

Opinion Dynamics and Learning in Social Networks

Daron Acemoglu · Asuman Ozdaglar

Published online: 1 October 2010
© Springer-Verlag 2010

Abstract We provide an overview of recent research on belief and opinion dynamics in social networks. We discuss both Bayesian and non-Bayesian models of social learning and focus on the implications of the form of learning (e.g., Bayesian vs. non-Bayesian), the sources of information (e.g., observation vs. communication), and the structure of social networks in which individuals are situated on three key questions: (1) whether social learning will lead to consensus, i.e., to agreement among individuals starting with different views; (2) whether social learning will effectively aggregate dispersed information and thus weed out incorrect beliefs; (3) whether media sources, prominent agents, politicians and the state will be able to manipulate beliefs and spread misinformation in a society.

Keywords Bayesian updating · Consensus · Disagreement · Learning · Misinformation · Non-Bayesian models · Rule of thumb behavior · Social networks

1 Introduction

Almost all social interactions are, at least in part, shaped by beliefs and opinions. Most of this we take for granted. To start with the most mundane example, few of us would have sampled all the different types of food or entertainment on offer around the world, but most of us have formed beliefs and opinions about which ones we would like, and every day we make decisions on the basis of these beliefs. More important and more interesting from the viewpoint of social science, most of us also have a set of (often complex and nuanced) beliefs about how others will act in different social situations, which guide our behavior in

We thank Stefana Stantcheva for excellent research assistance.

D. Acemoglu (✉)
Department of Economics, MIT, Cambridge, MA 02139, USA
e-mail: daron@mit.edu

A. Ozdaglar
Laboratory for Information and Decision Systems, Electrical Engineering and Computer Science
Department, MIT, Cambridge, MA 02139, USA
e-mail: asuman@mit.edu

social contexts ranging from business meetings to competitive sports. These beliefs are intimately intertwined with “social norms,” such as what we deem as acceptable behavior (to help somebody who has fallen down in the street and to not be part of activities that would harm others and so on). They also shape our political participation and attitudes. Political participation is certainly a “learned” attribute: those of us lucky enough to live in democracies learn about and believe in the need for and the virtues of political participation and often take part in elections and other political activities. And when we go to the polls, which candidate we support is again shaped by our beliefs and opinions about what is a just society, which candidates are more reliable and so on. Many who are born in repressive societies instead form beliefs about the dangers of such types of political participation, though some of them are also encouraged as much by their beliefs as by their friends and family, to take part in protests against regimes they deem as unjust or illegitimate. The importance of the beliefs we hold for how our daily lives and society in general function cannot be overstated.

Where do these beliefs and opinions come from? While, as evolutionary biology has taught us, certain phenotypic characteristics have biological and genetic bases, it is unlikely that any of our beliefs are imprinted on us by our genes. Instead, we acquire our beliefs and opinions through various types of learning experiences. Some of this learning takes place within families, when parents teach certain basic principles and beliefs to their children (e.g., [15, 16, 19, 20, 77]). Much of it, however, takes place through a process of “social learning,” whereby individuals obtain information and update their beliefs and opinions as a result of their own experiences, their observations of others’ actions and experiences, the communication with others about their beliefs and behavior, news from media sources, and propaganda and indoctrination from political leaders and the state. While the process of learning by an individual from his or her experience can be viewed as an “individual” learning problem, it also has an explicitly “social” character. This is in three related but distinct senses:

1. Learning is social because any given individual observes the behavior of or receives information through communication with a small subset of society—those we may want refer to as her *social network*, consisting of her friends, her coworkers and peers, her distant and close family members, and a certain group of leaders that she listens to and respects (e.g., village leaders, trusted politicians, trusted media sources and so on). We can thus not separate the process of learning and opinion formation from the specific social network in which an individual is situated.
2. Learning is also social because an individual will need to *interpret* the information that she obtains in a social context. She will inevitably trust some information more than others and she will have to form conjectures about the sources of the experiences and the intentions of members of her social network in communicating certain information to her (below, we will refer to this as her conjectures about the *strategies* of others).
3. Learning is also social because these interactions will lead to *dynamics* in learning and opinion formation. Once an individual obtains a piece of information from a specific peer in her social network, then she may pass on this information or some version of it to other members of her social network. The latter will typically have their own social networks that do not overlap with hers, so that this information, regardless of whether it is accurate or not, may spread both within and beyond the initial social network in which it originated.¹

¹Another related social aspect, which has received less attention, is that the social network of the individual might change dynamically and endogenously as a result of the information that she receives. For example,

The social aspect of belief and opinion formation—*social learning* for short—will be our central focus in this paper. We will investigate how the structure of social relationships in society, the (potentially selective) trust that individuals have toward others and their conjectures about others' behavior and intentions impact the formation of their beliefs and opinions. We will refer to these influences as *the impact of the social network* on opinion formation. For concreteness, we will study the impact of the social network in the context of three specific questions, which have received the bulk of the attention from the literature:

- a. Will social learning lead different individuals to hold beliefs that are in agreement, even though they might start with different views (priors) and observe the actions of and engage in communication with different subsets of the society? Put differently, will a *consensus* form among different individuals?
- b. Will social learning effectively aggregate dispersed information about an underlying state that represents some social or economic situation? Will there be *asymptotic learning* (learning in the long run) of the underlying state which is the source of uncertainty? For example, in many situations there will be sufficient information in society to settle a question such as whether a particular economic, social or political action is desirable, but different parts of this information will be held by different agents rather than by a single entity. In this situation, whether social learning through observation and communication will be able to aggregate those dispersed parts becomes central. A corollary to this question is the following: will social learning guarantee that incorrect beliefs (which can be refuted on the basis of the available evidence) disappear?
- c. Will media sources, “prominent agents,” politicians and officers of the state be able to manipulate the beliefs of individuals, indoctrinate them and convince them of views that may not be justifiable by the data and evidence? Put differently, how much room is there for *belief manipulation* and *misinformation*?

While most of these questions are asked as simple “yes or no” questions, the answers we will present, on the basis of a large body of research over the past several decades, will often provide additional conditions and insights under which the answer will be yes. We will also attempt, whenever we can, to relate these conditions to the impact of the social network (in particular to structural properties of the social network) in which agents are situated. In some cases, these mathematical conditions will be straightforward to map to reality, though in many instances more work is necessary to either sharpen these conditions or to create a better bridge between the mathematical results and the reality that they are supposed to represent.²

The issue of whether a group of agents who hold dispersed information will be able to aggregate this information and reach a consensus, and in fact, a correct consensus, has been the focus of a large body of mathematical and philosophical work throughout the last several centuries. The seminal work by Marquis de Condorcet, which is the basis of what is now referred to as the *Condorcet's Jury Theorem*, is particularly noteworthy. de Condorcet

because of some new information that she obtains, she may decide to no longer trust some of her friends or previously trusted information sources.

²Like almost all of the literature in this area, in investigating this question, we will focus on the behavior of opinions and beliefs in a society in the “long run,” meaning after a sufficiently long time has elapsed. This will enable us to develop sharper mathematical insights, though it should be borne in mind throughout that the relevant “long run” might be longer than the lifespan of a single individual and might arrive so slowly that non-long-run behavior might be of greater interest. We view the development of more powerful mathematical models and more insightful analysis of short-run opinion and belief dynamics as an exciting and important area for future work.

[24] observed that truthful reporting of information by a large group of individuals with each holding a belief or piece of information correlated with some underlying state θ is sufficient for aggregation of information. A similar perspective, but based both on theoretical reasoning and empirical backing, was developed a century later by the British scientist (and arguably one of the founders of modern statistics), Francis Galton. In a famous *Nature* article published in 1907, Galton espoused the view that a group of relatively uninformed individuals would collectively have much more knowledge than any single one of them. Galton visited an agricultural fair in Plymouth in 1906 to investigate a specific application of this idea. Participants at the fair were asked to guess the weight of an ox (after the animal was slaughtered and dressed, meaning the head and other parts were removed). Eight hundred people apparently took part in this contest. Galton was well aware that these were no experts and in fact, in the *Nature* article that he wrote in the following year, he stated: “[. . . most of the contestants were. . .] as well fitted for making a just estimate of the dressed weight of an ox, as an average voter is judging the merits of most political issues on which he votes.” The remarkable thing reported in the *Nature* article was that, when Galton looked at the 787 valid entries, he found that the median estimate was extremely close to the actual weight of the ox. This estimate was 1197 pounds, while the actual weight was 1198. Galton concluded [42]:

“The result seems more creditable to the trustworthiness of a democratic judgment than might have been expected.”

Despite these famous arguments, it is now generally believed that this type of aggregation of dispersed information is neither a theoretical nor an empirical necessity. Instead, in some situations this type of aggregation may take place (at least approximately), while in many others it will not. In fact, disagreement on many economic, social and political phenomena, such as the scientific standing of evolution, the likelihood that US health care reforms will increase overall spending, whether Iranian elections in 2009 did produce a majority for Ahmadinejad, whether the current trend of emissions will lead to significant climate change, whether a particular government or politician is competent, and so on, is ubiquitous, despite the availability of a reasonable amount of data bearing on these questions. This suggests that useful models of learning should not always predict consensus, and certainly not the weeding out of incorrect beliefs. Instead, these should merely be possible outcomes among others, depending on the nature of disagreement (e.g., whether some people are likely to support an incompetent government because of a specific subset of the policies they adopt), the informativeness of the information agents possess and receive (e.g., whether there is new and relatively precise information arriving on an issue, or whether social learning is mostly a question of aggregating already existing priors and information), and the structure of the social network leading to the exchange of information (e.g., whether there are relatively isolated clusters of individuals not communicating with those outside their clusters).

When consensus and the weeding out of incorrect beliefs are not guaranteed, there will also be room for systematic indoctrination and spread of misinformation by certain “prominent” agents, media sources, politicians and the state. The ability of authoritarian regimes such as those in China and Iran to indoctrinate significant subsets of their populations with nationalistic or religious propaganda underscores this possibility (but the fact that these governments place great emphasis on controlling all news sources, and especially the Internet and foreign news, also highlights that such indoctrination is not easy or automatic).

More specific evidence that this type of influence on beliefs exists and can be important comes from DellaVigna and Kaplan [29], who provide an illustration of this by studying the expansion of the Fox News cable channel across US towns. As is well-documented and accepted, the coverage in Fox News is more right-wing and pro-Republican than that of other

TV stations and most other media sources. If media sources can indeed have a significant influence on opinions and beliefs, for example, by spreading misinformation, or more benignly by providing information that others would not and blocking some information that others would provide, then we might expect that exposure to Fox News may increase the support for the Republican Party. DellaVigna and Kaplan exploit the fact that the Fox News channel was introduced in October 1996, and then spread across different towns slowly, reaching several towns before 2000, and many after 2000, depending on its agreements and the decisions of local cable companies. DellaVigna and Kaplan then compare the change in the vote share of the Republican candidate between the 1996 and the 2000 elections. The 1996 elections were before the introduction of Fox News, so differential change between towns with and without access to this channel can be interpreted as the impact of Fox News. They find that the Republican candidate gained between 0.4 to 0.7 percentage points in towns where the population had access to the Fox News cable channel compared to similar towns without such access. The available evidence also suggests that there are certain marked differences in opinions (often very strongly held) between individuals consistently watching the Fox News vs. those obtaining their news from CNN or CNBC.

This discussion highlights the need for developing models in which certain subset of agents can have a major influence on the opinions of others (possibly by spreading misinformation) and also models in which even in the very long run, consensus may not arise. These models, though generally in their infancy, provide a potential framework for the study of persistent disagreements and indoctrination, and might also generate insights about what types of societies (partly based on their social networks) might be able to develop some “robustness” to indoctrination and misinformation.

We next provide a brief outline of the rest of this paper. In Sect. 2, we start by discussing two alternative approaches to social learning and opinion formation, one based on Bayesian updating of beliefs, and the other one based on non-Bayesian reasonable “rules of thumb” on how people form their opinions on the basis of evidence and social influences. The Bayesian learning approaches assume that individuals update their beliefs optimally (from a statistical point of view) given an underlying model of the world. We emphasize in this section that Bayesian approaches make several demanding requirements from the agents. For example, they require that the agents have a reliable “model of the world” enabling them to assign priors to all possible events and that they can update their beliefs by forming complex conjectures on the behavior of others in society. They also put considerable structure on the updating problem by ruling out many states as “zero probability events”. These features also imply that Bayesian approaches might end up putting too much structure, making issues such as indoctrination and spread of misinformation more difficult, almost impossible, to model. Finally, the inference problem facing Bayesian agents, particularly when they are situated in social networks in which information travels in complex ways, is quite challenging. While non-Bayesian *rule-of-thumb* learning models avoid some of these difficulties and might provide better approximations to the behavior of most agents, we do not currently have sufficient empirical evidence to distinguish between several different types of non-Bayesian models. This suggests that Bayesian approaches might be a useful benchmark for understanding the implications of several different types of deviations from Bayesian updating and different rules-of-thumb behaviors. In the end, both Bayesian and non-Bayesian approaches provide useful insights and contrasting them is often instructive. Ultimately, which class of models is most appropriate should depend on the specific question and social context. For example, questions related to belief manipulation and misinformation appear to be more naturally treated in a non-Bayesian context.

In Sect. 3, we present two models of Bayesian social learning by a set of agents observing others’ actions or communicating with each other over a social network. We also briefly

discuss the role of markets in the formation of opinions. We emphasize several important themes in this section. First, we illustrate why the Bayesian models are a natural benchmark, but also why they involve a high degree of sophisticated and complex reasoning on the part of the agents. Second, we show that, even though agents are assumed to be highly rational and have the correct model of the world on which they base their belief updates, strategic interactions that inevitably exist in such Bayesian environments place endogenous limits on the aggregation of dispersed information. We will illustrate this by showing the emergence of a phenomenon often referred to as “herding,” in which Bayesian rational agents follow others’ actions or opinions even when these are only imperfectly informative (and in the process, prevent the aggregation of information from which others would have also benefited). Third, we also show that even when aggregation of information breaks down and social learning leads to incorrect beliefs, there are strong forces toward consensus in Bayesian models. They thus do not present a natural framework for understanding persistent disagreements. Finally, there will also be natural limits on the spread of misinformation in Bayesian models. While herding is possible, herding will never happen on an action that is likely to be a “bad action”.

In Sect. 4, we present several non-Bayesian models of learning. We start with a simple and widely used model, the so-called DeGroot model of belief updating. We emphasize why this is a simple and tractable model, but also why it illustrates some of the shortcomings of non-Bayesian models. We then present a variant of this model which avoids some of these shortcomings and also enables a first attempt at modeling the possible spread of misinformation (propagated by a set of “prominent agents” which may include community leaders as well as media outlets). However, in this benchmark model, even though misinformation might spread, persistent disagreement is not possible and opinions will converge to a consensus (though this consensus is stochastic and cannot be known in advance). We then present an extended model in which both misinformation and persistent disagreement can coexist, and highlight possible research directions in this area.

Section 5 concludes with a brief discussion of future research. Throughout, we have tried to ground the discussion sufficiently by presenting baseline models and some relevant details of the analysis that will be useful in future work, but we have economized on space by referring the reader to the original papers for the proofs and more involved results.

2 Bayesian and Non-Bayesian Perspectives

Opinion formation is directly or indirectly about learning. An individual starts with some views (priors) about a subject, which will affect her economic, political or social decisions, and then updates them according to some process. Thus overall, opinion formation has three key components:

1. *Priors.* Any model of opinion formation has to start with some type of prior opinions for an individual or group of individuals. For example, an individual might start with “diffuse priors” meaning that she can be easily swayed by the information she receives. Or an individual might have such strong views that her posterior opinion will not be very different from her prior opinion even after she receives a substantial amount of new data and communications.
2. *Sources of information.* An individual will update her prior based on new information that she receives. This might come from her own experiences, from observing others’ actions and experiences, or from communication with others. Any model of opinion formation will implicitly or explicitly specify the sources of information, which will in general be

at least partly affected by an individual’s “social network”. For example, both when it comes to observing others’ experiences and to communication, an individual is much more likely to learn from and communicate with some people than others—typically, family members, friends, coworkers and other peers are much more likely to influence an individual by communicating with her or by providing her information on the basis of their experiences.

3. *Method of information processing.* The third and key part of the process of opinion formation is how the individual will combine her priors and the information she receives. By Bayesian models, we refer to those in which individuals use the Bayes rule to form the “best” mathematical estimate of the relevant unknowns given their priors and understanding of the world. Non-Bayesian models are defined as all those that are not Bayesian. They include information updating processes that similarly combine priors and information to yield a posterior, but they could also include various approaches that appear “non-informational” at first sight, for example, those in which an individual might change her opinion in a manner similar to being “infected” by a disease.

In this section, we first discuss the basic approach of Bayesian models and why they might be a useful starting point. We then discuss several cognitive difficulties that Bayesian models face and discuss some alternatives. Finally, we conclude with an analysis demonstrating how Bayesian updating does not always lead to consensus and information aggregation as is sometimes presumed.

2.1 Bayesian Approaches

The Bayes rule is familiar and simple in its abstract form. It states that for two probabilistic events A and B , the probability that A is true conditional on B being true, $\mathbb{P}(A | B)$, is given by

$$\mathbb{P}(A | B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} = \frac{\mathbb{P}(B | A) \cdot \mathbb{P}(A)}{\mathbb{P}(B)}, \tag{1}$$

i.e., by the ratio of the probability of the event that both A and B are true, $\mathbb{P}(A \cap B)$, and the unconditional probability of event B , $\mathbb{P}(B)$. Since a similar relationship holds for $\mathbb{P}(B | A)$ (i.e., $\mathbb{P}(B | A) = \mathbb{P}(A \cap B)/\mathbb{P}(A)$), the second expression follows. This formula is very powerful when we apply it to the issue of social (and for that matter individual) learning.

Consider a situation in which an individual is trying to form an opinion about some *underlying state* $\theta \in \Theta$. The state could correspond to some economic variable, such as potential earnings in an occupation or profitability of a line of business, or to a social or political variable, such as whether a politician is to be trusted or a certain ideology is useful or beneficial. As we noted in the previous subsection, he would first have to start with some “priors”. We can capture this by a function $\mathbb{P}(\theta)$, which gives the prior belief of the individual about the likelihood of each possible value of θ in Θ . The second key ingredient of learning is the information that the individual will receive. Let us represent this information by some *signal* $s \in S$. This could correspond to some observation concerning θ (e.g., how much do others in this occupation earn? How much profit do they make from this line of business? Is there any evidence that the politician in question is corrupt?).

The Bayesian approach then posits that the individual will update her prior after observing s according to a version of (1), in particular:

$$\mathbb{P}(\theta | s) = \frac{\mathbb{P}(s | \theta) \cdot \mathbb{P}(\theta)}{\mathbb{P}(s)}. \tag{2}$$

This equation implies that, provided that the individual knows $\mathbb{P}(s \mid \theta)$, $\mathbb{P}(s)$ and $\mathbb{P}(\theta)$, she can compute the probability that the true state is θ given the signal s , $\mathbb{P}(\theta \mid s)$. That the individual knows $\mathbb{P}(\theta)$ implies that she has priors on each possible value of θ . That she knows $\mathbb{P}(s \mid \theta)$ implies that she has an understanding of what types of signals to expect when the true value is θ .

These are not innocuous requirements. We can think of these two requirements together as positing that the individual has a reliable *model of the world*, meaning that her views about the likelihood of different underlying states and the distribution of the signals conditional on the states are accurate or at the very least, “reliable”. As we will see, such models of the world play a central role in Bayesian approaches.

More specifically, the first requirement implies that the individual has a complete set of priors. When θ takes on a few values, this may be a reasonable requirement. But when the set of possible values of θ , Θ , is large, this requirement becomes quite demanding. From a practical point of view, it implies that the individual needs to have a fairly complex view of the world, assigning priors (probabilities) to each possible event. From a mathematical point of view, this also requires some care, since if Θ is uncountable, one would have to define priors (and of course probabilities) only for measurable subsets of Θ , and this implies that many possible states must receive zero probability according to any well-defined prior or not even have well-defined priors (see, e.g., [14, 82]). We will see in Sect. 2.3 that the presence of “zero probability” events plays a crucial (and sometimes subtle) role in Bayesian updating.

The second requirement, that the individual should know, or should have well formed opinions about, $\mathbb{P}(s \mid \theta)$, is equally stringent. This quantity captures the conditional probability law of the signal s for each possible value of θ . Even in the simplest situations, this might be too much information for an individual to process or to have reliable knowledge about. And again, if she does not have reliable knowledge about $\mathbb{P}(s \mid \theta)$, (2) will not imply accurate posteriors and learning that follows from Bayesian updating may be unreliable. However, the real challenging implication of the requirement that the individual should have reliable knowledge about $\mathbb{P}(s \mid \theta)$ arises in social, rather than individual, learning situations. Suppose, for example, that s is (or includes) information that an individual obtains from the actions of others. This might be the choice of somebody with experience, or information about their performance, or something that they directly communicate to the individual. In all of these cases, there would be no unique function $\mathbb{P}(s \mid \theta)$ describing the relationship between s and θ independent of the behavioral rule or *strategy* of the other players.

Let us consider a specific example to elaborate on this point. Consider the problem of an individual, “the observer,” learning about some underlying state θ , which, for example, corresponds to how profitable a certain line of business is. In doing so, she will observe the success y of some other agent. Let us posit that there is a relationship between success denoted by y and the underlying state θ , but this relationship also depends on the effort that an agent exerts, e . Thus we can write this relationship as $\mathbb{P}(y \mid \theta, e)$. Suppose, putting aside the previous difficulties we have raised, that the observer fully understands and has reliable knowledge on this relationship. But this would be useless to her unless she also knows e . This means that beyond understanding the relevant “physical” or “technological” relationships between the underlying state, effort and success, the observer must also have a good understanding of the “strategic” relationships and form the right conjectures about e . This problem becomes much more complex when e itself is chosen by each agent as a function of their own beliefs about θ and their own observations. These issues will be central in our analysis in the next section when we study Bayesian models of observational learning. For now, our purpose has been to highlight the complexities that this might involve.

Where do the reliable models of the world for Bayesian agents come from? There is no good answer to this question, and in most Bayesian models, it is assumed, as we will see, that individuals have beliefs about $\mathbb{P}(\theta)$ and $\mathbb{P}(s \mid \theta)$ that coincide with the true data generating process and with each other's beliefs, and in fact, there is *common knowledge* that they all share the same priors. The only uncertainty is about the specific value of θ ; there is no uncertainty or doubt about the underlying model of the world, and this plays a central role in the implications of Bayesian models. While mathematically tractable and powerful, this approach raises further questions. Part of our interest in opinion and belief dynamics is motivated by our desire to understand where a complex set of beliefs that individuals hold come from. But assuming that they already start with common beliefs on certain key parts of the model that correspond to the true data generating process and that there is common knowledge on this, is potentially quite restrictive. Some, following a view, forcefully articulated by LJ Savage in the *Foundations of Statistics* [81] and the statement that Bayesian individuals receiving informative signals about some underlying state θ should eventually learn and agree on θ , maintain that common priors and common knowledge are justified or at the very least are good approximations to reality. However, we will see in Sect. 2.3 below that the Bayesian foundations of such common priors are not necessarily very strong. This argument raises further challenges for the Bayesian approach, at least in the case that starts with a model of the world that is the same for all agents and is common knowledge.

Overall, the requirements of the Bayesian approach may be quite challenging for most individuals when it comes to complex issues. For this reason, some, like Gilboa and Schmeidler [45], Gilboa, Postlewaite and Schmeidler [43] and Gintis [46], argue that a Bayesian perspective is often too restrictive. Gilboa and Schmeidler [45], for example, suggest that an approach for prior formulation based on “empirical frequency” would be much more realistic and fruitful. Even if we can specify such priors for individuals, posteriors implied by (2) would only be reliable if these priors are reliable. This, at least in part, shifts the burden to another requirement: individual should have “reliable” priors, which as we will see in Sect. 2.3, is not a trivial requirement either.³

Despite all of these complexities and somewhat unrealistic assumptions about what individuals need to know and what they should form conjectures on, we believe that the Bayesian approach is a useful benchmark. This is not because the Bayesian approach will necessarily make good predictions. Indeed, if the degree of complexity necessary for informing Bayesian posteriors, for forming conjectures about what the underlying model of the world is and how others behave, is too high, the predictions of the Bayesian approach are unlikely to be realistic. Nevertheless, how agents learn and form their opinions when they are acting in the Bayesian manner gives us an obvious reference point to which non-Bayesian models can be compared. This is particularly important when there are several alternative non-Bayesian models for which one wishes to understand where certain predictions come from and how plausible they may be.

2.2 Non-Bayesian Approaches

Non-Bayesian approaches start by specifying simple *rules of thumb*, which are either motivated by how people behave or appear to behave in simple situations, including in laboratory experiments, or are justified because they are both simple and lead to a process of

³As is well appreciated, trying to replace such priors by something akin to Laplace's “principle of insufficient reason” is hopeless. In some situations, priors will not matter much because they will be overwhelmed by data. We will discuss below how Bayesian updating and abundant data may not be sufficient for this, however, and priors may not be overwhelmed by data in many relevant situations.

opinion formation that has desirable properties. The desirable properties might be “normative,” meaning that they lead to the formation of accurate opinions or enable agents to learn the underlying state, or they may be “positive,” meaning that opinion dynamics might match some data we have available.

The literature has considered several non-Bayesian approaches. The simplest ones, used both in economics and in other fields, involve some type of *imitation*. For instance, each individual could, with some probability, adopt the behavior or beliefs of some others she knows or has observed, or alternatively, she may use a combination of their behavior or beliefs to fashion her own. A specific example of this type of learning is the DeGroot model which we discuss in Sect. 4.1, where each individual updates her beliefs as a weighted average of the beliefs of her social neighbors, with weights given by the “trust” she has for those neighbors (see also [30–32]). Another example will be one in which individuals meet one person at a time from their social neighborhood (friends, coworkers or peers) and update their opinion to a weighted average of their initial opinion and the opinion of the person they have just met (see Sect. 4.2).

In economics, game theory and biology, several other alternatives are also used, although for space reasons, we will not be able to discuss them in great detail in the current paper. The most important idea in many of these alternatives is that individuals should change their views in the direction of beliefs and actions that have been more successful. Put in terms of the discussion of imitation above, these models still involve some amount of imitation, but observations are accompanied by some information about the performance of different alternatives. Then the alternatives that have exhibited a better performance are more likely to be imitated or receive greater weight. A particularly simple version is the so-called *replicator dynamics* inspired by evolutionary biology. In replicator dynamics, just as in genetic evolution, particular actions or beliefs that have performed better in the past are more likely to replicate and their fraction in the population increases (e.g., [79, 80, 92]). Suppose, for example, that an individual is updating her beliefs about some variable θ corresponding to whether a certain line of business is profitable. She observes the entry behavior of some past agents and also receives some information about how successful they have been. Simple imitation would involve a behavior based on the fraction of agents, among those he has observed, that have chosen to enter into this line of business. Replicator dynamics type updates, on the other hand, would involve the individual imitating one of the past agents that has been “successful”. Other ideas have been used, particularly in the context of learning in games where agents learn not only an underlying state but also the actions of other players in the game. These include fictitious play (e.g., [38, 78]) where individuals form beliefs about future play based on past patterns and regret-based update rules (e.g., [36, 37, 56]), where individuals choose actions that perform well in the sense of minimizing regret, relative to the history of play.

A more general and flexible approach is adopted in the case-based decision theory developed by Gilboa and Schmeidler [44, 45] (see also [74, 84]). In this approach, beliefs are formed according to *empirical similarity*, meaning that individuals form beliefs about a situation (decision problem) based on their experiences in similar situations in the past, where similarities are defined by means of an exogenously specified *similarity function*. The empirical similarity approach, as well as the ideas based on replicator dynamics, fictitious play and regret matching, are more flexible than the simple imitation-based rules of thumbs we will discuss below, but they have similar implications (though also important differences). In particular, the tendency toward consensus and the possibility that dispersed information may not aggregate despite consensus are common features of these approaches.

In sociology and physics, several other forms of rule-of-thumb behavior are also used. Most work in this area relies on classical models of interacting particle systems whose motivation comes from statistical mechanics (see [65] and references therein). In these models, opinions are represented by either finitely many discrete values (as in the Ising model, introduced by Glauber [48] and the voter model, introduced independently by Clifford and Sudbury [23] and by Holley and Liggett [60]), or continuous values (as in the DeGroot model [28] and the bounded confidence models of Krause [63] (further developed by Hegselmann and Krause [59], Deffuant, Neau, Amblard and Weisbuch [27] and Weisbuch, Kirman and Herreiner [93])). The opinions of the agent evolve dynamically over time as a function of their neighbors' opinions. A number of recent papers use these models to study various social phenomena, including opinion formation, information transmission, and effects of opinion leaders, using an analysis that is largely based on simulations (see [9, 35, 69, 96]).

This array of choices illustrates the limits of our knowledge in this area. Each of these rules leads to different types of behavior (and often, at least some of the implications also depend on the specific parameters one chooses within a general class of rules). It is conceivable that future work will bring additional information from empirical studies in order to eliminate some possibilities and put more structure on remaining ones. However, until this happens, there is a high degree of arbitrariness. This often motivates researchers to use either the Bayesian models or a benchmark where all of the relevant information is learned or aggregated at least in the limit (“asymptotically”) as a way of evaluating non-Bayesian models. We believe that both Bayesian and non-Bayesian approaches are useful and generate instructive insights. Which type of approach is appropriate is likely to depend on the specific question being investigated. For example, we will argue below that issues of persistent disagreements, misinformation and belief manipulation might be better analyzed using non-Bayesian approaches. In fact, it would be difficult to understand the spread of misinformation and persistent disagreements among individuals with access to similar information if we insist that all “reasonable” learning models should be strictly Bayesian and/or lead to the efficient aggregation of information.

One might still conjecture that non-Bayesian updating rules that lead to opinions that are very different from Bayesian benchmarks or do not aggregate the available information should not arise or survive. While comparison to these benchmarks is useful, deviations from either the Bayesian benchmark or the benchmark in which relevant information is aggregated effectively should not be viewed as potential shortcomings of a model (since there is no reason to expect that reality is either Bayesian or “efficient”). In fact, as we will see in the next section, aggregation of dispersed information is not even a feature of most Bayesian models (even though individuals are highly sophisticated and are assumed to have the correct model of the world).

2.3 Learning and Disagreement By Bayesian Agents

In this subsection, we first present the benchmark situation in which, in line with Savage's arguments, Bayesian updating leads to learning in a single agent problem and to consensus when several agents learn from the same or similar information, provided that some standard assumptions on priors and informativeness of signals are adopted. However, the main message of this subsection will be more nuanced. We will illustrate why the standard assumptions on priors and informativeness of signals are less innocuous than commonly presumed. We will in fact show that under reasonable assumptions, Bayesian updating will imply lack of learning and lack of consensus (even before we bring in social learning aspects). The material in this section borrows heavily from [6].

2.3.1 Bayesian Learning and Agreement

Consider two individuals, denoted by $i = 1$ and $i = 2$, who observe a sequence of signals $\{s_t\}_{t=0}^n$ where $s_t \in \{a, b\}$. Throughout this section, we assume that both agents are Bayesian. The underlying state is $\theta \in \{A, B\}$, and agent i assigns ex ante probability $\pi^i \in (0, 1)$ to $\theta = A$ (the generalization of the ideas presented here to more agents and more states is straightforward). The individuals believe that, given θ , the signals are exchangeable, i.e., they are independently and identically distributed with an unknown distribution. More specifically, the probability of $s_t = a$ given $\theta = A$ is an unknown number p_A ; likewise, the probability of $s_t = b$ given $\theta = B$ is an unknown number p_B —as shown in the following table:

	A	B
a	p_A	$1 - p_B$
b	$1 - p_A$	p_B

Standard models, for example as in [81] can be thought of as special cases of this model in which p_A and p_B are given and known. Here we will relax this assumption and assume that there is potentially some uncertainty about p_θ (where $\theta \in \{A, B\}$), and we will capture this uncertainty facing individual i by his *subjective probability distribution* denoted by \mathbb{F}_θ^i . Savage’s standard model is then the special case where \mathbb{F}_θ^i is degenerate and puts probability 1 on some \hat{p}_θ^i .

Consider any infinite sequence $s \equiv \{s_t\}_{t=1}^\infty$ of signals and write S for the set of all such sequences. The posterior belief of individual i about θ after observing the first n signals $\{s_t\}_{t=1}^n$ is given by Bayes’s rule as

$$\phi_n^i(s) \equiv \mathbb{P}^i(\theta = A \mid \{s_t\}_{t=1}^n),$$

where $\mathbb{P}^i(\theta = A \mid \{s_t\}_{t=1}^n)$ denotes the posterior probability that $\theta = A$ given a sequence of signals $\{s_t\}_{t=1}^n$ under prior π^i and subjective probability distribution \mathbb{F}_θ^i . Since the sequence of signals, s , is generated by an exchangeable process, the order of the signals does not matter for the posterior. The latter only depends on

$$r_n(s) \equiv \#\{t \leq n \mid s_t = a\},$$

i.e., on the number of times $s_t = a$ out of the first n signals. By the strong law of large numbers, $r_n(s)/n$ converges to some $\rho(s) \in [0, 1]$ almost surely, for both individuals. Defining the set

$$\bar{S} \equiv \{s \in S : \lim_{n \rightarrow \infty} r_n(s)/n \text{ exists}\}, \tag{3}$$

this observation implies that $\mathbb{P}^i(s \in \bar{S}) = 1$ for $i = 1, 2$. We will often state our results for all sample paths s in \bar{S} , which equivalently implies that these statements are true almost surely or with probability 1. Now, a straightforward application of the Bayes rule gives

$$\phi_n^i(s) = \frac{1}{1 + \frac{1 - \pi^i}{\pi^i} \frac{\mathbb{P}^i(r_n \mid \theta = B)}{\mathbb{P}^i(r_n \mid \theta = A)}}, \tag{4}$$

where $\mathbb{P}^i(r_n \mid \theta)$ is the probability of observing the signal $s_t = a$ exactly r_n times out of n signals with respect to the distribution \mathbb{F}_θ^i .

The following theorem presents a slight generalization of the standard result, for example, as formulated by Savage [81].

Theorem 1 Assume that each \mathbb{F}_θ^i puts probability 1 on \hat{p}_θ for some $\hat{p}_\theta > 1/2$, i.e., $\mathbb{F}_\theta^i(\hat{p}_\theta) = 1$ and $\mathbb{F}_\theta^i(p) = 0$ for each $p < \hat{p}_\theta$. Then, for each $i = 1, 2$:

1. There is asymptotic learning of the underlying state, in the sense that $\mathbb{P}^i(\lim_{n \rightarrow \infty} \phi_n^i(s) = 1 | \theta = A) = 1$.
2. There is asymptotic agreement between the two agents, in the sense that $\mathbb{P}^i(\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| = 0) = 1$.

This standard result states that when the individuals know the conditional distributions of the signals (and hence they agree what those distributions are), they will learn the truth from experience and observation (almost surely as $n \rightarrow \infty$) and two individuals observing the same sequence will necessarily come to agree what the underlying state, θ , is. A simple intuition for this result is that the underlying state θ is *fully identified* from the limiting frequencies, so that both individuals can infer the underlying state from the observation of the limiting frequencies of signals.

However, there is more to this theorem than this simple intuition. Each individual is sure that they will be confronted either with a limiting frequency of a signals equal to \hat{p}_A , in which case they will conclude that $\theta = A$, or they will observe a limiting frequency of $1 - \hat{p}_B$, and they will conclude that $\theta = B$; and they attach zero probability to the events that they will observe a different asymptotic frequency. What happens if an individual observes a frequency ρ of signals different from \hat{p}_A and $1 - \hat{p}_B$ in a large sample of size n ? The answer to this question will provide the intuition for some of the results that we will present next. Observe that this event has zero probability under the individual’s beliefs at the limit $n = \infty$. However, for $n < \infty$ he will assign a strictly positive (but small) probability to such a frequency of signals resulting from *sampling variation*. Moreover, it is straightforward to see that there exists a unique $\hat{\rho}(\hat{p}_A, \hat{p}_B) \in (1 - \hat{p}_B, \hat{p}_A)$ such that when $\rho > \hat{\rho}(\hat{p}_A, \hat{p}_B)$, the required sampling variation that leads to ρ under $\theta = B$ is *infinitely greater* (as $n \rightarrow \infty$) than the one under $\theta = A$. This cutoff value $\hat{\rho}(p_A, p_B)$ is clearly the solution to the equation $p_A^\rho(1 - p_A)^{1-\rho} = p_B^{1-\rho}(1 - p_B)^\rho$, given by

$$\hat{\rho}(p_A, p_B) \equiv \frac{\log(p_B/(1 - p_A))}{\log(p_B/(1 - p_A)) + \log(p_A/(1 - p_B))} \in (1 - p_B, p_A). \tag{5}$$

Consequently, when $\rho > \hat{\rho}(\hat{p}_A, \hat{p}_B)$, the individual will asymptotically assign probability 1 to the event that $\theta = A$. Conversely, when $\rho < \hat{\rho}(\hat{p}_A, \hat{p}_B)$, he will assign probability 1 to $\theta = B$.

2.3.2 Lack of Learning and Disagreement

It is clear that this theorem relies on the feature that $\mathbb{F}_\theta^i(1/2) = 0$ for each $i = 1, 2$ and each θ . This implies that both individuals attach zero probability to a range of possible models of the world—i.e., they are certain that p_θ cannot be less than $1/2$. There are two reasons for considering situations in which this is not the case. First, the preceding discussion illustrates why assigning zero probability to certain models of the world is important; it enables individuals to ascribe any frequency of signals that are unlikely under these models to sampling variability. This kind of inference may be viewed as somewhat unreasonable, since individuals are reaching very strong conclusions based on events that have vanishingly small probabilities (since sampling variability vanishes as $n \rightarrow \infty$). Second, once we take into account uncertainty about the underlying models that individuals have, it may also be natural to allow them to attach positive (albeit small) probabilities to all possible values of p_θ . This latter feature will lead to very different consequences as shown by the next theorem.

Theorem 2 *Suppose that for each i and θ , \mathbb{F}_θ^i has a continuous, non-zero and finite density f_θ^i over $[0, 1]$. Then,*

1. *There is no asymptotic learning, i.e., $\mathbb{P}^i(\lim_{n \rightarrow \infty} \phi_n^i(s) \neq 1 | \theta = A) = 1$ for $i = 1, 2$.*
2. *There is no asymptotic agreement between the two agents, i.e., $\mathbb{P}^i(\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| \neq 0) = 1$ whenever $\pi^1 \neq \pi^2$ and $F_\theta^1 = F_\theta^2$ for each $\theta \in \{A, B\}$.*

The proof of this theorem is given in Acemoglu, Chernozhukov, and Yildiz [6].

Let us also provide a brief sketch of the arguments leading to the main results of no learning and asymptotic agreement. When \mathbb{F}_θ^i has a continuous, non-zero and finite density f_θ^i over $[0, 1]$, it can be shown that for $s \in \bar{S}$,

$$\phi_\infty^i(\rho(s)) \equiv \lim_{n \rightarrow \infty} \phi_n^i(s) = \frac{1}{1 + \frac{1-\pi^i}{\pi^i} R^i(\rho(s))}, \tag{6}$$

where $\rho(s) = \lim_{n \rightarrow \infty} r_n(s)/n$, and for all $\rho \in [0, 1]$,

$$R^i(\rho) \equiv \frac{f_B^i(1 - \rho)}{f_A^i(\rho)} \tag{7}$$

is the asymptotic likelihood ratio. From this result, one can prove Theorem 2 readily, as $R^i(\rho)$ does not vanish or diverge to infinity: any asymptotic frequency of different signals can be generated both under $\theta = A$ and under $\theta = B$. This implies that some residual uncertainty will remain, and priors will be important in shaping asymptotic beliefs.

In fact, in the presence of learning under this type of uncertainty, two Bayesian agents may end up disagreeing more after receiving exactly the same infinite sequence of signals—again a feature that it is never possible in the standard model (without uncertainty). We state this result in the following theorem.

Theorem 3 *Suppose that for each i and θ , \mathbb{F}_θ^i has a continuous, non-zero and finite density f_θ^i over $[0, 1]$ and that there exists $\epsilon > 0$ such that $|R^1(\rho) - R^2(\rho)| > \epsilon$ for each $\rho \in [0, 1]$. Then, there exists an open set of priors π^1 and π^2 , such that for all $s \in \bar{S}$,*

$$\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| > |\pi^1 - \pi^2|;$$

in particular,

$$\mathbb{P}^i \left(\lim_{n \rightarrow \infty} |\phi_n^1(s) - \phi_n^2(s)| > |\pi^1 - \pi^2| \right) = 1.$$

Intuitively, even a small difference in priors ensures that individuals will interpret signals differently, and if the original disagreement is relatively small, after almost all sequences of signals, the disagreement between the two individuals grows. Consequently, the observation of a common sequence of signals causes an initial difference of opinion between individuals to widen (instead of the standard merging of opinions under certainty). Theorem 3 also shows that both individuals are certain ex ante that their posteriors will diverge after observing the same sequence of signals, because they understand that they will interpret the signals differently. This result strengthens the general message that under rich enough priors, Bayesian updating does not guarantee learning or consensus. Instead, for some priors individuals will “agree to eventually disagree even more”.

Acemoglu, Chernozhukov, and Yildiz [6] in fact establish a much stronger result: agreement by two Bayesian agents is *fragile* in the sense that if we take the limit in the above model such that \mathbb{F}_θ^i 's become Dirac (concentrated around single points), then even though this limit is arbitrarily close to the standard case where asymptotic agreement is guaranteed, Bayesian updating does not necessarily lead to agreement. In fact, there necessarily exist some priors such that the limiting sequence of beliefs (as \mathbb{F}_θ^i 's become Dirac) differs from limiting beliefs (which do involve consensus). These results, therefore, strengthen the points already highlighted in this section that Bayesian updating gets a lot of mileage by restricting priors. When these types of restrictions on priors are reasonable and the cognitive requirements of Bayesian updating are not excessive in terms of complex computations and reasoning, Bayesian approaches to learning may be a natural starting point. Else, it may have more limited relevance for real world learning problems.⁴

3 Bayesian Social Learning In Networks

In the previous section we saw how both Bayesian and non-Bayesian approaches to opinion formation have strengths and shortcomings. We therefore think that there is much to learn from both types of approaches and from their contrast. In this and the next section, we will outline several mathematical models of opinion formation and dynamics. In this section, we focus on Bayesian models. We start with Bayesian models of *observational learning*, meaning models in which individuals learn from the observation of others' actions. We then turn to models of communication learning, where individuals learn through communication. In both cases our focus will be on interactions in the context of relatively general social networks and on understanding what classes of learning problems and what types of social networks will lead to "good outcomes," that is, to an equilibrium in which the social group will be able to aggregate the available dispersed information. In the next section, we turn to non-Bayesian models.

We will summarize the relevant literature in the context of the models we present below. But a high-level summary of the general lessons from these models is useful. We will find that even in very simple environments, Bayesian learning need not lead to the formation of accurate opinions and beliefs that aggregate all of the dispersed information, though it does typically lead to consensus or "quasi-consensus" (meaning only small differences in opinion remaining among the agents). The structure of the network—which encodes information on such things as who observes whose actions and who communicates with whom—influences opinion dynamics, though some of the effects of network structure are on the speed of learning rather than on asymptotic outcomes, and the analysis of speed of learning is much more challenging and is generally an open question for future research.

3.1 Bayesian Observational Learning

We start with models of Bayesian learning from observations of past actions.

⁴See also [26] on how common knowledge might fail to arise from Bayesian observations (our discussion here emphasizes lack of common priors), and [54] for another critique of the common priors.

3.1.1 Problem

Consider the following problem first studied by Banerjee [11] and Bikhchandani, Hirshleifer, and Welch [13]. A countably infinite number of agents (individuals), indexed by $n \in \mathbb{N}$, sequentially make a single decision each. The payoff of agent n depends on an underlying, payoff-relevant state of the world, θ , and on her decision. Learning and opinion formation in this society will be about θ .

Suppose, for simplicity, that both the underlying state and decisions are binary. In particular, the decision of agent n is denoted by $x_n \in \{0, 1\}$ and the underlying state is taken to be $\theta \in \{0, 1\}$. The payoff of agent n is

$$u_n(x_n, \theta) = \begin{cases} 1 & \text{if } x_n = \theta, \\ 0 & \text{if } x_n \neq \theta. \end{cases} \quad (8)$$

To simplify notation, we assume that both values of the underlying state are equally likely, so that $\mathbb{P}(\theta = 0) = \mathbb{P}(\theta = 1) = 1/2$.

As a specific example, $\theta = 1$ might denote whether a particular product is high quality or whether a certain line of business is profitable (but may also correspond to more abstract issues such as whether a particular ideology provides the right perspective on political, social or economic events). The variable x would then correspond to purchasing or entry decisions (or to whether an individual subscribes to a certain ideology).

Assume also that each individual receives an independent binary signal $s \in \{0, 1\}$ such that $s = \theta$ has probability $q > 1/2$. The signals are *private information* and are thus referred to as *private signals*. In addition, each individual observes all past actions (i.e., individual n observes all of x_1, \dots, x_{n-1} in addition to s_n). Since this is a dynamic game of incomplete information, we focus on the standard equilibrium concept (weak) *Perfect Bayesian Equilibrium*, PBE, which simply requires that each individual chooses the action that maximizes their utility given their beliefs and beliefs are formed by Bayesian updating which correctly conditions on the equilibrium strategies of all other agents (e.g., [39]). There are two important features of this equilibrium concept, which are entirely standard in game theory: first, there will be Bayesian updating (this is similar to the example of Bayesian updating we have already discussed in the previous section). Second, again as we emphasized in the previous section, each individual will understand and correctly condition on the equilibrium strategies of other agents. The second feature, though simple, is quite important and subtle: individuals will not simply look at what actions others choose, but will try to infer what their signals must have been from their actions (based on their understanding or views about others' strategies). This clearly requires more complex reasoning than just Bayesian updating.

The key result in [13] and in [11] is the following striking *herding* result. Focus on an equilibrium in which the first two individuals choose the action in line with their signal (i.e., for $n = 1, 2$, $x_n = 1$ if and only if $s_n = 1$). First observe that this behavior is in fact a PBE. Indeed, it is a strict best response for the first agent to choose $x_n = s_n$ given that the signals are informative ($q > 1/2$) and the prior is that the two states are equally likely to start with. What about the second agent? Suppose that she observes $x_1 = 1$. Given the behavior of the first agent (and here conditioning on the strategies of others is crucial), she can infer perfectly that the first agent must have received a signal of $s_1 = 1$. If in addition she also receives a signal of $s_2 = 1$, it is again a strict best response for her to choose $x_2 = 1$. What if she receives a signal of $s_2 = 0$. Now given that she knows that $s_1 = 1$ and since $s = \theta$ has probability $q > 1/2$ regardless of the underlying state θ , the posterior of this agent, defined

as the probability that θ is equal to 1 given his information, will be exactly the same as the initial prior, i.e., $1/2$. Therefore, it is a weak best response for her to choose $x_2 = 0$ if $s_2 = 0$. Thus we have verified that there exists a PBE in which both of the first two agents follow their signals.

What about agent $n = 3$? She has observed $x_1 = x_2 = 1$ and suppose that she receives a signal of $s_3 = 0$. Conditioning on the equilibrium in which both agents 1 and 2 make choices in line with their signals, she knows (i.e., she can infer from their behavior) that $s_1 = s_2 = 1$. She has effectively access to three signals $s_1 = s_2 = 1$ (from the behavior of the first two agents) and her private signal, $s_3 = 0$. But then, given the signal structure, her posterior will be strictly less than the initial prior, $1/2$, for any $q > 1/2$. This means that even if $s_3 = 0$, she will choose $x_3 = 1$. Of course, if $s_3 = 1$, she will choose $x_3 = 1$ a fortiori. Hence we have obtained the result that it is a strict best response for her to herd on the behavior of the first two agents, entirely ignoring her own information.⁵

What about agent $n = 4$? Again, conditioning correctly on others' strategies, he will infer from $x_1 = x_2 = 1$ that $s_1 = s_2 = 1$, and also understand that agent 3 has ignored her information. But this means that he is in exactly the same situation as agent 3, and in the same way as it was a strict best response for $n = 3$ to herd, it is a strict best response for $n = 4$ to do so. By induction, the same conclusion applies to any $n \geq 4$. We have thus obtained the result that following $x_1 = x_2 = 1$, all agents will ignore their private signals and they will all choose the same action, $x = 1$. Clearly, $x_1 = x_2 = 1$ is entirely consistent with $\theta = 0$. In particular, it can happen with probability $(1 - q)^2$ when $\theta = 0$, which can be quite close to $1/2$. This implies that a *mistaken herd*, in which all agents choose the incorrect action and there is almost no aggregation of the dispersed private information, can happen with reasonably high probability (and it is straightforward to see that a similar herd also arises if three of the first four agents, or five of the first eight agents and so on, choose $x = 1$). This example therefore sharply illustrates how dispersed information will fail to be aggregated even though all agents would like to know this information and are all acting in a fairly sophisticated Bayesian (and in fact game theoretic) manner. It is also important to emphasize another feature here: it is no coincidence that a mistaken herd cannot happen with probability greater than $1/2$. Since agents are Bayesian and have the correct model of the world, their beliefs at any point in time are "accurate," meaning that if they believed the true state to be $\theta = 1$ with probability p , then this is indeed the probability that an outside observer with access to exactly the same data would assign to the event that $\theta = 1$. But since individuals will only choose action $x = 1$ if they believe $\theta = 1$ is more likely than $\theta = 0$, they can never be more likely to make a mistake than taking the correct action. This fact, though simple, has important implications, particularly concerning the possibility of misinformation and indoctrination. There will be limits to how much misinformation and belief manipulation there can be in a Bayesian world, at least unless we relax the assumption that individuals start with a reliable model of the world (e.g., with accurate priors). Motivated by this feature, in the next section we will use non-Bayesian approaches to develop our benchmark models of spread of misinformation and belief manipulation.

3.1.2 Related Literature

There is now a sizable literature building on the insights we have just discussed characterizing the conditions under which there is aggregation of dispersed information through

⁵The same argument also shows that herding would occur even if agent 2 mixed between $x_2 = 0$ and $x_2 = 1$ conditional on $s_2 = 0$ and $x_1 = 1$. Thus our specific selection of equilibrium does not affect the conclusion.

a process of social learning. Most relevant for this subsection are the works focusing on Bayesian updating, game theoretic analysis and observational learning. This literature is also voluminous and includes, among others, the important work by Smith and Sørensen [86], who generalize the environment we have just outlined to include a richer set of private signals. They define the notion of unbounded private signals, which will be introduced below, and show that with *unbounded private signals*, which make it possible for individuals to receive highly informative signals, the type of herding we just identified does not happen. However, with bounded private signals, it is a robust phenomenon which necessarily occurs.

Welch [94], Lee [64], Chamley and Gale [21], and Vives [90] consider various other extensions. These papers maintain the assumption that all past actions are observed. Using language from the analysis of networks, we can say that they focus on the *full observation network topology*. Two papers that have gone beyond this “full observation network topology” are [12, 85]. Both of these papers study social learning with random sampling of past actions.

Thus much of the literature has not tackled the issue of learning taking place in the context of a social network, where each individual only observes the actions of a subset of the other agents (those with whom he is more *connected* through friendship, work or other social engagements). We will next present the model and the analysis from [5], which introduces a network structure to capture these interactions (as well as allowing more general signal structures). Allowing for a network structure leads to several challenges. The first one is that, in contrast to models with *full observation network topology*, one cannot use martingale properties, which imply that the expectation of future beliefs conditional on publicly available information is equal to current beliefs. The second is that the inference problem that individuals face is much more complex, because a given action might be driven not by the private signals that an individual observes, but by the actions that they have observed from the past (and we may be uncertain as to what these actions are or about whether or not the individual in question has indeed observed others).

In the analysis that follows, we will maintain three assumptions that are worth highlighting. First, agents will be assumed to be Bayesian and understand the structure of the game and the strategies of others. Relaxing this assumption is important and some ways of doing so will be considered in the next section when we discuss non-Bayesian models. Second, we will assume that there is no heterogeneity in the preferences of the agents in this social network, i.e., they all have preferences given by (8). Acemoglu, Dahleh, Lobel, and Ozdaglar [4] extend this analysis to incorporate heterogeneity in the preferences toward different types of actions. Finally, we will maintain the assumption that each individual takes only a single decision. Repeated decisions introduce additional challenges. The most important one of those is that agents will no longer choose the action that maximizes their static payoff, but may choose to *strategically experiment* in order to induce others to choose actions that will be more informative in the future. The analysis of strategic experimentation is unfortunately very challenging. Bala and Goyal [10] and Gale and Kariv [40] study models of social learning with repeated interactions, but both of these of papers ignore the strategic experimentation issue ([10] in fact use a non-Bayesian model of learning, though their approach has many parallels with Bayesian models).

3.1.3 Learning in Social Networks

We now outline the model and the results from [5]. We continue to assume that the unknown state θ is binary, $\{0, 1\}$ (and that $\mathbb{P}(\theta = 0) = \mathbb{P}(\theta = 1) = 1/2$). Each agent $n \in \mathbb{N}$ still has utility given by (8) and forms beliefs about this state from a *private signal* $s_n \in \bar{\mathcal{S}}$ (where $\bar{\mathcal{S}}$

is a metric space or simply a Euclidean space) and from his observation of the actions of other agents. Conditional on the state of the world θ , the signals are independently generated according to a probability measure \mathbb{F}_θ . We refer to the pair of measures $(\mathbb{F}_0, \mathbb{F}_1)$ as the *signal structure* of the model. We assume that \mathbb{F}_0 and \mathbb{F}_1 are *absolutely continuous* with respect to each other, which immediately implies that no signal is fully revealing about the underlying state. We also assume that \mathbb{F}_0 and \mathbb{F}_1 are not identical, so that some signals are *informative*.

As already indicated, in contrast to the rest of the literature on social learning discussed above, we assume that agents do not necessarily observe all previous actions. Instead, they observe the actions of other agents according to the structure of a *social network*. To introduce the notion of a social network, let us first define a *neighborhood*. Each agent n observes the decisions of the agents in his (stochastically-generated) neighborhood, denoted by $B(n)$. Since agents can only observe actions taken previously, $B(n) \subseteq \{1, 2, \dots, n - 1\}$. Each neighborhood $B(n)$ is generated according to an arbitrary probability distribution \mathbb{Q}_n over the set of all subsets of $\{1, 2, \dots, n - 1\}$. We impose no special assumptions on the sequence of distributions $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ except that the draws from each \mathbb{Q}_n are independent from each other for all n and from the realizations of private signals. The sequence $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ is the *network topology* of the social network formed by the agents. The network topology is common knowledge, whereas the realized neighborhood $B(n)$ and the private signal s_n are the private information of agent n . We say that $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ is a *deterministic* network topology if the probability distribution \mathbb{Q}_n is a degenerate (Dirac) distribution for all n . Otherwise, $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ is a *stochastic* network topology. A *social network* consists of a network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ and a signal structure $(\mathbb{F}_0, \mathbb{F}_1)$.

Notice that our framework is general enough to nest the majority of social network models studied in the literature, including the popular preferential attachment and small-world network structures. For example, the preferential attachment model can be nested by choosing a stochastic network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ with a collection of subsets S_1, \dots, S_k of agents such that agents in S_1 have a very high probability of being in each $B(n)$, those in S_2 also have a high probability, but lower than the corresponding probability for those in S_1 , of being in each $B(n)$, and so on. The small-world network structure can be nested by choosing a partition $\{S_j\}$ of \mathbb{N} such that for each $n \in S_j$, the probability that any agent m in S_j with $m < n$ is also in $B(n)$ is very high, while the probability that an agent m who is not in S_j is in $B(n)$ is low but positive. More generally, any network structure can be represented by a judicious choice of $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ provided that we keep the assumption that the realizations of $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ are independent, which is adopted to simplify the analysis. The independence assumption on the neighborhoods does not impose a restriction on the degree distribution (cardinality) of the agents nor on their degree of clustering.

We next introduce the definitions of equilibrium and asymptotic learning, and we provide a characterization of equilibrium strategies. In particular, we show that equilibrium decision rules of individuals can be decomposed into two parts: one that only depends on an individual’s private signal, and the other that is a function of the observations of past actions. We also show why a full characterization of individual decisions is non-trivial and motivate an alternative proof technique, relying on developing bounds on improvements in the probability of the correct decisions, that will be used in the rest of our analysis.

Given the description above, it is evident that the information set I_n of agent n is given by her signal s_n , her neighborhood $B(n)$, and all decisions of agents in $B(n)$, that is,

$$I_n = \{s_n, B(n), x_k \text{ for all } k \in B(n)\}. \tag{9}$$

The set of all possible information sets of agent n is denoted by \mathcal{I}_n . A strategy for individual n is a mapping $\sigma_n : \mathcal{I}_n \rightarrow \{0, 1\}$ that selects a decision for each possible information set.

A strategy profile is a sequence of strategies $\sigma = \{\sigma_n\}_{n \in \mathbb{N}}$. We use the standard notation $\sigma_{-n} = \{\sigma_1, \dots, \sigma_{n-1}, \sigma_{n+1}, \dots\}$ to denote the strategies of all agents other than n and also (σ_n, σ_{-n}) for any n to denote the strategy profile σ . Given a strategy profile σ , the sequence of decisions $\{x_n\}_{n \in \mathbb{N}}$ is a stochastic process and we denote the measure generated by this stochastic process by \mathbb{P}_σ .

Definition 1 A strategy profile σ^* is a pure-strategy Perfect Bayesian Equilibrium of this game of social learning if for each $n \in \mathbb{N}$, σ_n^* maximizes the expected payoff of agent n given the strategies of other agents σ_{-n}^* .

Given a strategy profile σ , the expected payoff of agent n from action $x_n = \sigma_n(I_n)$ is simply $\mathbb{P}_\sigma(x_n = \theta \mid I_n)$. Therefore, for any equilibrium σ^* , we have

$$\sigma_n^*(I_n) \in \operatorname{argmax}_{y \in \{0,1\}} \mathbb{P}_{(y, \sigma_{-n}^*)}(y = \theta \mid I_n). \tag{10}$$

We denote the set of equilibria of the game by Σ^* . It is clear that Σ^* is nonempty. Given the sequence of strategies $\{\sigma_1^*, \dots, \sigma_{n-1}^*\}$, the maximization problem in (10) has a solution for each agent n and each $I_n \in \mathcal{I}_n$. Proceeding inductively, and choosing either one of the actions in case of indifference determines an equilibrium.

Our main focus is whether equilibrium behavior will lead to information aggregation. This is captured by the notion of *asymptotic learning*, which is introduced next.

Definition 2 Given a signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ and a network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$, we say that asymptotic learning occurs in equilibrium σ if x_n converges to θ in probability (according to measure \mathbb{P}_σ), that is,

$$\lim_{n \rightarrow \infty} \mathbb{P}_\sigma(x_n = \theta) = 1.$$

Notice that asymptotic learning requires that the probability of taking the correct action converges to 1.⁶ Therefore, asymptotic learning will fail when, as the network becomes large, the limit inferior of the probability of all individuals taking the correct action is strictly less than 1. The following proposition characterizes optimal decisions by (Bayesian) individuals as a function of their observations and signal. The proof of this proposition, like the other ones in this subsection, is omitted and can be found in [5].

Proposition 1 Let $\sigma \in \Sigma^*$ be an equilibrium of the game. Let $I_n \in \mathcal{I}_n$ be an information set of agent n . Then, the decision of agent n , $x_n = \sigma(I_n)$, satisfies

$$x_n = \begin{cases} 1, & \text{if } \mathbb{P}_\sigma(\theta = 1 \mid s_n) + \mathbb{P}_\sigma(\theta = 1 \mid B(n), x_k, k \in B(n)) > 1, \\ 0, & \text{if } \mathbb{P}_\sigma(\theta = 1 \mid s_n) + \mathbb{P}_\sigma(\theta = 1 \mid B(n), x_k, k \in B(n)) < 1, \end{cases}$$

and $x_n \in \{0, 1\}$ otherwise.

This proposition establishes an additive decomposition in the equilibrium decision rule between the information obtained from the private signal of the individual and from the

⁶It is also clear that asymptotic learning is equivalent to the posterior beliefs converging to a distribution putting probability 1 on the true state.

observations of others’ actions (in his neighborhood). The next definition formally distinguishes between the two components of an individual’s information.

Definition 3 We refer to the probability $\mathbb{P}_\sigma(\theta = 1 \mid s_n)$ as the private belief of agent n , and the probability

$$\mathbb{P}_\sigma(\theta = 1 \mid B(n), x_k \text{ for all } k \in B(n)),$$

as the social belief of agent n .

Proposition 1 and Definition 3 imply that the equilibrium decision rule for agent $n \in \mathbb{N}$ is equivalent to choosing $x_n = 1$ when the sum of his private and social beliefs is greater than 1. Consequently, the properties of private and social beliefs will shape equilibrium learning behavior.

The private belief of an individual is a function of his private signal $s \in S$ and is not a function of the strategy profile σ since it does not depend on the decisions of other agents. We represent probabilities that do not depend on the strategy profile by \mathbb{P} . We use the notation p_n to represent the private belief of agent n , i.e.,

$$p_n = \mathbb{P}(\theta = 1 \mid s_n).$$

A straightforward application of Bayes’ rule implies that for any n and any signal $s_n \in \bar{S}$, the private belief p_n of agent n is given by

$$p_n = \left(1 + \frac{d\mathbb{F}_0}{d\mathbb{F}_1}(s_n) \right)^{-1}. \tag{11}$$

We next define the support of a private belief. The support of private beliefs plays a key role in asymptotic learning behavior. Since the p_n are identically distributed for all n (which follows by the assumption that the private signals s_n are identically distributed), in the following we will use agent 1’s private belief p_1 to define the support and the conditional distributions of private beliefs,

Definition 4 The support of the private beliefs is the interval $[\underline{\beta}, \bar{\beta}]$, where the end points of the interval are given by

$$\begin{aligned} \underline{\beta} &= \inf\{r \in [0, 1] \mid \mathbb{P}(p_1 \leq r) > 0\}, \quad \text{and} \\ \bar{\beta} &= \sup\{r \in [0, 1] \mid \mathbb{P}(p_1 \leq r) < 1\}. \end{aligned}$$

The signal structure has bounded private beliefs if $\underline{\beta} > 0$ and $\bar{\beta} < 1$ and unbounded private beliefs if $\underline{\beta} = 1 - \bar{\beta} = 1$.

When private beliefs are bounded, there is a maximum informativeness to any signal. When they are unbounded, agents may receive arbitrarily strong signals favoring either state (this follows from the assumption that $(\mathbb{F}_0, \mathbb{F}_1)$ are absolutely continuous with respect to each other).

The conditional distribution of private beliefs given the underlying state $j \in \{0, 1\}$ can be directly computed as

$$\mathbb{G}_j(r) = \mathbb{P}(p_1 \leq r \mid \theta = j). \tag{12}$$

The signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ can be equivalently represented by the corresponding *private belief distributions* $(\mathbb{G}_0, \mathbb{G}_1)$, and in what follows, it will typically be more convenient to work with $(\mathbb{G}_0, \mathbb{G}_1)$ rather than $(\mathbb{F}_0, \mathbb{F}_1)$. It is straightforward to verify that $\mathbb{G}_0(r)/\mathbb{G}_1(r)$ is non-increasing in r and $\mathbb{G}_0(r)/\mathbb{G}_1(r) > 1$ for all $r \in (\underline{\beta}, \bar{\beta})$.

We also introduce a key property of network topologies and signal structures that impacts asymptotic learning. Intuitively, for asymptotic learning to occur, the information that each agent receives from other agents should not only come from a bounded subset of agents. This property is established in the following definition. For this definition and throughout the paper, if the set $B(n)$ is empty, we set $\max_{b \in B(n)} b = 0$.

Definition 5 The network topology has expanding observations if for all $K \in \mathbb{N}$, we have

$$\lim_{n \rightarrow \infty} \mathbb{Q}_n \left(\max_{b \in B(n)} b < K \right) = 0.$$

If the network topology does not satisfy this property, then we say it has non-expanding observations.

Recall that the neighborhood of agent n is a random variable $B(n)$ (with values in the set of subsets of $\{1, 2, \dots, n - 1\}$) and distributed according to \mathbb{Q}_n . Therefore, $\max_{b \in B(n)} b$ is a random variable that takes values in $\{0, 1, \dots, n - 1\}$. The expanding observations condition can be restated as the sequence of random variables $\{\max_{b \in B(n)} b\}_{n \in \mathbb{N}}$ converging to infinity in probability. Similarly, it follows from the preceding definition that the network topology has non-expanding observations if and only if there exists some $K \in \mathbb{N}$ and some scalar $\epsilon > 0$ such that

$$\limsup_{n \rightarrow \infty} \mathbb{Q}_n \left(\max_{b \in B(n)} b < K \right) \geq \epsilon.$$

An alternative restatement of this definition might clarify its meaning. Let us refer to a finite set of individuals C as *excessively influential* if there exists a subsequence of agents who, with probability uniformly bounded away from zero, observe the actions of a subset of C . Then, the network topology has non-expanding observations if and only if there exists an excessively influential group of agents. Note also that if there is a minimum amount of arrival of new information in the network, so that the probability of an individual observing some other individual from the recent past goes to one as the network becomes large, then the network topology will feature expanding observations. This discussion therefore highlights that the requirement that a network topology has expanding observations is quite mild and most social networks, including all of those discussed above, satisfy this requirement.

Acemoglu, Dahleh, Lobel, and Ozdaglar [5] establish the following characterization results for asymptotic learning by this social network (society) of agents.

Theorem 4 Assume that the signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ has unbounded private beliefs and the network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ has expanding observations. Then, asymptotic learning occurs in every equilibrium $\sigma \in \Sigma^*$.

Theorem 5 There exists no equilibrium $\sigma \in \Sigma^*$ with asymptotic learning if either:

1. the network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ has non-expanding observations; or
2. the signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ has bounded private beliefs and the network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ satisfies one of the following three conditions:

- (a) $B(n) = \{1, \dots, n - 1\}$ for all n ,
- (b) $|B(n)| \leq 1$ for all n , or
- (c) there exists some constant M such that $|B(n)| \leq M$ for all n and

$$\lim_{n \rightarrow \infty} \max_{b \in B(n)} b = \infty \quad \text{with probability 1.}$$

Several points are worth noting. First, Theorem 4 is surprising. Each agent is solving a complex updating problem, and may directly have access to a very limited amount of information. For instance, an example of a network topology with expanding observations is one in which each agent observes only one other agent from the past (which may be some recent neighbor or somebody randomly drawn from those that have already acted). It might appear that the information transmitted by the action of this one agent is quite limited; not only is this just a single agent rather than a collection of many agents as much of the previous literature assumes, but also it will not be known who this agent himself has observed from the past, and thus it is difficult to infer this individual’s own signal from his action. The result in the theorem follows nonetheless because of a form of an *improvement principle*, whereby each agent can do as well as the one individual she observes by just copying him. In fact, with unbounded private signals, she will strictly do better than copying him by following her signals when they are “extreme”. This reasoning suggests that for any sequence of agents, expected payoffs are strictly increasing along the sequence, which, combined with the expanding observations assumption, ensures that this improvement sequence will continue until there is convergence to the right decision with probability 1, implying asymptotic learning. The role of the unbounded private signals in this result is related to but different from that in [86]. In Smith and Sørensen, the social belief is a martingale, and unbounded private signals ensure that this martingale cannot converge anywhere but to the correct belief. Here, in contrast, the unbounded private belief ensures that the above-mentioned improvement principle must be strict.

Second, first part of Theorem 5 shows that there are also several network structures that preclude learning even with unbounded private signals. These are network structures in which there are *excessively influential* agents who are the only source of all of the information for a non-trivial subset of society. While mathematically not surprising, this result is still useful in highlighting that social learning of the underlying state, and thus aggregation of dispersed information, is only possible if all but a trivial fraction of the agents have access to new information, meaning information coming from “recent” actions.

Third, Theorem 5 also shows that for many commonly-studied network topologies, bounded private signals preclude learning of the underlying state. Note that bounded private signals do not mean uninformative signals. The support $(\underline{\beta}, \bar{\beta})$ defined above could be quite close to $(0, 1)$. And there is always sufficient (dispersed) information to reveal the underlying state θ when a sufficiently large number of agents are considered. Therefore, the failure of the equilibrium to aggregate this dispersed information is not a forgone conclusion. It is, however, the equilibrium outcome because each agent stops using his own information, thus free riding on the information revealed by past actions, before there is sufficient aggregation of dispersed information on the underlying state. Thus there is a herding-related reason why social learning does not take place. This is also the reason why in many situations with bounded private beliefs the equilibrium not only involves lack of social learning, but also convergence to a consensus or more appropriately to a quasi-consensus, whereby beliefs across individuals are very similar (sufficiently similar so that they take the same action).

Fourth, Theorem 5 does not, however, state that social learning will necessarily fail with bounded private signals. In fact, this is not true, and Acemoglu, Dahleh, Lobel, and Ozdaglar

[5] identify a range of network topologies where social learning takes place even though the information structure may involve bounded private signals (and there are other structures where social learning takes place in some equilibria but not in others). In particular, Acemoglu, Dahleh, Lobel, and Ozdaglar [5] show that in a class of stochastic network topologies where a subset of the agents do not receive sufficient information from others and are thus forced to use their own signals and the remaining agents observe individuals from this subset with sufficiently high probability, social learning always takes place. This result illustrates the importance of the network topology, and in fact its specific stochastic properties, for long-run learning outcomes.

Fifth, despite the just-mentioned result on learning with bounded private signals, which illustrates how learning depends on the specific network topology, the network topology appears relatively unimportant in Theorems 4 and 5, beyond the issue of expanding vs. non-expanding observations. This is to some degree a shortcoming of the model, since informal intuition and some empirical work indicate that network interactions are important for information dissemination. In fact, this shortcoming is in part a consequence of our focus on long-run learning outcomes. It can be shown that the speed of learning is very different between network structures that allow social learning (with unbounded private signals). For example, when each individual observes only one person from the past, the speed of learning is much greater when this person is a recent neighbor than when he is chosen from the past (uniformly) randomly. Thus in this context, network effects might exhibit themselves more in the speed of learning than on whether there is learning asymptotically. However, a systematic analysis of the issue of speed of learning in general networks is much more challenging.

Sixth, social networks play an important role in opinion formation in practice because we tend to talk to people in our social network. This type of learning is not well captured by observational models, and instead requires a model of communication. This will be our next task.

Finally, returning to some of the themes raised in the previous section, Bayesian updating here is quite complex. It requires not only a standard application of Bayes rule, but also necessitates each individual to correctly conjecture what types of information each agent he has observed will have received (e.g., which ones of the many past actions he has observed, and which actions have the individuals observed by this agent observed themselves, and so on). This observation highlights that models of Bayesian observation learning in social networks might require a very high degree of sophistication from the agents and suggests that incorporating some aspects of non-Bayesian updating in this context might be necessary. This is a topic we will take up in Sect. 4.

3.2 Bayesian Communication Learning

3.2.1 Problem

Consider next a similar environment where information concerning an underlying state θ is held dispersely and agents will make a decision on the basis of θ . However, learning will not arise from observation of others' actions, but through communication. The difficulty with information aggregation in models of observational learning resulted from the selfish interests of the agents: they ignore their impact on others' learning. With communication, selfish interests are again a problem, but in a different way. If all of the agents have common interests and could communicate costlessly and rapidly, we would expect that information sharing would occur and this is the basis of Condorcet's Jury Theorem. There are two problems preventing this strong conclusions from emerging, however. First, not all agents may

have the same interests. If an agent suspects that another one might try to mislead him or her, communication will be hampered. In the game theory literature, this issue is typically addressed by using the so-called “cheap talk” models first proposed by Crawford and Sobel [25]. In these models, a sender, who has some information, will send a message to a receiver, who will use this information to make a decision. The problem is that the sender and the receiver may have different interests, and thus the sender might try to influence the receiver’s decision by communicating incorrect information. Crawford and Sobel [25], and the literature that follows them, characterizes the amount of information that can be transmitted in such an interaction (see, e.g., [22, 34, 70]).

Second, communication, even without strategic interactions, is time-consuming. An individual might prefer to make a decision sooner rather than later, even if this involves receiving less information. But in doing so, the individual will also be affecting the information that others obtain. Thus, selfish behavior, now in the form of trying to reach a decision quickly, will again affect the information that is available to others in a society or in a social network.

We next present a model based on Acemoglu, Bimpikis, and Ozdaglar [3] designed to address these issues. The literature on learning by communication in groups is somewhat smaller than the observational learning literature (there are several non-Bayesian models of communication, which we will discuss in the next section). Most closely related to the model we will present here are [41, 55], which look at one-shot cheap talk games superimposed over a network of interacting agents. Their main focus is on the conditions under which there will be truthful communication in a given network. Though they represent important advances in our understanding of the limits of communication in the context of social groups and networks, these papers do not provide an analysis of the aggregation of dispersed information, which is our focus here.

3.2.2 Model


We next study a simple model of learning with communication based on [3]. We consider n agents situated in a communication network represented by a directed graph $G^n = (\mathcal{N}^n, \mathcal{E}^n)$, where $\mathcal{N}^n = \{1, \dots, n\}$ is the set of agents and \mathcal{E}^n is the set of directed edges with which agents are linked. Agents make decisions over time and the payoff to each agent depends on her decision and an underlying unknown *state of the world* θ , which we assume to be binary, i.e., $\theta \in \{0, 1\}$. For simplicity, we again assume that both values of θ are equally likely, i.e., $\mathbb{P}(\theta = 0) = \mathbb{P}(\theta = 1) = 1/2$.

Agent i forms beliefs about θ from a *private signal* $s_i \in S_i$ (where S_i is a Euclidean space), as well as from information she obtains from other agents through the network G^n . At each time period, $t = 0, 1, \dots$, agent i can decide to take an irreversible *action*, 0 or 1, or wait for another time period. Her payoff is thus

$$u_i^n(x_i^n, \theta) = \begin{cases} \delta^t \pi & \text{if } x_{i,t}^n = \theta \text{ and } x_{i,t}^n = \text{“wait” for } t < \tau, \\ 0 & \text{otherwise,} \end{cases}$$

where $x_i^n = [x_{i,t}^n]_{t=0,1,\dots}$ denotes the sequence of agent i ’s actions ($x_{i,t}^n \in \{\text{“wait”}, 0, 1\}$). Here, $x_{i,t}^n = 0$ or 1 denotes agent i taking action 0 or 1 respectively, while “wait” designates the agent deciding to wait for that time period without taking an action; $\pi > 0$ is the payoff from the correct action. Without loss of generality, we normalize π to be equal to 1. We say that the agent “exits”, if she chose to take action 0 or 1. The discount factor $\delta \in (0, 1)$ implies that an earlier exit is preferred to a later one.

Fig. 1 The information set of Agent 1 under truthful communication



We say that *agent i sends information (or a message) to agent j* (or equivalently *j receives information from i*) if there is an edge from *i* to *j* in graph G^n , i.e., $(i, j) \in \mathcal{E}^n$. Let $I_{i,t}^n$ denote the *information set* of agent *i* at time *t* and $\mathcal{I}_{i,t}^n$ denote the set of all possible information sets. For any *i, j*, such that $(i, j) \in \mathcal{E}^n$, the messages that *i* can send to *j* at time *t* are defined through a mapping $m_{ij,t}^n : \mathcal{I}_{i,t}^n \rightarrow \mathcal{M}_{ij,t}^n$, where $\mathcal{M}_{ij,t}^n$ denotes the set of all possible messages. This mapping makes it clear that the messages that *i* can send to *j* could in principle depend on the information set of agent *i* as well as the identity of agent *j*. Importantly, we assume that the cardinality (“dimensionality”) of $\mathcal{M}_{ij,t}^n$ is no less than that of $\mathcal{I}_{i,t}^n$, so that communication can take the form of agent *i* sharing all her information with agent *j*. This has two key implications. First, an agent can communicate (indirectly) with a much larger set of agents than just her immediate neighbors, albeit with a time delay. As an example, an agent can communicate with the neighbors of her neighbors in two time periods (see Fig. 1). Second, mechanical duplication of information is avoided. For example, the second time agent *j* communicates with agent *i*, she can repeat her original signal, but this will not be recorded as an additional piece of information by agent *j*, since given the size of the message space $\mathcal{M}_{ij,t}^n$, each piece of information can be “tagged”. This ensures that under truthful communication, there need be no confounding of new information and previously communicated information.

The information set of agent *i* at time $t \geq 1$ is given by

$$I_{i,t}^n = \{s_i, m_{ji,\tau}^n, \text{ for all } 1 \leq \tau < t \text{ and } j \text{ such that } (j, i) \in \mathcal{E}^n\}$$

and $I_{i,0}^n = \{s_i\}$. In particular, the information set of agent *i* at time $t \geq 1$ consists of her private signal and all the messages her neighbors sent to *i* in previous time periods. Agent *i*’s action $x_{i,t}^n$ at time *t* is a mapping from her information set to the set of actions, i.e.,

$$\sigma_{i,t}^n : \mathcal{I}_{i,t}^n \rightarrow \{\text{“wait”}, 0, 1\}.$$

The tradeoff between taking an action (0 or 1) and waiting, should be clear at this point. An agent would wait, in order to communicate with a larger set of agents and potentially choose the correct action with higher probability. On the other hand, the future is discounted, therefore, delaying is costly. Moreover, an agent may delay her decision so that other agents, whom she receives information from, delay their exit decisions. These strategic considerations introduce a game among the agents, which we refer to as the *information exchange game*.

We will need the notion of the neighborhood of an agent in the following analysis. In particular, the *neighborhood* of agent i at time t is defined as

$$B_{i,t}^n = \{j \neq i \mid \exists \text{ a directed path from } j \text{ to } i \text{ with at most } t \text{ links in } G^n\},$$

i.e., $B_{i,t}^n$ consists of all agents that are at most t links away from agent i in graph G^n . Intuitively, if agent i waits for t periods and all of the intervening agents receive and send information truthfully, i will have access to all of the signals of the agents in the set $B_{i,t}^n$.

We next define the equilibria of the information exchange game.⁷

Definition 6 An action strategy profile $\sigma^{n,*}$ is a *pure-strategy Perfect Bayesian Equilibrium* of the information exchange game if for every $i \in \mathcal{N}^n$ and time t , $\sigma_{i,t}^{n,*}$ maximizes the expected payoff of agent i given the strategies of other agents $\sigma_{-i}^{n,*}$, i.e.,

$$\sigma_{i,t}^{n,*} \in \arg \max_{y \in \{\text{“wait”}, 0, 1\}} \mathbb{E}_{((y, \sigma_{i,-t}^{n,*}), \sigma_{-i}^{n,*})} (u_i(x_i^n, \theta) \mid I_{i,t}^n).$$

Let us consider a sequence of communication networks $\{G^n\}_{n=1}^\infty$ and refer to it as a *society*. We use the term *equilibrium* to denote a sequence of equilibria $\sigma = \{\sigma^n\}_{n=1}^\infty$ of information exchange games. The next definition introduces asymptotic learning for a given society. For any fixed $n \geq 1$ and any equilibrium of the information exchange game σ^n , we introduce the indicator variable:

$$M_{i,t}^n = \begin{cases} 1 & \text{if agent } i \text{ takes the correct decision by time } t, \\ 0 & \text{otherwise.} \end{cases} \tag{13}$$

Definition 7 We say that asymptotic learning occurs in society $\{G^n\}_{n=1}^\infty$ along equilibrium σ if for every $\epsilon > 0$, we have

$$\lim_{n \rightarrow \infty} \lim_{t \rightarrow \infty} \mathbb{P}_\sigma \left(\left[\frac{1}{n} \sum_{i=1}^n (1 - M_{i,t}^n) \right] \epsilon \right) = 0.$$

This definition states that asymptotic learning occurs with all but a negligible fraction of the agents taking the correct action (as the society grows infinitely large).

Let us also simplify the analysis and the exposition by assuming that private signals are binary. In particular, we assume that $s_i \in \{0, 1\}$ for all i . Let $L(x)$ denote the likelihood ratio for private signal x , i.e., $L(x) = \frac{\mathbb{P}(x \mid \theta=1)}{\mathbb{P}(x \mid \theta=0)}$. We assume that $L(1) = \frac{\beta}{1-\beta}$ and $L(0) = \frac{1-\beta}{\beta}$ (with $1/2 > \beta$). Moreover, the (common) *precision* of the private signals β is taken to be less than the discount factor δ , i.e., $\beta < \delta$. We will also further simplify the analysis by assuming

⁷Note that $\sigma_{i,-t}^n$ denotes the vector of actions of agent i at all times except t .

that communication is truthful (thus, until she exists, an agent truthfully reports all of her information). These results are extended to the case with strategic communication in [3].

We next introduce the concepts that are instrumental for asymptotic learning: *the minimum observation radius* and *k-radius sets*. We define the *minimum observation radius* of agent i as the following stopping time:

Definition 8 The *minimum observation radius* of agent i is defined as d_i^n , where

$$d_i^n = \arg \min_t \min_{I_{i,t}^n \in \{0,1\}^{|N_{i,t}^n|}} \{x_{i,t}^{n,*}(I_{i,t}^n) \in \{0, 1\}\}.$$

In particular, the minimum observation radius of agent i can be simply interpreted as the minimum number of time periods that agent i will wait before she takes an irreversible action 0 or 1, given that all other agents do not exit, over any possible realization of the private signals. Given the notion of a minimum observation radius, we define *k-radius sets* as follows.

Definition 9 Let V_k^n be defined as

$$V_k^n = \{i \in \mathcal{N} \mid |B_{i,d_i^n}^n| \leq k\}.$$

We refer to V_k^n as the *k-radius set*.

Intuitively, V_k^n includes all agents that may take an action before they receive signals from more than k other individuals—the size of their (indirect) neighborhood by the time their minimum observation radius d_i^n is reached is no greater than k . Equivalently, agent i belongs to set V_k^n if the number of agents that lie at distance less than d_i^n from i are at most k . From Definition 9 it follows immediately that

$$i \in V_k^n \Rightarrow i \in V_{k'}^n \quad \text{for all } k' > k. \tag{14}$$

The following theorem provides a necessary and sufficient condition for asymptotic learning to occur in a society under the assumption that communication is truthful.

Theorem 6 *Suppose communication is truthful. Then, asymptotic learning occurs in society $\{G^n\}_{n=1}^\infty$ (in any equilibrium σ) if and only if*

$$\lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{n} \cdot |V_k^n| = 0. \tag{15}$$

This result states that asymptotic learning is precluded if there exists a significant fraction of the society that will take an action before seeing a large set of signals, because in this case there will be a positive probability of each individual making a mistake, since her decision is based on a small set of signals. Intuitively, this condition requires that most agents are a short distance away from *information hubs*, defined as agents that have a very large (in the limit, infinite) number of connections. This motivates two different types of information hubs as conduits of asymptotic learning (both of these labels are inspired by [47]). The first are *information mavens*, which have a large in-degree, enabling them to aggregate information. If most agents are close to an information maven, asymptotic learning is guaranteed. The second type of hubs are *social connectors*, which have large out-degree, enabling them

to communicate their information to a large number of agents. Social connectors are only useful for asymptotic learning if they are close to mavens, so that they can distribute their information. Thus, asymptotic learning is also obtained if most agents are close to a social connector, who is in turn a short distance away from a maven.

In [3], we study an environment, in which individuals may (strategically) misreport their information if they have an incentive to do so. We show that individuals may in general choose to misreport their information in order to delay the action of their neighbors, thus obtaining more information from them in the future. Nevertheless, we establish that whenever truthful communication leads to asymptotic learning, it is an ϵ -equilibrium of the strategic communication game to report truthfully. Interestingly, the converse is not necessarily true: strategic communication may lead to asymptotic learning in some special cases in which truthful communication precludes learning.

The most important implication of this analysis is that even if truthful communication can be guaranteed, the conditions for learning are quite stringent. In particular, Theorem 6 shows that such learning will only take place when there are mavens and social connectors that can effectively aggregate almost all of the dispersed information. While this might sometimes be a good approximation to media sources or other central individuals that play the role of information aggregation in society, it does not provide a good description of most social networks and societies, where information flows are much more local. What happens when the dispersed information cannot be aggregated through communication? Because the underlying model here corresponds to a strongly connected network, even when all of the dispersed information cannot be aggregated, there are strong forces here toward consensus. In particular, a common situation would be one in which highly connected individuals will not wait for all of the possible information that they can gather, because of the cost of delay, and will instead make a decision at some point. But since they are highly connected, this choice will be observed by others and will influence their beliefs. Because there will be no more information forthcoming from these individuals, in most situations the rest of the society might also immediately follow the actions of these highly connected individuals. In this situation, even though all individuals may not hold exactly the same beliefs, they will still have beliefs that are very similar, and in particular, they will all choose the same action. Therefore, there will be a form of “herding” in this environment as well, now based on (partially informative) communication rather than observation of past actions. The reason why this information does not aggregate is therefore not a lack of consensus, but rather the fact that consensus forms too soon.

This discussion therefore highlights that both in Bayesian models of observational learning and communication, there are reasons to expect opinions not to effectively aggregate dispersed information, though the result will often be some type of consensus. Therefore, Bayesian models highlight the difficulties of aggregating dispersed information, while also emphasizing that it is rare for individuals observing each others’ actions or communicating together to arrive at very different beliefs.⁸

3.3 Learning in Markets

Even though the underlying state θ may be related to economic decisions, the models we have discussed so far allow for social and economic interactions but not for market-based interactions. A different perspective, pioneered by the famous economist Frederich von Hayek,

⁸As we saw in the previous section, this type of differences in long-run beliefs can arise when there is more uncertainty about the underlying state and the process through which signals are generated.

also argues that dispersed information can be aggregated effectively through markets, instead of through individuals' observation of and communication with their network. Hayek reached a similar conclusion to that of Galton on the collective intelligence of groups because he argued that markets and prices were effective at aggregating dispersed information. For example, Hayek wrote ([58], p. 526):

“The mere fact that there is a one [market] price for any commodity. . . brings about the solution which. . . might have been arrived at by one single mind possessing all the information which is in fact dispersed among all people involved in the process.”

Hayek in fact believed that while markets could successfully aggregate dispersedly-held complex information, no single individual or government could play the same role because the cost of computation and processing it would be prohibitive. Thus in Hayek's vision, social learning mediated by the market is superior to what a centralized authority could achieve even if it had access to the same information.

The economics literature has attempted to construct models to formalize this perspective. One approach has been to augment the standard Arrow-Debreu competitive equilibrium analysis with dispersed information. The key assumption is that while each individual might have dispersed and even private information, he still behaves in a competitive manner, as a price taker, meaning he takes prices as fixed and independent of his actions (e.g., [8, 52, 75, 76]). Under some (sometimes quite stringent) assumptions, this approach shows that market prices might effectively aggregate all the relevant information. This is a striking result, though there are several challenges in mapping it to reality.

First, the result follows only when in the full information case (after dispersed information has been effectively aggregated), there is a one-to-one mapping between observed prices and the underlying state θ . This may not be the case in many situations, particularly if the underlying state θ has a high dimension.

Second, as noted by the seminal paper by Grossman and Stiglitz [53], if all information is reflected in prices, then there may not be incentives for individuals to acquire information, and this may lead to the disappearance (or nonexistence) of this type of fully revealing equilibrium.

Third and perhaps most important, this entire approach has been based on the competitive equilibrium framework, and requires a Walrasian auctioneer to do the price setting and adjustment. This is particularly important, since full revelation arises because individuals can submit fully contingent demands to the Walrasian auctioneer. Hayek's intuition, in contrast, is based on decentralized trading. In fact, as we have emphasized, Hayek believed that markets were achieving something that a central authority with all of this information in its hands could not. Therefore, a systematic investigation of the role of markets in social learning requires more microfounded models, where trade takes place in a decentralized manner, possibly over a social network rather than in a centralized Walrasian market. It would also be necessary to investigate interactions between non-market learning and information revealed by prices in markets.⁹

4 Non-Bayesian Learning

In this section, we present non-Bayesian models of belief formation with learning through communication and show that the forces toward the emergence of incorrect beliefs in the

⁹Wolinsky [95], Ostrovsky [73] and Golosov, Lorenzoni and Tsyvinski [49] investigate the revelation of information by prices in markets in which buyers and sellers match and bargain over prices.

long run are even stronger. We first present a classical model introduced in [28], which is a simple model of belief and consensus formation over social networks (see also [30, 50]).¹⁰ We will see, however, that the specific assumptions it makes on how beliefs are updated may have certain non-desirable implications. Motivated by this, we will consider a variant of this model that avoids some of these properties and will then enrich it by introducing issues related to the spread of misinformation, belief manipulation and persistent disagreements between agents.

4.1 DeGroot Model of Belief Updating

We consider a set $\mathcal{N} = \{1, \dots, n\}$ of agents interacting in a social network. Each agent i starts with an initial belief about an underlying state, which we denote by $x_i(0) \in \mathbb{R}$. Agents exchange information about their beliefs with their neighbors. We assume that agents update their beliefs at discrete time instances. At any time $k \geq 0$, agent i updates his belief according to the relation

$$x_i(k + 1) = \sum_{j=1}^n T_{ij}x_j(k). \tag{16}$$

Here the nonnegative scalar T_{ij} indicates the weight that agent i puts on agent j 's belief. Letting $x(k) = [x_1(k), \dots, x_n(k)]'$ denote the vector of beliefs, the evolution of the beliefs can be expressed as

$$x(k + 1) = Tx(k) \quad \text{for all } k \geq 0,$$

where the weight matrix $T = [T_{ij}]_{i,j \in \mathcal{N}}$ represents the social network of interactions, i.e., $T_{ij} = 0$ implies that agent i does not get direct information from agent j regarding his belief, or equivalently, there is no directed link from agent i to j in the underlying social network.¹¹ We assume that the weight matrix T is a (row) *stochastic matrix*, i.e., the sum of entries across each row is equal to one. Hence, at each time instance agents update their beliefs to a convex combination of their current beliefs and the beliefs of their neighbors.¹²

¹⁰The DeGroot model and variations have also been studied extensively in the cooperative control literature as natural algorithms for achieving cooperative behavior with local information in networked-systems (see [18, 33, 62, 67, 71, 72, 87–89]).

¹¹The appealing interpretation is that this type of averaging would be optimal at time $t = 0$ if all agents had independent beliefs drawn from a normal distribution with mean equal to the underlying state and weights adjusted according to the precision of the distribution. At times $t > 0$, the weights assigned to the neighbors need to be updated to account for the new information they obtained from their neighbors. Therefore, this type of updating can be viewed as a boundedly rational version of the optimal processing of information where the weights are kept constant over time. This interpretation is discussed in detail in [30] in a related context.

¹²The literature has also considered versions of the DeGroot model with time varying weights and belief-dependent weights (the weight matrix T is a function of agent beliefs $x(k)$). Time varying weights have been studied both in social models of belief formation and in the cooperative control literature and capture the natural setting in which an individual changes self and/or neighbor weights as he gets more information or as the underlying network of interactions changes over time. Belief-dependent weights represent situations in which the underlying communication patterns are affected by the current beliefs of agents. A natural belief-dependent weight model is due to [63], which allows an agent to pay attention to beliefs that do not differ too much from his own (see also [27, 59, 93]). The convergence of beliefs in these models is studied in [68], which under mild conditions on the weight matrix $T(x(k))$, shows that the set of agents can be partitioned into groups such that each group reaches a consensus (see also [17] for a proof of convergence on Krause's model and properties of limiting beliefs, and [66] for application of a belief-dependent communication model within a multi-agent optimization framework).

This is a simple update rule, capturing the “imitation” aspect of non-Bayesian models we discussed in Sect. 2. When we think of a one-step update, this rule is also quite reasonable. For example, the entries of the matrix T can be interpreted as the trust that an individual places on another (with $T_{ij} = 0$ corresponding to the case in which agent i does not trust j), and it may appear natural that an individual should update her beliefs to be closer to those of the agents whom she trusts. However, when this update rule is applied dynamically, it may not be as compelling. To see this, note that an individual will update her beliefs at each date according to the beliefs of the agents that she trusts. Take a special case in which $T_{ij} > 0$ (agent i trusts agent j) and $T_{jk} = 0$ for all $k \neq j$, so that agent j does not update her beliefs. In this situation, after the first period, agent i has already taken all the relevant information from j . But according to the DeGroot update rule, she will keep on updating her own information according to the unchanging information of agent j , creating an extreme form of *duplication of information*. In most practical situations, we would imagine that after repeated interactions, agent i should realize that agent j has an unchanging opinion and there is no more point in updating her views because of agent j 's different beliefs. While this special case is extreme (because agent j 's opinions are not changing), it illustrates a general feature that the DeGroot update rule might be too myopic, especially in the context of individuals that are interacting in the same manner in each period.

The attractive feature of this model, on the other hand, is that the analysis of consensus is quite straightforward. In particular, standard results from Markov Chain Theory can be used to establish sufficient conditions that ensure convergence of the beliefs to a stationary distribution, which here will represent consensus among all of the agents, i.e., there will exist $\bar{x} \in \mathbb{R}$ such that $x(k)$ converges to $e'\bar{x}$ as k goes to infinity, where e is the vector of all ones. This is because the matrix T has been assumed to be a (row) stochastic matrix as noted above (this feature ensures that the stationary distribution can be represented as $e'\bar{x}$).

In particular, it can be shown that if the weight matrix T is such that the Markov chain with transition matrix T is *irreducible* and *aperiodic*, then agent beliefs reach a consensus in the limit. The irreducibility of the Markov chain is equivalent to the underlying social network being strongly connected so that there is a directed path from every node to every other node. Most of the literature in this area guarantees aperiodicity by assuming that $T_{ii} > 0$ for some or all i so that some or all agents assign positive weight to their own belief in the update relation (16).¹³

This discussion also implies that consensus will arise under fairly weak assumptions in this model (in particular, to avoid consensus, one needs to assume that the society is not “strongly connected,” meaning that there is no communication between two or more subsets, which would not be a realistic description of any society, however fragmented). Therefore, certain aspects of this model need to be generalized for the study of persistent disagreements, the spread of misinformation and belief manipulation. In the next subsection, we will present a variant of this model, which deals with several of these problems.

4.2 Spread of Misinformation

We next consider a variation on the DeGroot model, developed in [1], which enables us to study the effect of prominent agents on the beliefs of the society. We consider a similar setup as in the previous section with a set $\mathcal{N} = \{1, \dots, n\}$ of agents interacting over a social network with initial beliefs $x_i(0) \in \mathbb{R}$ about some underlying state $\theta \in \mathbb{R}$. We assume, in a way

¹³It is not necessary to have $T_{ii} > 0$ for even a single i to ensure consensus. Necessary and sufficient conditions for reaching a consensus in this model are presented in [51].

that parallels the assumptions in the Bayesian models, that there is sufficient information about θ held in a dispersed manner among all of the agents in the society. In particular, let us suppose that the average of initial beliefs in the society is equal to θ , i.e.,

$$\frac{1}{n} \sum_{i=1}^n x_i(0) = \theta.$$

We also assume that there are two types of agents; *regular and forceful*. Regular agents exchange information with their neighbors (when they meet). In contrast, forceful agents influence others disproportionately.

We use an asynchronous continuous-time model to represent meetings between agents. In particular, we assume that each agent meets and communicates with other agents at instances defined by a rate one Poisson process independent of other agents. This implies that the meeting instances (over all agents) occur according to a rate n Poisson process at times t_k , $k \geq 1$. Note that in this model, by convention, at most one node is active (i.e., is meeting another) at a given time. We discretize time according to meeting instances (since these are the relevant instances at which the beliefs change), and refer to the interval $[t_k, t_{k+1})$ as the k^{th} time slot. On average, there are n meeting instances per unit of absolute time. Suppose that at time (slot) k , agent i is chosen to meet another agent (probability $1/n$). In this case, agent i will meet agent $j \in \mathcal{N}$ with probability $p_{ij} \geq 0$ (we assume that $\sum_{j=1}^n p_{ij} = 1$ for all i). We let P denote the matrix with entries p_{ij} , which is again a (row) stochastic matrix.¹⁴

Following a meeting between i and j , there is a potential exchange of information. Throughout, we assume that all events that happen in a meeting are *independent of any other event that happened in the past*. Let $x_i(k)$ denote the belief of agent i about the underlying state at time k . The agents update their beliefs according to one of the following three possibilities.

- (i) Agents i and j reach pairwise consensus and the beliefs are updated according to

$$x_i(k + 1) = x_j(k + 1) = \frac{x_i(k) + x_j(k)}{2}.$$

We denote the conditional probability of this event (conditional on i meeting j) as β_{ij} .

- (ii) Agent j influences agent i , in which case for some $\epsilon \in (0, 1/2]$, beliefs change according to

$$x_i(k + 1) = \epsilon x_i(k) + (1 - \epsilon)x_j(k), \quad \text{and} \quad x_j(k + 1) = x_j(k). \tag{17}$$

In this case beliefs of agent j do not change.¹⁵ We denote the conditional probability of this event as α_{ij} , and refer to it as the *influence probability*. Note that we allow $\epsilon = 1/2$, so that agent i may be treating agent j just as a regular agent, except that agent j himself does not change his beliefs.

¹⁴There is a natural parallel between the matrix T indeed DeGroot model and the matrix P here. However, we will see that it will in fact be an augmented version of this matrix that will play a mathematically similar role to that of T into the DeGroot model.

¹⁵We could allow the own-belief weight ϵ to be different for each agent i . This generality does not change the results or the economic intuitions, so for notational convenience, we assume this weight to be the same across all agents.

(iii) Agents i and j do not agree and stick to their beliefs, i.e.,

$$x_i(k + 1) = x_i(k), \quad \text{and} \quad x_j(k + 1) = x_j(k).$$

This event has probability $\gamma_{ij} = 1 - \beta_{ij} - \alpha_{ij}$.

The first advantage of this framework is that the problem of duplication of information highlighted in the context of the DeGroot model, though still present, is now less severe. In particular, ignoring the events in which one agent influences the other (events with probability α_{ij} described in (ii) above), two agents who match again after having communicated in the recent past will have either reached a consensus or will not have exchanged any information, and thus one more round of updating is either inconsequential or would not run into the problem that there is duplication of information. There will be communication, instead, if one of them has met with somebody else and has updated her beliefs in the process, and in this case, there is relevant information to be exchanged, though there will still be some amount of replication of information.

The second advantage is that we can introduce, in a natural way, the possibility that some agents are “prominent” and influence others, without themselves being influenced to the same degree. This will then enable us to model some agents trying to manipulate the beliefs of others or spreading misinformation. In particular, our description above makes it clear that when $\alpha_{ij} > 0$ agent j will have a special influence on agent i . To emphasize this, any agent j for whom the influence probability $\alpha_{ij} > 0$ for some $i \in \mathcal{N}$ is referred to as a *forceful agent*. Moreover, the directed link (j, i) is referred to as a *forceful link*.¹⁶

We can interpret forceful agents in multiple different ways. First, forceful agents may correspond to community leaders or news media, who have a disproportionate effect on the beliefs of their followers. In such cases, it is natural to consider ϵ small and the leaders or media not updating their own beliefs as a result of others listening to their opinion. Second, forceful agents may be indistinguishable from regular agents, and thus regular agents engage in what they think is information exchange, but forceful agents, because of stubbornness or some other motive, do not incorporate the information of these agents in their own beliefs. In this case, it may be natural to think of ϵ as equal to 1/2. The results that follow remain valid with either interpretation.

The influence structure described above will determine the evolution of beliefs in the society. Below, we will give a more precise separation of this evolution into two components, one related to the underlying social network, and the other to influence patterns.

We next state our assumptions on the belief evolution model among the agents. Our first assumption is about the connectivity of the agents in the social network. Consider the directed graph $(\mathcal{N}, \mathcal{E})$, where \mathcal{E} is the set of directed links induced by the positive meeting probabilities p_{ij} , i.e.,

$$\mathcal{E} = \{(i, j) \mid p_{ij} > 0\}. \tag{18}$$

We assume that the graph $(\mathcal{N}, \mathcal{E})$ is *strongly connected*, i.e., for all $i, j \in \mathcal{N}$, there exists a directed path connecting i to j with links in the set \mathcal{E} . This assumption ensures that every agent “communicates” with every other agent (possibly through multiple links). This is not an innocuous assumption, since otherwise the graph $(\mathcal{N}, \mathcal{E})$ (and the society that it represents) would segment into multiple non-communicating parts. Though not innocuous, this assumption is also natural for two reasons. First, the evidence suggests that most subsets of

¹⁶We refer to directed links/edges as links and undirected ones as edges.

the society are not only connected, but are connected by means of several links (e.g., [61, 91]), and the same seems to be true for indirect linkages via the Internet. Second, if the society is segmented into multiple non-communicating parts, the insights here would apply, with some modifications, to each of these parts.

Let us also use d_{ij} to denote the length of the shortest path from i to j and d to denote the *maximum shortest path length* between any $i, j \in \mathcal{N}$, i.e.,

$$d = \max_{i, j \in \mathcal{N}} d_{ij}. \tag{19}$$

Since $(\mathcal{N}, \mathcal{E})$ is strongly connected, these are all well defined. Finally, we also impose the following *no-man-is-an-island* assumption, that there is positive probability that every agent (even if he is forceful) receives some information from an agent in his neighborhood. In particular, for all $(i, j) \in \mathcal{E}$, the sum of the averaging probability β_{ij} and the influence probability α_{ij} is positive, i.e.,

$$\beta_{ij} + \alpha_{ij} > 0 \quad \text{for all } (i, j) \in \mathcal{E}.$$

The assumption that the network is strongly connected ensures that there is a path from any forceful agent to other agents in the network, implying that for any forceful agent i , there is a link $(i, j) \in \mathcal{E}$ for some $j \in \mathcal{N}$. Then the no-man-is-an-island assumption guarantees that even the forceful agents at some point adopt information from the other agents in the network.¹⁷ This is a central assumption for the analysis in this subsection, and we will see in the next subsection that the implications of the model are enriched considerably (though the analysis also becomes more involved) when this assumption is relaxed.

We can express the preceding belief update model compactly as follows. Let $x(k) = (x_1(k), \dots, x_n(k))$ denote the vector of agent beliefs at time k . The agent beliefs are updated according to the relation

$$x(k + 1) = W(k)x(k), \tag{20}$$

where $W(k)$ is a random matrix given by

$$W(k) = \begin{cases} A^{ij} \equiv I - \frac{(e_i - e_j)(e_i - e_j)'}{2} & \text{with probability } p_{ij}\beta_{ij}/n, \\ J^{ij} \equiv I - (1 - \epsilon)e_i(e_i - e_j)' & \text{with probability } p_{ij}\alpha_{ij}/n, \\ I & \text{with probability } p_{ij}\gamma_{ij}/n, \end{cases} \tag{21}$$

for all $i, j \in \mathcal{N}$. The preceding belief update model implies that the matrix $W(k)$ is a stochastic matrix for all k , and is independent and identically distributed over all k .

Given our assumptions, we have for some nonnegative matrix \tilde{W} ,

$$\mathbb{E}[W(k)] = \tilde{W} \quad \text{for all } k \geq 0. \tag{22}$$

The matrix, \tilde{W} , which we refer to as the *mean interaction matrix*, represents the evolution of beliefs in the society. It incorporates elements from both the underlying social network (which determines the meeting patterns) and the influence structure. In what follows, it will be useful to separate these into two components, both for our mathematical analysis and to

¹⁷This assumption is stated for all $(i, j) \in \mathcal{E}$, thus a forceful agent i receives some information from any j in his “neighborhood”. This is without any loss of generality, since we can always set $p_{ij} = 0$ for those j ’s that are in i ’s neighborhood but from whom i never obtains information.

clarify the intuitions. For this purpose, let us use the belief update model (20)–(21) and write the mean interaction matrix \tilde{W} as follows:¹⁸

$$\begin{aligned} \tilde{W} &= \frac{1}{n} \sum_{i,j} p_{ij} [\beta_{ij} A^{ij} + \alpha_{ij} J^{ij} + \gamma_{ij} I] \\ &= \frac{1}{n} \sum_{i,j} p_{ij} [(1 - \gamma_{ij}) A^{ij} + \gamma_{ij} I] + \frac{1}{n} \sum_{i,j} p_{ij} \alpha_{ij} [J^{ij} - A^{ij}], \end{aligned}$$

where A^{ij} and J^{ij} are matrices defined in (21), and the second equality follows from the fact that $\beta_{ij} = 1 - \alpha_{ij} - \gamma_{ij}$ for all $i, j \in \mathcal{N}$. We use the notation

$$T = \frac{1}{n} \sum_{i,j} p_{ij} [(1 - \gamma_{ij}) A^{ij} + \gamma_{ij} I], \quad D = \frac{1}{n} \sum_{i,j} p_{ij} \alpha_{ij} [J^{ij} - A^{ij}], \quad (23)$$

to write the mean interaction matrix, \tilde{W} , as

$$\tilde{W} = T + D. \quad (24)$$

Here, the symmetric matrix T only depends on meeting probabilities (matrix P) and on the probability that following a meeting no exchange takes place, γ_{ij} . We can therefore think of the matrix T as representing the underlying *social network* (friendships, communication among coworkers, decisions about which news outlets to watch, etc.), and refer to it as the *social network matrix*. It will be useful below to represent the social interactions using an undirected (and weighted) graph induced by the social network matrix T . This graph is given by $(\mathcal{N}, \mathcal{A})$, where \mathcal{A} is the set of undirected edges given by

$$\mathcal{A} = \{\{i, j\} \mid T_{ij} > 0\}, \quad (25)$$

and the weight w_e of edge $e = \{i, j\}$ is given by the entry $T_{ij} = T_{ji}$ of the matrix T . We refer to this graph as the *social network graph*.

The matrix D , on the other hand, can be thought of as representing the *influence structure* in the society, and is hence called the *influence matrix*. It incorporates information about which individuals and links are forceful, i.e., which types of interactions will lead to one individual influencing the other without updating his own beliefs. It is also useful to note for interpreting the mathematical results below that T is a doubly stochastic matrix, while D is not. Therefore, (24) gives a decomposition of the mean connectivity matrix \tilde{W} into a doubly stochastic and a remainder component, and enables us to use tools from matrix perturbation theory.

The next result shows that (stochastic) consensus will emerge despite the possibility of misinformation and stochastic communication. We refer to this as stochastic consensus because the consensus value of beliefs itself is a random variable that depends on the initial beliefs and the random sequence of matrices $\{W(k)\}$.

Theorem 7 *The sequences $\{x_i(k)\}, i \in \mathcal{N}$, generated by (20) converge to a consensus belief, i.e., there exists a scalar random variable \bar{x} such that*

$$\lim_{k \rightarrow \infty} x_i(k) = \bar{x} \quad \text{for all } i \text{ with probability one.}$$

¹⁸In the sequel, the notation $\sum_{i,j}$ will be used to denote the double sum $\sum_{i=1}^n \sum_{j=1}^n$.

Moreover, the random variable \bar{x} is a convex combination of initial agent beliefs, i.e.,

$$\bar{x} = \sum_{j=1}^n \pi_j x_j(0),$$

where $\pi = [\pi_1, \dots, \pi_n]$ is a random vector that does not depend on the initial beliefs, and satisfies $\pi_j \geq 0$ for all j , and $\sum_{j=1}^n \pi_j = 1$.

The key implication of this result is that, despite the presence of forceful agents, the society will ultimately reach a consensus. Though surprising at first, this result is intuitive in light of our no-man-is-an-island assumption. Note, however, that, in contrast to the DeGroot model discussed in the previous subsection, this consensus value here is a random variable even conditional on initial beliefs. In particular, the consensus value will depend on the order in which meetings have taken place. The main role of this result for us is that we can now conduct our analysis on quantifying the extent of the spread of misinformation by looking at this consensus value of beliefs.

The next theorem characterizes $\mathbb{E}[\bar{x}]$ in terms of the limiting behavior of the matrices \tilde{W}^k as k goes to infinity.

Theorem 8 *Let \bar{x} be the limiting random variable of the sequences $\{x_i(k)\}$, $i \in \mathcal{N}$ generated by (20) (cf. Theorem 7). Then we have:*

(a) *The matrix \tilde{W}^k converges to a stochastic matrix with identical rows $\bar{\pi}$ as k goes to infinity, i.e.,*

$$\lim_{k \rightarrow \infty} \tilde{W}^k = e\bar{\pi}'.$$

(b) *The expected value of \bar{x} is given by a convex combination of the initial agent values $x_i(0)$, where the weights are given by the components of the probability vector $\bar{\pi}$, i.e.,*

$$\mathbb{E}[\bar{x}] = \sum_{i=1}^n \bar{\pi}_i x_i(0) = \bar{\pi}'x(0).$$

Combining Theorem 7 and Theorem 8(a) (and using the fact that the results hold for any $x(0)$), we have $\bar{\pi} = \mathbb{E}[\pi]$. The stationary distribution $\bar{\pi}$ is crucial in understanding the formation of opinions since it encapsulates the weight given to each agent (forceful or regular) in the limiting mean consensus value of the society. We refer to the vector $\bar{\pi}$ as the *consensus distribution* corresponding to the mean interaction matrix \tilde{W} and to its component $\bar{\pi}_i$ as the *weight* of agent i .

It is also useful at this point to highlight how consensus will form around the correct value in the absence of forceful agents. Let $\{x(k)\}$ be the belief sequence generated by the belief update rule of (20). When there are no forceful agents, i.e. $\alpha_{ij} = 0$ for all i, j , then the interaction matrix $W(k)$ for all k is either equal to an averaging matrix A^{ij} for some i, j or equal to the identity matrix I ; hence, $W(k)$ is a *doubly stochastic matrix* (i.e., it has both row and column sums equal to 1). This implies that the average value of $x(k)$ remains constant at each iteration and is given by

$$\frac{1}{n} \sum_{i=1}^n x_i(k) = \frac{1}{n} \sum_{i=1}^n x_i(0) \quad \text{for all } k \geq 0.$$

But since, by assumption, $\frac{1}{n} \sum_{i=1}^n x_i(0) = \theta$, we have from Theorem 7 the following simple but important corollary:

Corollary 1 *Assume that there are no forceful agents, i.e., $\alpha_{ij} = 0$ for all $i, j \in \mathcal{N}$. We have*

$$\lim_{k \rightarrow \infty} x_i(k) = \frac{1}{n} \sum_{i=1}^n x_i(0) = \theta \quad \text{with probability one.}$$

Therefore, in the absence of forceful agents, the society is able to aggregate information effectively. Theorem 8 then also implies that in this case $\bar{\pi}_i = 1/n$ for all i (i.e., beliefs converge to a deterministic value), so that no individual has excess influence. These results no longer hold when there are forceful agents.

We next study the effect of the forceful agents and the structure of the social network on the extent of misinformation and excess influence of individuals. As a measure of the extent of misinformation, we consider the expected value of the difference between the consensus belief \bar{x} (cf. Theorem 7) and the true underlying state, θ (or equivalently the average of the initial beliefs), i.e.,

$$\mathbb{E}[\bar{x} - \theta] = \mathbb{E}[\bar{x}] - \theta = \sum_{i \in \mathcal{N}} \left(\bar{\pi}_i - \frac{1}{n} \right) x_i(0), \quad (26)$$

(cf. Theorem 8).

The next theorem provides a key result on characterizing the extent of misinformation and establishes an upper bound on the l_∞ -norm of the difference between the stationary distribution $\bar{\pi}$ and the uniform distribution $\frac{1}{n}e$, which, from (26), also provides a bound on the deviation between expected beliefs and the true underlying state, θ . The proof uses results from Markov Chain Theory, which enable us to decompose the mean interaction matrix \tilde{W} in (24) into a component given by the social network matrix T , which is doubly stochastic, and an *influence matrix* D , which is the source of deviation of $\mathbb{E}\bar{x}$ from θ (see, in particular, [57, 83]).

Theorem 9

(a) *Let $\bar{\pi}$ denote the consensus distribution. The l_∞ -norm of the difference between $\bar{\pi}$ and $\frac{1}{n}e$ is bounded by*

$$\left\| \bar{\pi} - \frac{1}{n}e \right\|_\infty \leq \frac{1}{1 - \delta} \frac{\sum_{i,j} p_{ij} \alpha_{ij}}{2n},$$

where δ is a constant defined by

$$\delta = (1 - n\chi^d)^{\frac{1}{d}},$$

$$\chi = \min_{(i,j) \in \mathbb{E}} \left\{ \frac{1}{n} \left[p_{ij} \frac{1 - \gamma_{ij}}{2} + p_{ji} \frac{1 - \gamma_{ji}}{2} \right] \right\},$$

and d is the maximum shortest path length in the graph $(\mathcal{N}, \mathcal{E})$ [cf. (19)].

(b) Let \bar{x} be the limiting random variable of the sequences $\{x_i(k)\}$, $i \in \mathcal{N}$ generated by (20) (cf. Theorem 7). We have

$$\left| \mathbb{E}[\bar{x}] - \frac{1}{n} \sum_{i=1}^n x_i(0) \right| \leq \frac{1}{1-\delta} \frac{\sum_{i,j} p_{ij} \alpha_{ij}}{2n} \|x(0)\|_\infty.$$

Before providing the intuition for the preceding theorem, we provide a related bound on the l_2 -norm of the difference between $\bar{\pi}$ and the uniform distribution $\frac{1}{n}e$ in terms of the second largest eigenvalue of the social network matrix T .

Theorem 10 Let $\bar{\pi}$ denote the consensus distribution (cf. Theorem 8). The l_2 -norm of the difference between $\bar{\pi}$ and $\frac{1}{n}e$ is given by

$$\left\| \bar{\pi} - \frac{1}{n}e \right\|_2 \leq \frac{1}{1-\lambda_2(T)} \frac{\sum_{i,j} p_{ij} \alpha_{ij}}{n},$$

where $\lambda_2(T)$ is the second largest eigenvalue of the matrix T defined in (23).

Theorem 10 characterizes the variation of the stationary distribution in terms of the average influence, $\frac{\sum_{i,j} p_{ij} \alpha_{ij}}{n}$, which captures the importance of forceful agent in the society, and in terms of structural properties of the social network as represented by the matrix T —in particular, its second largest eigenvalue $\lambda_2(T)$. As is well known, the difference $1 - \lambda_2(T)$, also referred to as the *spectral gap*, governs the rate of convergence of the Markov Chain induced by the social network matrix T to its stationary distribution. In particular, the larger $1 - \lambda_2(T)$ is, the faster the k^{th} power of the transition probability matrix converges to the stationary distribution matrix. When the Markov chain converges to its stationary distribution rapidly, we say that the Markov chain is *fast mixing*.

In this light, Theorem 10 shows that, in a fast-mixing graph, given a fixed average influence $\frac{\sum_{i,j} p_{ij} \alpha_{ij}}{n}$, the consensus distribution is “closer” to the underlying $\theta = \frac{1}{n} \sum_{i=1}^n x_i(0)$ and the extent of misinformation is limited. This is intuitive. In a fast-mixing social network graph, there are several connections between any pair of agents. Now for any forceful agent, consider the set of agents who will have some influence on his beliefs. This set itself is connected to the rest of the agents and thus obtains information from the rest of the society. Therefore, in a fast-mixing graph (or in a society represented by such a graph), the beliefs of forceful agents will themselves be moderated by the rest of the society before they spread widely. In contrast, in a slowly-mixing graph, we can have a high degree of clustering around forceful agents, so that forceful agents get their (already limited) information intake mostly from the same agents that they have influenced. If so, there will be only a very indirect connection from the rest of the society to the beliefs of forceful agents and forceful agents will spread their information widely before their opinions also adjust. As a result, the consensus is more likely to be much closer to the opinions of forceful agents, potentially quite different from the true underlying state θ .

This discussion also gives the intuition for Theorem 9 since the constant δ in that result is closely linked to the mixing properties of the social network matrix and the social network graph. In particular, Theorem 9 clarifies that δ is related to the maximum shortest path and the minimum probability of (indirect) communication between any two agents in the society. These two notions also crucially influence the spectral gap $1 - \lambda_2(T_n)$, which plays a key role in Theorem 10.

There are several important implications of the model presented in this subsection and its analysis. First, the model has introduced a tractable variant of the DeGroot model, with an arguably easier interpretation and more limited duplication of information. Second, it has incorporated forceful agents, which give us a way of introducing belief manipulation and spread of misinformation. Third, because of the no-man-is-an-island assumption, despite the presence of forceful agents and the possibility of misinformation being spread, beliefs converge to a consensus in the long run, albeit a stochastic consensus. This is both a convenient implication, because it makes further analysis feasible and relatively straightforward, and a negative implication, because it does not allow the emergence of persistent disagreement. Finally, the analysis in this subsection has shown how one can go considerably beyond the question of whether there is (stochastic or nonstochastic) consensus and provide a full characterization of the divergence between consensus beliefs and a benchmark corresponding to a simple notion of aggregation of dispersed information. In particular, this analysis has shown that this measure of divergence depends on the presence and importance of forceful agents and structural properties of the social network (in particular, on the second largest eigenvalue, which captures how fast mixing the Markov chain induced by the social network of the society is).

4.3 Persistent Disagreement

The no-man-is-an-island assumption played a crucial role in the previous subsection in ensuring that the beliefs of the forceful agents are affected (even if infrequently) by the beliefs of the rest of the society. This feature then underpinned the emergence of (stochastic) consensus, which enabled the rest of the analysis to be conducted in a relatively simple manner. While “no-man-is-an-island” is a plausible assumption for forceful agents, it may not be a good description of how those trying to manipulate the beliefs of others or spread misinformation may act. For example, it may be realistic to assume this feature when one of the agents in a village or in a community has a leadership role but still listens to the rest of the group, but it may be highly implausible when we want to have a stylized description of how state media in authoritarian regimes such as Iran or China might try to manipulate the beliefs of its citizens. In this subsection, we relax this assumption. In addition to providing a generalization of the learning model presented in the previous subsection, this will enable us to have a framework in which disagreement among agents may be a long-run phenomenon despite the fact that the society is strongly connected.

More specifically, following [2], we consider a model with (fully) *stubborn agents* which are similar to the forceful agents in the previous subsection but never update their opinions and continuously influence those of the rest of the society. This model is developed and studied in [4]. We show that the presence of these agents leads to persistent disagreements among the rest of the society—because different individuals are within the “sphere of influence” of distinct stubborn agents and are influenced to varying degrees.

Consider a population $\mathcal{N} = \{1, \dots, n\}$ of agents, communicating and exchanging information. As in the previous two sections, each agent starts with an opinion $x_i(0) \in \mathbb{R}$ and is then “recognized” according to a rate-1 independent Poisson process in continuous time. Following this event, she meets one of the agents in her neighborhood according to a stochastic matrix P . We shall identify agents with the vertices of a (directed) graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, representing an underlying *social network*, where $(i, j) \in \mathcal{E}$ if and only if $P_{ij} > 0$.

As noted above, stubborn agents are those that never change or update their opinions, and we think of them as typically few in number. The (nonempty) set of stubborn agents is denoted by $S \subseteq \mathcal{N}$. The remaining agents are referred to as *regular agents*, and their set is

denoted by $\mathcal{A} \equiv \mathcal{N} \setminus \mathcal{S}$. These regular agents update their beliefs to some weighted average of their pre-meeting belief and the belief of the agent they met. Specifically, we assume that, if agent $a \in \mathcal{A}$ is recognized at time $t \geq 0$, and meets agent $v \in \mathcal{N}$, then her belief jumps from its current value $x_a(t^-)$ to

$$x_a(t) = (1 - \theta)x_a(t^-) + \theta x_v(t^-),$$

where $x_v(t^-)$ denotes the limit $\lim_{u \uparrow t} x_v(u)$. All other agents' beliefs remain constant. The parameter $\theta \in (0, 1)$ represents the trust that regular agent a puts on agent v 's belief (here assumed to be constant over all agents and over time for notational simplicity).

We denote the vector of beliefs at time t by $x(t) \in \mathbb{R}^{\mathcal{N}}$ and study its long-run behavior, under the assumption that the graph \mathcal{G} is strongly connected, which implies that the induced continuous-time Markov chain (i.e., the Markov chain with set of states given by \mathcal{N} and transition probability matrix P) admits a unique stationary probability distribution, denoted by π , supported over all \mathcal{N} .

The next theorem shows that, in contrast to the models we have seen so far, opinions no longer converge to a consensus. Instead, they exhibit both persistent disagreement and persistent fluctuations.

Theorem 11 *Assume that $x_s(0) \neq x_{s'}(0)$ for some $s, s' \in \mathcal{S}$. Then, with probability one, $x(t)$ does not converge.*

From a substantive point of view, this is one of the central results of the model. In particular, it shows that with probability one, opinions fail to converge. Thus persistent disagreement will remain, and in fact, all opinions will fluctuate even in the long run. Notably, there is persistent disagreement despite the fact that the social network in the society is strongly connected and thus the opinion of each regular agent can be (indirectly) influenced by that of any other regular agent. Disagreement arises because of the constant pull of the society in different directions by the influence of stubborn agents.

Nevertheless, we can show that regardless of the initial values of regular agents' beliefs, the belief vector $x(k)$ is convergent in distribution to a random vector X . Recall that this means that the probability law of $x(t)$, to be denoted by $\mathcal{L}(x(t))$, converges, according to the weak-* topology, to $\mathcal{L}(X)$, the probability law of an $\mathbb{R}^{\mathcal{N}}$ -valued random variable X , i.e.,

$$\lim_{t \rightarrow +\infty} \mathbb{E}[\varphi(x(t))] = \mathbb{E}[\varphi(X)],$$

for all bounded and continuous test functions $\varphi : \mathbb{R}^{\mathcal{N}} \rightarrow \mathbb{R}$.

Theorem 12 *For every value of the stubborn agents' beliefs $\{x_s(0)\} \in \mathbb{R}^{\mathcal{S}}$, there exists an $\mathbb{R}^{\mathcal{N}}$ -valued random variable X , such that*

$$\lim_{t \rightarrow +\infty} x(t) = X,$$

in distribution.

As in the previous subsection, we next derive results linking the social network structure to the distribution of asymptotic beliefs. We focus on the empirical average of beliefs:

$$\alpha_v \equiv \lim_{k \rightarrow +\infty} \frac{1}{k} \int_0^k x_v(u) du, \quad v \in \mathcal{N}.$$

Standard ergodic theorems for Markov chains imply that agent beliefs are ergodic, therefore the preceding limit is given by the expected value of the random vector X , i.e., $\alpha_v = \mathbb{E}[X_v]$ for all $v \in \mathcal{N}$, independent of the distribution of $x(0)$. We refer to $\alpha = [\alpha_v]_{v \in \mathcal{N}}$ as the *expected asymptotic belief vector*. The next theorem provides an explicit characterization of α .

Theorem 13 *The expected asymptotic belief vector α is the unique solution of the following Laplace equation*

$$((P - I)\alpha)_a = 0 \quad \text{for all } a \in \mathcal{A},$$

with boundary conditions $\alpha_s = x_s(0)$ for all $s \in \mathcal{S}$.

Here $((P - I)\alpha)_a$ refers to the a th element of the vector $((P - I)\alpha)$. The reason why we do not have $((P - I)\alpha) = 0$ is that this vector also includes the stubborn agents, whereas the requirement $((P - I)\alpha)_a = 0$ is only for $a \in \mathcal{A}$, i.e., only for regular agents. This theorem provides us with an explicit expression for the expected asymptotic belief vector α . In fact, α admits the following standard representation (see [7], Chap. 2, Lemma 27):

$$\alpha_v = \sum_{s \in \mathcal{S}} \mathbb{P}_v(\tau_{\mathcal{S}} = \tau_s) x_s(0), \quad \text{for all } v \in \mathcal{N},$$

where $\mathbb{P}_v(\cdot)$ denotes the probability measure associated with a continuous-time random walk $V(t)$, with initial state $V(0) = v$, transition rates P , while $\tau_{\mathcal{W}} \equiv \inf\{t \geq 0 : V(t) \in \mathcal{W}\}$ denotes the hitting time of an arbitrary subset $\mathcal{W} \subseteq \mathcal{N}$. This representation enables us to compute α exactly for certain social network topologies. In cases when the expected asymptotic belief vector α cannot be explicitly computed in a simple way, it is possible to provide bounds on its dispersion.

Theorem 14 *Assume that the stochastic matrix P is reversible. Then, for all $\epsilon > 0$, we have*

$$\frac{1}{n} \left| \left\{ v : \left| \alpha_v - \sum_{u \in \mathcal{N}} \pi_u \alpha_u \right| \geq \Delta_* \epsilon \right\} \right| \leq \frac{2}{\epsilon} \log(2e^2/\epsilon) \frac{\tau}{n\pi_* \mathbb{E}_\pi[\tau_{\mathcal{S}}]},$$

where $\Delta_* \equiv \max_{s,s' \in \mathcal{S}} x_s(0) - x_{s'}(0)$, $\pi_* \equiv \min_u \pi_u$, $\mathbb{E}_\pi[\cdot]$ denotes the expectation of the random walk $V(t)$ with initial distribution π , and τ is the (variational distance) mixing time of the random walk $V(t)$, i.e.,

$$\tau \equiv \inf\{t \geq 0 \mid \|\mathbb{P}_v(V(t) = \cdot) - \mathbb{P}_{v'}(V(t) = \cdot)\|_{TV} \leq e^{-1}, \forall v, v' \in \mathcal{N}\}.$$

Theorem 14 therefore provides a bound on the stationary distribution of beliefs in terms of the beliefs and the belief differences of the stubborn agents and structural properties of the social network, in particular as captured by the mixing time τ (in the same way that mixing times were important in the previous subsection). Intuitively, the theorem states that if the Markov chain $V(t)$ mixes in time faster than the expected hitting time of the stubborn agent set \mathcal{S} , then it will eventually hit any of them with approximately equal probability, and thus the expected asymptotic opinions do not vary much over the network.

5 Concluding Remarks

In this paper, we have provided an overview of recent work on opinion dynamics and learning in social networks. We emphasized the insights and the shortcomings of both Bayesian

and non-Bayesian approaches. Our focus has been on mathematical models linking the dynamics of opinions to the distribution of prior beliefs, the form of updating (e.g., Bayesian vs. non-Bayesian) and crucially to the structural properties of the social network in which agents are situated. In particular, we highlighted the importance of these factors on three sets of questions:

- a. *Consensus*: will social learning lead different individuals, who might start with different views (priors) and observe the actions of and engage in communication with different subsets of the society, to hold beliefs that are in agreement?
- b. *Asymptotic learning*: will social learning effectively aggregate dispersed information?
- c. *Misinformation*: will media sources, “prominent agents,” politicians and officers of the state be able to manipulate the beliefs of individuals, indoctrinate them and convince them of views that may not have the backing of data and evidence?

We have emphasized that even when agents are Bayesian and start with fairly accurate models of the world, that is, the correct understanding of how signals are generated and reasonable priors, asymptotic learning may not emerge because selfish behavior by each agent need not lead to the aggregation of dispersed information. Even though they do not guarantee asymptotic learning, Bayesian models create a strong tendency toward consensus and they limit the extent to which misinformation can arise (because Bayesian agents are relatively difficult to fool).

We then showed that some benchmark models of non-Bayesian learning also lead to consensus and similarly are unlikely to lead to asymptotic learning. The forces that lead to consensus also preclude the emergence of persistent disagreements and this feature might put limits on the extent of misinformation. We then outlined some recent work on the spread of misinformation and persistent long-run disagreement and linked these features to the structural properties of social networks in which communication and learning take place.

Our overview has been purposefully partial. A large amount of work on learning, opinion formation and communication has either been mentioned only in passing or has been ignored, because it is less directly related to our focus. We have also concentrated on mathematical models that enable the study of links between the structure of social networks and opinion dynamics. There are several areas upon which we touched only briefly and some which we could not mention because of space constraints that are important areas for future research. We end the paper by a brief discussion of some of these.

1. To enable sharp mathematical characterization, we have throughout focused on long-run properties of opinion dynamics. For example, rather than asking whether incorrect beliefs can persist for a long time, we have investigated whether there is asymptotic learning, meaning whether in the very long-run society will arrive at the correct beliefs and actions. Opinion dynamics away from this very long-run limit are often important and interesting, but more difficult to study. Moreover, even if there is asymptotic learning, the rate at which different societies arrive at this might be very different (as briefly mentioned in Sect. 3). The development of more powerful mathematical tools to study opinion dynamics away from the long-run stationary distribution and to investigate the rates at which the long-run distribution is reached constitutes an important area for future research.
2. Also to facilitate analysis, we have focused on models in which actions are either implicit (as in the models in Sect. 4) or each agent only takes a single action (as in the models in Sect. 3). An important area for future research is to consider models in which individuals interact with others repeatedly and update their information dynamically taking into account implications of this information both today and in the future.

3. We have emphasized repeatedly that non-Bayesian models provides several useful insights, but that they also have several ad hoc features and properties such as the “duplication of information” of the DeGroot model, which are unattractive (and which, in particular, we might expect not to survive in the long run). An interesting area of study is to develop more adaptive non-Bayesian models in which such easily identifiable myopic behavior is avoided by changes in the relevant rules of thumbs, at least in the long run.
4. We have seen that individuals’ conjectures about others’ behavior plays a crucial role in Bayesian models. These considerations are entirely absent in non-Bayesian models. It would also be interesting to extend non-Bayesian models to incorporate some of these concerns. For example, even if individuals are not Bayesian, they might worry about how to draw inferences from the behavior of others and about whether some other agents are trying to mislead them.
5. Models of misinformation and persistent disagreement are very much in their infancy. We have presented only one approach to these questions. Given the ubiquity of these issues in practice, more work is necessary to understand how misinformation spreads to some parts of the society and how individuals that communicate and share the same sources of information might nonetheless disagree significantly even in the very long run. It would be particularly important to see how such long-run disagreement is possible even when more Bayesian features are introduced or non-Bayesian models are made more adaptive.
6. In the area of persistent disagreement, it is also interesting to study whether persistent disagreement will lead to a situation in which there are clusters of relatively unchanging opinions within different parts of the social network, or whether there is greater fluidity and different types of opinions spread and retreat throughout the network at different points.
7. More work is also necessary on modeling indoctrination and belief manipulation. While the spread of misinformation in social networks gives us some clues about these issues, indoctrination in practice is often carried out by political movements or by the state using several instruments. Control of schools and the media might be particularly important. Currently, we know relatively little on this important topic.
8. Issues of misinformation also open up another area of study: if there will periodically be new sources of misinformation, either from parties that are purposefully trying to manipulate beliefs or because some agents and community leaders are stubborn and will not change their opinion even if these are not strongly based on the facts, then it becomes important to understand what types of societies and social structures are “robust” to the spread of misinformation and what can be done, from a design perspective, to make society more robust and opinions more stable or less subject to manipulation.

References

1. Acemoglu D, Ozdaglar A, ParandehGheibi A (2010) Spread of (mis)information in social networks. *Games Econ Behav*, forthcoming
2. Acemoglu D, Como G, Fagnani F, Ozdaglar A (2010) Persistent disagreement in social networks. Working paper
3. Acemoglu D, Bimpikis K, Ozdaglar A (2010) Communication information dynamics in (endogeneous) networks. LIDS report 2813, working paper
4. Acemoglu D, Dahleh M, Lobel I, Ozdaglar A (2010) Heterogeneity and social learning. Working paper
5. Acemoglu D, Dahleh M, Lobel I, Ozdaglar A (2009) Bayesian learning in social networks. LIDS report 2780
6. Acemoglu D, Chernozhukov V, Werning I (2007) Learning and disagreement in an uncertain world. Working paper

7. Aldous D, Fill J (2002) Reversible Markov chains and random walks on graphs. Monograph in preparation. <http://www.stat.berkeley.edu/aldous/RWG/book.html>
8. Allen B (1981) Generic existence of completely revealing equilibria for economies with uncertainty when prices convey information. *Econometrica* 49(5):1173–1199
9. Deffuant G, Amblard F, Weisbuch G, Faure T (2002) How can extremism prevail? A study based on the relative agreement interaction model. *J Artif Soc Soc Simul* 5(4)
10. Bala V, Goyal S (1998) Learning from neighbours. *Rev Econ Stud* 65(3):595–621
11. Banerjee A (1992) A simple model of herd behavior. *Q J Econ* 107:797–817
12. Banerjee A, Fudenberg D (2004) Word-of-mouth learning. *Games Econ Behav* 46:1–22
13. Bikhchandani S, Hirshleifer D, Welch I (1992) A theory of fads, fashion, custom, and cultural change as information cascades. *J Polit Econ* 100:992–1026
14. Binmore K (2008) *Rational decisions*. Princeton University Press, Princeton
15. Bisin A, Verdier T (2001) The economics of cultural transmission and the dynamics of preferences. *J Econ Theory* 97(2):298–319
16. Bisin A, Verdier T (2000) Beyond the melting pot: cultural transmission, marriage, and the evolution of ethnic and religious traits. *Q J Econ* 115(3):955–988
17. Blondel VD, Hendrickx JM, Tsitsiklis JN (2009) On Krause’s multi-agent consensus model with state-dependent connectivity. *IEEE Trans Autom Control* 54(11):2586–2597
18. Boyd S, Ghosh A, Prabhakar B, Shah D (2005) Gossip algorithms: design, analysis, and applications. In: *Proceedings of IEEE INFOCOM*
19. Boyd R, Richerson PJ (1985) *Culture and the evolutionary process*. The University of Chicago Press, Chicago
20. Cavalli-Sforza LL, Feldman MW (1981) *Cultural transmission and evolution: a quantitative approach*. Princeton University Press, Princeton
21. Chamley C, Gale D (1994) Information revelation and strategic delay in a model of investment. *Econometrica* 62:1065–1086
22. Chen Y, Kartik N, Sobel J (2008) Selecting cheap talk equilibria. *Econometrica* 76(1):117–136
23. Clifford P, Sudbury A (1973) A model for spatial conflict. *Biometrika* 60:581–588
24. de Condorcet NC (1785) *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, Paris
25. Crawford V, Sobel J (1982) Strategic information transmission. *Econometrica* 50(6):1431–1451
26. Cripps M, Ely J, Mailath G, Samuelson L (2008) Common learning. *Econometrica* 76:909–933
27. Deffuant G, Neau D, Amblard F, Weisbuch G (2000) Mixing beliefs among interacting agents. *Adv Complex Syst* 3:87–98
28. DeGroot MH, Reaching a consensus. *J Am Stat Assoc* 69:118–121
29. DellaVigna S, Kaplan E (2007) The Fox News effect: media bias and voting. *Q J Econ* 122:1187–1234
30. DeMarzo PM, Vayanos D, Zwiebel J (2003) Persuasion bias, social influence, and unidimensional opinions. *Q J Econ* 118(3):909–968
31. Ellison G, Fudenberg D (1993) Rules of thumb for social learning. *J Polit Econ* 101(4):612–643
32. Ellison G, Fudenberg D (1995) Word-of-mouth communication and social learning. *Q J Econ* 110:93–126
33. Fagnani F, Zampieri S (2008) Randomized consensus algorithms over large scale networks. *IEEE J Sel Areas Commun*, forthcoming
34. Farrell J, Rabin M (1996) Cheap talk. *J Econ Perspect* 10(3):103–118
35. Fortunato S, Stauffer D (2006) Computer simulations of opinions and their reactions to extreme events. In: *Albeverio S, Jentsch V, Kantz H (eds) Extreme events in nature and society, 2005*. Springer, Berlin
36. Foster D, Vohra RV (1997) Calibrated learning and correlated equilibrium. *Games Econ Behav* 21:40–55
37. Foster D, Vohra RV (1999) Regret in the online decision problem. *Games Econ Behav* 29:7–35
38. Fudenberg D, Levine DK (1998) *The theory of learning in games*. MIT Press, Cambridge
39. Fudenberg D, Tirole J (1991) *Game theory*. MIT Press, Cambridge
40. Gale D, Kariv S (2003) Bayesian learning in social networks. *Games Econ Behav* 45(2):329–346
41. Galeotti A, Ghiglino C, Squintani F (2010) Strategic information transmission in networks. Working paper
42. Galton F (1907) *Vox populi*. *Nature* 75:450–451
43. Gilboa I, Postlewaite A, Schmeidler D (2007) Rationality of belief. Or: why Savage’s axioms are neither necessary nor sufficient for rationality. Working paper
44. Gilboa I, Schmeidler D (1995) Case-based decision theory. *Q J Econ* 110:605–639
45. Gilboa I, Schmeidler D (2001) *A theory of case-based decisions*. Cambridge University Press, Cambridge

46. Gintis H (2009) *The bounds of reason: game theory and the unification of the behavioral sciences*. Princeton University Press, Princeton
47. Gladwell M (2000) *The tipping point: how little things can make a big difference*. Little Brown
48. Glauber RJ (1963) Time-dependent statistics of the Ising model. *J Math Phys* 4:294–307
49. Golosov M, Lorenzoni G, Tsyvinski A (2009) Decentralized trading with private information. Working paper
50. Golub B, Jackson MO (2010) How homophily affects diffusion and learning in networks. Unpublished manuscript
51. Golub B, Jackson MO (2007) Naïve learning in social networks and the wisdom of crowds. *Am Econ J, Microecon* 2(1):112–149
52. Grossman SJ (1977) The existence of future markets, and my special expectations and informational externalities. *Rev Econ Stud* 44(2):431–449
53. Grossman S, Stiglitz JE (1980) On the impossibility of informationally efficient markets. *Am Econ Rev* 70(3):393–408
54. Gul F (1998) A comment on Aumann's Bayesian view. *Econometrica* 66(4):923–927
55. Hagenbach J, Koessler F (2010) Strategic communication networks. Working paper
56. Hart S, Mas-Colell A (2000) A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68(5):1127–1150
57. Haviv M, Van Der Heyden L (1984) Perturbation bounds for the stationary probabilities of a finite Markov chain. *Adv Appl Probab* 16(4):804–818
58. Hayek FA (1945) The use of knowledge in society. *Am Econ Rev* 35(4):519–530
59. Hegselmann R, Krause U (2002) Opinion dynamics and bounded confidence models, analysis, and simulations. *J Artif Soc Soc Simul* 5:1–33
60. Holley R, Liggett TM (1975) Ergodic theorems for weakly interacting systems and the voter model. *Ann Probab* 3:643–663
61. Jackson MO (2008) *Social and economic networks*. Princeton University Press, Princeton
62. Jadbabaie A, Lin J, Morse S (2003) Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Trans Autom Control* 48(6):988–1001
63. Krause U (2000) A discrete nonlinear and nonautonomous model of consensus formation. In: Elaydi S, Ladas G, Pospodja J, Rakowski J (eds) *Communications in difference equations*. Gordon and Breach, Amsterdam
64. Lee I (1993) On the convergence of informational cascades. *J Econ Theory* 61:395–411
65. Liggett TM (1985) *Interacting particle systems*. Springer, New York
66. Lobel I, Ozdaglar A, Feijer D (2010) Distributed multi-agent optimization with state-dependent communication. LIDS report 2834
67. Lobel I, Ozdaglar A (2010) Distributed subgradient methods for convex optimization over random networks. LIDS report 2800. *IEEE Trans Autom Control*, to appear
68. Lorenz J (2005) A stabilization theorem for dynamics of continuous opinions. *Physica A* 355:217–223
69. Mobilia M, Petersen A, Redner S (2007) On the role of zealotry in the voter model. *J Stat Mech Theory Exp* 128:447–483
70. Morgan J, Stocken P (2008) Information aggregation in polls. *Am Econ Rev* 98(3):864–896
71. Nedic A, Ozdaglar A (2009) Distributed subgradient methods for multi-agent optimization. *IEEE Trans Autom Control* 54(1):48–61
72. Olshevsky A, Tsitsiklis JN (2009) Convergence speed in distributed consensus and averaging. *SIAM J Control Optim* 48(1):33–55
73. Ostrovsky M (2009) Information aggregation in dynamic markets with strategic traders. Working paper
74. Quine WV (1969) Natural kinds. In: Rescher N (ed) *Essays in honor of Carl G. Hempel*. Reidel, Dordrecht
75. Radner R (1972) Existence of equilibrium of plans, prices and price expectations in a sequence of markets. *Econometrica* 40(2):289–303
76. Radner R (1982) Equilibrium under uncertainty. In: Arrow KJ, Intiligator MD (eds) *Handbook of mathematical economics*, vol 2. Elsevier, Amsterdam, pp 923–1006
77. Richerson PJ, Boyd R (2005) *Not by genes alone: how culture transformed human evolution*. The University of Chicago Press, Chicago
78. Robinson J (1951) An iterative method of solving a game. *Ann Math* 54:296–301
79. Samuelson L (1997) *Evolutionary games and equilibrium selection*. MIT Press, Cambridge
80. Sandholm W (2010) *Population games and evolutionary dynamics*. MIT Press, Cambridge
81. Savage LJ (1954) *The foundations of statistics*. Wiley, New York
82. Schmeidler D (1989) Subjective probability and expected utility without additivity. *Econometrica* 57(2):571–587
83. Schweitzer PJ (1968) Perturbation theory and finite Markov chains. *J Appl Probab* 5(2):401–413

84. Simon HA (1957) *Models of man*. Wiley, New York
85. Smith L, Sorensen P (1998) Rational social learning with random sampling. Unpublished manuscript
86. Smith L, Sorensen P (2000) Pathological outcomes of observational learning. *Econometrica* 68(2):371–398
87. Tahbaz-Salehi A, Jadbabaie A (2008) A necessary and sufficient condition for consensus over random networks. *IEEE Trans Autom Control* 53(3):791–795
88. Tsitsiklis JN (1984) Problems in decentralized decision making and computation. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology
89. Tsitsiklis JN, Bertsekas DP, Athans M (1986) Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Trans Autom Control* 31(9):803–812
90. Vives X (1997) Learning from others: a welfare analysis. *Games Econ Behav* 20:177–200
91. Watts D (2003) *Six degrees: the science of a connected age*. Norton, New York
92. Weibull J (1995) *Evolutionary game theory*. MIT Press, Cambridge
93. Weisbuch G, Kirman A, Herreiner D (2000) Market organization. *Economica* 110:411–436
94. Welch I (1992) Sequential sales, learning and cascades. *J. Finance* 47:695–732
95. Wolinsky A (1990) Information revelation in a market with pairwise meetings. *Econometrica* 58(1):1–23
96. Wu F, Huberman B (2008) How public opinion forms. In: WINE' 08