i in the DeGroot model at time t . We are assuming that the DeGroot model is defined using the same $w(i, j)$ and that the private signals are identical for the voter and DeGroot model. Under these assumption it is easy to verify by induction on i and t that

$$\mathbb{P}[A_t^i = 1] = D_t^i.$$

Thus

$$\mathbb{P}[A_\infty^i = 1] = D_\infty^i = \sum_{i \in V} \alpha_i W_i,$$

as needed. \square

In the next section we will discuss a variant of the voter model that does lead to learning.

We next briefly discuss the question of the convergence rate of the voter model. Here again the connection to the Markov chain of the DeGroot model is paramount (see, e.g., Holley and Liggett [24]). We will not discuss this beautiful theory in detail. Instead, we will just discuss the case of undirected graphs where all the weights are 1.

Exercise. Consider the voter model on an undirected graph with n vertices. This is equivalent to letting $w(i, j) = 1/d_i$ for all i , where $d_i = |\partial i|$.

- Show that $X_t = \sum d_i A_t^i$ is a martingale.
- Let T be the stopping time where $A_t^i = 0$ for all i or $A_t^i = 1$ for all i . Show that $\mathbb{E}[X_T] = \mathbb{E}[X_0]$ and use this to deduce that

$$\mathbb{P}[A_\infty = 1|W] = \frac{\sum_{i \in V} d_i W_i}{\sum_{i \in V} d_i}$$

- Let $d = \max_i d_i$. Show that

$$\mathbb{E}[(X_t - X_{t-1})^2 | t < T] \geq 1/(2d).$$

Use this to conclude that

$$\mathbb{E}[T]/(2d) \leq \mathbb{E}[(X_T - X_0)^2] \leq n^2,$$

so

$$\mathbb{E}[T] \leq 2dn^2.$$

2.2.4. A variant of the voter model

As we just saw, the voter model does not lead to learning even on large egalitarian networks. It is natural to ask if there are variants of the model that do. We will now describe such a variant (see e.g. [5, 31]). For simplicity we consider an undirected graph $G = (V, E)$ and the following asynchronous dynamics.

- At time $t = 0$, let $A_i^0 = (W_i, 1)$.
- At each time $t \geq 1$ choose an edge $e = (i, j)$ of the graph at random and continue as follows:
 - For all $k \notin \{i, j\}$, let $A_k^t = A_k^{t-1}$.
 - Denote $(a_i, w_i) = A_i^{t-1}$ and $(a_j, w_j) = A_j^{t-1}$.
 - If $a_i \neq a_j$ and $w_i = w_j = 1$, let $a'_i = a_i, a'_j = a_j$ and $w'_i = w'_j = 0$.
 - If $a_i \neq a_j$ and $w_i = 1 > w_j = 0$, let $a'_i = a'_j = a_i$ and $w'_i = w_i$ and $w'_j = w_j$.

- Similarly, if $a_i \neq a_j$ and $w_j = 1 > w_i = 0$, let $a'_i = a'_j = a_j$ and $w'_i = w_i$ and $w'_j = w_j$.
- if $a_i \neq a_j$ and $w_j = w_i = 0$, let $a'_i = a'_j = 0(1)$ with probability $1/2$ each. Let $w'_i = w'_j = 0$.
- Otherwise, if $a_i = a_j$, let $a'_i = a_i, a'_j = a_j, w_i = w'_i, w_j = w'_j$.
- With probability $1/2$ let $A_i^t := (a'_i, w'_i)$ and $A_j^t := (a'_j, w'_j)$. With probability $1/2$ let $A_i^t := (a'_j, w'_j)$ and $A_j^t := (a'_i, w'_i)$.

Here is a useful way to think about this dynamics. The n players all begin with opinions given by W_i . Moreover these opinions are all strong (this is indicated by the second coordinate of the action being 1). At each round a random edge is chosen and the two agents sharing the edge declare their opinions regarding S . If their opinions are identical, then nothing changes except that with probability $1/2$ the agents swap their location on the edge. If the opinions regarding S differ and one agent is strong (second coordinate is 1) while the second one is weak (second coordinate is 0) then the weak agent is convinced by the strong agent. If the two agents are strong, then they keep their opinion but become weak. If the two of them are weak, then they both choose the same opinion at random. At the end of the exchange, the agents again swap their positions with probability $1/2$. We leave the following as an exercise:

Proposition 2.6. *Let $A_i^t = (X_i^t, Y_i^t)$. Then a.s.*

$$\lim X_i^t = X,$$

where

- $X = 1$ if $\sum_i W_i > n/2$,
- $X = 0$ if $\sum_i W_i < n/2$ and
- $\mathbb{P}[X = 1] = 1/2$ if $\sum_i W_i = n/2$.

Thus this variant of the voter model yields optimal learning.

2.3. Deterministic iterated dynamics

A natural deterministic model of discrete opinion exchange dynamics is **majority dynamics**, in which each agent adopts, at each time period, the opinion of the majority of its neighbors. This is a model that has been studied since the 1940's in such diverse fields as biophysics [27], psychology [10] and combinatorics [21].

2.3.1. Definition

In this section, let A_0^i take values in $\{-1, +1\}$, and let

$$A_{t+1}^i = \operatorname{sgn} \sum_{j \in \partial i} A_t^j.$$

we assume that $|\partial i|$ is odd, so that there are never cases of indifference and $A_t^i \in \{-1, +1\}$ for all t and i . We assume also that the graph is undirected.

A classical combinatorial result (that has been discovered independently repeatedly; see discussion and generalization in [21]) is the following.

Theorem 2.7. *Let $G = (V, E)$ be a finite undirected graph. Then*

$$A_{t+1}^i = A_{t-1}^i$$

for all i , for all $t \geq |E|$, and for all initial opinion sets $\{A_0^j\}_{j \in V}$.

That is, each agent (and therefore the entire dynamical system) eventually enters a cycle of period at most two. We prove this below.

A similar result applies to some infinite graphs, as discovered by Moran [28] and Ginosar and Holzman [20]; see also [44, 6]. Given an agent i , let $n_r(G, i)$ be the number of agents at distance exactly r from i in G . Let $\mathfrak{g}(G)$ denote the asymptotic growth rate of G given by

$$\mathfrak{g}(G) = \limsup_r n_r(G, i)^{1/n}.$$

This can be shown to indeed be independent of i . Then

Theorem 2.8 (Ginosar and Holzman, Moran). *If G has degree at most d and $\mathfrak{g}(G) < \frac{d+1}{d-1}$ then for each initial opinion set $\{A_0^j\}_{j \in V}$ and for each $i \in V$ there exists a time T_i such that*

$$A_{t+1}^i = A_{t-1}^i$$

for all $t \geq T_i$.

That is, each agent (but *not* the entire dynamical system) eventually enters a cycle of period at most two. We will not give a proof of this theorem.

In the case of graphs satisfying $\mathfrak{g}(G) < (d+1)/(d-1)$, and in particular in finite graphs, we shall denote

$$A_\infty^i = \lim_t A_{2t}^i.$$

This exists surely, by Theorem 2.8 above.

In this model we shall consider a state of the world $S \in \{-1, +1\}$ with conditionally i.i.d. Bernoulli private signals in $\{-1, +1\}$, so that $\mathbb{P}[W_i = S] = \frac{1}{2} + \delta$. As above, we set $A_0^i = W_i$.

2.3.2. Questions and answers

We ask the usual questions with regards to this model.

1. **Convergence.** While it is easy to show that agents' opinions do not necessarily converge in the usual sense, they do converge to sequences of period at most two. Hence we will consider the limit action $A_\infty^i = \lim_t A_{2t}^i$ as defined above to be the action that agent i converges to.

2. **Agreement.** This is easily *not* the case in this model that $A_\infty^i = A_\infty^j$ for all $i, j \in V$. However, in [29] it is shown that agreement is reached, with high probability, for good enough expander graph.²
3. **Learning.** Since we do not have agreement in this model, we will consider a different notion of learning. This notion may actually be better described as *retention of information*. We define it below. Condorcet’s Jury Theorem [12], in an early version of the law of large numbers, states that given n conditionally i.i.d. private signals, one can estimate S correctly, except with probability that tends to zero with n . The question of retention of information asks whether this still holds when we introduce correlations “naturally” by the process of majority dynamics. Let G be finite, undirected graphs. Let

$$\hat{S} = \operatorname{argmax}_{s \in \{-1, +1\}} \mathbb{P} \left[S = s \mid A_\infty^1, \dots, A_\infty^{|V|} \right].$$

This is the maximum a-posteriori (MAP) estimator of S , given the limit actions. Let

$$\iota(G, \delta) = \mathbb{P} \left[\hat{S} \neq S \right],$$

where G and δ appear implicitly in the right hand side. This is the probability that the best possible estimator of S , given the limit actions, is *not* equal to S .

Finally, let $\{G_n\}_{n \in \mathbb{N}}$ be a sequence of finite, undirected graphs. We say that we have *retention of information* on the sequence $\{G_n\}$ if $\iota(G_n, \delta) \rightarrow_n 0$ for all $\delta > 0$. This definition was first introduced, to the best of our knowledge, in Mossel, Neeman and Tamuz [29].

Is information retained on all sequences of growing graphs? The answer, as we show below, is no. However, we show that information is retained on sequences of *transitive graphs* [29].

2.3.3. Convergence

To prove convergence to period at most two for finite graphs, we define the *Lyapunov* functional

$$L_t = \sum_{(i,j) \in E} (A_{t+1}^i - A_t^j)^2.$$

We prove Theorem 2.7 by showing that L_t is monotone decreasing, that $A_{t+1}^i = A_{t-1}^i$ whenever $L_t - L_{t-1} = 0$, and that $L_t = L_{t-1}$ for all $t > |E|$. This proof appears (for a more general setting) in Goles and Olivos [21]. For this we will require the following definitions:

$$J_t^i = (A_{t+1}^i - A_{t-1}^i) \sum_{j \in \partial i} A_t^j$$

²We do not define expander graphs formally here; informally, they are graphs that resemble random graphs.

and

$$J_t = \sum_{i \in V} J_t^i.$$

Claim 2.9. $J_t^i \geq 0$ and $J_t^i = 0$ iff $A_{t+1}^i = A_{t-1}^i$.

Proof. This follows immediately from the facts that

$$A_{t+1}^i = \operatorname{sgn} \sum_{j \in \partial i} A_t^j,$$

and that $\sum_{j \in \partial i} A_t^j$ is never zero. □

It follows that

Corollary 2.10. $J_t \geq 0$ and $J_t = 0$ iff $A_{t+1}^i = A_{t-1}^i$ for all $i \in V$.

We next show that L_t is monotone decreasing.

Proposition 2.11. $L_t - L_{t-1} = -J_t$.

Proof. By definition,

$$L_t - L_{t-1} = \sum_{(i,j) \in E} (A_{t+1}^i - A_t^j)^2 - \sum_{(i,j) \in E} (A_t^i - A_{t-1}^j)^2.$$

Opening the parentheses and canceling identical terms yields

$$L_t - L_{t-1} = -2 \sum_{(i,j) \in E} A_{t+1}^i A_t^j + 2 \sum_{(i,j) \in E} A_t^i A_{t-1}^j.$$

Since the graph is undirected we can change variable on the right sum and arrive at

$$\begin{aligned} L_t - L_{t-1} &= -2 \sum_{(i,j) \in E} A_{t+1}^i A_t^j - A_t^j A_{t-1}^i \\ &= -2 \sum_{(i,j) \in E} (A_{t+1}^i - A_{t-1}^i) A_t^j. \end{aligned}$$

Finally, applying the definitions of J_t^i and J_t yields

$$L_t - L_{t-1} = - \sum_{i \in V} J_t^i = -J_t. \quad \square$$

Proof of Theorem 2.7. Since $L_0 \leq |E|$, $L_t \leq L_{t-1}$ and L_t is integer, it follows that $L_t \neq L_{t-1}$ at most $|E|$ times. Hence, by Proposition 2.11, $J_t > 0$ at most $|E|$ times. But if $J_t = 0$, then the state of the system at time $t+1$ is the same as it was at time $t-1$, and so it has entered a cycle of length at most two. Hence $J_t = 0$ for all $t > |E|$, and the claim follows. □

2.3.4. Retention of information

In this section we prove that

1. There exists a sequence of finite, undirected graphs $\{G_n\}_{n \in \mathbb{N}}$ of size tending to infinity such that $\iota(G, \delta)$ does not tend to zero for any $0 < \delta < \frac{1}{2}$.
2. Let $\{G_n\}_{n \in \mathbb{N}}$ be a sequence of finite, undirected, connected *transitive* graphs of size tending to infinity. Then $\iota(G_n, \delta) \rightarrow_n 0$, and, furthermore, if we let G_n have n vertices, then

$$\iota(G_n, \delta) \leq C n^{-\frac{C\delta}{\log(1/\delta)}}.$$

for some *universal* constant $C > 0$.

A *transitive graph* is a graph for which, for every two vertices i and j there exists a *graph homomorphism* σ such that $\sigma(i) = j$. A graph homomorphism h is a permutation on the vertices such that $(i, j) \in E$ iff $(\sigma(i), \sigma(j)) \in E$. Equivalently, the group $\text{Aut}(G) \leq S_{|V|}$ acts transitively on V .

Berger [8] gives a sequence of graphs $\{H_n\}_{n \in \mathbb{N}}$ with size tending to infinity, and with the following property. In each $H_n = (V, E)$ there is a subset of vertices W of size 18 such that if $A_t^i = -1$ for some t and all $i \in W$ then $A_\infty^j = -1$ for all $j \in V$. That is, if all the vertices in W share the same opinion, then eventually all agents acquire that opinion.

Proposition 2.12. $\iota(H_n, \delta) \geq (1 - \delta)^{18}$.

Proof. With probability $(1 - \delta)^{18}$ we have that $A_0^i = -S$ for all $i \in W$. Hence $A_\infty^j = -S$ for all $j \in V$, with probability at least $(1 - \delta)^{18}$. Since the MAP estimator \hat{S} can be shown to be a symmetric and monotone function of A_∞^j , it follows that in this case $\hat{S} = -S$, and so

$$\iota(H_n, \delta) = \mathbb{P} \left[\hat{S} \neq S \right] \geq (1 - \delta)^{18}. \quad \square$$

We next turn to prove the following result

Theorem 2.13. *Let G a finite, undirected, connected transitive graph with n vertices, n odd. then*

$$\iota(G, \delta) \leq C n^{-\frac{C\delta}{\log(1/\delta)}}.$$

for some universal constant $C > 0$.

Let $\hat{S} = \text{sgn} \sum_{i \in V} A_\infty^i$ be the result of a majority vote on the limit actions. Since n is odd then \hat{S} takes values in $\{-1, +1\}$. Note that \hat{S} is measurable in the initial private signals W_i . Hence there exists a function $f : \{-1, +1\}^n \rightarrow \{-1, +1\}$ such that

$$\hat{S} = f(W_1, \dots, W_n).$$

Claim 2.14. *f satisfies the following conditions.*

1. **Symmetry.** For all $x = (x_1, \dots, x_n) \in \{-1, +1\}^n$ it holds that $f(-x_1, \dots, -x_n) = -f(x_1, \dots, x_n)$.
2. **Monotonicity.** $f(x_1, \dots, x_n) = 1$ implies that $f(x_1, \dots, x_{i-1}, 1, x_{i+1}, \dots, x_n) = 1$ for all $i \in [n]$.

3. **Anonymity.** *There exists a subgroup $G \leq S_n$ that acts transitively on $[n]$ such that $f(x_{\sigma(1)}, \dots, x_{\sigma(n)}) = f(x_1, \dots, x_n)$ for all $x \in \{-1, +1\}^n$ and $\sigma \in G$.*

This claim is straightforward to verify, with anonymity a consequence of the fact that the graph is transitive.

Influences, Russo's formula, the KKL theorem and Talagrand's theorem. To prove Theorem 2.13 we use *Russo's formula*, a classical result in probability that we prove below.

Let X_1, \dots, X_n be random variables taking values in $\{-1, +1\}$. For $-\frac{1}{2} < \delta < \frac{1}{2}$, let \mathbb{P}_δ be the distribution such that $\mathbb{P}_\delta[X_i = +1] = \frac{1}{2} + \delta$ independently. Let $g : \{-1, +1\}^n \rightarrow \{-1, +1\}$ be a monotone function (as defined above in Claim 2.14). Let $Y = g(X)$, where $X = (X_1, \dots, X_n)$.

Denote by $\tau_i : \{-1, +1\}^n \rightarrow \{-1, +1\}^n$ the function given by $\tau_i(x_1, \dots, x_n) = (x_1, \dots, x_{i-1}, -x_i, x_{i+1}, \dots, x_n)$. We define the *influence* I_i^δ of $i \in [n]$ on Y as the probability that i is *pivotal*:

$$I_i^\delta = \mathbb{P}_\delta[g(\tau_i(X)) \neq g(X)].$$

That is I_i^δ is the probability that the value of $Y = g(X)$ changes, if we change X_i .

Theorem 2.15 (Russo's formula).

$$\frac{d\mathbb{P}_\delta[Y = +1]}{d\delta} = \sum_i I_i^\delta,$$

Proof. Let $\mathbb{P}_{\delta_1, \dots, \delta_n}$ be the distribution on X such that

$$\mathbb{P}_{\delta_1, \dots, \delta_n}[X_i = +1] = \delta_i.$$

We prove the claim by showing that

$$\frac{\partial \mathbb{P}_{\delta_1, \dots, \delta_n}[Y = +1]}{\partial \delta_i} = \mathbb{P}_{\delta_1, \dots, \delta_n}[g(\tau_i(X)) \neq g(X)],$$

and noting that $\mathbb{P}_{\delta, \dots, \delta} = \mathbb{P}_\delta$, and that for general differentiable $h : \mathbb{R}^n \rightarrow \mathbb{R}$ it holds that

$$\frac{\partial h(y, \dots, y)}{\partial y} = \sum_i \frac{\partial h(x_1, \dots, x_n)}{\partial x_i}(y).$$

Indeed, if we denote $\mathbb{E} = \mathbb{E}_{\delta_1, \dots, \delta_n}$ and $\mathbb{P} = \mathbb{P}_{\delta_1, \dots, \delta_n}$, then

$$\frac{\partial}{\partial \delta_i} \mathbb{P}[Y = +1] = \frac{\partial}{\partial \delta_i} \frac{1}{2} \mathbb{E}[g(X)].$$

Given $x = (x_1, \dots, x_n) \in \mathbb{R}^n$, denote $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$. Then

$$\mathbb{E}[g(X)] = \sum_x \mathbb{P}[X_{-i} = x_{-i}, X_i = x_i] g(x)$$

$$= \sum_x \mathbb{P}[X_{-i} = x_{-i}] \mathbb{P}[X_i = x_i] g(x),$$

where the second equality follows from the independence of the X_i 's. Hence

$$\begin{aligned} \frac{\partial}{\partial \delta_i} \mathbb{P}_{\delta_1, \dots, \delta_n} [Y = +1] &= \frac{\partial}{\partial \delta_i} \frac{1}{2} \sum_x \mathbb{P}[X_{-i} = x_{-i}] \mathbb{P}[X_i = x_i] g(x) \\ &= \frac{1}{2} \sum_x \mathbb{P}[X_{-i} = x_{-i}] x_i g(x), \end{aligned}$$

where the second equality follows from the fact that $\mathbb{P}[X = +1] = \delta_i$ and $\mathbb{P}[X = -1] = 1 - \delta_i$. Now, $\sum_{x_i} x_i g(x)$ is equal to zero when $g(\tau_i(x)) = g(x)$, and to two otherwise, since g is monotone. Hence

$$\begin{aligned} \frac{\partial}{\partial \delta_i} \mathbb{P}_{\delta_1, \dots, \delta_n} [Y = +1] &= \sum_x \mathbb{P}[X_{-i} = x_{-i}] \mathbb{1}(g(\tau_i(x)) \neq g(x)) \\ &= \mathbb{P}[g(\tau_i(X)) \neq g(X)]. \quad \square \end{aligned}$$

Kahn, Kalai and Linial [25] prove a deep result on Boolean functions on the hypercube (i.e., functions from $\{-1, +1\}^n$ to $\{-1, +1\}$), which was later generalized by Talagrand [43]. Their theorem states that there must exist an i with influence at least $O(\log n/n)$.

Theorem 2.16 (Talagrand). *Let $\epsilon_\delta = \max_i I_i^\delta$ and $q_\delta = \mathbb{P}_\delta [Y = 1]$. Then*

$$\sum_i I_i^\delta \geq K \log(1/\epsilon_\delta) q_\delta (1 - q_\delta).$$

for some universal constant K .

Using this result, the proof of Theorem 2.13 is straightforward, and we leave it as an exercise to the reader.

3. Bayesian models

In this chapter we study *Bayesian agents*. We call an agent Bayesian when its actions maximize the expectation of some utility function. This is a model which comes from Economics, where, in fact, its use is the default paradigm. We will focus on the case in which an agent's utility depends only on the state of the world S and on its actions, and is the same for all agents and all time periods.

3.0.5. Toy model: Continuous actions

Before defining general Bayesian models, we consider the following simple model on an undirected connected graph. Let $S \in \{0, 1\}$ be a binary state of the world, and let the private signals be i.i.d. conditioned on S .

We denote by H_t^i the information available to agent i at time t . This includes its private signal, and the actions of its neighbors in the previous time periods:

$$H_t^i = \left\{ W_i, A_{t'}^j : j \in \partial i, t' < t \right\}. \quad (3.1)$$

The actions are given by

$$A_t^i = \mathbb{P} [S = 1 | H_t^i]. \quad (3.2)$$

That is, each agent's action is its *belief*, or the probability that it assigns to the event $S = 1$, given what it knows.

For this model we prove the following results:

- **Convergence.** The actions of each agent converge almost surely to some A_∞^i . This is a direct consequence of the observation that $\{\sigma(H_t^i)\}_{t \in \mathbb{N}}$ is a filtration, and so $\{A_t^i\}_{t \in \mathbb{N}}$ is a bounded martingale. Note that this does not use the independence of the signals.
- **Agreement.** The limit actions A_∞^i are almost surely the same for all $i \in V$. This follows from the fact that if i and j are connected then $A_\infty^i + A_\infty^j \in H_\infty^i \cap H_\infty^j$ and if A_∞^i and A_∞^j are not a.s. equal then:

$$\mathbb{E} \left[\left(\frac{1}{2}(A_\infty^i + A_\infty^j) - S \right)^2 \right] < \max \left(\mathbb{E} [(A_\infty^i - S)^2], \mathbb{E} [(A_\infty^j - S)^2] \right).$$

Note again that this argument does not use the independence of the signals. We will show this in further generality in Section 3.2 below. This is a consequence of a more general agreement theorem that applies to all Bayesian models, which we prove in Section 3.1.

- **Learning.** When $|V| = n$, we show in Section 3.4 that $A_\infty^i = \mathbb{P} [S = 1 | W_1, \dots, W_n]$. This is the strongest possible learning result; the agents' actions are the same as they would be if each agent knew all the others' private signals. In particular, it follows that $\mathbb{P} [\text{round}(A_\infty^i) \neq S]$ is exponentially small in n . This result crucially relies on the independence of the signals as the following example shows.

Example 3.1. Consider two agents 1, 2 with $W_i = 0$ or 1 with probability 1/2 each and independently, and $S = W_1 + W_2 \pmod{2}$. Note that here $A_t^i = 1/2$ for $i = 1, 2$ and all t , while it is trivial to recover S from W_1, W_2 .

3.0.6. Definitions and some observations

Following our general framework (see Section 1.5) we shall (mostly) consider a state of the world $S \in \{0, 1\}$ chosen from the uniform distribution, with conditionally i.i.d. private signals. We will consider both discrete and continuous actions, and each shall correspond to a different *utility function*. A utility function will simply be a continuous map $u : \{0, 1\} \times [0, 1] \rightarrow [0, 1]$. The quantity $u(S, a)$ represents what the agent gains when choosing action a when the state is

S . In a sense that we will soon define formally, agents will be *utility maximizers*: they will choose their actions so as to maximize their utilities.

More precisely, we shall denote by $U_t^i = u(S, A_t^i)$ agent i 's utility at time t , and study *myopic agents*, or agents who strive to maximize, at each period t , the expectation of U_t^i .

As in the toy model above, we denote by H_t^i the information available to agent i at time t , including its private signal, and the actions of its neighbors in the previous time periods:

$$H_t^i = \left\{ W_i, A_{t'}^j : j \in \partial i, t' < t \right\}. \quad (3.3)$$

Given a utility function $U_t^i = u(S, A_t^i)$, a Bayesian agent will choose

$$A_t^i = \operatorname{argmax}_s \mathbb{E} [u(S, s) | H_t^i]. \quad (3.4)$$

Equivalently, one can define A_t^i as a random variable which, out of all $\sigma(H_t^i)$ -measurable random variables, maximizes the expected utility:

$$A_t^i = \operatorname{argmax}_{A \in \sigma(H_t^i)} \mathbb{E} [u(S, A)]. \quad (3.5)$$

We assume that in cases of indifference (i.e., two actions that maximize the expected utility) the agents chooses one according to some known deterministic rule.

We consider two utility functions; a discrete one that results in discrete actions, and a continuous one that results in continuous actions. The first utility function is

$$U_t^i = \mathbb{1}(A_t^i = S). \quad (3.6)$$

Although this function is not continuous as a function from $[0, 1]$ to $[0, 1]$, we will, in this case, consider the set of allowed actions to be $\{0, 1\}$, and so $u : \{0, 1\} \times \{0, 1\} \rightarrow \mathbb{R}$ will be continuous again.

To maximize the expectation of U_t^i conditioned on H_t^i , a myopic agent will choose the action

$$A_t^i = \operatorname{argmax}_{s \in \{0, 1\}} \mathbb{P} [S = s | H_t^i], \quad (3.7)$$

which will take values in $\{0, 1\}$.

We will also consider the following utility function, which corresponds to continuous actions:

$$U_t^i = 1 - (A_t^i - S)^2. \quad (3.8)$$

To maximize the expectation of this function, an agent will choose the action

$$A_t^i = \mathbb{P} [S = 1 | H_t^i]. \quad (3.9)$$

This action will take values in $[0, 1]$.

An important concept in the context of Bayesian agents is that of *belief*. We define agent i 's belief at time t to be

$$B_t^i = \mathbb{P}[S = 1 | H_t^i]. \quad (3.10)$$

This is the probability that $S = 1$, conditioned on all the information available to i at time t . It is easy to check that, in the discrete action case, the action is the rounding of the belief. In the continuous action case the action equals the belief.

An important distinction is between *bounded* and *unbounded* private signals [42]. We say that the private signal W_i is **bounded** when there exists an $\epsilon > 0$ such the private belief $B_0^i = \mathbb{P}[S = 1 | W_i]$ is supported on $[\epsilon, 1 - \epsilon]$. We will say that it is unbounded when the private belief $B_0^i = \mathbb{P}[S = 1 | W_i]$ can be arbitrarily close to both 1 and 0; formally, when the convex closure of the support of B_0^i is equal to $[0, 1]$.

Unbounded private signals can be thought of as being “unboundedly strong”, and therefore could be expected to promote learning. This is indeed the case, as we show below.

The following claim follows directly from the fact that the sequence of sigma-algebras $\sigma(H_t^i)$ is a filtration.

Claim 3.2. *The sequence of beliefs of agent i , $\{B_t^i\}_{t \in \mathbb{N}}$, is a bounded martingale.*

It follows that a limiting belief almost surely exists, and we can define

$$B_\infty^i = \lim_{t \rightarrow \infty} B_t^i. \quad (3.11)$$

Furthermore, if we let $H_\infty^i = \cup_t H_t^i$, then

$$B_\infty^i = \mathbb{P}[S = 1 | H_\infty^i]. \quad (3.12)$$

We would like to also define the limiting *action* of agent i . However, it might be the case that the actions of an agent do not converge. We therefore define A_t^i to be an *action set*, given by the set of accumulation points of the sequence A_t^i . In the case that A_∞^i is a singleton $\{x\}$, we denote $A_\infty^i = x$, in a slight abuse of notation. Note that in the case that actions take values in $\{0, 1\}$ (as we will consider below), A_∞^i is either equal to 1, to 0, or to $\{0, 1\}$.

The following claim is straightforward.

Claim 3.3. *Fix a continuous utility function u . Then*

$$\lim_t \mathbb{E}[u(S, A_t^i) | H_t^i] = \mathbb{E}[u(S, a) | H_\infty^i] \geq \mathbb{E}[u(S, b) | H_\infty^i]$$

for all $a \in A_\infty^i$ and all b .

That is, any action in A_∞^i is optimal (that is, maximizes the expected utility), given what the agent knows at the limit $t \rightarrow \infty$. It follows that

$$\mathbb{E}[u(S, a) | H_\infty^i] = \mathbb{E}[u(S, b) | H_\infty^i]$$

for all $a, b \in A_\infty^i$. It also follows that in the case of actions in $\{0, 1\}$, $A_\infty^i = \{0, 1\}$ only if i is asymptotically indifferent, or expects the same utility from both 0 and 1.

We will show that an oft-occurring phenomenon in the Bayesian setting is agreement on limit actions, so that A_∞^i is indeed a singleton, and $A_\infty^i = A_\infty^j$ for all $i, j \in V$. In this case we can define A_∞ as the common limit action.

3.1. Agreement

In this section we show that regardless of the utility function, and, in fact, regardless of the private signal structure, Bayesian agents always reach agreement, except in cases of indifference. This theorem originated in the work of Aumann [2], with contributions by Geanakoplos and others [19, 41]. It first appeared as below in Gale and Kariv [17]. Rosenberg, Solan and Vieille [38] correct an error in the proof and extend this result to the even more general setting of *strategic agents* (which we will not discuss), as is done in [34].

Theorem 3.4 (Gale and Kariv). *Fix a utility function $U_t^i = u(S, A_t^i)$, and consider $(i, j) \in E$. Then*

$$\mathbb{E}[u(S, a^i) | H_\infty^i] = \mathbb{E}[u(S, a^j) | H_\infty^i]$$

for any $a^i \in A_\infty^i$ and $a^j \in A_\infty^j$.

That is, any action in A_∞^j is optimal, given what i knows, and so has the same expected utility as any action in A_∞^i . Note that this theorem applies even when private signals are not conditionally i.i.d., and when S is not necessarily binary.

Note that (3.5) is a particularly useful way to think of the agents' actions, as the proof of the following claim shows.

Claim 3.5. *For all $(i, j) \in E$ it holds that*

1. $\mathbb{E}[U_{t+1}^i] \geq \mathbb{E}[U_t^i]$.
2. $\mathbb{E}[U_{t+1}^i] \geq \mathbb{E}[U_t^j]$.

Proof. 1. Since $\sigma(H_t^i)$ is included in $\sigma(H_{t+1}^i)$, the maximum in (3.5) is taken over a larger space for A_{t+1}^i than it is for A_t^i , and therefore a value at least as high is achieved.

2. Since A_t^j is $\sigma(H_{t+1}^i)$ -measurable, it follows from (3.5) that $\mathbb{E}[u(S, A_{t+1}^i)] \geq \mathbb{E}[u(S, A_t^j)]$. \square

Exercise. Prove the following corollary.

Corollary 3.6. *For all $i, j \in V$,*

$$\lim_t \mathbb{E}[U_t^i] = \lim_t \mathbb{E}[U_t^j]$$

Exercise. Prove Theorem 3.4 using Corollary 3.6 and Claim 3.3.

3.2. Continuous utility models

As mentioned above, in the case that the utility function is

$$U_t^i = 1 - (A_t^i - S)^2,$$

it follows readily that

$$A_t^i = B_t^i = \mathbb{P}[S = 1 | H_t^i],$$

and so, by Claim 3.2, the actions of each agent form a martingale, and furthermore each converge to a singleton A_∞^i . Aumann's celebrated *Agreement Theorem* from the paper titled "Agreeing to Disagree" [2], as followed-up by Geanakoplos and Polemarchakis in the paper titled "We can't disagree forever" [19], implies that all these limiting actions are equal. This follows from Theorem 3.4.

Theorem 3.7. *In the continuous utility model*

$$A_\infty^i = \mathbb{P}[S = 1 | H_\infty^i]$$

and furthermore

$$A_\infty^i = A_\infty^j$$

for all $i, j \in V$.

Note again that this holds also for private signals that are not conditionally i.i.d.

Proof. As was mentioned above, since the actions A_t^i are equal to the beliefs B_t^i , they are a bounded martingale and therefore converge. Hence $A_\infty^i = B_\infty^i$ and, by (3.12),

$$A_\infty^i = \mathbb{P}[S = 1 | H_\infty^i].$$

Assume $(i, j) \in E$. By Theorem 3.4 we have that

$$\mathbb{E}[u(S, A_\infty^i) | H_\infty^i] = \mathbb{E}[u(S, A_\infty^j) | H_\infty^i].$$

It hence follows from Claim 3.3 that both A_∞^i and A_∞^j maximize $\mathbb{E}[u(S, \cdot) | H_\infty^i]$. But the unique maximizer is $\mathbb{P}[S = 1 | H_\infty^i]$, and so $A_\infty^i = A_\infty^j$. For general i and j , the claim now follows from the fact that the graph is connected. \square

3.3. Bounds on number of rounds in finite probability spaces

In this section we consider the case of a **finite** probability space. Let S be binary, and let the private signals $\bar{W} = (W_1, \dots, W_{|V|})$ be chosen from an arbitrary (not necessarily conditionally independent) distribution over a finite joint probability space of size M . Consider general utility functions $U_t^i = u(S, A_t^i)$.

The following theorem is a strengthening of a theorem by Geanakoplos [18], using ideas from [30].

Theorem 3.8 (Geanakoplos). *Let d be the diameter of the graph G . Then the actions of each agent converge after at most $M \cdot |V|$ time periods:*

$$A_t^i = A_{t'}^i,$$

for all $i \in V$ and all $t, t' \geq M \cdot |V|$. Furthermore, the number of time periods t such that $A_{t+1}^i \neq A_t^i$ is at most M .

The key observation is that each sigma-algebra $\sigma(H_t^i)$ is generated by some subset of the set of random variables $\{\mathbb{1}(\overline{W} = m)\}_{m \in \{1, \dots, M\}}$.

Proof. By (3.5), if $\sigma(H_t^i) = \sigma(H_{t'}^i)$ then $A_t^i = A_{t'}^i$. It remains to show, then, that $\sigma(H_t^i) = \sigma(H_{t'}^i)$ for all $t, t' \geq M \cdot |V|$, and that $\sigma(H_t^i) \neq \sigma(H_{t+1}^i)$ at most M times.

Now, every sub-sigma-algebra of $\sigma(\overline{W})$ (such as $\sigma(H_t^i)$) is simply a partition of the finite space $\{1, \dots, M\}$. Furthermore, for every i , the sequence $\sigma(H_t^i)$ is a filtration, so that each $\sigma(H_{t+1}^i)$ is a refinement of $\sigma(H_t^i)$. A simple combinatorial argument shows that any such sequence has at most M unique partitions, and so $\sigma(H_t^i) \neq \sigma(H_{t+1}^i)$ at most M times.

Finally, note that if $\sigma(H_t^i) = \sigma(H_{t+1}^i)$ for all $i \in V$ at some time t , then this is also the case for all later time periods. Hence, as long as the process hasn't ended, it must be that $\sigma(H_t^i) \neq \sigma(H_{t+1}^i)$ for some agent i . It follows that the process ends after at most $M \cdot |V|$ time periods. \square

3.4. From agreement to learning

This section is adapted from Mossel, Sly and Tamuz [32].

In this section we prove two very general results that relate agreement and learning in Bayesian models. As in our general framework, we consider a binary state of the world $S \in \{0, 1\}$ chosen from the uniform distribution, with conditionally i.i.d. private signals. We do not define actions, but only study what can be said when, at the end of the process (whatever it may be) the agents reach agreement.

Formally, consider a finite set of agents of size n , or a countably infinite set of agents, each with a private signal W_i . Let \mathcal{F}_i be the sigma-algebra that represents what is known by agent i . We require that W_i is \mathcal{F}_i measurable (i.e., each agent knows its own private signal), and that each \mathcal{F}_i is a sub-sigma-algebra of $\sigma(W_1, W_2, \dots)$. Let agent i 's belief be

$$B_i = \mathbb{P}[S = 1 | \mathcal{F}_i],$$

and let agent i 's action be

$$A_i = \operatorname{argmax}_{s \in \{0, 1\}} \mathbb{P}[S = s | \mathcal{F}_i].$$

We let $A_i = \{0, 1\}$ when both maximize $\mathbb{P}[S = s | \mathcal{F}_i]$.

We say that agents agree on beliefs when there exists a random variable B such that almost surely $B_i = B$ for all agents i . Likewise, we say that agents

agree on actions when there exists a random variable A such that almost surely $A_i = A$ for all agents i . Such agreement arises often as a result of repeated interaction of Bayesian agents.

We show below that agreement on beliefs is a sufficient condition for learning, and in fact implies the strongest possible type of learning. We also show that when private signals are unbounded beliefs then agreement on actions is also a condition for learning.

3.4.1. Agreement on beliefs

The following theorem and its proof is taken from [32]. This theorem also admits a proof as a corollary of some well known results on *rational expectation equilibria* (see, e.g., [14, 36]), but we will not delve into this topic.

Theorem 3.9. *Let the private signals (W_1, \dots, W_n) be independent conditioned on S , and let the agents agree on beliefs. Then*

$$B = \mathbb{P}[S = 1 | W_1, \dots, W_n].$$

That is, if the agents have exchanged enough information to agree on beliefs, they have exchanged *all the relevant information*, in the sense that they have the same belief that they would have had they shared *all the information*.

Proof. Denote agent i 's private log-likelihood ratio by

$$Z_i = \log \frac{d\mu_1^i}{d\mu_0^i}(W_i).$$

Since $\mathbb{P}[S = 1] = \mathbb{P}[S = 0] = 1/2$ it follows that

$$Z_i = \log \frac{\mathbb{P}[S = 1 | W_i]}{\mathbb{P}[S = 0 | W_i]}.$$

Denote $Z = \sum_{i \in [n]} Z_i$. Then, since the private signals are conditionally independent, it follows by Bayes' rule that

$$\mathbb{P}[S = 1 | W_1, \dots, W_n] = \text{logit}(Z), \quad (3.13)$$

where $\text{logit}(z) = e^z / (e^z + e^{-z})$.

Since

$$B = \mathbb{P}[S = 1 | B] = \mathbb{E}[\mathbb{P}[S = 1 | B, W_1, \dots, W_n] | B]$$

then

$$B = \mathbb{E}[\text{logit}(Z) | B], \quad (3.14)$$

since, given the private signals (W_1, \dots, W_n) , further conditioning on B (which is a function of the private signals) does not change the probability of the event $S = 1$.

Our goal is to show that $B = \mathbb{P}[S = 1|W_1, \dots, W_n]$. We will do this by showing that conditioned on B , Z and $\text{logit}(Z)$ are uncorrelated. It will follow that conditioned on B , Z is constant, so that $Z = Z(B)$ and

$$B = \mathbb{P}[S = 1|B] = \mathbb{P}[S = 1|Z(B)] = \mathbb{P}[S = 1|W_1, \dots, W_n].$$

By the law of total expectation we have that

$$\mathbb{E}[Z_i \cdot \text{logit}(Z)|B] = \mathbb{E}[\mathbb{E}[Z_i \text{logit}(Z)|B, Z_i]|B].$$

Note that $\mathbb{E}[Z_i \cdot \text{logit}(Z)|B, Z_i] = Z_i \mathbb{E}[\text{logit}(Z)|B, Z_i]$ and so we can write

$$\mathbb{E}[Z_i \cdot \text{logit}(Z)|B] = \mathbb{E}[Z_i \mathbb{E}[\text{logit}(Z)|B, Z_i]|B].$$

Since Z_i is \mathcal{F}_i measurable, and since, by (3.14), $B = \mathbb{E}[\text{logit}(Z)|\mathcal{F}_i] = \mathbb{E}[\text{logit}(Z)|B]$, then $B = \mathbb{E}[\text{logit}(Z)|B, Z_i]$ and so it follows that

$$\mathbb{E}[Z_i \cdot \text{logit}(Z)|B] = \mathbb{E}[Z_i B|B] = B \cdot \mathbb{E}[Z_i|B] = \mathbb{E}[\text{logit}(Z)|B] \cdot \mathbb{E}[Z_i|B]. \quad (3.15)$$

where the last equality is another substitution of (3.14). Summing this equation (3.15) over $i \in [n]$ we get that

$$\mathbb{E}[Z \cdot \text{logit}(Z)|B] = \mathbb{E}[\text{logit}(Z)|B] \mathbb{E}[Z|B]. \quad (3.16)$$

Now, since $\text{logit}(Z)$ is a monotone function of Z , by Chebyshev's sum inequality we have that

$$\mathbb{E}[Z \cdot \text{logit}(Z)|B] \geq \mathbb{E}[\text{logit}(Z)|B] \mathbb{E}[Z|B] \quad (3.17)$$

with equality only if Z (or, equivalently $\text{logit}(Z)$) is constant. Hence Z is constant conditioned on B and the proof is concluded. \square

3.4.2. Agreement on actions

In this section we consider the case that the agents agree on actions, rather than beliefs. The boundedness of private beliefs plays an important role in the case of agreement on actions. When private beliefs are bounded then agreement on actions does not imply learning, as shown by the following example, which is reminiscent of Bala and Goyal's [3] *royal family*. However, when private beliefs are unbounded then learning does occur with high probability, as we show below.

Example 3.10. *Let there be $n > 100$ agents, and call the first hundred "the Senate". The private signals are bits that are independently equal to S with probability $2/3$. Let*

$$A_S = \text{argmax}_a \mathbb{P}[S = a|W_1, \dots, W_{100}],$$

and let $\mathcal{F}_i = \sigma(W_i, A_S)$.

This example describes the case in which the information available to each agent is the decision of the senate - which aggregates the senators' private information optimally - and its own private signal. It is easy to convince oneself that $A_i = A_S$ for all $i \in [n]$, and so actions are indeed agreed upon. However, the probability that $A_S \neq S$ - i.e., the Senate makes a mistake - is constant and does not depend on the number of agents n . Hence the probability that the agents choose the wrong action does *not* tend to zero as n tends to infinity. This cannot be the case when private beliefs are unbounded, as Mossel, Sly and Tamuz [32] show.

Theorem 3.11 (Mossel, Sly and Tamuz). *Let the private signals (W_1, \dots, W_n) be i.i.d. conditioned on S , and have unbounded beliefs. Let the agents agree on actions. Then there exists a sequence $q(n) = q(n, \mu_0, \mu_1)$, depending only on the conditional private signal distributions μ_1 and μ_0 , such that $q(n) \rightarrow 1$ as $n \rightarrow \infty$, and*

$$\mathbb{P}[A = S] \geq q(n).$$

In particular,

$$q(n) \leq \min_{\epsilon > 0} \max \left\{ \frac{2\epsilon}{1-\epsilon}, \frac{4}{n\mathbb{P}[B_i < \epsilon | S = 0]} \right\}.$$

For the case of a countably infinite set of agent, we prove (using an essentially identical technique) the following similar statement.

Theorem 3.12. *Identify the set of agents with \mathbb{N} , let the private signals (W_1, W_2, \dots) be i.i.d. conditioned on S , and have unbounded beliefs. Let all but a vanishing fraction of the agents agree on actions. That is, let there exist a random variable A such that almost surely*

$$\limsup_n \frac{1}{n} |\{i \in \mathbb{N} : A_i \neq A\}| = 0.$$

Then $\mathbb{P}[A = S] = 1$.

Recall that B_0^i denoted the probability of $S = 1$ given agent i 's private signal:

$$B_0^i = \mathbb{P}[S = 1 | W_i].$$

The condition of unbounded beliefs can be equivalently formulated to be that for any $\epsilon > 0$ it holds that $\mathbb{P}[B_0^i < \epsilon] > 0$ and $\mathbb{P}[B_0^i > 1 - \epsilon] > 0$.

We shall need two standard lemmas to prove this theorem.

Lemma 3.13. $\mathbb{P}[S = 0 | B_0^i < \epsilon] > 1 - \epsilon$.

Proof. Since B_0^i is a function of W_i then

$$\mathbb{P}[S = 1 | B_0^i = b_i] = \mathbb{E}[\mathbb{P}[S = 1 | W_i] | B_0^i(W_i) = b_i] = \mathbb{E}[B_0^i | B_0^i = b_i] = b_i,$$

and so $\mathbb{P}[S = 1 | B_0^i] = B_0^i$. It follows that $\mathbb{P}[S = 0 | B_0^i] = 1 - B_0^i$, and so $\mathbb{P}[S = 0 | B_0^i < \epsilon] > 1 - \epsilon$. \square

Lemma 3.14 below is a version of Chebyshev's inequality, quantifying the idea that the expectation of a random variable Z , conditioned on some event A , cannot be much lower than its unconditional expectation when A has high probability.

Exercise. Prove the following lemma.

Lemma 3.14. *Let Z be a real random variable with finite variance, and let A be an event. Then*

$$\mathbb{E}[Z] - \sqrt{\frac{\text{Var}[Z]}{\mathbb{P}[A]}} \leq \mathbb{E}[Z|A] \leq \mathbb{E}[Z] + \sqrt{\frac{\text{Var}[Z]}{\mathbb{P}[A]}}$$

We are now ready to prove Theorem 3.12.

Proof of Theorem 3.12. Consider a set of agents \mathbb{N} who agree (except for a vanishing fraction) on the action. Assume by contradiction that $q = \mathbb{P}[A \neq 0 | S = 0] > 0$.

Recall that $B_i = \mathbb{P}[S = 1 | \mathcal{F}_i]$. Since $\mathbb{P}[S = 1 | B_0^i] = B_0^i$,

$$\mathbb{E}[B_i | B_0^i] = \mathbb{E}[\mathbb{P}[S = 1 | \mathcal{F}_i] | B_0^i] = \mathbb{P}[S = 1 | B_0^i] = B_0^i.$$

Applying Markov's inequality to B_i we have that $\mathbb{P}[B_i \geq \frac{1}{2} | B_0^i < \epsilon] < 2\epsilon$, and in particular

$$\mathbb{P}[A_i \neq 0, S = 0 | B_0^i < \epsilon] = \mathbb{P}[B_i \geq \frac{1}{2}, S = 0 | B_0^i < \epsilon] < 2\epsilon$$

so

$$\mathbb{P}[A_i \neq 0, S = 0, B_0^i < \epsilon] \leq 2\epsilon \mathbb{P}[B_0^i < \epsilon]$$

Denote

$$K(n) = \frac{1}{n} \sum_{i \in [n]} \mathbb{1}(B_0^i < \epsilon) = \frac{1}{n} \sum_{i \in [n]} \mathbb{1}(B_0^i < \epsilon, A_i = 0) + \frac{1}{n} \sum_{i \in [n]} \mathbb{1}(B_0^i < \epsilon, A_i \neq 0) \quad (3.18)$$

Let $K_1(n)$ denote the first sum and $K_2(n)$ denote the second sum. From our assumption that a vanishing fraction of agents disagree it follows that a.s.

$$\begin{aligned} & \limsup \mathbb{E}[K_1(n) | A \neq 0, S = 0] \\ & \leq \frac{1}{q} \limsup \mathbb{E}[K_1(n) | A \neq 0] \\ & \leq \frac{1}{q} \limsup \mathbb{E} \left[\frac{1}{n} \sum_{i \in [n]} \mathbb{1}(A_i = 0) \middle| A \neq 0 \right] = 0. \end{aligned}$$

It also follows that for all n

$$\mathbb{E}[K_2(n) | A \neq 0, S = 0] \leq \frac{1}{q} \mathbb{E}[K_2(n), A \neq 0, S = 0] \leq \frac{2\epsilon \mathbb{P}[B_0^i < \epsilon]}{q}.$$

Thus

$$\limsup_n \mathbb{E}[K(n)|A \neq 0, S = 0] \leq \frac{2\epsilon \mathbb{P}[B_0^i < \epsilon]}{q}.$$

We hence bound $\mathbb{E}[K|A \neq 0, S = 0]$ from above. We will now bound it from below to obtain a contradiction.

Applying lemma 3.14 to K and the event $\{A \neq 0\}$ (under the conditional measure $S = 0$) yields that

$$\mathbb{E}[K(n)|A \neq 0, S = 0] \geq \mathbb{E}[K(n)|S = 0] - \sqrt{\frac{\text{Var}[K(n)|S = 0]}{q}}.$$

Since the agents' private signals (and hence their private beliefs) are independent conditioned on $S = 0$, K (conditioned on S) is the average of n i.i.d. variables. Hence $\text{Var}[K(n)|S = 0] = n^{-1} \text{Var}[\mathbb{1}(B_0^i < \epsilon)|S = 0]$ and $\mathbb{E}[K(n)|S = 0] = \mathbb{P}[B_0^i < \epsilon|S = 0]$. Thus we have that

$$\mathbb{E}[K(n)|A \neq 0, S = 0] \geq \mathbb{P}[B_0^i < \epsilon|S = 0] - n^{-1/2} \sqrt{\frac{\text{Var}[\mathbb{1}(B_0^i < \epsilon)|S = 0]}{q}}. \quad (3.19)$$

and so

$$\liminf_n \mathbb{E}[K(n)|A_i \neq 0, S = 0] \geq \mathbb{P}[B_0^i < \epsilon|S = 0]$$

Joining the lower bound with the upper bound we obtain that

$$\mathbb{P}[B_0^i < \epsilon|S = 0] \leq \frac{2\epsilon \mathbb{P}[B_0^i < \epsilon]}{q},$$

and applying Bayes rule we obtain

$$q < \frac{\epsilon}{\mathbb{P}[S = 0|B_0^i < \epsilon]}.$$

Since by Lemma 3.13 above we know that $\mathbb{P}[S = 0|B_0^i < \epsilon] > 1 - \epsilon$, then

$$q < \frac{\epsilon}{1 - \epsilon}.$$

Since this holds for all ϵ , we have shown that $q = 0$, which is a contradiction. \square

3.5. Sequential models

In this section we consider a classical class of learning models called **sequential models**. We retain a binary state of the world S and conditionally i.i.d. private signals, but relax two assumption.

- We no longer assume that the graph G is strongly connected. In fact, we consider the particular case that the set of agents is countably infinite,

identify it with \mathbb{N} , and let $(i, j) \in E$ iff $j < i$. That is, the agents are ordered, and each agent observes the actions of its predecessors.

- We assume that each agent acts once, after observing the actions of its predecessors. That is, agent i acts only once, at time i .

In this section, we denote agent i 's (single) action by A_i . Hence agent i 's information when taking its action, which we denote by H_i , is

$$H_i = \{W_i, A_j : j < i\}.$$

We likewise denote agent i 's belief at time i by $B_i = \mathbb{P}[S = 1|H_i]$. We assume discrete utilities, so that

$$A_i = \operatorname{argmax}_{s \in \{0,1\}} \mathbb{P}[S = s|H_i],$$

and let $A_i = 1$ when $\mathbb{P}[S = 1|H_i] = 1/2$.

Since each agent acts only once, we explore a different notion of learning in this section. The question we consider is the following: when is it the case that $\lim_{i \rightarrow \infty} A_i = S$ with probability one? Since the graph is fixed, the answer to this question depends only on the private signal distributions μ_0 and μ_1 .

This model (in a slightly different form) was introduced independently by Bikhchandani, Hirshleifer and Welch [9], and Banerjee [4]. A significant later contribution is that of Smith and Sørensen [42].

An interesting phenomenon that arises in this model is that of an *information cascade*. An information cascade is said to occur if, given an agent i 's predecessor's actions, i 's action does not depend on its private signal. This happens if the previous agents' actions present such compelling evidence towards the event that (say) $S = 1$, that any realization of the private signal would not change this conclusion. Once this occurs - that is, once one agent's action does not depend on its private signal - then this will also hold for all the agents who act later.

3.5.1. The external observer at infinity

An important tool in the analysis of this model is the introduction of an external observer x that observes all the agents' actions but none of their private signals. We denote by $H_i^x = \{A_j : j < i\}$ the information available to x at time i , and denote by

$$B_i^x = \mathbb{P}[S = 1|H_i^x]$$

and

$$B_\infty^x = \lim_i B_i^x = \mathbb{P}[S = 1|H_\infty^x]$$

the beliefs of x at times t and infinity respectively, where, as before, $H_\infty^x = \cup_i H_i^x$. The same martingale argument used above can also be used here to show that the limit B_∞^x indeed exists and satisfies the equality above.

Exercise. Show that the likelihood ratio

$$L_i^x = \frac{1 - B_i^x}{B_i^x}$$

is also a martingale, *conditioned on* $S = 1$.

The martingale $\{B_i^x\}$ converges almost surely to B_∞^x in $[0, 1]$, and conditioned on $S = 1$, B_∞^x has support $\subseteq (0, 1]$. The reason that $B_\infty^x \neq 0$ when conditioning on $S = 1$, is the fact that $\mathbb{P}[S = 1|B_\infty^x] = B_\infty^x$, and so $\mathbb{P}[S = 1|B_\infty^x = 0] = 0$.

We also define actions for x , given by

$$A_i^x = \operatorname{argmax}_{s \in \{0,1\}} \mathbb{P}[S = s|H_i^x] = \operatorname{round}(B_i^x).$$

We again assume that in cases of indifference, the action 1 is chosen.

Claim 3.15. $A_{i+1}^x = A_i$

That is, the external observer simply copies, at time $t + 1$, the action of agent t . This follows immediately from the fact that A_i is $\sigma(H_i)$ -measurable, and so $H_{i+1}^x \subseteq H_i$. It follows that $\lim_i A_i = \lim_i A_i^x$, and so we have learning - in the sense we defined above for this section by $\lim_i A_i = S$ - iff the external observer learns in the usual sense of $\lim_i A_i^x = S$.

3.5.2. The agents' calculation

We write out each agent's calculation of its belief B_i , from which follows its action A_i . This is more easily done by calculating the likelihood ratio

$$L_i = \frac{1 - B_i}{B_i}.$$

By Bayes' law, since $\mathbb{P}[S = 1] = \mathbb{P}[S = 0] = \frac{1}{2}$, and since $H_i = (H_i^x, W_i)$

$$L_i = \frac{\mathbb{P}[S = 0|H_i]}{\mathbb{P}[S = 1|H_i]} = \frac{\mathbb{P}[H_i|S = 0]}{\mathbb{P}[H_i|S = 1]} = \frac{\mathbb{P}[H_i^x, W_i|S = 0]}{\mathbb{P}[H_i^x, W_i|S = 1]}.$$

Since the private signals are conditionally i.i.d., W_i is conditionally independent of H_i^x , and so

$$L_i = \frac{\mathbb{P}[H_i^x|S = 0]}{\mathbb{P}[H_i^x|S = 1]} \cdot \frac{\mathbb{P}[W_i|S = 0]}{\mathbb{P}[W_i|S = 1]}.$$

We denote by P_i the *private likelihood ratio* $\mathbb{P}[W_i|S = 0]/\mathbb{P}[W_i|S = 1]$, so that

$$L_i = L_i^x \cdot P_i. \quad (3.20)$$

3.5.3. The Markov chain and the martingale

Another useful observation is that $\{B_i^x\}_{i \in \mathbb{N}}$ is not only a martingale, but also a Markov chain. We denote this Markov chain on $[0, 1]$ by \mathcal{M} . To see this, note that conditioned on S , the private likelihood ratio P_i is independent of B_j^x , $j < i$, and so its distribution conditioned on $B_i^x = \mathbb{P}[S = 1|H_i^x]$ is the same as its distribution conditioned on (B_0^x, \dots, B_i^x) , which are $\sigma(H_i^x)$ -measurable.

3.5.4. Information cascades, convergence and learning

An information cascade is the event that, for some i , conditioned on H_i^x , A_i is independent of W_i . That is, an information cascade is the event that the observer at infinity knows, at time i , which action agent i is going to take, even though it only knows the actions of i 's predecessors and does not know i 's private signal. Equivalently, an information cascade occurs when A_i is $\sigma(H_i^x)$ -measurable. It is easy to see that it follows that A_j will also be $\sigma(H_i^x)$ -measurable, for all $j \geq i$.

Claim 3.16. *An information cascade is the event that B_i^x is a fixed point of \mathcal{M} .*

By ‘‘fixed point of \mathcal{M} ’’ we mean that a.s. $B_{i+1}^x = B_i^x$.

Proof of Claim 3.16. If A_i is $\sigma(H_i^x)$ measurable then $\sigma(H_i^x) = \sigma(H_i^x, A_i) = \sigma(H_{i+1}^x)$. It follows that

$$B_i^x = \mathbb{P}[S = 1 | H_i^x] = \mathbb{P}[S = 1 | H_{i+1}^x] = B_{i+1}^x.$$

Conversely, if $B_i^x = B_{i+1}^x$ w.p. one, then $A_i^x = A_{i+1}^x$ with probability one, and it follows that $A_i = A_{i+1}^x$ is $\sigma(H_i^x)$ -measurable. \square

Theorem 3.17. *The limit $\lim_i A_i$ exists almost surely.*

Proof. As noted above, $A_i = A_{i+1}^x$. Assume by contradiction that A_{i+1}^x takes both values infinitely often. Since $A_i^x = \mathbb{1}(B_i^x \geq \frac{1}{2})$, and since B_i^x converges to B_∞^x , it follows that $B_\infty^x = \frac{1}{2}$.

Note that by the Markov chain nature of $\{B_i^x\}$,

$$B_{i+1}^x = f(B_i^x, A_i) \tag{3.21}$$

for $f : [0, 1] \times \{0, 1\} \rightarrow [0, 1]$ independent of i and given by

$$f(b, a) = \mathbb{E}[B_i | B_i^x = b, A_i = a].$$

Since $A_i = \mathbb{1}(B_i \geq \frac{1}{2})$, it follows that $B_i = |B_i - \frac{1}{2}|(2A_i - 1) + \frac{1}{2}$, and so

$$f(b, a) = \mathbb{E}[|B_i - \frac{1}{2}| | B_i^x = b, A_i = a] (2a - 1) + 1/2.$$

Hence f is continuous at $(1/2, 1)$ and $(1/2, 0)$, even if $B_i = \frac{1}{2}$ with positive probability. It follows by taking the limit of (3.21) that if $\lim_i B_i^x = 1/2$ then $f(1/2, 1) = f(1/2, 0)$. But then B_i^x would equal $f(1/2, \cdot)$ for all i , since $B_0^x = 1/2$, and $A_i^x = 1$ for all i , which is a contradiction. \square

Since $\lim_i A_i$ exists almost surely we can define

$$A = \lim_i A_i.$$

Since $A_i \neq A$ for only a finite number of agents, we can directly apply Theorem 3.12 to arrive at the following result.

Theorem 3.18. *When private signals are unbounded then $A = S$ w.p. one.*

When private signals are bounded then information cascades occur with probability one, and A is no longer almost surely equal to S .

Theorem 3.19. *When private signals are bounded then $\mathbb{P}[A = S] < 1$.*

Proof. When private signals are bounded then the convex closure of the support of P_i is equal to $[\epsilon, M]$ for some $\epsilon, M > 0$. It follows then from (3.20) that if $L_i^x \leq 1/M$ then a.s. $L_i \leq 1$, and so $A_i = 1$. Likewise, if $L_i^x > 1/\epsilon$ then a.s. $A_i = 0$. Hence $[0, 1/M]$ and $(1/\epsilon, \infty)$ are all fixed points of \mathcal{M} .

Note that $\mathbb{P}[A_i^x = S | H_i^x] = \max\{B_i^x, 1 - B_i^x\}$. Hence we can prove the claim by showing that $B_\infty^x = \lim_i B_i^x$ is in $(0, 1)$, since then it would follow that $\lim_i \mathbb{P}[A_i^x = S] < 1$, and in particular $\mathbb{P}[\lim_i A_i = S] < 1$.

Indeed, condition on $S = 1$, and assume by contradiction that $\lim_i B_i^x = 1$. Then L_i^x will equal some $\delta \in (0, 1/M)$ for i large enough. But δ is a fixed point of \mathcal{M} , and so L_j^x will equal δ hence and B_i^x will not converge to one. The same argument applies if we condition on $S = 0$ and argue that L_i^x will equal some $N \in (1/\epsilon, \infty)$ for i large enough. \square

3.6. Learning from discrete actions on networks

This section is adapted from Mossel, Sly and Tamuz [33].

In this section we study asymptotic learning on general (undirected) social networks. We here choose to dive more deeply into the proofs - as compared to the previous sections of this survey - in order to showcase the various techniques needed to tackle this problem. These techniques include graph limits (Section 3.6.2), a notion of δ -independence (Section 3.6.4) and more. Indeed, the proof of the main result of this section, Theorem 3.21, does not (as far as we know) admit a short intuitive explanation, but rather requires the introduction and digestion of some abstract ideas, and in particular the topology on rooted graphs that we define and analyze in Section 3.6.2.

We study general social networks that are **undirected**, and consider both the finite and the countably infinite case. We consider agents who maximize, at each time t , the utility function (see (3.6))

$$U_t^i = \mathbb{1}(A_t^i = S).$$

Hence they choose actions using (3.7):

$$A_t^i = \operatorname{argmax}_{s \in \{0,1\}} \mathbb{P}[S = s | H_t^i].$$

We ask the following questions:

1. **Agreement.** Do the agents reach agreement? In this model we say that i and j agree if $A_\infty^i = A_\infty^j$. This happens under a weak condition on the private signals.

2. **Learning.** When the agents do agree on some limit action A_∞ , does this action equal S ? The answer to this question depends on the graph, and that for undirected graphs indeed $A_\infty = S$ with high probability (for large finite graphs) or with probability one (for infinite graphs).

The condition on private signals that implies agreement on limit actions is the following. By the definition of beliefs, $B_0^i = \mathbb{P}[S = 1|W_i]$. We say that the private signals *induce non-atomic beliefs* when the distribution of B_0^i is non-atomic. The rationale behind this definition is that it precludes the possibility of *indifference* or *ties*.

Theorem 3.20. *Let (μ_0, μ_1) induce non-atomic beliefs. Then there exists a random variable A_∞ such that almost surely $A_\infty^i = A_\infty$ for all i .*

We refer the reader to [33] for a proof of this Theorem. In Section 3.6.6 we give an example that shows that this claim indeed does not necessarily hold when private signals are atomic.

The following theorem states that when such agreement is guaranteed then the agents learn the state of the world with high probability, when the number of agents is large. This phenomenon is known as *asymptotic learning*.

Theorem 3.21 (Mossel, Sly and Tamuz). *Let μ_0, μ_1 be such that for every connected, undirected graph G there exists a random variable A_∞ such that almost surely $A_\infty^i = A_\infty$ for all $u \in V$. Then there exists a sequence $q(n) = q(n, \mu_0, \mu_1)$ such that $q(n) \rightarrow 1$ as $n \rightarrow \infty$, and $\mathbb{P}[A_\infty = S] \geq q(n)$, for any choice of undirected, connected graph G with n agents.*

Informally, when agents agree on limit action sets then they necessarily learn the correct state of the world, with probability that approaches one as the number of agents grows. This holds uniformly over all possible connected and undirected social network graphs.

The following theorem is a direct consequence of the two theorems above, since the property proved by Theorem 3.20 is the condition required by Theorem 3.21.

Theorem 3.22. *Let μ_0 and μ_1 induce non-atomic beliefs. Then there exists a sequence $q(n) = q(n, \mu_0, \mu_1)$ such that $q(n) \rightarrow 1$ as $n \rightarrow \infty$, and $\mathbb{P}[A_\infty^i = S] \geq q(n)$, for all agents i and for any choice of undirected, connected G with n agents.*

Before delving into the proof of Theorem 3.21 we introduce additional definitions in 3.6.1 and prove some general lemmas in 3.6.2, 3.6.3 and 3.6.4.

3.6.1. Additional general notation

We denote the actions of the neighbors of i up to time t by

$$I_t^i = \{A_{t'}^j : j \in \partial i, t' < t\},$$

and let I_∞^i denote all the actions of i 's neighbors:

$$I_\infty^i = \{A_{[0, \infty)}^j : j \in \partial i\} = \{A_{t'}^j : j \in \partial i, t' \geq 0\}.$$

We denote the probability that i chooses the correct action at time t by

$$p_t^i = \mathbb{P}[A_t^i = S].$$

and accordingly

$$p_\infty^i = \lim_{t \rightarrow \infty} p_t^i.$$

For a set of vertices $U \subseteq V$ we denote by $W(U)$ the private signals of the agents in U .

3.6.2. Sequences of rooted graphs and their limits

In this section we define a topology on undirected, connected *rooted graphs*. We call convergence in this topology *convergence to local limits*, and use it repeatedly in the proof of Theorem 3.21. The core of the proof of Theorem 3.21 is the topological Lemma 3.25, which we prove here. This lemma is a claim related to *local graph properties*, which we also introduce here.

Let $G = (V, E)$ be an undirected, connected, finite or countably infinite graph, and let $i \in V$ be a vertex in G . We denote by (G, i) the **rooted graph** G with root i .

Let $G = (V, E)$ and $G' = (V', E')$ be graphs. $h : V \rightarrow V'$ is a **graph isomorphism** between G and G' if $(i, j) \in E \Leftrightarrow (h(i), h(j)) \in E'$.

Let (G, i) and (G', i') be rooted graphs. Then $h : V \rightarrow V'$ is a **rooted graph isomorphism** between (G, i) and (G', i') if h is a graph isomorphism and $h(i) = i'$.

We write $(G, i) \cong (G', i')$ whenever there exists a rooted graph isomorphism between the two rooted graphs.

Given a graph $G = (V, E)$ and two vertices $i, j \in V$, the graph distance $d(i, j)$ is equal to the length in edges of a shortest (directed) path between i and j . We denote by $B_r(G, i)$ the ball of radius r around the vertex i in the graph $G = (V, E)$: Let V' be the set of vertices j such that $d(i, j)$ is at most r . Let $E' = \{(i, j) \in E : i, j \in V'\}$. Then $B_r(G, i)$ is the rooted graph with vertices V' , edges E' and root i .

We next define a topology on (undirected, connected) rooted graphs (or rather on their isomorphism classes; we shall simply refer to these classes as graphs). A natural metric between rooted graphs is the following (see Benjamini and Schramm [7], Aldous and Steele [1]). Given (G, i) and (G', i') , let

$$D((G, i), (G', i')) = 2^{-R},$$

where

$$R = \sup\{r : B_r(G, i) \cong B_r(G', i')\}.$$

This is indeed a metric: the triangle inequality follows immediately, and a standard diagonalization argument is needed to show that if $D((G, i), (G', i')) = 0$ then $B_\infty(G, i) \cong B_\infty(G', i')$ and so $(G, i) \cong (G', i')$.

This metric induces a topology that will be useful to us. As usual, the basis of this topology is the set of balls of the metric; the ball of radius 2^{-R} around the graph (G, i) is the set of graphs (G', i') such that $B_R(G, i) \cong B_R(G', i')$. We refer to convergence in this topology as convergence to a *local limit*, and provide the following equivalent definition for it.

Let $\{(G_r, i_r)\}_{r=1}^\infty$ be a sequence of rooted graphs. We say that the sequence converges if there exists a rooted graph (G', i') such that

$$B_r(G', i') \cong B_r(G_r, i_r),$$

for all $r \geq 1$. We then write

$$(G', i') = \lim_{r \rightarrow \infty} (G_r, i_r),$$

and call (G', i') the **local limit** of the sequence $\{(G_r, i_r)\}_{r=1}^\infty$.

Let \mathcal{G}_d be the set of rooted graphs with degree at most d .

Exercise. Show that \mathcal{G}_d is *compact*, and deduce from that the following lemma:

Lemma 3.23. *Let $\{(G_r, i_r)\}_{r=1}^\infty$ be a sequence of rooted graphs in \mathcal{G}_d . Then there exists a subsequence $\{(G_{r_n}, i_{r_n})\}_{n=1}^\infty$ with $r_{n+1} > r_n$ for all n , such that $\lim_{n \rightarrow \infty} (G_{r_n}, i_{r_n})$ exists.*

We next define *local properties* of rooted graphs. Let P be property of rooted graphs or a Boolean predicate on rooted graphs. We write $(G, i) \in P$ if (G, i) has the property, and $(G, i) \notin P$ otherwise.

We say that P is a **local property** if, for every $(G, i) \in P$ there exists an $r > 0$ such that if $B_r(G, i) \cong B_r(G', i')$, then $(G', i') \in P$. Let r be such that $B_r(G, i) \cong B_r(G', i') \Rightarrow (G', i') \in P$. Then we say that (G, i) **has property P with radius r** , and denote $(G, i) \in P^{(r)}$. That is, if (G, i) has a local property P then there is some r such that knowing the ball of radius r around i in G is sufficient to decide that (G, i) has the property P .

An alternative name for a local property would therefore be a *locally decidable* property. In our topology, local properties are nothing but *open sets*: the definition above states that if $(G, i) \in P$ then there exists an element of the basis of the topology that includes (G, i) and is also in P . This is a necessary and sufficient condition for P to be open.

We use this fact to prove the following lemma. Let \mathcal{B}_d be the set of infinite, connected, undirected graphs of degree at most d , and let \mathcal{B}_d^r be the set of \mathcal{B}_d -rooted graphs

$$\mathcal{B}_d^r = \{(G, i) : G \in \mathcal{B}_d, i \in G\}.$$

Exercise. Prove the following lemma.

Lemma 3.24. \mathcal{B}_d^r is compact.

We now state and prove the main lemma of this section. Note that the set of graphs \mathcal{B}_d satisfies the conditions of this lemma.

Lemma 3.25. *Let \mathcal{A} be a set of infinite, connected graphs, let \mathcal{A}^r be the set of \mathcal{A} -rooted graphs*

$$\mathcal{A}^r = \{(G, i) : G \in \mathcal{A}, i \in G\},$$

and assume that \mathcal{A} is such that \mathcal{A}^r is compact.

Let P be a local property such that for each $G \in \mathcal{A}$ there exists a vertex $j \in G$ such that $(G, j) \in P$. Then for each $G \in \mathcal{A}$ there exist an r_0 and infinitely many distinct vertices $\{j_n\}_{n=1}^\infty$ such that $(G, j_n) \in P^{(r_0)}$ for all n .




FIG 1. Schematic diagram of the proof of lemma 3.25. The rooted graph (G', i') is a local limit of (G, i_r) . For $r \geq R$, the ball $B_R(G', i')$ is isomorphic to the ball $B_R(G, i_r)$, with $w' \in G'$ corresponding to $j_r \in G$.

Proof. Let G be an arbitrary graph in \mathcal{A} . Consider a sequence $\{k_r\}_{r=1}^\infty$ of vertices in G such that for all $r, s \in \mathbb{N}$ the balls $B_r(G, k_r)$ and $B_s(G, k_s)$ are disjoint.

Since \mathcal{A}^r is compact, the sequence $\{(G, k_r)\}_{r=1}^\infty$ has a converging subsequence $\{(G, k_{r_n})\}_{n=1}^\infty$ with $r_{n+1} > r_n$. Write $i_r = k_{r_n}$, and let

$$(G', i') = \lim_{r \rightarrow \infty} (G, i_r).$$

Note that since \mathcal{A}^r is compact, $(G', i') \in \mathcal{A}^r$ and in particular $G' \in \mathcal{A}$ is an infinite, connected graph. Note also that since $r_{n+1} > r_n$, it also holds that the balls $B_r(G, i_r)$ and $B_s(G, i_s)$ are disjoint for all $r, s \in \mathbb{N}$.

Since $G' \in \mathcal{A}$, there exists a vertex $j' \in G'$ such that $(G', j') \in P$. Since P is a local property, $(G', j') \in P^{(r_0)}$ for some r_0 , so that if $B_{r_0}(G', j') \cong B_{r_0}(G, j)$ then $(G, j) \in P$.

Let $R = d(i', j') + r_0$, so that $B_{r_0}(G', j') \subseteq B_R(G', i')$. Then, since the sequence (G, i_r) converges to (G', i') , for all $r \geq R$ it holds that $B_R(G, i_r) \cong B_R(G', i')$. Therefore, for all $r > R$ there exists a vertex $j_r \in B_R(G, j_r)$ such that $B_{r_0}(G, j_r) \cong B_{r_0}(G', j')$. Hence $(G, j_r) \in P^{(r_0)}$ for all $r > R$ (see Fig 1). Furthermore, for $r, s > R$, the balls $B_R(G, i_r)$ and $B_R(G, i_s)$ are disjoint, and so $j_r \neq j_s$.

We have therefore shown that the vertices $\{j_r\}_{r>R}$ are an infinite set of distinct vertices such that $(G, j_r) \in P^{(r_0)}$, as required. \square

3.6.3. Coupling isomorphic balls

This section includes technical claims that we will use later. Their spirit is that everything that happens to an agent up to time t depends only on the state of the world and a ball of radius t around it. We leave their proofs as an exercise.

Lemma 3.26. *Consider two processes with identical private signal distributions (μ_0, μ_1) , on different graphs $G = (V, E)$ and $G' = (V', E')$.*

Let $t \geq 1$, $i \in V$ and $i' \in V'$ be such that there exists a rooted graph isomorphism $h : B_t(G, i) \rightarrow B_t(G', i')$.

Let M be a random variable that is measurable in $\sigma(H_t^i)$. Then there exists an M' that is measurable in $H_t^{i'}$ such that the distribution of (M, S) is identical to the distribution of (M', S') .

In particular, we use this lemma in the case where M is an estimator of S . Then this lemma implies that the probability that $M = S$ is equal to the probability that $M' = S'$.

Recall that $p_t^i = \mathbb{P}[A_t^i = S] = \max_{A \in \sigma(H_t^i)} \mathbb{P}[A = S]$. Hence we can apply this lemma (3.26) above to A_t^i and $A_t^{i'}$:

Corollary 3.27. *If $B_t(G, i)$ and $B_t(G', i')$ are isomorphic then $p_t^i = p_{i'}(t)$.*

3.6.4. δ -independence

To prove that agents learn S we will show that the agents must, over the duration of this process, gain access to a large number of measurements of S that are *almost* independent. To formalize the notion of almost-independence we define δ -independence and prove some easy results about it. The proofs in this section are again left as an exercise to the reader.

Let μ and ν be two measures defined on the same space. We denote the total variation distance between them by $d_{TV}(\mu, \nu)$. Let A and B be two random variables with joint distribution $\mu_{(A,B)}$. Then we denote by μ_A the marginal distribution of A , μ_B the marginal distribution of B , and $\mu_A \times \mu_B$ the product distribution of the marginal distributions.

Let (X_1, X_2, \dots, X_k) be random variables. We refer to them as **δ -independent** if their joint distribution $\mu_{(X_1, \dots, X_k)}$ has total variation distance of at most δ from the product of their marginal distributions $\mu_{X_1} \times \dots \times \mu_{X_k}$:

$$d_{\text{TV}}(\mu_{(X_1, \dots, X_k)}, \mu_{X_1} \times \dots \times \mu_{X_k}) \leq \delta.$$

Likewise, (X_1, \dots, X_l) are **δ -dependent** if the distance between the distributions is more than δ .

Claim 3.28. *Let A, B and C be random variables such that $\mathbb{P}[A \neq B] \leq \delta$ and (B, C) are δ' -independent. Then (A, C) are $2\delta + \delta'$ -independent.*

Claim 3.29. *Let (X, Y) be δ -independent, and let $Z = f(Y, B)$ for some function f and B that is independent of both X and Y . Then (X, Z) are also δ -independent.*

Claim 3.30. *Let $A = (A_1, \dots, A_k)$, and X be random variables. Let (A_1, \dots, A_k) be δ_1 -independent and let (A, X) be δ_2 -independent. Then (A_1, \dots, A_k, X) are $(\delta_1 + \delta_2)$ -independent.*

As an application of these claim we state the following lemma. The proof is again left as a (non-trivial) exercise.

Lemma 3.31. *For every $1/2 < p < 1$ there exist $\delta = \delta(p) > 0$ and $\eta = \eta(p) > 0$ such that if S and (X_1, X_2, X_3) are binary random variables with $\mathbb{P}[S = 1] = 1/2$, $1/2 < p - \eta \leq \mathbb{P}[X_i = S] < 1$, and (X_1, X_2, X_3) are δ -independent conditioned on S then $\mathbb{P}[a(X_1, X_2, X_3) = S] > p$, where a is the MAP estimator of S given (X_1, X_2, X_3) .*

In other words, one's odds of guessing S using three conditionally almost-independent bits are greater than using a single bit.

3.6.5. Asymptotic learning

In this section we prove Theorem 3.21. To prove this theorem we will need a number of intermediate results, which are given over the next few sections.

Estimating the limiting optimal action set A_∞ . We would like to show that although the agents have a common optimal action set A_∞ only at the limit $t \rightarrow \infty$, they can estimate this set well at a large enough time t .

The action A_t^i is agent i 's MAP estimator of S at time t . We likewise define K_t^i to be agent i 's MAP estimator of A_∞ , at time t :

$$K_t^i = \operatorname{argmax}_{K \in \{0,1\}} \mathbb{P}[A_\infty = K | H_t^i]. \quad (3.22)$$

We show that the sequence of random variables K_t^i converges to A_∞ for every i , or that alternatively $K_t^i = A_\infty$ for each agent i and t large enough:

Lemma 3.32. $\mathbb{P}[\lim_{t \rightarrow \infty} K_t^i = A_\infty] = 1$ for all $i \in V$.

Lemma 3.32 follows by direct application of the more general Lemma 3.33 which we leave as an exercise. Note that a consequence is that $\lim_{t \rightarrow \infty} \mathbb{P}[K_t^i = A_\infty] = 1$.

Exercise. Prove the following lemma.

Lemma 3.33. *Let $\mathcal{K}_1 \subseteq \mathcal{K}_2, \dots$ be a filtration of σ -algebras, and let $\mathcal{K}_\infty = \cup_t \mathcal{K}_t$. Let K be a random variable that takes a finite number of values and is measurable in \mathcal{K}_∞ . Let $M(t) = \operatorname{argmax}_k \mathbb{P}[K = k | \mathcal{K}(t)]$ be the MAP estimator of K given \mathcal{K}_t . Then*

$$\mathbb{P}\left[\lim_{t \rightarrow \infty} M(t) = K\right] = 1.$$

We would like at this point to provide the reader with some more intuition on A_t^i, K_t^i and the difference between them. Assuming that $A_\infty = 1$ then by definition, from some time t_0 on, $A_t^i = 1$, and from Lemma 3.32, $K_t^i = 1$. The same applies when $A_\infty = 0$. However, when $A_\infty = \{0, 1\}$ then A_t^i takes both values 0 and 1 infinitely often, but K_t^i will eventually equal $\{0, 1\}$. That is, agent i will realize at some point that, although it thinks at the moment that 1 is preferable to 0 (for example), it is in fact the most likely outcome that its belief will converge to 1/2. In this case, although it is not optimal, a *uniformly random* guess of which is the best action may not be so bad. Our next definition is based on this observation.

Based on K_t^i , we define a second “action” C_t^i . Let C_t^i be picked uniformly from K_t^i : if $K_t^i = 1$ then $C_t^i = 1$, if $K_t^i = 0$ then $C_t^i = 0$, and if $K_t^i = \{0, 1\}$ then C_t^i is picked independently from the uniform distribution over $\{0, 1\}$.

Note that we here extend our probability space by including in I_t^i (the observations of agent i up to time t) an extra uniform bit that is independent of all else and S in particular. Hence this does not increase i ’s ability to estimate S , and if we can show that in this setting i learns S then i can also learn S without this bit. In fact, we show that asymptotically it is as good an estimate for S as the best estimate A_t^i :

Claim 3.34.

$$\lim_{t \rightarrow \infty} \mathbb{P}[C_t^i = S] = \lim_{t \rightarrow \infty} \mathbb{P}[A_t^i = S] = p$$

for all i .

Exercise. Prove Claim 3.34.

The probability of getting it right. Recall that $p_t^i = \mathbb{P}[A_t^i = S]$ and $p_\infty^i = \lim_{t \rightarrow \infty} p_t^i$ (i.e., p_t^i is the probability that agent i takes the right action at time t). We state here a few easy related claims that will later be useful to us. The next claim is a rephrasing of the first part of Claim 3.5.

Claim 3.35. $p_{t+1}^i \geq p_t^i$.

The following claim is a rephrasing of Corollary 3.6.

Claim 3.36. *There exists a $p \in [0, 1]$ such that $p_\infty^i = p$ for all i .*

We make the following definition in the spirit of these claims:

$$p = \lim_{t \rightarrow \infty} \mathbb{P} [A_t^i = S].$$

In the context of a specific social network graph G we may denote this quantity as $p(G)$.

For time $t = 1$ the next standard claim follows from the fact that the agents' signals are informative.

Claim 3.37. $p_t^i > 1/2$ for all i and t .

Recall that $|\partial i|$ is the out-degree of i , or the number of neighbors that i observes. The next lemma states that an agent with many neighbors will have a good estimate of S already at the second round, after observing the first action of its neighbors.

Lemma 3.38. *There exist constants $C_1 = C_1(\mu_0, \mu_1)$ and $C_2 = C_2(\mu_0, \mu_1)$ such that for any agent i it holds that*

$$p_1^i \geq 1 - C_1 e^{-C_2 \cdot |\partial i|}.$$

Intuitively, this follows from the fact that i 's neighbors will provide him with $|\partial i|$ independent signals. We leave the proof as an exercise.

The following claim is a direct consequence of the previous lemmas of this section.

Claim 3.39. *Let $d(G) = \sup\{|\partial i|\}$ be the out-degree of the graph G ; note that for infinite graphs it may be that $d(G) = \infty$. Then there exist constants $C_1 = C_1(\mu_0, \mu_1)$ and $C_2 = C_2(\mu_0, \mu_1)$ such that*

$$p(G) \geq 1 - C_1 e^{-C_2 \cdot d(G)}.$$

Proof. Let i be an arbitrary vertex in G . Then by Lemma 3.38 it holds that

$$p_1^i \geq 1 - C_1 e^{-C_2 \cdot \partial i},$$

for some constants C_1 and C_2 . By Lemma 3.35 we have that $p_{t+1}^i \geq p_t^i$, and therefore

$$p_\infty^i = \lim_{n \rightarrow \infty} p_n^i \geq 1 - C_1 e^{-C_2 \cdot \partial i}.$$

Finally, $p(G) = p_\infty^i$ by Lemma 3.36, and so

$$p_\infty^i \geq 1 - C_1 e^{-C_2 \cdot \partial i}.$$

Since this holds for an arbitrary vertex i , the claim follows. \square

Local limits and pessimal graphs. We now turn to apply local limits to our process. We consider here and henceforth the same model as applied, with the same private signals, to different graphs. We write $p(G)$ for the value of p on the process on G , $A_\infty(G)$ for the value of A_∞ on G , etc.

Lemma 3.40. *Let $(G, i) = \lim_{r \rightarrow \infty} (G_r, i_r)$. Then $p(G) \leq \liminf_r p(G_r)$.*

Proof. Since $B_r(G_r, i_r) \cong B_r(G, i)$, by Lemma 3.27 we have that $p^{i_r} = p_r^{i_r}$. By Claim 3.35 $p_r^{i_r} \leq p(G_r)$, and therefore $p_r^i \leq p(G_r)$. The claim follows by taking the limit inferior of both sides. \square

Recall that \mathcal{B}_d denotes the set of infinite, connected, undirected graphs of degree at most d . Let

$$\mathcal{B} = \bigcup_d \mathcal{B}_d.$$

Let

$$p^* = p^*(\mu_0, \mu_1) = \inf_{G \in \mathcal{B}} p(G)$$

be the probability of learning in the pessimal graph.

Note that by Claim 3.37 we have that $p^* > 1/2$. We show that this infimum is in fact attained by some graph:

Lemma 3.41. *There exists a graph $H \in \mathcal{B}$ such that $p(H) = p^*$.*

Proof. Let $\{G_r = (V_r, E_r)\}_{r=1}^\infty$ be a series of graphs in \mathcal{B} such that $\lim_{r \rightarrow \infty} p(G_r) = p^*$. Note that $\{G_r\}$ must all be in \mathcal{B}_d for some d (i.e., have uniformly bounded degrees), since otherwise the sequence $p(G_r)$ would have values arbitrarily close to 1 and its limit could not be p^* (unless indeed $p^* = 1$, in which case our main Theorem 3.21 is proved). This follows from Lemma 3.38.

We now arbitrarily mark a vertex i_r in each graph, so that $i_r \in V_r$, and let (H, i) be the limit of some subsequence of $\{G_r, i_r\}_{r=1}^\infty$. Since \mathcal{B}_d is compact (Lemma 3.24), (H, i) is guaranteed to exist, and $H \in \mathcal{B}_d$.

By Lemma 3.40 we have that $p(H) \leq \liminf_r p(G_r) = p^*$. But since $H \in \mathcal{B}$, $p(H)$ cannot be less than p^* , and the claim is proved. \square

Independent bits. We now show that on infinite graphs, the private signals in the neighborhood of agents that are “far enough away” are (conditioned on S) almost independent of A_∞ (the final consensus estimate of S).

Lemma 3.42. *Let G be an infinite graph. Fix a vertex i_0 in G . Then for every $\delta > 0$ there exists an r_δ such that for every $r \geq r_\delta$ and every vertex i with $d(i_0, i) > 2r$ it holds that $W(B_r(G, i))$, the private signals in $B_r(G, i)$, are δ -independent of A_∞ , conditioned on S .*

Here we denote graph distance by $d(\cdot, \cdot)$.

Proof. Fix i_0 , and let i be such that $d(i_0, i) > 2r$. Then $B_r(G, i_0)$ and $B_r(G, i)$ are disjoint, and hence independent conditioned on S . Hence $K_r^{i_0}$ is independent of $W(B_r(G, i))$, conditioned on S .

Lemma 3.32 states that $\mathbb{P}[\lim_{r \rightarrow \infty} K_r^{i_0} = A_\infty] = 1$, and so there exists an r_δ such that for every $r \geq r_\delta$ it holds that $\mathbb{P}[K_r^{i_0} = A_\infty] > 1 - \frac{1}{2}\delta$.

Recall Claim 3.28: for any A, B, C , if $\mathbb{P}[A = B] = 1 - \frac{1}{2}\delta$ and B is independent of C , then (A, C) are δ -independent.

Applying Claim 3.28 to A_∞ , $K_r^{i_0}$ and $W(B_r(G, i))$ we get that for any r greater than r_δ it holds that $W(B_r(G, i))$ is δ -independent of A_∞ , conditioned on S . \square

We will now show, in the lemmas below, that in infinite graphs each agent has access to any number of “good estimators”: δ -independent measurements of S that are each almost as likely to equal S as p^* , the minimal probability of estimating S on any infinite graph.

We say that agent $i \in G$ has k (δ, ϵ) -good estimators if there exists a time t and estimators M_1, \dots, M_k such that $(M_1, \dots, M_k) \in H_t^i$ and

1. $\mathbb{P}[M_i = S] > p^* - \epsilon$ for $1 \leq i \leq k$.
2. (M_1, \dots, M_k) are δ -independent, conditioned on S .

The proof of the next claim is straightforward.

Claim 3.43. *Let P denote the property of having k (δ, ϵ) -good estimators. Then P is a local property of the rooted graph (G, i) . Furthermore, if $u \in G$ has k (δ, ϵ) -good estimators measurable in H_t^u then $(G, i) \in P^{(t)}$, i.e., (G, i) has property P with radius t .*

We are now ready to prove the main lemma of this subsection:

Lemma 3.44. *For every $d \geq 2$, $G \in \mathcal{B}_d$, $\epsilon, \delta > 0$ and $k \geq 0$ there exists a vertex i , such that i has k (δ, ϵ) -good estimators.*

Informally, this lemma states that if G is an infinite graph with bounded degrees, then there exists an agent that eventually has k almost-independent estimates of S with quality close to p^* , the minimal probability of learning.

Proof. In this proof we use the term “independent” to mean “independent conditioned on S ”.

We choose an arbitrary d and prove by induction on k . The basis $k = 0$ is trivial. Assume the claim holds for k , any $G \in \mathcal{B}_d$ and all $\epsilon, \delta > 0$. We shall show that it holds for $k + 1$, any $G \in \mathcal{B}_d$ and any $\delta, \epsilon > 0$.

By the inductive hypothesis for every $G \in \mathcal{B}_d$ there exists a vertex in G that has k $(\delta/100, \epsilon)$ -good estimators (M_1, \dots, M_k) .

Now, having k $(\delta/100, \epsilon)$ -good estimators is a *local property* (Claim 3.43). We now therefore apply Lemma 3.25: since every graph $G \in \mathcal{B}_d$ has a vertex with k $(\delta/100, \epsilon)$ -good estimators, any graph $G \in \mathcal{B}_d$ has a time t_k for which infinitely many distinct vertices $\{j_r\}$ have k $(\delta/100, \epsilon)$ -good estimators measurable at time t_k .

In particular, if we fix an arbitrary $i_0 \in G$ then for every r there exists a vertex $j \in G$ that has k $(\delta/100, \epsilon)$ -good estimators and whose distance $d(i_0, j)$ from i_0 is larger than r .

We shall prove the lemma by showing that for a vertex j that is far enough from i_0 which has $(\delta/100, \epsilon)$ -good estimators (M_1, \dots, M_k) , it holds that for a time t_{k+1} large enough $(M_1, \dots, M_k, C_{t_{k+1}}^j)$ are (δ, ϵ) -good estimators.

By Lemma 3.42 there exists an r_δ such that if $r > r_\delta$ and $d(i_0, j) > 2r$ then $W(B_r(G, j))$ is $\delta/100$ -independent of A_∞ . Let $r^* = \max\{r_\delta, t_k\}$, where t_k is such

that there are infinitely many vertices in G with k good estimators measurable at time t_k .

Let j be a vertex with k $(\delta/100, \epsilon)$ -good estimators (M_1, \dots, M_k) at time t_k , such that $d(i_0, j) > 2r^*$. Denote

$$\bar{M} = (M_1, \dots, M_k).$$

Since $d(i_0, j) > 2r_\delta$, $W(B_{r^*}(G, j))$ is $\delta/100$ -independent of A_∞ , and since $B_{t_k}(G, j) \subseteq B_{r^*}(G, j)$, $W(B_{t_k}(G, j))$ is $\delta/100$ -independent of A_∞ . Finally, since $\bar{M} \in \sigma(H_{t_k}^j)$, \bar{M} is a function of $W(B_{t_k}(G, j))$, and so by Claim 3.29 we have that \bar{M} is also $\delta/100$ -independent of A_∞ .

For t_{k+1} large enough it holds that

- $K_{t_{k+1}}^j$ is equal to A_∞ with probability at least $1 - \delta/100$, since

$$\lim_{t \rightarrow \infty} \mathbb{P} \left[K_t^j = A_\infty \right] = 1,$$

by Claim 3.32.

- Additionally, $\mathbb{P} \left[C_{t_{k+1}}^j = S \right] > p^* - \epsilon$, since

$$\lim_{t \rightarrow \infty} \mathbb{P} \left[C_t^j = S \right] = p \geq p^*,$$

by Claim 3.34.

We have then that (\bar{M}, A_∞) are $\delta/100$ -independent and $\mathbb{P} \left[K_{t_{k+1}}^j \neq A_\infty \right] \leq \delta/100$. Claim 3.28 states that if (A, B) are δ -independent $\mathbb{P} [B \neq C] \leq \delta'$ then (A, C) are $\delta + 2\delta'$ -independent. Applying this here we get that $(\bar{M}, K_{t_{k+1}}^j)$ are $\delta/25$ -independent.

It follows by application of Claim 3.30 that $(M_1, \dots, M_k, K_{t_{k+1}}^j)$ are δ -independent. Since $C_{t_{k+1}}^j$ is a function of $K_{t_{k+1}}^j$ and an independent bit, it follows by another application of Claim 3.29 that $(M_1, \dots, M_k, C_{t_{k+1}}^j)$ are also δ -independent.

Finally, since $\mathbb{P} \left[C_{t_{k+1}}^j = S \right] > p^* - \epsilon$, j has the $k+1$ (δ, ϵ) -good estimators $(M_1, \dots, C_{t_{k+1}}^j)$ and the proof is concluded. \square

Asymptotic learning. As a tool in the analysis of finite graphs, we would like to prove that in infinite graphs the agents learn the correct state of the world almost surely.

Theorem 3.45. *Let $G = (V, E)$ be an infinite, connected undirected graph with bounded degrees (i.e., G is a general graph in \mathcal{B}). Then $p(G) = 1$.*

Note that an alternative phrasing of this theorem is that $p^* = 1$.

Proof. Assume the contrary, i.e. $p^* < 1$. Let H be an infinite, connected graph with bounded degrees such that $p(H) = p^*$, such as we have shown exists in Lemma 3.41.

By Lemma 3.44 there exists for arbitrarily small $\epsilon, \delta > 0$ a vertex $w \in H$ that has access at some time T to three δ -independent estimators (conditioned on S), each of which is equal to S with probability at least $p^* - \epsilon$. By Claims 3.31 and 3.37, the MAP estimator of S using these estimators equals S with probability higher than p^* , for the appropriate choice of low enough ϵ, δ . Therefore, since j 's action A_t^j is the MAP estimator of S , its probability of equaling S is $\mathbb{P}[A_t^j = S] > p^*$ as well, and so $p(H) > p^*$ - contradiction. \square

Using Theorem 3.45 we can now prove Theorem 3.21, which is the corresponding theorem for finite graphs:

Proof of Theorem 3.21. Assume the contrary. Then there exists a series of graphs $\{G_r\}$ with r agents such that $\lim_{r \rightarrow \infty} \mathbb{P}[A_\infty(G_r) = S] < 1$, and so also $\lim_{r \rightarrow \infty} p(G_r) < 1$.

By the same argument of Theorem 3.45 these graphs must all be in \mathcal{B}_d for some d , since otherwise, by Lemma 3.39, there would exist a subsequence of graphs $\{G_{r_d}\}$ with degree at least d and $\lim_{d \rightarrow \infty} p(G_{r_d}) = 1$. Since \mathcal{B}_d is compact (Lemma 3.24), there exists a graph $(G, i) \in \mathcal{B}_d$ that is the limit of a subsequence of $\{(G_r, i_r)\}_{r=1}^\infty$.

Since G is infinite and of bounded degree, it follows by Theorem 3.45 that $p(G) = 1$, and in particular $\lim_{r \rightarrow \infty} p_\infty^i(r) = 1$. As before, $p_{i_r}(r) = p_\infty^i(r)$, and therefore $\lim_{r \rightarrow \infty} p_{i_r}(r) = 1$. Since $p(G_r) \geq p_{i_r}(r)$, $\lim_{r \rightarrow \infty} p(G_r) = 1$, which is a contradiction. \square

3.6.6. Example of non-atomic private beliefs leading to non-learning

We sketch an example in which private beliefs are atomic and asymptotic learning does not occur.

Example 3.46. *Let the graph G be the undirected chain of length n , so that $V = \{1, \dots, n\}$ and (i, j) is an edge if $|i - j| = 1$. Let the private signals be bits that are each independently equal to S with probability $2/3$. We choose here the tie breaking rule under which agents defer to their original signals³.*

We leave the following claim as an exercise to the reader.

Claim 3.47. *If an agent i has at least one neighbor with the same private signal (i.e., $W_i = W_j$ for j a neighbor of i) then i will always take the same action $A_t^i = W_i$.*

Since this happens with probability that is independent of n , with probability bounded away from zero an agent will always take the wrong action, and so asymptotic learning does not occur. It is also clear that optimal action sets do not become common knowledge, and these fact are indeed related.

³We conjecture that changing the tie-breaking rule does not produce asymptotic learning, even for randomized tie-breaking.

References

- [1] D. Aldous and J. Steele. The objective method: Probabilistic combinatorial optimization and local weak convergence. *Probability on Discrete Structures (Volume 110 of Encyclopaedia of Mathematical Sciences)*, ed. H. Kesten, 110:1–72, 2003. [MR2023650](#)
- [2] R. J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976. [MR0433654](#)
- [3] V. Bala and S. Goyal. Learning from neighbours. *Review of Economic Studies*, 65(3):595–621, July 1998.
- [4] A. V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.
- [5] F. Benezit, P. Thiran, and M. Vetterli. Interval consensus: from quantized gossip to voting. In *ICASSP 2009*, pages 3661–3664, 2009.
- [6] I. Benjamini, S.-O. Chan, R. O’Donnell, O. Tamuz, and L.-Y. Tan. Convergence, unanimity and disagreement in majority dynamics on unimodular graphs and random graphs. *Stochastic Processes and their Applications*, 126(9):2719–2733, 2016. [MR3522298](#)
- [7] I. Benjamini and O. Schramm. Recurrence of distributional limits of finite planar graphs. *Selected Works of Oded Schramm*, pages 533–545, 2011. [MR2883381](#)
- [8] E. Berger. Dynamic monopolies of constant size. *Journal of Combinatorial Theory, Series B*, 83(2):191–200, 2001. [MR1866395](#)
- [9] S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, pages 992–1026, 1992.
- [10] D. Cartwright and F. Harary. Structural balance: a generalization of heider’s theory. *Psychological review*, 63(5):277, 1956.
- [11] P. Clifford and A. Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973. [MR0343950](#)
- [12] J.-A.-N. Condorcet. *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. De l’Imprimerie Royale, 1785.
- [13] M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- [14] P. DeMarzo and C. Skiadas. On the uniqueness of fully informative rational expectations equilibria. *Economic Theory*, 13(1):1–24, 1999. [MR1658993](#)
- [15] J. L. Doob. *Stochastic Processes*. John Wiley and Sons, 1953. [MR0058896](#)
- [16] J. L. Doob. *Classical potential theory and its probabilistic counterpart*, volume 262. Springer, 2001. [MR1814344](#)
- [17] D. Gale and S. Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 45(2):329–346, November 2003. [MR2023667](#)
- [18] J. Geanakoplos. Common knowledge. *Handbook of game theory with economic applications*, 2:1437–1496, 1994. [MR1313236](#)
- [19] J. Geanakoplos and H. Polemarchakis. We can’t disagree forever. *Journal of Economic Theory*, 28(1):192–200, 1982.

- [20] Y. Ginosar and R. Holzman. The majority action on infinite graphs: strings and puppets. *Discrete Mathematics*, 215(1–3):59–72, 2000. [MR1746448](#)
- [21] E. Goles and J. Olivos. Periodic behaviour of generalized threshold functions. *Discrete Mathematics*, 30(2):187–189, 1980. [MR0566436](#)
- [22] B. Golub and M. O. Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010.
- [23] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963. [MR0144363](#)
- [24] R. A. Holley and T. M. Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. *The annals of probability*, pages 643–663, 1975. [MR0402985](#)
- [25] J. Kahn, G. Kalai, and N. Linial. The influence of variables on boolean functions. In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, pages 68–80, 1988.
- [26] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. AMS Bookstore, 2009. [MR2466937](#)
- [27] W. S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4):115–133, 1943. [MR0010388](#)
- [28] G. Moran. On the period-two-property of the majority operator in infinite graphs. *Transactions of the American Mathematical Society*, 347(5):1649–1668, 1995. [MR1297535](#)
- [29] E. Mossel, J. Neeman, and O. Tamuz. Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems*, pages 1–22, 2013.
- [30] E. Mossel, N. Olsman, and O. Tamuz. Efficient bayesian learning in social networks with gaussian estimators. In *Proceedings of the 54th annual Allerton conference on Communication, control, and computing*. IEEE Press, 2016.
- [31] E. Mossel and G. Schoenebeck. Reaching consensus on social networks. In *Proceedings of 1st Symposium on Innovations in Computer Science*, pages 214–229, 2010.
- [32] E. Mossel, A. Sly, and O. Tamuz. On agreement and learning. Preprint at <http://arxiv.org/abs/1207.5895>, 2012.
- [33] E. Mossel, A. Sly, and O. Tamuz. Asymptotic learning on bayesian social networks. *Probability Theory and Related Fields*, pages 1–31, 2013. [MR3152782](#)
- [34] E. Mossel, A. Sly, and O. Tamuz. Strategic learning and the topology of social networks. *Econometrica*, 83(5):1755–1794, 2015. [MR3414192](#)
- [35] M. J. Osborne and A. Rubinstein. *A course in game theory*. MIT press, 1994. [MR1301776](#)
- [36] M. Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–2647, 2012. [MR3001136](#)

- [37] L. Page, S. Brin, R. Motwani, and T. Winograd. The PageRank citation ranking: bringing order to the web. Stanford InfoLab, 1999.
- [38] D. Rosenberg, E. Solan, and N. Vieille. Informational externalities and emergence of consensus. *Games and Economic Behavior*, 66(2):979–994, 2009. [MR2543312](#)
- [39] A. Rubinstein. *Economic fables*. Open Book Publishers, 2012.
- [40] L. Saloff-Coste. Lectures on finite markov chains. In P. Bernard, editor, *Lectures on Probability Theory and Statistics*, volume 1665 of *Lecture Notes in Mathematics*, pages 301–413. Springer Berlin Heidelberg, 1997. [MR1490046](#)
- [41] J. Sebenius and J. Geanakoplos. Don't bet on it: Contingent agreements with asymmetric information. *Journal of the American Statistical Association*, 78(382):424–426, 1983. [MR0711118](#)
- [42] L. Smith and P. Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000. [MR1748010](#)
- [43] M. Talagrand. On Russo's approximate zero-one law. *The Annals of Probability*, 22(3):1576–1587, 1994. [MR1303654](#)
- [44] O. Tamuz and R. J. Tessler. Majority dynamics and the retention of information. *Israel Journal of Mathematics*, 206(1):483–507, 2015. [MR3319649](#)