Verifiable Communication on Networks*

Germán Gieczewski†

August 2021

Abstract

This paper models the diffusion of verifiable information on a network populated by biased agents. Some agents, who are exogenously informed, choose whether to inform their neighbors. Informing a neighbor affects her behavior, but also enables her to inform others. Agents cannot lie; they can, however, feign ignorance. The model yields three main results. First, unless a large set of agents is initially informed, learning is incomplete. Second, full learning is more likely for moderate than for extreme states of the world. Third, when agents are forward-looking, concerns about learning cascades lead to an endogenous division of the population into like-minded groups that do not communicate with each other.

JEL codes: D83, D85

Keywords: Verifiable information; Networks; Social learning; Strategic communication; Learning cascades

1 Introduction

The Internet has given us unprecedented access to endless sources of information, and made it easier to share this information with others through social networks. Nevertheless, wild differences of opinion persist around politically or emotionally charged

*I am thankful to Daron Acemoglu, Glenn Ellison, Drew Fudenberg, Robert Gibbons, Catherine Hafer, Matías Iaryczower, Mihai Manea, Jacopo Perego, Juuso Toikka, Robert Townsend, Alex Wolitzky, Muhamet Yildiz and seminar participants at MIT for helpful comments and suggestions. Declarations of interest: none.

†Department of Politics, Princeton University. Address: 311 Fisher Hall, Princeton NJ 08540, USA. Phone number: 6172330642. E-mail: germang@princeton.edu.

topics, even those for which definitive evidence-based answers are available. For example, significant minorities of Americans believe that climate change is not real or not man-made (Leiserowitz et al., 2019; Howe et al., 2015; Funk et al., 2015); that vaccines (Frankovic, 2021; Moore, 2015; Omer et al., 2012) or bio-engineered foods (Funk et al., 2015) are dangerous; or that Barack Obama was not born in the United States (Morales, 2011; Jardina and Traugott, 2019). These and similar examples bear out two stylized facts. First, beliefs about charged issues are partly driven by political affiliation or other "deep" preferences, with self-identified liberals and conservatives often having divergent views (Frankovic, 2021; Dimock et al., 2015). Second, misconceptions about a substantive issue are accompanied by misconceptions about the existing body of evidence: the number of climate change skeptics is similar to the number of people who believe that scientific opinion is divided on the issue (Funk et al., 2015). Thus, while biases may play a role, failures to converge to a common truth appear to be tied to breakdowns in the transmission of information. If such breakdowns are possible even in these stark examples, the hope of reaching a consensus is even weaker when it comes to politically charged debates that have no clear-cut answer—for instance, the effects of minimum wage increases, healthcare reform, or gun control.

In this paper, I propose *motivated communication* as a cause of persistent disagreement in social discourse, and aim to characterize the extent to which it can limit the diffusion of information. I define motivated communication as the act of strategically using valid evidence to support one's position, even if said position was chosen for self-serving reasons. For instance, a climate change skeptic may cite scientific studies—perhaps outliers—to support her claims, but her position may be better explained by the fact that she works for an oil company.

More concretely, this paper studies a model of information transmission underpinned by the following broad assumptions, motivated by our setting. There is a large population of agents connected by a social network. Information is verifiable and reproducible, but initially only held by some agents. Agents have differing biases, but are able to discuss and share objective evidence if they wish to do so. When deciding what information to share with their acquaintances, they are driven by the desire to bring the views of others into alignment with their own.

Here is a concrete example. The citizens of a country in political turmoil share news on an online social network. Citizens differ in their ulterior motives: some are connected with the ruling party or are favored by its policies, while others are routinely mistreated. In addition, citizens may have hard information about the quality of the regime (e.g., a link to a video showing the president taking bribes) acquired directly from Nature—say, from reading a news website—or from other agents. Citizens have random opportunities to talk to others, and decide what links to share; a link that is shared can, in turn, be passed on by the recipient. Citizens want to bias others towards their own views, either for public or private reasons—for example, to topple the government, or so that they can spend time with their friends at a rally.

Since information consists of links to a verifiable source, agents cannot lie, but their biases tempt them to lie by omission, that is, to only share evidence that bolsters their position. For example, government supporters will avoid mentioning corruption scandals even if they know the allegations to be true. This self-censorship stifles the free diffusion of information. Moreover, in the face of this behavior, the beliefs of biased agents who remain ignorant will end up reinforcing their biases: a dissident will think her position vindicated if she meets a government supporter who fails to produce evidence supporting the government.

The model yields three main insights. The first one concerns the extent of information diffusion, that is, how close the population comes to learning the state of the world. I show that full diffusion is in general possible only when the set of initially informed agents contains agents with diverse biases. More precisely, if the true state of the world is high, then the truth spreads to all agents only if some high-bias agents are initially informed, and vice versa.

The second insight concerns the kind of information that proliferates. I show that full diffusion is more likely to obtain when the state of the world is *moderate*—roughly speaking, when the realized state is close to the mean of its distribution. More precisely, if the state of the world is symetrically distributed around 0, with support $[-1, 1]$, there is an interval $I = (-\epsilon, \epsilon) \subset [-1, 1]$ such that, if the state of the world is in I, then everyone will learn this as long as any one agent is initially informed. On the other hand, *extreme* states of the world, far from 0—paradoxically, the ones that would be more valuable to learn—will remain less known. (When the state is extreme, learning is more valuable because, so long as an agent remains uninformed, her actions will be badly mismatched with the true state.) The logic behind this result is that extreme information is only communicated in a predictable direction, which stifles its propagation: for instance, when the state is high, high-bias agents inform lower-bias agents to increase their beliefs, but low-bias agents are unlikely to pass on this information. On the other hand, news of a moderate state can cycle through the population by traveling in both directions. This implies that misinformation is more likely to persist around issues for which a conclusive answer would shift everyone's posterior beliefs in the same direction, rather than those for which the truth lies in the middle. More generally, it means that the existence of moderate evidence promotes the diffusion of information, so that conversely, if all news are extreme—or if even balanced news can always be divided into small, good-or-bad nuggets to be shared selectively—communication becomes more difficult.

The third insight concerns the *structure* of information diffusion when learning is incomplete. I show that, when the network is densely connected, failures of information diffusion are driven by concerns about learning cascades. These concerns lead to an endogenous division of the population into like-minded groups, or *segments*, that communicate with each other only in one direction. The intuition here is that, in a dense network—that is, a network in which every agent has a number of friends, some with similar biases—myopic incentives would drive agents to share their information with at least some neighbors in such a way that everyone would eventually become informed. However, as forward-looking agents understand the implications of their behavior, in equilibrium, they choose not to inform neighbors who are close in bias but who would then share information with the wrong group. Hence, in dense networks, higher sophistication leads to less information being transmitted. In sparse networks, the effect of sophistication can go either way: forward-looking agents may fear learning cascades, but they may also seek alliances of convenience, that is, they may be willing to use a disliked neighbor as a conduit to reach a like-minded player further down the network.

This paper contributes to a growing literature on strategic communication on networks. Like the present paper, the prototypical model in this literature (Galeotti et al., 2013) considers biased agents on a network who choose whether to communicate; each agent has a different bias and wants everyone's actions to match her own state-contingent bliss point. However, the most common variant of this framework (Galeotti et al., 2013; Hagenbach and Koessler, 2010; Dewan et al., 2015; Dewan and Squintani, 2018) assumes that communication is non-verifiable (cheap talk) and that there is a single round of communication, that is, agents are not allowed to pass on signals that other players have shared with them. In contrast, in this paper, agents exchange hard evidence and multiple rounds of communication are possible. As a result, more information is transmitted in my setting: while cheap talk is informative only when the difference between the agents' biases is small enough (Crawford and Sobel, 1982), verifiable communication gives rise to one-sided revelation strategies, whereby agents reveal convenient facts and hide the rest. More importantly, the learning cascades that are at the heart of my results are, by design, absent from the aforementioned papers.

Within this literature, the two papers closest to this one are Squintani (2019) and Bloch et al. (2018). They both consider verifiable messages that can be re-shared, but differ from this paper in other respects. Squintani (2019) aims to characterize optimal and/or pairwise stable networks, and finds conditions under which these networks are either a star or an ordered line. In contrast, I take the network as given and characterize the equilibrium for a large class of networks. Another difference is that, in Squintani (2019), there is a single decision-maker whose beliefs all agents are concerned with; this assumption limits the complexity of learning cascades, as there is no trade-off between informing wanted and unwanted agents. In contrast, in this paper, each agent cares about the beliefs of all other agents.

In Bloch et al. (2018), there are two types of players: unbiased types that want to match the state, and biased types whose bliss point is independent of the state. In contrast, I consider agents of differing biases who care equally about the state. The main difference, however, lies in the communication protocol. In Bloch et al. (2018), agents share information with either every neighbor, or none; and information-sharing strategies are static (that is, players cannot adjust their messages depending on when they are informed, or by whom). In contrast, I assume agents can share information with only some of their neighbors, and their strategies are allowed to depend on the history (and they typically do).

More broadly, the paper is related to the social learning literature, in which agents directly observe beliefs and update naively (DeMarzo et al., 2003; Golub and Jackson, 2010), or they learn by observing others' actions (Banerjee, 1992; Smith and Sørensen, 2000; Acemoglu et al., 2011). In either case, the choice of what to transmit to others is typically devoid of strategic considerations. A general lesson from this paper is that, when agents are patient, use information differently and care about each others' actions, concerns about spreading information to unwanted agents dampen learning. Similar insights may extend to a model of observational learning, such as Acemoglu et al. (2011), with strategic considerations.

The paper is also connected to a literature on explanations for the persistence of disagreement. For example, agents may interpret information differently (Acemoglu et al., 2016); they may be unable to separate others' priors from their information (Sethi and Yildiz, 2012); the acquisition and transmission of information may be costly (Calvó-Armengol et al., 2015; Perego and Yuksel, 2016); or agents' beliefs may be slightly misspecified (Frick et al., 2020).

Finally, in the two-player case, the communication game in this paper is similar to existing models of communication with verifiable information (Milgrom and Roberts, 1986; Glazer and Rubinstein, 2006). A common result in this literature (Milgrom and Roberts, 1986; Hagenbach et al., 2014) is that unraveling leads to full revelation. In this paper, there is only partial revelation because the amount of information held by the sender is uncertain, which allows a sender with "bad" information to feign ignorance without much being inferred by the receiver. My model of verifiable communication differs more markedly from the Bayesian persuasion literature (Gentzkow and Kamenica, 2011), in which the sender is allowed to commit to a revelation strategy *before* learning the state. In my setting, this assumption would be unnatural, as the sender knows her available information before meeting the next receiver.

The paper proceeds as follows. Section 2 presents the model. Section 3 considers the case of myopic agents as a benchmark. Section 4 analyzes the case of forward-looking agents. Section 5 discusses the results. Appendix A contains additional analysis of the model in Section 4. Appendix B contains all the proofs.

2 The Model

Preliminaries

There is a set $N = \{1, \ldots, n\}$ of players connected by an undirected network $G \subseteq N \times N$. Each agent i has a type $b_i \in \mathbb{R}$ denoting her bias. Without loss of generality, we assume $b_1 \leq \ldots \leq b_n$. We denote the set of i's neighbors by N_i.

There is a state of the world θ that is fixed over time, and distributed according to a c.d.f. F with support contained in $[-1, 1]$ and mean 0. At the beginning of the game, Nature draws θ and shows it to a subset S of players with probability $\gamma(S)$, where $\gamma : 2^N \to [0,1]$ is a probability distribution.[1] We denote by $\gamma_0 = 1 - \gamma(\emptyset)$ the probability that Nature informs a non-empty set of players, and assume $\gamma_0 \in (0,1)$. The assumption that $\gamma_0 < 1$ guarantees that full unraveling in the style of Milgrom and Roberts (1986) cannot obtain, as the players do not infer too much from being uninformed. It can be taken as reflecting a world in which opportunities to communicate strategically with neighbors do not arrive too often and are thus somewhat unexpected.

Later in the paper we will vary γ_0 as other parameters are held constant. This should be taken to mean that $\gamma(S)|_{S \neq \emptyset}$ varies proportionally with γ_0.

Actions

After Nature informs the players in S, a game is played over discrete and infinite time: $t = 0, 1, \ldots$. In each period t, agent i chooses an action a_{it}. i's utility function is

$$U_i(\theta, a_i, a_{-i}) = -\sum_{t=0}^{\infty} \delta^t \left[\sum_{j \in N} (a_{jt} - \theta - b_i)^2 \right].$$

This functional form, which is common in the literature (Galeotti et al., 2013), ensures that i's optimal action at time t is the expected value of $\theta + b_i$ given her information; and, moreover, that i wants to induce each player $j \neq i$ to play according to i's own bias, that is, i wants to induce $a_{jt} = \theta + b_i$, whereas j would choose $a_{jt} = \theta + b_j$ if fully informed.[2] The discount factor $\delta \in (0,1)$ is the same for all agents.

Communication

In each period t, with probability $\alpha \in (0,1)$, Nature creates an opportunity to communicate. More precisely, Nature picks a player, each with probability $\frac{\alpha}{n}$, as the potential sender; with probability $1 - \alpha$, no one is chosen. If i is chosen, she is given an opportunity to message a neighbor j; Nature picks each neighbor to be the

[1] All the results go through if Nature instead shows a noisy signal s of θ, as long as all informed players see the same realization. The conditional expectation $E(\theta|s)$ would then take the role of θ in the rest of the paper.

[2] The assumption that i's own action and the actions of other players have the same weight in i's utility function is immaterial to the results. Heterogeneous weights over the actions of other players can be incorporated into the model without complications.

potential receiver with probability $\frac{1}{|N_i|}$.[3] When i meets j, i can either share θ (if she is informed) or share nothing, potentially as a function of her private history: $m(\theta, i, j, t, h_i^t) = 1$ or 0, where 1 denotes a message that reveals θ and 0 denotes no message. If j receives θ she learns θ as well as the fact that i sent the message at time t, but if i shares nothing, j does not know that a meeting has occurred. In particular, meetings are one-way: when i has an opportunity to talk to j, j is not able to talk back to i. We also allow i to meet j even if j is already informed.

Intuitively, we can imagine i having a one-way meeting with j to mean that i has "logged on" to the social network and considered messaging j, but j is unaware of this unless she gets a message. In other settings (for example, face-to-face meetings) it would be more natural to assume that meetings are two-way. The only difference this would introduce is that j's posterior would jump when she meets another player i who sends no message, but would stay constant otherwise. In contrast, with one-way meetings, posteriors evolve independently of when meetings happen (since j's posterior under ignorance responds to the *expected* number of meetings in which she's been the receiver), which simplifies the analysis.

Finally, we assume that actions a_{it} are not observable by other players and payoffs are not observable, so j can only learn from i through i's messages. The ramifications of learning from other players' behavior (Banerjee, 1992) or outcomes (Wolitzky, 2018) have been studied elsewhere; they are orthogonal to the point of this paper.

Timing

The game proceeds as follows. Before period 0, Nature determines the value of θ and informs a set of players S. Then, at each time $t \geq 0$, Nature generates a meeting (or not); a message is sent (or not); and the players choose their actions a_{it}.

Network Properties

A *path* is a sequence of players such that every two consecutive elements share a link. A network is *connected* if there is a path between any two players. A network is *complete* if there is a direct link between any two players. A *forest* is a network containing no cycles. A *tree* is a connected forest.

[3] We may allow meeting probabilities that vary across links, as well as multiple meetings in the same period, without substantially altering the results.

Given an open interval $I = (a, b) \subseteq [b_1, b_n]$, we define an *interval* (N_I, G_I) of the graph (N, G) as the induced subgraph with vertex set $N_I = \{i \in N : b_i \in I\}$. Given a network (N, G) and a value of $\gamma_0 > 0$, we say (N, G) is *well-connected* if all its intervals of size $\frac{1-\gamma_0}{2-\gamma_0}$ are nonempty and connected.

Remark 1. If (N, G) is well-connected, then it is connected, and all its intervals of size larger than $\frac{1-\gamma_0}{2-\gamma_0}$ are connected.

The concept of well-connectedness, which will be central to our results, captures the notion that players with close biases are able to share information with each other without needing to use highly dissimilar agents as intermediaries. While a stronger condition than connectedness, it is more likely to hold if links between ideologically similar agents are disproportionately common, that is, if the network displays homophily (McPherson et al., 2001).

Equilibrium Concept

Our solution concept is sequential equilibrium. We denote a general strategy profile by $\sigma = (a(\theta, i, t), m(\theta, i, j, t, h_i^t))$, where $a(\theta, i, t)$ denotes the action chosen by agent i at time t, given her observation of $\theta \in [-1, 1] \cup \{\emptyset\}$ ($\theta = \emptyset$ denotes that i is uninformed at time t);[4] and $m(\theta, i, j, t, h_i^t)$ is the probability that i shares her information with agent j if able to, conditional on her observation of the state θ and her history h_i^t, with the restriction that $m(\emptyset, \cdot) \equiv 0$, i.e., no message is sent if i is uninformed.[5]

We denote the beliefs held by an agent i at a history h_i^t by $\mu(i, t, h_i^t)$. Since we will mainly be interested in i's beliefs about θ, rather than about the entire history of play, we also define $p(\theta, i, t)$ as the probability that player i is informed at the end of period t, if the true state is θ. (In general $p(\theta, i, t)$ depends on θ because the players' communication strategies condition on the state.) We denote the limit of $p(\theta, i, t)$ as $t \to \infty$ by $p(\theta, i)$. In addition, we denote i's expectation about θ under ignorance by $\overline{\theta}(i, t)$, i.e., $\overline{\theta}(i, t) = \frac{\int_{-1}^{1} \theta(1-p(\theta,i,t)dF(\theta)}{\int_{-1}^{1}(1-p(\theta,i,t))dF(\theta)}$, and we denote its limit as $t \to \infty$ by $\overline{\theta}(i)$.[6]

We can immediately solve for the players' equilibrium actions: it follows from the definition of the utility functions that each agent i should optimally match her action

[4] In principle a_{it} could depend on h_i^t, but it will depend only on θ, i and t in any equilibrium, since a_{it} is not observed by anyone else, and (θ, i, t) is a sufficient statistic for i's posterior.

[5] Unlike a_{it}, m_{it} may depend on h_i^t, since h_i^t may contain information about who else is informed.

[6] $p(\theta, i)$ is well-defined since $p(\theta, i, t) \leq 1$ is weakly increasing in t. $\overline{\theta}(i)$ is pinned down by Bayes' rule, as the assumption $\gamma_0 < 1$ guarantees that i remains uninformed with probability at least $1 - \gamma_0$.

a_{it} to the state of the world—or her expectation of it, if uninformed—plus her bias.

Remark 2. In any equilibrium, $a_{it} = \theta + b_i$ if i is informed at time t, and $a_{it} = \overline{\theta}(i,t) + b_i$ otherwise.

3 Benchmark: Myopic Preferences

As a prelude to analyzing the full model, we solve the case of myopic preferences ($\delta = 0$) as a benchmark. The analysis is much simpler in this case: under myopic preferences, when i meets j, i cares only about j's current action, as opposed to the current and future actions of all other players. Hence, i shares θ with j at time t if and only if doing so increases $E_{it}(-(\theta + b_i - a_{jt})^2 | \theta, h_i^t)$.

There are two reasons for studying the myopic case. First, the assumption of myopic preferences is descriptively accurate in some applications. Agents may care about the private effects of actions with public consequences: for example, if i and j are parents with children in the same school, i may want j to vaccinate his children so that i's children are safe while they are near j's. Similarly, in the context of a protest, i may want to change j's mind so that j will attend a rally with i; i may care more about the consumption value of the rally than its political consequences. Alternatively, agents may be truly myopic and fail to anticipate the downstream consequences of their messages—a real concern in such a complex model. Finally, in practice, communication is not always instrumental; instead, agents may be simply driven by a desire to "win arguments".

Second, the myopic case serves as an instructive bridge between my setting and the extant literature on strategic communication on networks, most of which assumes a single round of communication. In a model with myopic preferences and multi-round communication, learning cascades are possible but the players do not take them into account—rather, they behave *as if* the current round of communication were the last.

Proposition 1 characterizes the agents' equilibrium behavior under myopic preferences.

Proposition 1 (Myopic Equilibrium). *Assume that $\delta = 0$. In any equilibrium, if i is informed at time t, $m(\theta, i, j, t) = 1$ if either $\theta > \overline{\theta}(j,t)$ and $2(b_j - b_i) < \theta - \overline{\theta}(j,t)$, or $\theta < \overline{\theta}(j,t)$ and $2(b_i - b_j) < \overline{\theta}(j,t) - \theta$.*

In addition, if F admits a continuous density f and α is low enough, then the equilibrium is unique.

The first part of the Proposition reflects the following intuition: if i can inform j at time t, she faces a choice between revealing θ and letting j retain his posterior under ignorance, which has mean $\overline{\theta}(j,t)$. If $b_i > b_j$, i will want to reveal θ when doing so increases j's action. But θ may also be shared when it has the opposite effect if $|b_i - b_j|$ is small relative to $|\theta - \overline{\theta}(j,t)|$: in that case, i's incentive to bias j is trumped by her desire to prevent him from choosing an action badly mismatched with the state of the world. An important special case is if $\theta = 1$ and γ_0 is close to 0 (that is, information is scarce, which guarantees that posteriors under ignorance are close to 0): then i informs j if $b_j < b_i + \frac{1}{2}$. Conversely, for $\theta = -1$, i informs j if $b_j > b_i - \frac{1}{2}$. Finally, taking α to be low enough guarantees equilibrium uniqueness by ensuring that $\overline{\theta}(i,t)$ is not too sensitive to the strategies of other players in period t.

Corollary 1 characterizes the extent of information transmission resulting from these equilibrium strategies, in two instructive cases.

Figure 1: $p(\theta, j, t)$ on the equilibrium path for $n = 4$, $j = 2$, $\theta \sim U[-1, 1]$, $\delta = 0$, $\gamma_0 = 0.75$, $\alpha = 0.3$, and $b_{i+1} - b_i > \frac{1}{2-\gamma_0}$ for all i

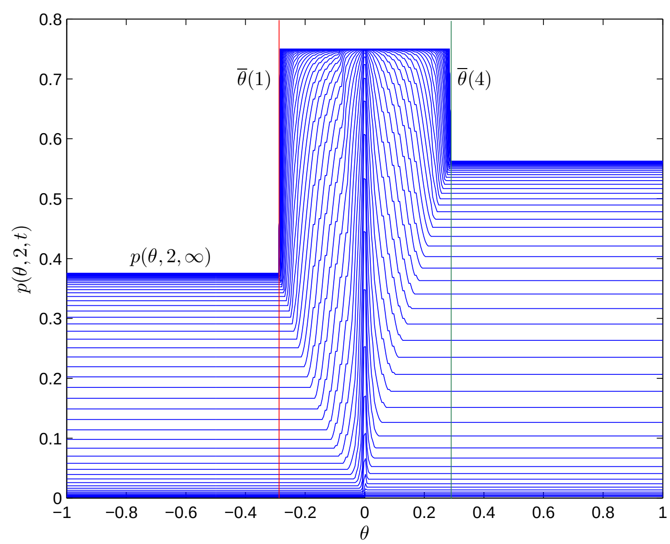

Corollary 1.

(a) Assume that $\delta = 0$. Suppose (N, G) is a well-connected network, and the state is binary: $P(\theta = 1) = P(\theta = -1) = 0.5$. Then there is full diffusion: $p(\theta, i) = \gamma_0$ for all θ, i in any equilibrium.

Figure 2: Posterior beliefs for $i = 1, 2, 3, 4$

(a) $p(\theta, i, \infty)$ (b) $\overline{\theta}(i, t)$

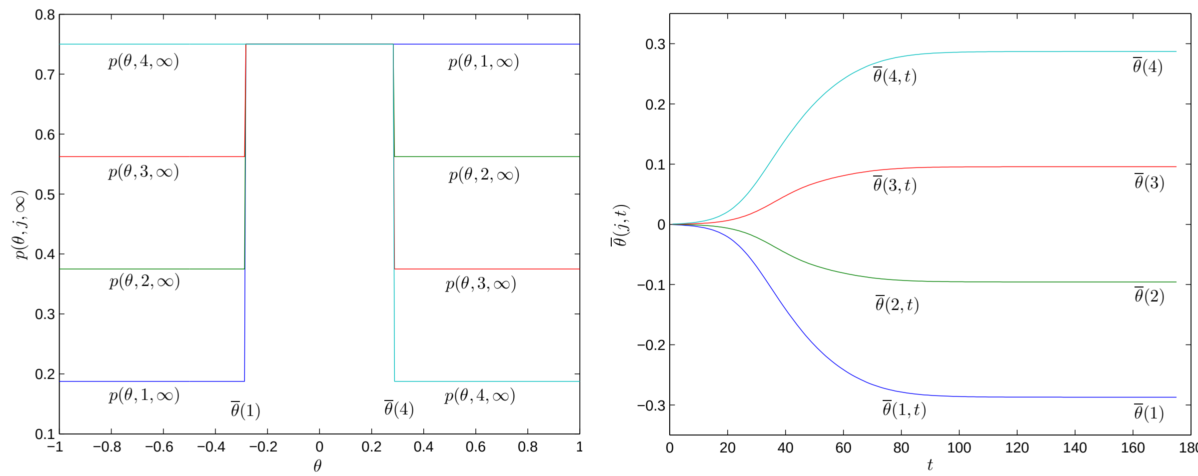

(b) On the other hand, suppose (N, G) is such that $b_{i+1} - b_i > \frac{1}{2-\gamma_0}$ for all i, i is linked to $i+1$ for all i, and 1 and n are linked to everyone. Then, in any equilibrium,

(i) there is full diffusion for moderate states: $\overline{\theta}(1) < \overline{\theta}(n)$, and $p(\theta, j) = \gamma_0$ for $\theta \in \big(\overline{\theta}(1), \overline{\theta}(n)\big)$.

(ii) There is partial diffusion for extreme states: for $\theta < -\frac{\gamma_0}{2-\gamma_0}$, j is informed iff S is nonempty and $\min(S) \leq j$. For $\theta > \frac{\gamma_0}{2-\gamma_0}$, j is informed iff S is nonempty and $\max(S) \geq j$.

The intuition behind Corollary 1 is as follows. If the network is well-connected, then for every player i, her closest neighbors—in either direction—are close enough that she would like to inform them no matter what the state is. But the receivers will also find close neighbors to inform, and so on. Thus information spreads through the whole network.

Part (b) presents a contrasting case in which players are separated by large bias gaps (note that, if $b_{i+1} - b_i > \frac{1}{2-\gamma_0} > \frac{1-\gamma_0}{2-\gamma_0}$ for any i, the network cannot be well-connected). In this case, information is only shared in a self-serving fashion: $i > j$ only informs j in order to increase her posterior belief, and vice versa. Both partial and full diffusion are possible, depending on the state. Figure 1 illustrates this case in

Figure 3: Myopic transmission with large gaps

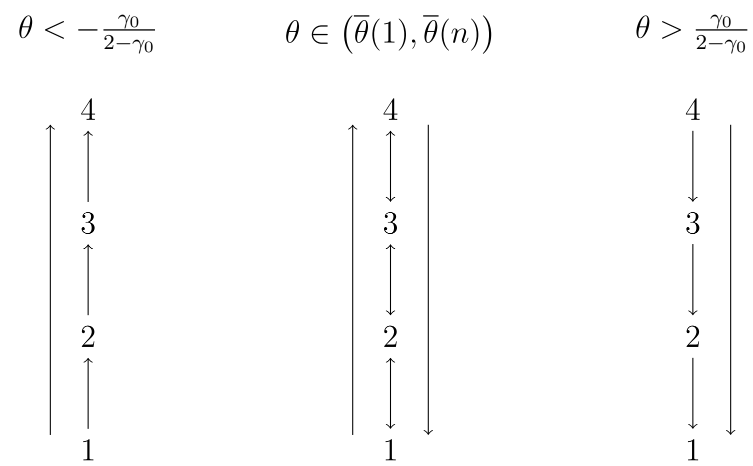

a complete 4-player network. The graph shows player 2's probability of learning the state over time. This probability is increasing in t, but converges to different values depending on θ. If θ is low, only player 1 will tell her; if θ is high, only players 3 and 4 will tell her; if θ is intermediate, she will be told no matter who is initially informed.

The transmission of moderate information in the case of large gaps is enabled by the following fact: over time, the posteriors under ignorance of extreme types tend to reinforce their biases. In particular, $\overline{\theta}(n,t) > \overline{\theta}(1,t)$ and $\overline{\theta}(n,t)$ is increasing in t, while $\overline{\theta}(1,t)$ is decreasing in t, as shown in Figure 2b. The reason is that a low-bias player, for instance, expects others to send her messages that will drive her beliefs up, and so expects to learn the state with higher likelihood if it is high (Figure 2a). If she is not informed, she guesses that θ is low and that is why no one has told her.

In particular, this means that n's long-run mean posterior under ignorance is higher than 1's: $\overline{\theta}(n) > \overline{\theta}(1)$. When θ lies between these two values, every player wants to share θ with both 1 and n, in order to moderate their actions. And once 1 and n are both informed, they inform everyone else: for any player i between them, either 1 or n will be willing to tell i. This is represented diagramatically in Figure 3.

Finally, note that the variability of posterior beliefs increases with information abundance: the higher γ_0 is, the more uninformed players can infer from their lack of information, which prompts other players to share more values of θ. As a result, full diffusion is attained in the limit as $\gamma_0 \to 1$, due to an argument related to the classic unraveling result under verifiable communication (Milgrom and Roberts, 1986). This is the subject of Corollary 2.

Corollary 2. *As* $\gamma \to 0$, $p(\theta, i) \to_{|| ||_\infty} 0$ *and* $\overline{\theta}(i) \to 0$ *for all* i*. Conversely, assume* (N, G) *is such that* 1 *and* n *are linked to everyone, and* F *has full support in* $[-1, 1]$*. Then, as* $\gamma_0 \to 1$, $p(\theta, i) \to_{|| ||_1} 1$ *for all* i*.*

4 Forward-Looking Players

Learning Cascades

We now consider the case of forward-looking players who care about the future actions of the entire population. Two new considerations appear, both related to *learning cascades*. First, an agent i may want to inform j, even if this negatively affects i's payoff from j's action, so that j can then inform a third player k. Second, i may want to hide information from j, even when it improves j's action, to prevent j from informing k.

These incentives are illustrated in Figures 4 and 5. In Figure 4, we assume that $(b_1, b_2, b_3) = (0, 1, 1.6)$, $\gamma(\{2\}) = \gamma_0$ (only player 2 may be initially informed), and $\gamma_0 > 0$ is small. Suppose $\theta = 1$ is realized and 2 is informed. Clearly, if 3 becomes informed, he will inform 1. While 2 has a myopic incentive not to inform 3, she does wants to inform 1, and the latter incentive dominates. Hence, 2 informs 3.

Figure 4: Learning cascades encourage communication

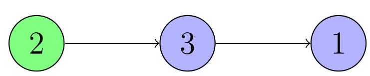

In Figure 5, assume that $(b_1, b_2, b_3) = (0, 0.4, 0.8)$; $\gamma(\{1\}) = \gamma_0$ is small; and $\theta = 1$ is realized. Then agent 1 has a slight incentive to inform 2, and 2 has a slight incentive to inform 3, but 1 is opposed to 3 becoming informed. Hence, if 1 were myopic, she would inform 2; but, being far-sighted, she chooses not to.

Figure 5: Learning cascades discourage communication

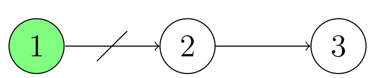

These examples show that learning cascades may result in either more or less communication. However, there is an asymmetry between the two. In Figure 4, the result depends on 2 *not* being linked to 1 (as otherwise 2 could message 1 directly and bypass 3); in Figure 5, the result is robust to adding links. This insight applies more broadly. In a sparse network, the first motive can dominate, as in order to reach faraway agents a player may be forced to use undesirable neighbors as intermediaries. On the other hand, in a dense network, the first motive vanishes (as direct communication is possible) but the second remains, hence discouraging communication.

In arbitrary networks, the set of equilibria of the communication game with forward-looking players may be very complex. For instance, there may be multiple equilibria that are Pareto-ordered; multiple equilibria that are not Pareto-ordered; equilibria in which two informed players want to be the first to send a message (preemption), or the last (war of attrition); equilibria in which a player's communication strategy depends, even in the long run, on whom she has been informed by; and there may not be a pure strategy equilibrium. A detailed discussion of these possibilities is found in Appendix A.

What I show next is that, on a large class of networks, much of this complexity can be ruled out. Indeed, a reasonable condition on the network structure (well-connectedness), combined with assumptions about the communication protocol (*mutual observability* and an *activity rule*) guarantee the existence of a tractable equilibrium with several appealing properties. Briefly, this equilibrium is coalition-proof in a certain sense; it is robust to small changes in the network structure; and the extent of information diffusion it induces is uniquely determined and admits a simple description.

Communication on Well-Connected Networks

Next, we define a class of strategy profiles that the equilibrium of the game will belong to. For ease of exposition, assume a binary state of the world, with $P(\theta = -1) = P(\theta = 1) = 0.5$.

A sequence of bias profiles $\hat{b}_1 < \hat{b}_2 < \ldots < \hat{b}_l \subseteq [b_1, b_n]$ partitions the population N into *segments* $\hat{M}_1, \hat{M}_2, \ldots, \hat{M}_{l+1}$, where $\hat{M}_i = N_{[\hat{b}_{i-1}, \hat{b}_i)}$.[7] A strategy profile σ is *segmented* if there are two sequences $\underline{b}_1 < \ldots < \underline{b}_{\underline{k}}$ and $\overline{b}_1 < \ldots < \overline{b}_{\overline{k}}$, defining respective partitions $\underline{M}_1, \ldots, \underline{M}_{\underline{k}+1}$ and $\overline{M}_1, \ldots, \overline{M}_{\overline{k}+1}$ such that:

(i) if $\theta = -1$ and $\min(S) \in \underline{M}_i$, then all the players in $\underline{M}_i$ and higher segments become informed with probability converging to 1, and all the players in lower

[7] Implicitly we have taken $\hat{b}_0 = -\infty$, $\hat{b}_{l+1} = \infty$.

segments remain uninformed.

(ii) if $\theta = 1$ and $\max(S) \in \overline{M}_j$, then all the players in $\overline{M}_j$ and lower segments become informed with probability converging to 1, and all the players in higher segments remain uninformed.

The essence of a segmented strategy profile is that the players are split into groups according to their biases. Groups share information internally, but information only travels between groups in a single direction: when the state is high, higher-bias groups only talk to lower groups, and the opposite happens if the state is low. This is illustrated in Figure 6. There, if $\theta = 1$ and 6 is initially informed, information spreads in 6's segment and the one below it, but not the one above it.

Figure 6: Outcome of a segmented strategy profile with $S = \{6\}$ and $\theta = 1$

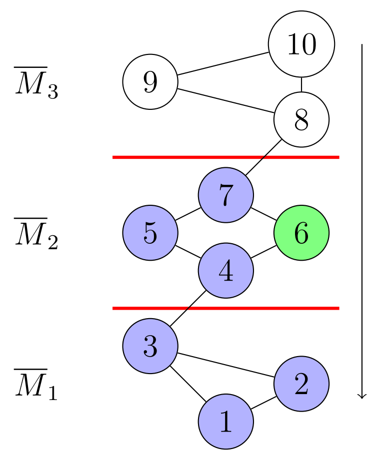

We call a segmented strategy profile *natural* if, for $\theta = 1$ and γ_0 low enough, each $i \in \overline{M}_l$ i prefers that all the members of $\overline{M}_l$ above her be informed (compared to none of them being informed), and for any higher segment $\overline{M}_{l'}$ ($l' > l$) prefers that no members of $\overline{M}_{l'}$ be informed (compared to all being informed); plus the analogous condition for $\theta = -1$.

Remark 3. For generic values of $(b_1, \ldots, b_n)$, there is a unique set of segments compatible with a natural strategy profile, determined recursively. For $\theta = 1$,

(i) Start with player n who is in the top segment. For each $i = n-1, n-2, \ldots$ let $A_i = \{i+1, \ldots, n\}$, and let b_{A_i} be the average bias of players in A_i. Then, if

$$b_{A_i} < b_i + \frac{1}{2},$$

i is in the top segment and we proceed to $i-1$. If $b_{A_i} > b_i + \frac{1}{2}$, the top segment is $\{i+1, \ldots, n\}$, and we denote $i = n_2$. Proceed to step 2.

(ii) In step k ($k \geq 2$), n_k is in the k-th highest segment. For $i = n_k - 1, n_k - 2, \ldots$ let $A_i = \{i+1, \ldots, n_k\}$. If $b_{A_i} < b_i + \frac{1}{2}$, i is in the k-th highest segment and we proceed to $i-1$. If $b_{A_i} > b_i + \frac{1}{2}$, the k-th segment is $\{i+1, \ldots, n_k\}$, and we denote $i = n_{k+1}$. Proceed to step $k+1$.

The algorithm for $\theta = -1$ is analogous.

Before presenting our main results, we make two additions to the game defined in Section 2. We say a communication game with an *activity rule* is one in which, if—since the last time anyone was informed—each informed player has declined at least K opportunities to inform each of her uninformed neighbors, then all meetings stop forever.[8] A communication game with *mutual observability* (and an activity rule) is one in which, at all times, informed players automatically observe the set of other informed players (and the status of the activity rule) in addition to knowing θ.[9]

Proposition 2 provides an equilibrium characterization for the communication game with forward-looking players, mutual observability and an activity rule.

Proposition 2. *Suppose that (N, G) is well-connected. Then, for generic $(b_1, \ldots, b_n)$, for γ_0 low enough and δ close enough to 1:*

(i) The game has a unique equilibrium which is natural and in pure strategies.

[8]For γ_0 low, the activity rule serves purely as an equilibrium selection device. Indeed, for every equilibrium of the game with an activity rule, there is a payoff-equivalent equilibrium of the game without it, as follows: while the activity rule would not have bound, play the same strategies. After the activity rule would have bound, no more messages are sent. If anyone deviates after the activity rule would have bound, play in the continuation as if this deviation had happened at the last opportunity before the activity rule bound.

[9]Even without the assumption of mutual observability, players still make inferences about the state of the game: e.g., if i informs j, j knows that i is informed and the activity rule has reset. Moreover, informed players could share additional information through the timing of superfluous messages, i.e., they could have have a "Morse code" that re-sends θ at specific times to indicate information about the state of the game, or manipulates the timing of the first message to the same effect. The extent of learning by informed players about the state of the game is thus hard to control in general. Assuming mutual observability amounts to assuming that the players exogenously receive all the information about the state of the game that they could stand to gain from such communication, which makes any attempts at further learning irrelevant. See Appendix A for partial results without this assumption.

(ii) The threshold sequences satisfy $\underline{b}_{i+1} - \underline{b}_i \geq \frac{1}{2}$ *and* $\overline{b}_{i+1} - \overline{b}_i \geq \frac{1}{2}$ *for all* i*. In addition, if the biases* $b_1, \ldots, b_n$ *are roughly uniformly distributed in* $[b_1, b_n]$*, then* $\underline{b}_{i+1} - \underline{b}_i \approx 1$ *and* $\overline{b}_{i+1} - \overline{b}_i \approx 1$ *for all* i*.*[10]

The intuition behind the equilibrium is as follows. Suppose that $\theta = 1$ and n is informed. Clearly n wants to tell everyone; if the network has enough links, everyone becomes informed in the long run. Now suppose that $n-1$ is informed. $n-1$ wants to inform all the lower-bias players, and also wants to inform n if $b_n - b_{n-1}$ is small enough. Hence everyone becomes informed. If we look successively at lower-bias players, eventually some $n' < n$ is reached that would rather not inform n. However, n' understands that telling anyone in $\{n'+1, \ldots, n\}$ leads to everyone learning, so the choice is whether or not to inform the whole group. Since n' is still willing to inform most members of the group, everyone learns. Eventually, we reach some $n'' < n'$ that would rather not inform $\{n''+1, \ldots, n\}$. $\{n''+1, \ldots, n\}$ then becomes the top segment. Now $n''-1$ is happy to inform n'' because she knows that information will not flow upwards from n'', so the process restarts. This is the logic giving rise to natural strategies.

Proposition 2 has several implications. First, for any two networks (N, G), (N, G'), both well-connected and defined on the same set of agents, the set of agents informed in equilibrium is identical, because the natural segments are uniquely determined by the distribution of biases. Therefore fine details of the network structure are irrelevant in the long run, as long as the network has enough links to be well-connected.

Second, if we compare Proposition 2 to Corollary 1.(a), the clear takeaway is that sophisticated agents transmit less information: even though their myopic incentives would lead them to fully disseminate the state, in equilibrium, their behavior leads to only partial diffusion whenever the distribution of biases is wide enough that there are multiple segments. Instead, the results in Proposition 2 hew more closely to the "large gaps" case of Corollary 1.(b), even when the bias gaps between neighbors are small. The reason is that forward-looking agents think about informing groups, rather than individuals, and large enough groups always have very different average biases.

Finally, there is a sense in which the equilibrium is coalition-proof. Say all the informed players are members of a segment $\hat{M}_l$. By construction, these players have a common interest in informing all the lower segments and none of the upper segments. In equilibrium, this is exactly what they do.

[10] A rigorous statement is given in the Appendix.

Figure 7: A pathological non-segmented equilibrium

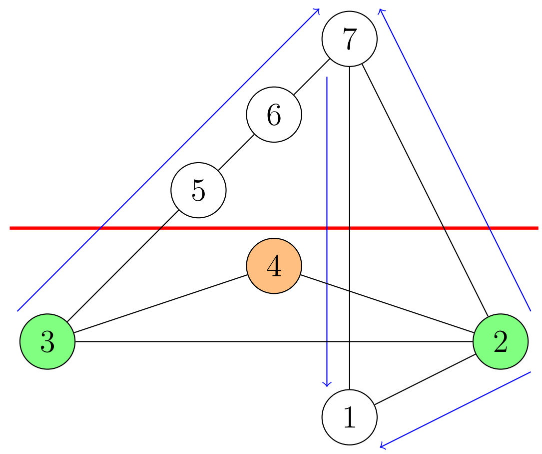

Under some additional conditions, the existence of a natural equilibrium can also be guaranteed in the communication game without mutual observability or an activity rule (see Proposition 5 in Appendix A). However, other equilibria may arise that are not segmented or not natural. An example is given in Figure 7. Assume that $\gamma(\{2,3\}) > 0$ is small and $\gamma(S) = 0$ otherwise, and $\theta = 1$. The dividing red line denotes that $\{5,6,7\}$ communicate with each other, but $\{1,2,3,4\}$ would rather not share information with them. That is, in a natural segmented equilibrium, $\{1,2,3,4\}$ and $\{5,6,7\}$ are separate segments.

Without mutual observability or an activity rule, there may be a non-natural equilibrium in which 3 talks to 5 and 2 talks to 7, that is, $\{1,2,\ldots,7\}$ forms a single segment. The reason is that, if 2 expects 3 to share information with the upper segment, then 2's own message to 7 has no marginal impact on long-run information diffusion, and vice versa. Instead, such messages only have a short-term effect, which may be positive: for instance, 3 expects that 7 will become informed through 2 sooner than through the path $3 \to 5 \to 6 \to 7$, so informing 5 mainly amounts to informing $\{5,6\}$ a bit sooner, which may be desirable. Similarly, if 2 has not yet been able to inform 1, she may inform 7 to speed up transmission to 1.

There may also be a non-segmented equilibrium if we instead assume $\gamma(\{4\}) > 0$ in the same example. The reason is that, counterintuitively, if 4 is the only informed player it may be optimal to inform nobody. Indeed, although *a priori* 4 should be more willing to share information than 2 and 3, she knows that informing either of them will trigger a continuation that is bad for all of them. Finally, note that these equilibria are not ex post coalition-proof: $\{1,2,3,4\}$ would all prefer to move to a natural segmented equilibrium, but cannot necessarily do so. In a sense, then, they constitute a coordination failure.

Such pathological equilibria rely on the players' uncertainty about who else is informed, or about whom other informed players *will* inform. Mutual observability removes the first kind of uncertainty, while an activity rule removes the second by allowing the players to "lock in" a desirable information distribution by staying silent for a number of periods.

Diffusion of Moderate States

For the sake of tractability, Section 4 has so far focused on the case of a binary state. We now discuss the case of a continuous state, and recover results regarding the diffusion of moderate states similar to those in Section 3.

Assume now that F admits a density f, symmetric around 0 and with full support in $[-1,1]$, and that the distribution of biases is also symmetric around 0. As before, we consider the communication game with mutual observability and an activity rule.

Proposition 3. *Let (N,G) be a network such that i is linked to $i+1$ for all i. Then, for generic $(b_1,\ldots,b_n)$, there is a function $m(\gamma)$ with $m(\gamma) \underset{\gamma\to 0}{\longrightarrow}$ such that, for γ_0 low enough and δ close enough to 1, the game has an equilibrium such that:*

(i) Restricted to all θ except for a set of measure up to $m(\gamma_0)$, the equilibrium is segmented and natural. For such θ, if $\theta<0$, i becomes informed iff she is in $\min(S)$'s segment or higher. If $\theta>0$, i becomes informed iff she is in $\max(S)$'s segment or lower.

(ii) There is $\epsilon>0$ such that, for $\theta\in(-\epsilon,\epsilon)$, everyone becomes informed.

The intuition behind this equilibrium structure is as follows. First, as in Proposition 1, the players' posteriors under ignorance reinforce their biases, so that $\overline{\theta}(1)<\ldots<\overline{\theta}(n)$. When $\theta>\overline{\theta}(n)$, players always want to inform lower-bias neighbors, and only higher-bias neighbors who are close enough; this leads to the same general results as in Proposition 2 when $\theta=1$. The outcome when $\theta<\overline{\theta}(1)$ is analogous. However, values of θ lying in $[\overline{\theta}(1),\overline{\theta}(n)]$ create a different set of incentives. Suppose $\overline{\theta}(i)<\theta<\overline{\theta}(i+1)$ for some i. Then both i and $i+1$ want to inform everyone, as the actions of higher-bias players will be lowered after learning θ, while the actions of lower-bias players will be increased. Consider now $i-1$, who wants to inform everyone except for i; and $i+2$, who wants to inform everyone except for $i+1$. Since i and $i+1$ talk to each other, both $i-1$ and $i+2$ have effectively only two options: inform both i and $i+1$ or neither. However, depending on whether sharing θ increases or decreases the average of i and $i+1$'s actions, it will always be optimal for at least one of $\{i-1, i+2\}$ to inform them. Hence, if any player in $\{i-1, i, i+1, i+2\}$ is informed, everyone becomes informed.

This argument can be iterated: given an interval of players who inform everyone, at least one of the adjacent players on either end will be willing to inform that set, and thus everyone. The argument ends when we run out of players on one side, pinning down $\hat{b}_1(\theta)$ or $\hat{b}_{k_\theta}(\theta)$. After that, the remaining players have monotonic incentives, so their behavior is as in Proposition 2. However, for θ close to zero, the set of players who inform everyone turns out to be the entire population. That is, just as in the myopic setting, moderate states of the world are communicated to everyone.

Well-Connectedness in Random Graphs

The condition of well-connectedness is satisfied only by networks with a relatively high number of links. However, it is not too restrictive, and is plausibly satisfied by real social networks. To illustrate, take an Erdõs-Rényi random graph $G(n, p)$,[11] with n equal to 320 million. Assume γ_0 is small, biases are uniformly distributed and $b_n - b_1 = 4$ (so that, per part (ii) of Proposition 2, there are 4 segments in equilibrium). Then, if we assume an average degree of $np = 404$, at least 99% of the resulting networks are well-connected (see Proposition 6 in Appendix B). This bound is not tight, so in fact even a somewhat lower average degree would be sufficient.

That aside, an average degree of 404 is approximately in line with online social networks: for example, the average Facebook user in the U.S. has 338 friends (Smith, 2014). In addition, well-connectedness can be satisfied for much lower average degrees if the network displays homophily. Indeed, a high global average degree is not required for well-connectedness—rather, it suffices for the average degree, restricted to ideologically similar players, to be relatively high. To illustrate, in our example, a person whose bias is in the 50th percentile must have an expected 101 friends with

[11]That is, a graph with n nodes in which any two nodes are linked with probability p.

biases between the 38th and 62nd percentiles.

Equilibrium Analysis for Trees

All of our previous results rely on the network having a relatively high number of links. The other extreme—a network that is sparse enough—is also tractable but yields very different results. Namely, in the case of a tree with a single source of information, the problem of learning cascades can be solved recursively, starting with the peripheral players who are upstream of only one other player. We will not assume mutual observability or an activity rule, as they make no difference here.

Proposition 4. *Assume that the state of the world is binary (*$P(\theta = -1) = P(\theta = 1) = 0.5$*),* (N, G) *is a tree, and* $\gamma(S) > 0$ *only for singleton* S*. Then, for* δ *is close enough to* 1 *and* γ_0 *low enough, there is generically a unique sequential equilibrium, characterized as follows.*

If $S = \{i_0\}$*, we say* j *is* downstream *of* i *if the shortest path connecting* i_0 *to* j *passes through* i*. Let* $l(i)$ *be* i*'s path distance from* i_0*. Let* $\bar{l} = \max_i l(i)$*.*

(i) Players at distance $\bar{l}$ *only have one neighbor. If informed, they cannot inform anyone else.*

(ii) For $k = 1, \ldots, \bar{l}$*, In the* k*th round, consider all players at distance* $\bar{l} - k$*. Each such player* i *can only be informed by her unique neighbor at distance* $\bar{l} - k - 1$*. For each of her other neighbors* j*, if* $\theta = 1$*,* i *informs* j *at the first opportunity if*

$$b_A < b_i + \frac{1}{2},$$

where A *is the set of players downstream of* j*, including* j*, who are informed in equilibrium if* j *is informed, and* b_A *is the average bias of players in this set. Else, if* $b_A > b_i + \frac{1}{2}$*,* i *never informs* j*.*

Analogously, if $\theta = -1$*,* i *informs* j *iff* $b_{A'} > b_i - \frac{1}{2}$*, where* A' *is the set of downstream players informed in equilibrium.*

An example is given in Figure 8. Assume that $(b_1, \ldots, b_6) = (0.95, 1, 1.4, 1.6, 1.8, 2.2)$, $\theta = 1$, and 2 is informed by Nature. She chooses to inform 4 but not 3. These decisions are influenced by their downstream consquences: 4 will inform 1 but not 6,

Figure 8: Communication on a tree

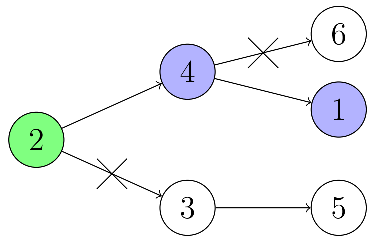

while 3 would inform 5. If, for instance, b_6 was lowered to 2.05, 4 would now inform 6, and 2 would no longer inform 4.

We finish with two observations. First, the assumption of a single source is crucial: if two players are initially informed, multiple equilibria can arise (see Appendix A). Second, even though well-connected networks and trees both admit tractable equilibria, they result in different patterns of information transmission. Whereas on a well-connected network the players break up into segments based on their biases, on a tree, communication is influenced by the network structure as much as by the distribution of biases. This can lead to either more or less communication, and to outcomes that are more sensitive to the parameters.

We illustrate this in two ways. First, consider a tree of the following form: i_0 is linked to n; i is linked to $i+1$ for all $i \in \{1, \ldots, i_0 - 2\} \cup \{i_0 + 1, \ldots, n-1\}$; and $i_0 - 1$ is linked to $i_0 + 1$. Suppose i_0 is the only player informed by Nature. How far does information spread when $\theta = 1$? Clearly, if n is informed, everyone is; else no one (besides i_0) is informed. Moreover, there is a threshold $\bar{b}$ such that, holding everything but b_{i_0} constant, if $b_{i_0} > \bar{b}$, then everyone is informed, while if $b_{i_0} < \bar{b}$, then no one is. Thus, the outcome can depend dramatically on fine details about the agents' preferences.

Second, imagine that a planner could design the network, with the goal of having a certain set of players become informed when $\theta = 1$. If i_0 is the only player informed by Nature, then the planner can guarantee that any set $A \subseteq N - \{i_0\}$ of players becomes informed, subject to the condition that i_0 prefers all of A to be informed rather than no one, i.e., $b_A < b_i + \frac{1}{2}$. Moreover, the planner can always achieve this with a forest: she can connect the players in A in a descending line, with i_0 connected to the highest-bias member of A, and leave all other players isolated. On the other hand, any set A that violates the condition could not be the set of informed agents resulting from *any* network, as i_0 would rather message no one than allow this outcome on the equilibrium path.

5 Discussion

Welfare Implications

Ex post (conditional on θ), agents often want others to remain uninformed. For example, when θ is high, agent 1 may want agent n to remain uninformed, so that n chooses a lower action. However, this does not imply that low levels of information transmission are optimal. Indeed, ignorance can only lead n to choose lower actions if n is also sometimes uninformed when θ is low, and in those cases n's action under ignorance would be too high. Thus, ex ante, n being less informed cannot push her average action closer to 1's preferences, and only results in a worse fit between her actions and the state. This logic implies that equilibria leading to more information transmission are preferred ex ante by every agent.[12] Formally:

Remark 4. Consider two message strategy profiles m, $\tilde{m}$. Let $p(\theta, j, t)$, $\tilde{p}(\theta, j, t)$ denote j's probability of being informed at time t when the state is θ, under each respective message strategy. Suppose that $p(\theta, j, t) \geq \tilde{p}(\theta, j, t)$ for all θ, j, t. Denote i's ex ante payoffs generated in each case by U_i, $\tilde{U}_i$, if actions are optimal given messages. Then $U_i \geq \tilde{U}_i$ for all i.

The model thus matches the conventional wisdom that assigns a negative connotation to breakdowns in communication. An immediate implication of Remark 4 is that an equilibrium with full diffusion, if it exists, is ex ante optimal for all agents (if they are patient). Therefore, all agents would be better off if they could collectively commit to playing the game myopically (compare Proposition 2 to Corollary 1).

Ex post, equilibria with partial diffusion lead to distributions of actions that may be more or less dispersed than those obtained under full information. For instance, under the conditions of Proposition 2, it is straightforward to show that an agent's long-run mean posterior under ignorance increases in her bias: $-1 < \overline{\theta}(1) \leq \ldots \leq \overline{\theta}(n) < 1$. This means that, when Nature informs no one ($S = \emptyset$), actions in the long run are more dispersed than under full information. When a strict subset $Y \subset N$ of players eventually becomes informed, the actions of players in $N - Y$ are still more dispersed with respect to one another, but may be closer to the actions of players in

[12] A similar result is shown in Corollary 2 of Galeotti et al. (2013).

Y: for instance, if $\theta = 1$ and $Y = \{1, \ldots, i\}$, the actions of agents above i will be lower than they would be under full information, while the actions of agents below i will fully take into account that $\theta = 1$.

Alternative Models

Several possible changes to the model deserve discussion. One alternative concerns the communication protocol: in many settings players cannot personalize their messages (e.g., on Twitter). Embedding this assumption in the model may lead to more communication if the network displays homophily and agents are myopic, as they would then be willing to share information with their mostly-similar neighbors, without being able to hide it from the more biased ones. But it may also lead to less communication if agents are forward-looking, as they would be fearful of information spilling into the wrong hands, and would potentially have to censor all their communication to prevent this.

Another alternative involves changing the payoff structure: in some settings, each player i may care about the average of the others' actions, rather than everyone's individual actions. (For example, i may care about the vote share in an election rather than individual votes.) The distinction is subtle but important: a focus on average actions tends to lead to coarser segments. For instance, if $\theta = 1$, players in the bottom half of the bias distribution (who perceive the average action as too high) may hide the state not just from players above them, but from *all* players. Conversely, players in the top half (who perceive the average action as too low) may want to inform everyone. Note, however, that an objective function like the one in the paper is recovered if players care about average actions, but also penalize variance in the distribution of actions. (For instance, more dispersed votes may lead to gridlock or uncertain political outcomes.)

Another potential extension concerns endogenous network formation and information acquisition. In the model, biased players are more likely to see information contradicting, rather than reinforcing, their biases; beliefs reinforce biases in the *absence* of information, a form of "negative echo chambers". While puzzling, this result is a natural consequence of our assumption that the players are active as senders but passive as receivers (i.e., they cannot choose which of their neighbors to listen to, nor change their probability of being informed by Nature). Giving players more control over information sources may lead to "positive echo chambers", but not necessarily more learning; instead, if players care about influencing others, they may choose to spend most of their time listening to like-minded agents, who are more likely to share useful ammunition for arguing with other groups.

Finally, this paper focuses on the case where information is all-or-nothing, that is, an informative message fully reveals the state. In reality, multiple partially informative signals may be available. If anything, a model with this feature would tend to generate less communication: if players are exposed to many signals, so that each one has a small impact on beliefs, there would never be a motive to share a high signal with a high-bias player in order to correct a large mistake in her beliefs.

References

Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar, "Bayesian Learning in Social Networks," *The Review of Economic Studies*, 2011, *78* (4), 1201–1236.

_ , Victor Chernozhukov, and Muhamet Yildiz, "Fragility of Asymptotic Agreement under Bayesian Learning," *Theoretical Economics*, 2016, *11* (1), 187–225.

Banerjee, Abhijit V, "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics*, 1992, *107* (3), 797–817.

Bloch, Francis, Gabrielle Demange, and Rachel Kranton, "Rumors and Social Networks," *International Economic Review*, 2018, *59* (2), 421–448.

Boguná, Marián, Romualdo Pastor-Satorras, Albert Díaz-Guilera, and Alex Arenas, "Models of Social Networks Based on Social Distance Attachment," *Physical Review E*, 2004, *70* (5), 056122.

Calvó-Armengol, Antoni, Joan De Martí, and Andrea Prat, "Communication and Influence," *Theoretical Economics*, 2015, *10* (2), 649–690.

Crawford, Vincent P and Joel Sobel, "Strategic Information Transmission," *Econometrica*, 1982, pp. 1431–1451.

DeMarzo, Peter M, Dimitri Vayanos, and Jeffrey Zwiebel, "Persuasion Bias, Social Influence, and Unidimensional Opinions," *The Quarterly Journal of Economics*, 2003, *118* (3), 909–968.

Dewan, Torun and Francesco Squintani, "Leadership with Trustworthy Associates," *American Political Science Review*, 2018, *112* (4), 844–859.

_ , Andrea Galeotti, Christian Ghiglino, and Francesco Squintani, "Information Aggregation and Optimal Structure of the Executive," *American Journal of Political Science*, 2015, *59* (2), 475–494.

Dimock, Michael, Jocelyn Kiley, Scott Keeter, Carroll Doherty, and Alec Tyson, "Beyond Red vs. Blue: The Political Typology," *Pew Research Center*, 2015. http://www.people-press.org/files/2014/06/6-26-14-Political-Typology-release1.pdf.

Erdos, Paul and Alfréd Rényi, "On the Evolution of Random Graphs," *Publ. Math. Inst. Hung. Acad. Sci*, 1960, *5* (1), 17–60.

Frankovic, Kathy, "Why Won't Americans Get Vaccinated?," *Retrieved on 07/23/2021*, 2021. https://today.yougov.com/topics/politics/articles-reports/2021/07/15/why-wont-americans-get-vaccinated-poll-data.

Frick, Mira, Ryota Iijima, and Yuhta Ishii, "Misinterpreting Others and the Fragility of Social Learning," *Econometrica*, 2020, *88* (6), 2281–2328.

Funk, C, L Rainie, and D Page, "Public and Scientists' Views on Science and Society," *Pew Research Center*, 2015. http://www.pewinternet.org/files/2015/01/PI_ScienceandSociety_Report_012915.pdf.

Galeotti, Andrea, Christian Ghiglino, and Francesco Squintani, "Strategic Information Transmission Networks," *Journal of Economic Theory*, 2013, *148* (5), 1751–1769.

Gentzkow, Matthew and Emir Kamenica, "Bayesian Persuasion," *American Economic Review*, 2011, *101* (6), 2590–2615.

Glazer, Jacob and Ariel Rubinstein, "A Study in the Pragmatics of Persuasion: a Game Theoretical Approach," *Theoretical Economics*, 2006, *1* (4), 395–410.

Golub, Benjamin and Matthew O. Jackson, "Naive Learning in Social Networks and the Wisdom of Crowds," *American Economic Journal: Microeconomics*, 2010, *2* (1), 112–49.

Hagenbach, Jeanne and Frédéric Koessler, "Strategic Communication Networks," *The Review of Economic Studies*, 2010, *77* (3), 1072–1099.

_ , _ , and Eduardo Perez-Richet, "Certifiable Pre-Play Communication: Full Disclosure," *Econometrica*, 2014, *82* (3), 1093–1131.

Howe, Peter D, Matto Mildenberger, Jennifer R Marlon, and Anthony Leiserowitz, "Geographic Variation in Opinions on Climate Change at State and Local Scales in the USA," *Nature Climate Change*, 2015, *5* (6), 596–603.

Jardina, Ashley and Michael Traugott, "The Genesis of the Birther Rumor: Partisanship, Racial Attitudes, and Political Knowledge," *Journal of Race, Ethnicity and Politics*, 2019, *4* (1), 60–80.

Leiserowitz, A., E. Maibach, S. Rosenthal, J. Kotcher, P. Bergquist, M. Ballew, M. Goldberg, and A. Gustafson, "Climate change in the American mind: April 2019," *Yale Project on Climate Change Communication (Yale University and George Mason University, New Haven, CT)*, 2019.

McPherson, Miller, Lynn Smith-Lovin, and James M Cook, "Birds of a Feather: Homophily in Social Networks," *Annual Review of Sociology*, 2001, pp. 415–444.

Milgrom, Paul and John Roberts, "Relying on the Information of Interested Parties," *The RAND Journal of Economics*, 1986, *17* (1), 18–32.

Moore, Peter, "Young Americans most worried about vaccines," *Retrieved on 9/6/2015*, 2015. `https://today.yougov.com/news/2015/01/30/young-americans-worried-vaccines/`.

Morales, Lymari, "Obama's Birth Certificate Convinces Some, but Not All, Skeptics," *Retrieved on 9/6/2015*, 2011. `http://www.gallup.com/poll/147530/obama-birth-certificate-convinces-not-skeptics.aspx`.

Omer, Saad B, Jennifer L Richards, Michelle Ward, and Robert A Bednarczyk, "Vaccination Policies and Rates of Exemption from Immunization, 2005-2011," *New England Journal of Medicine*, 2012, *367* (12), 1170–1171. PMID: 22992099.

Perego, Jacopo and Sevgi Yuksel, "Searching for Information and the Diffusion of Knowledge," *Working paper*, 2016.

Sethi, Rajiv and Muhamet Yildiz, "Public Disagreement," *American Economic Journal: Microeconomics*, 2012, *4* (3), 57–95.

Smith, Aaron, "6 New Facts about Facebook," *Pew Research Center. Retrieved on 11/6/2015*, 2014. http://www.pewresearch.org/fact-tank/2014/02/03/6-new-facts-about-facebook/.

Smith, Lones and Peter Sørensen, "Pathological Outcomes of Observational Learning," *Econometrica*, 2000, *68* (2), 371–398.

Squintani, Francesco, "Information Transmission in Political Networks," *Working paper*, 2019.

Wolitzky, Alexander, "Learning from Others' Outcomes," *American Economic Review*, 2018, *108* (10), 2763–2801.

A Other Equilibria with Forward-Looking Players

In networks that are not well-connected, and without the assumptions of mutual observability and an activity rule, concerns about learning cascades can support a variety of complex or pathological equilibria. We illustrate the possibilities in several examples. In all of these, we assume that $\gamma_0 > 0$ is small and that $\theta = 1$ is realized.

Figure 9: Multiple Pareto-ordered equilibria

(a) Messages are substitutes

(b) Messages are complements

First, Figures 9a and 9b show games with multiple equilibria that are Pareto-ordered, in the sense that the informed agents agree on the best equilibrium but cannot necessarily coordinate on it. In Figure 9a, take $(b_1, b_2, b_3, b_4, b_5) = (0, 1, 1, \bar{b}, \bar{b} + \epsilon)$, with $\frac{3}{2} < \bar{b} < 2$; $\gamma(\{2, 3\})$ small and $\gamma(S) = 0$ otherwise; and $\epsilon > 0$ small. Then 4 and 5 talk to 1 if they are informed, while 1 talks to no one. Both 2 and 3 want to inform 1, but not 4 or 5. However, it is worth using one and only one of them as an intermediary. Hence, there is an equilibrium in which 2 talks to 4 while 3 does not talk to 5, and another in which the opposite happens. Although both 2 and 3 would rather inform 4, if the "wrong" equilibrium is being played, there is no way for 2 to inform 4 and tell 3 that informing 5 is no longer necessary.

In Figure 9b, $(b_1, b_2, b_3, b_4, b_5) = (1, 1, \bar{b}, \bar{b}, \bar{b} + 0.4)$, where $1 < \bar{b} < \frac{3}{2}$; $\gamma(\{2, 3\})$ is small, and $\gamma(S) = 0$ otherwise. (This example requires a directed network, where 5 can receive but not share information, but more complex examples can deliver similar results in an undirected network.) Here 3 and 4 will talk to 5 if informed. 1 and 2 are willing to inform 3 and 4, but this comes at the cost of informing 5. It is then optimal to either inform no one or everyone; and one of these options (depending on the value of $\bar{b}$) is preferred by both 1 and 2. However, for a range of values of $\bar{b}$, both are equilibrium outcomes. Indeed, if no one is informed in equilibrium, then a deviation by 1 can inform $\{3, 5\}$, but not 4: a one-player deviation does not capture the full value of switching to the informative equilibrium. Conversely, if everyone is informed in equilibrium, then a deviation by 1 leaves 3 uninformed while having no effect on 4 or 5, which is never profitable.

Figure 10: Multiple unordered equilibria

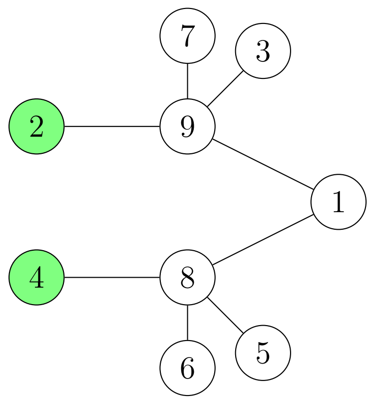

Figure 10 shows that there may be multiple equilibria that are not Pareto-ordered for the informed players. Suppose that $\gamma(\{2,4\})$ is small, and $\gamma(S) = 0$ otherwise. As in Figure 9a, 2 and 4 want to use either 8 or 9 as an intermediary to share information with 1. Hence, for appropriate bias vectors b, there are two pure-strategy equilibria. However, the choice of intermediary now results in other players learning, such as $\{3,7\}$ or $\{5,6\}$; since 2 and 4 in general differ in their preferences over which group to inform, examples can be constructed where 2 would rather use 9 as an intermediary, and 4 would rather use 8, or vice versa. Note that, if 2 and 4 could observe each other's actions, the first case would become a game of preemption (each player races to be the first to communicate) while the second would become a war of attrition (each player waits for the other to use her intermediary).

Figure 11: History-dependent strategies

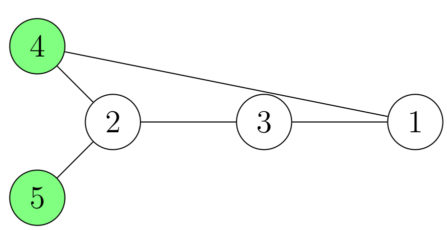

Figure 11 presents an example in which the equilibrium must condition on the players' histories in an important way. Assume that $(b_1, b_2, b_3, b_4, b_5) = (0, 1, 1.6, 3, 3)$, and $\gamma(\{4\}) = \gamma(\{5\}) = 0.1$ and $\gamma(S) = 0$ otherwise. Clearly 1 never talks, and 2 never talks to 4 or 5. 3 and 4 always inform 1 if they are informed, and 4 and 5 always inform 2. 3 would talk to 2 but never can (since 2 is 3's only potential source) and 3 talks to 1. The question then is: should 2 inform 3? The answer depends on how 2 became informed. If 2 was informed by 5, he should inform 3, but not if he was informed by 4.

The reason is that 2 wants to inform 1, but not 3. If 5 is the informed player, informing 3 amounts to also informing 1: 3 is used as an intermediary. However, if 4 is informed, then 1 would become informed no matter what 2 does.

This example can be altered to produce another in which players want to manipulate each other. Suppose that $\gamma(\{4\}) = 0.001$, $\gamma(\{5\}) = 0.1$, $\gamma(\{4, 5\}) = 0.009$ and $\gamma(S) = 0$ otherwise. Then the analysis proceeds as before, with one caveat: now 4 should *not* talk to 2. Indeed, 4 would rather risk a 10% chance that 2 will go uninformed (if 5 is uninformed) to ensure that, with a 90% chance, 2 will be informed by 5 and hence, thinking that 4 is likely uninformed, she will inform 3.

Figure 12: Non-existence of pure-strategy equilibrium

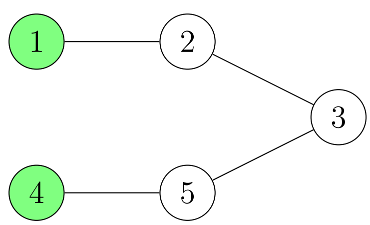

Finally, 12 shows that a pure-strategy equilibrium may not exist. Suppose that $(b_1, b_2, b_3, b_4, b_5) = (0, 0.45, 0.9, 2, 2.6)$; $\gamma(\{1, 4\})$ is small and $\gamma(S) = 0$ otherwise. (This example also assumes a directed network in which 3 is only a receiver.) Then 2 and 5 always inform 3. 1 wants to inform 2 but not 3, and would rather tell neither than both; 4 wants to inform 3 but not 5, and would rather tell both than neither. Then, if neither player is talking, 4 would deviate and inform 5. If 4 is talking and 1 is not, 1 would deviate and inform 2 (since 3 is already informed). If both 1 and 4 are talking, 4 would deviate and stop talking (since 3 is already informed). And if 1 is talking but not 4, 1 would deviate and stop talking.

Robustness of Natural Equilibrium

We finish with a survey of alternative sets of conditions which guarantee the existence of a natural equilibrium, without requiring the addition of mutual observability and an activity rule to the game. However, these conditions do not guarantee uniqueness, so they do not rule out the existence of other equilibria like the ones we have just described.

Proposition 5. *If* (N, G) *is well-connected, then generically, for* γ_0 *low enough and* $\delta = 1$,[13] *there is a natural equilibrium of the game without mutual observability or an activity rule, with the same segments* $\underline{M}_1, \ldots, \underline{M}_{\underline{k}+1}$ *and* $\overline{M}_1, \ldots, \overline{M}_{\overline{k}+1}$ *found in Proposition 2.*

In other words, a natural equilibrium always exists when $\delta = 1$ (but so do many other equilibria, as with $\delta = 1$ players are indifferent over any deviations that do not change the information distribution in the long run). Moreover, it has a simple structure: when $\theta = 1$, a player in $\overline{M}_l$ *always* informs players in segments $\overline{M}_l$ and higher, and *never* informs players in lower segments. (The case $\theta = -1$ is analogous.)

[13] This means that the players' ex ante payoffs are $U_i = \lim_{t\to\infty} E_0(u_{it})$, where u_{it} is i's flow payoff at time t. This is well-defined since $p(\theta, i, t)$ converges for all θ, i, so u_{it} converges.

B Proofs

Proof of Remark 1. Suppose (N, G) is well-connected but an interval (a, b) with $b - a > \frac{1-\gamma_0}{2-\gamma_0}$ is not connected. Let i_0 be the lowest member of $N_{(a,b)}$. Then, for some $i \in N_{(a,b)}$, there is no path between i_0 and i. Let i_1 be the lowest member of $N_{(a,b)}$ disconnected from i_0. Then there is a path between i_0 and $i_1 - 1$. Consider the interval N_I for $I = (b_{i_1-1} - \epsilon, b_{i_1-1} - \epsilon + \frac{1-\gamma_0}{2-\gamma_0})$ for small $\epsilon > 0$. If i_1 is in this interval for any $\epsilon \in (0, \frac{1-\gamma_0}{2-\gamma_0})$, then, by well-connectedness, there is a path from $i_1 - 1$ to i_1, a contradiction. If not, then the interval $\left(b_{i_1-1}, b_{i_1-1} + \frac{1-\gamma_0}{2-\gamma_0}\right)$ is empty, a contradiction. Finally, taking (a, b) large enough so that $N_{(a,b)} = N$, we conclude that N is connected. □

Proof of Remark 2. Each player i chooses a_{it} to maximize $E_{it}(-(a_{it} - \theta - b_i)^2)$. Since this is a concave function of a_{it}, the unique optimum is given by the FOC: $a_{it} = E_{it}(\theta) + b_i$, where $E_{it}(\theta) = \theta$ if i is informed and $E_{it}(\theta) = \overline{\theta}(i, t)$ if not, by definition. □

Proof of Proposition 1. The first part follows from the fact that i's marginal payoff from informing j at time t is

$$P\left[-(b_j - b_i)^2 + (b_j + \overline{\theta}(j,t) - b_i - \theta)^2\right] = P(\theta - \overline{\theta}(j,t))\left[(\theta - \overline{\theta}(j,t)) - 2(b_j - b_i)\right],$$

where $P > 0$ is the probability that j will be uninformed by the end of time t if i does not inform j.

The existence and uniqueness of equilibrium can be proven by solving the game forward. Indeed, suppose the result is true up to time t. Let $x(i)$ be a candidate value for agent i's mean posterior under ignorance at the end of period $t + 1$. The communication strategies at time $t+1$ are then uniquely determined by the first part of this Proposition. For $x \in [-1, 1]$, define $T(x)$ as

$$T(x) = \frac{\int_{-1}^{1} \theta(1 - \tilde{p}(\theta, i, t+1, x))f(\theta)d\theta}{\int_{-1}^{1}(1 - \tilde{p}(\theta, i, t+1))f(\theta)d\theta}, \text{ where} \tag{1}$$

$$1 - \tilde{p}(\theta, i, t+1, x) = (1 - p(\theta, i, t))\left(1 - \sum_{j \in N_i} \frac{\alpha p(\theta, j, t)}{n|N_j|}\mathbb{1}_{(\theta - x)(\theta - x - 2(b_i - b_j)) > 0}\right). \tag{2}$$

Here $\tilde{p}(\theta, i, t+1, x)$ is i's probability of being informed at the end of $t+1$, as a result of message strategies that depend on x. $\frac{\alpha p(\theta,j,t)}{n|N_j|}$ is the probability that j is chosen as the sender, i is chosen as the receiver, and j is informed. $\mathbb{1}_{(\theta-x)(\theta-x-2(b_i-b_j))>0}$ equals 1 iff j wants to inform i. Intuitively, $T(x)$ is i's equilibrium mean posterior if i's neighbors communicate *as if* her mean posterior is x. $x(i)$ is compatible with equilibrium iff $T(x(i)) = x(i)$. To finish the proof, we will show that, for $\alpha > 0$ small enough, T is a contraction. This will imply that the mean posteriors $\overline{\theta}(i, t+1)$ are uniquely determined, and hence that the equilibrium is uniquely determined up to time $t+1$.

Take two values $x > x'$. Equation 2 implies that, for any y,

$$1 - p(\theta, i, t) \geq 1 - \tilde{p}(\theta, i, t+1, y) \geq \left(1 - \frac{n-1}{n}\alpha\right)(1 - p(\theta, i, t)).$$

The upper bound obtains if no neighbors of i ever inform her; the lower bound obtains if every $j \neq i$ is connected to i and only to i, always informed, and wants to inform i. In particular, then, $|1 - \tilde{p}(\theta, i, t+1, x) - 1 + \tilde{p}(\theta, i, t+1, x')| < \alpha$. In addition, clearly $1 \geq 1 - \tilde{p}(\theta, i, t+1, y) \geq 1 - \gamma_0$ for all y.

In fact, $\tilde{p}(\theta, i, t+1, x)$ and $\tilde{p}(\theta, i, t+1, x')$ can only differ if, depending on whether $x(i) = x$ or x', at least one neighbor of i would change her strategy towards i in this period. In other words, for some $j \in N_i$, the expressions $(\theta - x)(\theta - x - 2(b_i - b_j))$ and $(\theta - x')(\theta - x' - 2(b_i - b_j))$ must have opposite signs. In particular, either $\theta - x$ and $\theta - x'$, or $\theta - x - 2(b_i - b_j)$ and $\theta - x' - 2(b_i - b_j)$, must have opposite signs (but not both). Equivalently $\theta \in [x', x] \triangle [x' + 2(b_i - b_j), x + 2(b_i - b_j)]$, a set of measure at most $2(x - x')$. Thus $\tilde{p}(\theta, i, t+1, x)$ and $\tilde{p}(\theta, i, t+1, x')$ only differ on a set of values of θ of measure at most $2(n-1)(x - x')$.

Let $\int_{-1}^{1} \theta(1 - \tilde{p}(\theta, i, t+1, x)) f(\theta) d\theta = A$, $\int_{-1}^{1} \theta(1 - \tilde{p}(\theta, i, t+1, x')) f(\theta) d\theta = A'$, $\int_{-1}^{1}(1 - \tilde{p}(\theta, i, t+1, x)) f(\theta) d\theta = B$, $\int_{-1}^{1}(1 - \tilde{p}(\theta, i, t+1, x')) f(\theta) d\theta = B'$. Our previous arguments imply that $|B|, |B'| \geq 1 - \gamma_0$ and $|A|, |B| \leq 1$. Let $\overline{f} = \max_\theta f(\theta)$. Then

$$\begin{aligned} |T(x) - T(x')| = \left|\frac{A}{B} - \frac{A'}{B'}\right| = \frac{|AB' - A'B|}{|BB'|} = \frac{|(A - A')B - A(B - B')|}{|BB'|} \leq \\ \leq \frac{|A - A'||B| + |A||B - B'|}{|B||B'|} \leq \\ \leq \frac{|A - A'| + |B - B'|}{(1 - \gamma_0)^2} \leq \frac{2 \times \alpha \times 2(n-1)(x - x')\overline{f}}{(1 - \gamma_0)^2}. \end{aligned}$$

For any α small enough that $\frac{4\overline{f}\alpha(n-1)}{(1-\gamma_0)^2} < 1$, T is a contraction, which completes the proof.[14] □

Proof of Corollary 1. (a) Consider the subgraph (N, G') given by deleting any links in (N, G) between players whose biases differ by $\frac{1-\gamma_0}{2-\gamma_0}$ or more. Clearly (N, G') is still well-connected. By Remark 1, (N, G') is connected. Note that $\overline{\theta}(i,t) \in [-\frac{\gamma_0}{2-\gamma_0}, \frac{\gamma_0}{2-\gamma_0}]$ for all i, t. (The upper bound is attained when $p(1,i,t)=0$, $p(-1,i,t)=\gamma_0$, and the lower bound when $p(1,i,t)=\gamma_0$, $p(-1,i,t)=0$.) Then, by Proposition 1, if $\theta = 1$ and i is informed, i will want to inform j whenever $b_j < b_i + \frac{1-\gamma_0}{2-\gamma_0}$. In particular, any two players at distance less than $\frac{1-\gamma_0}{2-\gamma_0}$ will inform each other, so information flows freely on all links in G'. Then, if any player is informed, everyone is informed eventually. Hence $p(\theta,i)=\gamma_0$ for all θ, i.
(b) We begin by noting that, if i has a meeting with j and $i > j$, i informs j iff $\theta > \overline{\theta}(j,t)$ (Proposition 1). Next, we argue that $\overline{\theta}(n,t)$ is increasing in t. Suppose for the sake of contradiction that $\overline{\theta}(n,t+1) \leq \overline{\theta}(n,t)$. Then, in equilibrium, informed players $i < n$ who meet n at time $t+1$ will inform n iff $\theta \leq \overline{\theta}(n,t+1) \leq \overline{\theta}(n,t)$. Hence $1-p(\theta,n,t+1)$ is strictly smaller[15] than $1-p(\theta,n,t)$ for $\theta \leq \overline{\theta}(n,t+1) \leq \overline{\theta}(n,t)$ and equal for $\theta > \overline{\theta}(n,t+1)$, whence $\overline{\theta}(n,t+1) > \overline{\theta}(n,t)$, a contradiction. By the same argument, $\overline{\theta}(1,t)$ is decreasing in t. In particular, this implies that $\overline{\theta}(1,t) < \overline{\theta}(n,t)$ for all t and hence $\overline{\theta}(1) < \overline{\theta}(n)$.

Next, we prove that there is full diffusion for states $\theta \in \left(\overline{\theta}(1), \overline{\theta}(n)\right)$. Take such a state, and let t_0 be such that $\theta \in \left(\overline{\theta}(1,t), \overline{\theta}(n,t)\right)$ for all $t \geq t_0$. Then, at any time $t \geq t_0$, all players other than 1 want to share θ with 1 (to increase her action), and all players other than n want to share θ with n (to decrease hers). Hence $p(\theta,1) = p(\theta,n) = \gamma_0$. But then, for any player $1 < i < n$, either 1 or n will share θ with i, depending on whether $\theta > \overline{\theta}(i,t)$ or vice versa. Hence $p(\theta,t) = \gamma_0$ for all t. On the other hand, if $\theta < -\frac{\gamma_0}{2-\gamma_0}$, then θ is only ever shared from lower types to higher types (as it is lower than anyone's posterior under ignorance), yielding the result. The case $\theta > \frac{\gamma_0}{2-\gamma_0}$ is analogous. The assumption that i is linked to $i+1$ for all i guarantees that if i is informed, every $k > i$ becomes informed as well. □

Proof of Corollary 2. As $\gamma_0 \to 0$, $p(\theta,i) \leq \gamma_0 \to 0$ for all θ and i. Then $\overline{\theta}(i,t) \to E(\theta)$, which equals 0 by assumption.

[14]When F does not admit a density, T may have discontinuities. But it can be shown that T is increasing, and moreover $T\left(-\frac{-\gamma_0}{2-\gamma_0}\right) > -\frac{-\gamma_0}{2-\gamma_0}$, $T\left(\frac{-\gamma_0}{2-\gamma_0}\right) < \frac{-\gamma_0}{2-\gamma_0}$, so existence is guaranteed.

[15]Here we use that n meets some informed player with positive probability, as n is connected to everyone.

Next consider the case $\gamma_0 \to 1$. Assume first that 1 and n are not always in the set of players informed by Nature (even conditional on these sets being nonempty). Note that, if $\overline{\theta}(1) \to -1$ and $\overline{\theta}(n) \to 1$ as $\gamma_0 \to 1$, we are done (by the proof of Corollary 1.(b)). For the sake of contradiction, then, suppose WLOG that $\overline{\theta}(1)$ does not converge to -1, so there is a sequence $\gamma_0^k \to 1$ such that $\overline{\theta}(1)^k \to \theta_0 > -1$. Then, as $k \to \infty$, the probability that 1 does *not* learn the state converges to 0 for $\theta > \theta_0$ but is bounded below by $1 - \sum_{1 \in S} \gamma(S) > 0$ for $\theta < -\frac{\gamma_0}{2-\gamma_0}$. But then $\overline{\theta}(1)^k$ must converge to some value in the interval $(-1, \theta_0)$, a contradiction.

If 1 and/or n are always informed by Nature conditional on S being nonempty, the result goes through for a different reason. If they are both always informed, one of them will always be willing to inform each other player i, depending on whether $\theta > \theta(i,t)$ or $\theta < \theta(i,t)$. If only one of them is always informed—say, 1—then our previous argument still implies $\overline{\theta}(n) \to 1$ as $\gamma_0 \to 1$. On the other hand, $\overline{\theta}(1,t) \equiv 0 = \overline{\theta}(i)$. For any player $1 < i < n$, $p(\theta, i) = \gamma_0$ for $\theta < \overline{\theta}(i)$ (as 1 would always inform i for these values of θ), while $p(\theta, i) \leq \gamma_0$ for $\theta \geq \overline{\theta}(i)$. This implies $\overline{\theta}(i) \geq 0$. Then 1 always informs i if $\theta < \overline{\theta}(i)$, while n informs i if $\theta \in (\overline{\theta}(i), \overline{\theta}(n)) \subseteq (\overline{\theta}(1), \overline{\theta}(n))$, where $\overline{\theta}(n) \to 1$. Hence $p(\theta, i) \to_{|| \, ||_1} 1$ as $\gamma_0 \to 1$ for all i. □

Proof of Remark 3. For $\gamma_0 = 0$, if $\theta = 1$, i's marginal payoff from a set of agents A being informed versus not is

$$\frac{1}{1-\delta} \sum_{j \in A} \left[-(b_j - b_i)^2 + (b_j - b_1 - 1)^2 \right] = \frac{1}{1-\delta} \sum_{j \in A} (2b_i - 2b_j + 1).$$

The result follows for $\gamma_0 = 0$. By a continuity argument, the outcome of the algorithm is the same for all $\gamma_0 > 0$ low enough, as long as none of the comparisons that must be checked yield equalities. □

Proof of Proposition 2. We will proceed as follows. First, we will prove that, for generic $(b_i)_i$, for all γ_0 low enough, the game has a unique equilibrium which is in pure strategies, and Markovian in a sense we will make precise. Second, we will show that, for generic $(b_i)_i$, for γ_0 low enough and δ close enough to 1, any pure strategy Markovian equilibrium must be natural. Finally we will prove part (ii).

The equilibrium is unique, pure and Markovian.

Preliminaries. Assume for now that $\gamma_0 = 0$. Denote the *state* of the game at a history h by $s(h) = (\theta, T, (k_{ij})_{i\in T, j\in N-T})$, where θ is the state of the world, T is the set of informed agents, and k_{ij} is the number of times i has declined to inform j since the last time anyone was informed (or K, whichever is lower), at the end of period t. Note that the game has a finite number of possible states.

Given a Markov strategy profile σ, for each state s, let $V_i(s, \sigma)$ be i's value function at time t if the state is s at the end of period t. Then, if i_0 is informed and is given an opportunity to send a message to an uninformed player j_0, her marginal payoff from doing so is

$$\Delta V_i(s, j, \sigma) = V_i(\theta, T \cup \{j\}, \mathbf{0}, \sigma) - V(\theta, T, (\tilde{k}_{ij}), \sigma),$$

where $\tilde{k}_{i_0 j_0} = \min(k_{i_0 j_0} + 1, K)$ and $\tilde{k}_{ij} = k_{ij}$ for all other ij. Say $s' \neq s$ *follows from* s if, the current state being s, the game could transition to s' in one period (i.e., s' has one more informed player and $(k_{ij}) = \mathbf{0}$, or s' has the same set of informed players and (k_{ij}) higher by 1 in one coordinate).

Equilibrium construction. We construct a candidate equilibrium by backward induction. Begin at terminal states (i.e., states with $k_{ij} = K$ for all i, j, or $T = N$). Call these 0-states. In this case, the game is over. Next, consider states s such that all states that follow from them are 0-states, and call these 1-states. At any such state s, it is optimal for i to inform j, given the chance, if $\Delta V_i(s, j, \sigma^*) \geq 0$, with indifference in case of equality. (Here σ^* represents the unique equilibrium continuation determined in the previous step.) Whenever $\Delta V_i(s, j, \sigma^*) \neq 0$, i's optimal action is unique, pure, and Markovian (i.e., the optimum is the same at all histories h that correspond to state s). In general, in the kth step, consider $(k-1)$-states, i.e., states s such that all states following from them are $(k-2)$-states or lower. Then, whenever $\Delta V_i(s, j, \sigma^*) \neq 0$, i's optimal action is unique, pure and Markovian.

Proceeding by induction on the set of states, if $\Delta V_i(s, j, \sigma^*) \neq 0$ for all i, j, s, then the equilibrium is unique, pure and Markovian. This algorithm terminates in at most $n^3(K+1)$ steps, as the state cannot advance more than $n^3(K+1)$ steps without reaching a terminal state. (T cannot grow more than $n-1$ times, and for each fixed T, each of the less than n^2 coordinates of (k_{ij}) cannot grow more than $K+1$ times.)

Genericity. Next, we will argue that the property $\Delta V_i(s, j, \sigma) \neq 0$ for all i, j, s, σ holds for a generic set of bias vectors $(b_1, \ldots, b_n)$. (Rather than check whether $\Delta V_i(s, j, \sigma^*) \neq 0$ for all i, j, s and the equilibrium continuation σ^*, we check the stronger condition that the optimal action be unique for *all* possible Markovian continuations.) The formula for $\Delta V_i(s, j, \sigma)$ depends not only on i, j and s, but also on the strategy profile σ followed in the continuation, whichever action i chooses. However, there is a finite number of pure Markovian strategy profiles σ. Indeed, such a profile is fully pinned down by specifying, for each state s (from a finite set), which informed players i would be willing to inform which uninformed players j, if given the opportunity. Hence, we need to check that a finite collection of expressions $(\Delta V_i(s,j,\sigma))_{i,j,s,\sigma}$ are all nonzero for generic $(b_i)_i$.

$V_i(s,\sigma)$ can be written as follows. For each j and each $t \in \mathbb{N}_0$, assuming WLOG that $\theta(s) = 1$, let $p(\theta, j, t, s, \sigma)$ be the probability that j will be informed by t periods from the present, if the current state is s, under the strategy profile σ. Then

$$\begin{aligned} V_i(s,\sigma) &= \sum_{j=1}^{n} \left[-P(\theta,j,s,\sigma)(b_j - b_i)^2 - \left(\frac{1}{1-\delta} - P(\theta,j,s,\sigma) \right)(b_j - b_i - 1)^2 \right] \\ &= \sum_{j=1}^{n} \left[-\frac{(b_j - b_i)^2}{1-\delta} + \left(\frac{1}{1-\delta} - P(\theta,j,s,\sigma) \right)(2(b_j - b_i) - 1) \right], \end{aligned}$$

where $P(\theta,j,s,\sigma) = \sum_t \delta^t p(\theta,j,t,s,\sigma)$. Since the $P(\theta,j,s,\sigma)$ are constants, ignoring the first term $-(b_j - b_i)^2$ for each j, $V_i(s,\sigma)$ is a linear function of the bias vector. In addition, if j is uninformed in state s, then $P(\theta,j,s,\sigma) \leq \frac{\delta}{1-\delta}$. On the other hand, if s' is the state following from s in which i has informed j, $P(\theta,j,s',\sigma) = 1$. Thus $\Delta V_i(s,j,\sigma)$ is a linear function of the bias vector, in which b_j has a nonzero coefficient. Thus, for each i,j,s,σ, $\Delta V_i(s,j,\sigma) = 0$ for an affine set of bias vectors of dimension $n-1$, hence measure zero.

The case $\gamma_0 > 0$. Next, we show that, for a generic set of values of the bias vector, the same equilibrium found above is the unique equilibrium for all $\gamma_0 > 0$ low enough. Given $(b_1, \ldots, b_n)$ for which $\Delta V_i(s,j,\sigma) \neq 0$ for all i,j,s,σ, let $M > 0$ be a uniform bound such that $|\Delta V_i(s,j,\sigma)| \geq M$ for all i,j,s,σ. Denote by $V_i(s,\sigma,\gamma_0,t)$ i's value function at time t when the state of the game is s and the Markovian profile σ is being played in the continuation, for any $\gamma_0 > 0$. (Calculating $V_i(s,\sigma,\gamma_0,t)$ explicitly involves calculating the mean posteriors under ignorance $\overline{\theta}(j,t)$, given σ and γ_0. V_i is generally a function of t in this case because $\overline{\theta}(j,t)$ is a function of t.) Then

$$\begin{aligned} V_i(s,\sigma,\gamma_0,t) &= \sum_{j=1}^{n} \sum_{t=0}^{\infty} \delta^t \left[-(b_j - b_i)^2 - (1 - p(\theta,j,t,s,\sigma))(b_j + \overline{\theta}(j,t) - b_i - 1)^2 \right] \\ \Longrightarrow & |V_i(s,\sigma,\gamma_0,t) - V_i(s,\sigma)| \leq \frac{1}{1-\delta}(2b_n - 2b_1 + 3)\gamma_0 \end{aligned}$$

for all t, owing to the fact that $|\overline{\theta}(j,t)| \leq \frac{\gamma_0}{2-\gamma_0} \leq \gamma_0 < 1$ and $|b_j - b_i| \leq b_n - b_1$.

Then, for $\gamma_0 \in \left(0, \frac{(1-\delta)M}{4b_n-4b_1+6}\right)$, $\Delta V_i(s,j,\sigma,\gamma_0,t)$ has the same sign as $\Delta V_i(s,j,\sigma)$ for all t. Then the same backward induction argument used before yields the result. (Rather than consider $(k-1)$-states in step k, consider all histories h corresponding to $(k-1)$-states in step k.)

The equilibrium is natural.

We will now show that, for generic $(b_i)_i$, for δ close enough to 1 and γ_0 low enough, any pure strategy Markovian equilibrium must be natural. Because the equilibrium is constant in a neighborhood of $\gamma_0 = 0$, it is enough to write the proof for $\gamma_0 = 0$. The proof relies on three key observations. First, if δ is close to 1, and the game reaches a terminal state quickly—so that every player that ever becomes informed does so quickly—then the players' payoffs depend mainly on the set of players who become informed in the long run, i.e., on the terminal state. Second, the activity rule guarantees that a terminal state will be reached quickly in expectation. Third, in any pure strategy Markovian profile, any terminal state reached with positive probability is reached with probability uniformly bounded away from zero. Formally, we have the following three lemmas.

Lemma 1. *Let p_t be the probability that an event A happens at time t, and $q_t = \sum_{t'=0}^{t} p_{t'}$ the probability that it happens by time t. Then, if the expected time $E(A)$ until A happens is finite, that is, $\sum_{t=0}^{\infty} tp_t = E < \infty$, then, for any $\delta \in (0,1)$, $\sum_{t=0}^{\infty} \delta^t(1-q_t) \leq E$.*

Proof. Note that, if $E < \infty$, $q_t \xrightarrow[t\to\infty]{} 1$. Then

$$\sum_{t=0}^{\infty} \delta^t(1-q_t) = \sum_{t=0}^{\infty} \delta^t \sum_{t'=t+1}^{\infty} p_{t'} = \sum_{t=0}^{\infty} p_t \sum_{t'=0}^{t-1} \delta^{t'} \leq \sum_{t=0}^{\infty} tp_t = E.$$

□

Lemma 2. *Let A be the event that the game reaches a terminal state. Then $E(A) \leq \frac{n^5(K+1)}{\alpha} < \infty$.*

Proof. In each period before a terminal state, Nature selects an informed sender and an uninformed receiver with probability at least $\frac{\alpha}{n(n-1)}$. (This lower bound would obtain if only one player j is uninformed, and only one other player i is linked to j.) Each time this happens, the state of the game s must advance to a state $s' \neq s$ that follows from s, that is, either T grows or T stays constant and (k_{ij}) grows. Thus the expected time until the state advances is at most $\frac{n(n-1)}{\alpha}$.

The state cannot advance more than $n^3(K+1)$ times without the game reaching a terminal state. Then $E(A) \leq \frac{n^5(K+1)}{\alpha} < \infty$. □

Lemma 3. *There is a fixed $P_0 > 0$ such that, if any terminal state s is reached with positive probability under some pure Markovian strategy profile σ, then it is reached with probability at least P_0.*

Proof. Suppose s can be reached under the strategy profile σ. Then, because σ is pure, there is a (finite) sequence of Nature moves $(z_1, \ldots, z_t)$ that *definitely* leads to s. In general, a Nature move takes the form $z_i \in N^2 \cup \emptyset$, where (i, j) means i is chosen as the sender and j as the receiver, and $\emptyset$ means no sender is chosen. Moreover, because σ is Markovian, any sequence of Nature moves that contains the same sequence of *nonempty* moves leads to s as well, as empty moves do not change the state. (For example, $((1,2), \emptyset, (1,3))$ leads to the same outcome as $((1,2),(1,3))$. In particular, if we remove all empty moves, we are left with a sequence of at most $n^3(K+1)$ moves that definitely leads to s. Any such sequence as probability at least $\min\left(1-\alpha, \frac{\alpha}{n^2}\right)^{n^3(K+1)} =: C$ of being realized. □

Rest of the proof. Finally we prove that the equilibrium σ must be natural. The proof is as follows. Denoting the set of segments for $\theta = 1$ by $\overline{M}_k$ for $k = 1, \ldots, \overline{k}+1$, we proceed by backward induction on k. At the k-th step, we show that, beginning at any state in which someone in $\overline{M}_k$ is informed (but no one in a higher segment is), in the terminal state, with probability 1, everyone in $\overline{M}_1, \ldots, \overline{M}_k$ is informed, and no one else is.

For $k = k+1$, suppose that the claim is false, i.e., it is possible that not everyone becomes informed, even if someone in the top segment, $\overline{M}_{\overline{k}+1}$, is informed. Take, then, a set $T \subseteq N$ which is extremal in the following way: T contains someone in $\overline{M}_{\overline{k}+1}$; beginning in state $(1, T, \mathbf{0})$, with positive probability, not everyone becomes informed; but for every $T' \supset T$, beginning in state $(1, T', \mathbf{0})$, everyone becomes informed with probability 1. Let $i_0 = \max(N - T)$. Suppose i_0 is linked to some $j \in T$ such that $|b_{i_0} - b_j| < \frac{1}{2}$. This leads to a contradiction because j would want to inform i_0. Indeed,

j's marginal payoff from deviating and informing i_0 at the first opportunity takes the following form:

$$(1-\delta)\Delta V_j(s, i_0, \sigma) = P \sum_{i \in N-T} \left[-(b_j - b_i)^2 + (b_j - b_i - 1)^2\right] + \mathcal{O}(1-\delta),$$

where $P \geq P_0$ is the probability that T becomes the terminal state on the equilibrium path,[16] and the term $\mathcal{O}(1-\delta)$ captures the residual payoffs generated on the path of play before the terminal state is reached, with or without a deviation. (That these payoffs are bounded by a fixed multiple of $1-\delta$ follows from Lemmas 1 and 2.) Note that the first term is positive since, by construction, $b_i < b_j + \frac{1}{2}$ for all $i \in N - T$, whence $b_{N-T} < b_j + \frac{1}{2}$. And for δ close enough to 1, the first term dominates, so j would deviate.

Hence no such j must be linked to i_0. If $b_n - b_{i_0} \geq \frac{1}{2}$, we obtain another contradiction, as the interval $\left(b_{i_0} - \epsilon, b_{i_0} - \epsilon + \frac{1}{2}\right)$ must be disconnected for $\epsilon > 0$ small enough. If $b_n - b_{i_0} < \frac{1}{2}$, then no agent in $\left(b_n - \frac{1}{2}, b_n\right]$ may be informed, as otherwise, by well-connectedness, there would be a path from an informed j in this interval to i_0, hence a direct link from some informed j' to some uninformed i' in this interval, leading to the same contradiction. Then, let $j_0 = \max(T)$. (By assumption, j_0 must be in the top segment.) By well-connectedness, $\left(b_{j_0} - \epsilon, b_{j_0} - \epsilon + \frac{1}{2}\right) \subseteq [b_1, b_n]$ must be connected for small $\epsilon > 0$, hence j_0 must be linked to some $i_1 > j_0$ who, by construction, is uninformed. j_0's payoff from deviating and informing i_1 takes the following form:

$$\begin{aligned}(1-\delta)\Delta V_{j_0}(s, i_1, \sigma) = P \textstyle\sum_{i > j_0} \left[-(b_j - b_i)^2 + (b_j - b_i - 1)^2\right] + \\ + P \textstyle\sum_{i \in N-T \text{ s.t. } i < j_0} \left[-(b_j - b_i)^2 + (b_j - b_i - 1)^2\right] + \mathcal{O}(1-\delta),\end{aligned}$$

where the first sum is positive by construction of the natural segments, the second sum is positive term by term, and the third term vanishes.

For $k \leq \overline{k}$, the argument has two steps. First, we argue that, if everyone in the segments $\overline{M}_1, \ldots, \overline{M}_k$ has become informed, then, with probability 1, no one else will be informed in equilibrium. We prove this claim with an extremal argument, working backwards from the states in which the activity rule binds. Suppose the claim is false,

[16] By construction, the informed set in the terminal state must be either T or N with probability 1.

and let $s = (1, T, k_{ij})$ be a state in which $T = \cup_{i=1}^{k} \overline{M}_i$; with positive probability, more players are informed in the continuation; and, for every state $s' = \left(1, T, k'_{ij}\right) \neq s$ with $k'_{ij} \geq k_{ij}$, no one else is informed in the continuation with probability 1. (Such a state s must exist because no one else is informed if the activity rule binds). Then, by construction of s and by the inductive hypothesis, if an informed player i is picked as a sender by Nature in state s and an uninformed j is the receiver, i's marginal payoff from informing j is

$$(1 - \delta)\Delta V_i(s, j, \sigma) = \sum_{i \in M} \left[-(b_j - b_i)^2 + (b_j - b_i - 1)^2\right] + \mathcal{O}(1 - \delta),$$

where $M = \cup_{i=k+1}^{k'} \overline{M}_i$ if $j \in \overline{M}_{k'}$. By construction of the natural segments, this payoff is negative for δ close to 1, a contradiction.

Second, we argue that, if someone in $\overline{M}_k$ is informed, everyone in this and lower segments will become informed, and no one in higher segments will. The proof is similar to the case $k = k+1$. Briefly, suppose not, and take $T \subseteq \overline{M}_1 \cup \ldots \cup \overline{M}_k$ to be extremal in the following way: beginning in state $(1, T, \mathbf{0})$, with positive probability, either some players in $\overline{M}_1 \cup \ldots \cup \overline{M}_k$ are never informed, or some players outside of this set are informed (or both); but, for any T' such that $T \subset T' \subseteq \overline{M}_1 \cup \ldots \cup \overline{M}_k$, beginning in state $(1, T', \mathbf{0})$, all players in $\overline{M}_1 \cup \ldots \cup \overline{M}_k$ are informed in the terminal state, and no one else is, w.p. 1. (Such a T must exist because, if the players in $\overline{M}_1 \cup \ldots \cup \overline{M}_k$ and only them are informed, no more players are ever informed.)

By construction, starting from $(1, T, \mathbf{0})$, the only sets of players that can be informed in the long run with positive probability are T or $\overline{M}_1 \cup \ldots \cup \overline{M}_{k'}$ for some $k' \geq k$. (If a player in $\overline{M}_1 \cup \ldots \cup \overline{M}_k - T$ is informed first, then the outcome is $\overline{M}_1 \cup \ldots \cup \overline{M}_k$, by definition of T; if a player in $\overline{M}_{k'}$ $(k' > k)$ is informed first, by the inductive step, the outcome is $\overline{M}_1 \cup \ldots \cup \overline{M}_{k'}$; if no one is informed, the outcome is T.)

By an analogous argument to the one given for $k = \overline{k}+1$, when the set of informed players is T, there must be an informed player i connected to a $j \in \overline{M}_1 \cup \ldots \cup \overline{M}_k - T$, such that either $i \in \overline{M}_k$ or $i > \max(\overline{M}_1 \cup \ldots \cup \overline{M}_k - T)$. In either case, i strictly prefers the outcome $\overline{M}_1 \cup \ldots \cup \overline{M}_k$ to all other possible outcomes. Hence, whenever she can message j, she must do so. Since i will always get a chance to do this before the activity rule binds, T cannot be the set of informed players in the terminal state with positive probability.

Now note that, for all $i \in \overline{M}_1 \cup \ldots \cup \overline{M}_k$, i's utility if $\overline{M}_1 \cup \ldots \cup \overline{M}_{k'}$ is informed in the long run is strictly decreasing in k' for $k' \geq k$. Thus, letting k^* be the maximal k' for which $\overline{M}_1 \cup \ldots \cup \overline{M}_{k'}$ is informed with positive probability, any player who is to inform someone in $\overline{M}_{k^*}$ on the equilibrium path would strictly prefer to deviate, a contradiction.

The argument for $\theta = -1$ is analogous.

Part (ii). The first half follows trivially from Remark 3.

For the second half, we will prove the following. Define a random variable X that takes the values $b_1, \ldots, b_n$ each with probability $\frac{1}{n}$, and denote its c.d.f. by F. Then, if there is $\zeta \leq 1$ such that for all $b' - b \geq \frac{1}{2}$ we have

$$\frac{1}{2\zeta}(b' - b) \geq E(X - b | X \in [b', b]) \geq \frac{\zeta}{2}(b' - b),$$

then the natural segment thresholds $\underline{b}_1 < \ldots < \underline{b}_{\underline{k}}$, $\overline{b}_1 < \ldots < \overline{b}_{\overline{k}}$ satisfy

$$\frac{1}{\zeta} \geq \overline{b}_{i+1} - \overline{b}_i \geq \zeta, \quad \frac{1}{\zeta} \geq \underline{b}_{i+1} - \underline{b}_i \geq \zeta.$$

The argument goes as follows. Suppose $\theta = 1$ is realized, $j+1$ is the lowest member of segment $[\overline{b}_i, \overline{b}_{i+1})$, and j is the highest member of the segment immediately below. It must be that j prefers not to inform the set of agents in $[\overline{b}_i, \overline{b}_{i+1})$ when γ_0 is low enough, in particular when $\gamma_0 = 0$. If $\gamma_0 = 0$, the net payoff from informing them is

$$\begin{aligned}
0 \geq \Delta &= -\int_{b_j}^{\overline{b}_{i+1}} (b' - b_j) dF(b') + \int_{b_j}^{\overline{b}_{i+1}} (b' - b_j - 1) dF(b') \\
&= \int_{b_j}^{\overline{b}_{i+1}} [1 - 2(b' - b_j)]\, dF(b') \quad \Longrightarrow 1 \geq 2E(X - b_j | X \in (b_j, \overline{b}_{i+1}]) \\
&\Longrightarrow \frac{1}{\zeta} \geq \overline{b}_{i+1} - b_j \quad \Longrightarrow \frac{1}{\zeta} \geq \overline{b}_{i+1} - \overline{b}_i.
\end{aligned}$$

Conversely, $j+1$ must prefer to inform the set of members of $[\overline{b}_i, \overline{b}_{i+1}]$ above him.

When $\gamma_0 = 0$, the net payoff from informing them is

$$\begin{aligned}
0 \leq \tilde{\Delta} &= -\int_{b_{j+1}}^{\bar{b}_{i+1}} (b' - b_{j+1}) dF(b') + \int_{b_{j+1}}^{\bar{b}_{i+1}} (b' - b_{j+1} - 1) dF(b') \\
&= \int_{b_{j+1}}^{\bar{b}_{i+1}} \left[1 - 2(b' - b_{j+1})\right] dF(b') \quad \Longrightarrow 1 \leq 2E(X - b_{j+1} | X \in (b_{j+1}, \bar{b}_{i+1}]) \\
&\Longrightarrow \bar{b}_{i+1} - b_{j+1} \geq \zeta \quad \Longrightarrow \bar{b}_{i+1} - \bar{b}_i \geq \zeta.
\end{aligned}$$

The proof for the thresholds $\underline{b}_i$ is analogous.

□

Proof of Proposition 3. To begin, construct a natural equilibrium when $\gamma_0 = 0$. This can be done in exactly the same fashion as in Proposition 2: since informed players have common knowledge of θ, and the beliefs of uninformed players are independent of the exact distribution of θ as long as $E(\theta) = 0$, the continuations after Nature picks each possible value of θ can be solved separately. Indeed, the case of any $\theta > 0$ is equivalent to $\theta = 1$ in the binary case if we multiply all the biases by $\frac{1}{\theta}$, and an analogous argument applies to $\theta < 0$. Finally, for $\theta = 0$, everyone is indifferent at all times. Denote such a natural equilibrium for the case $\gamma_0 = 0$ by σ_0.[17]

Next we move on to the case $\gamma_0 > 0$. First we show that, for small $\gamma_0 > 0$, we have $\overline{\theta}(1, \gamma_0) < \ldots < \overline{\theta}(n, \gamma_0)$, where we are making the dependence on γ_0 explicit. The argument goes as follows. Calculate the long-run mean posteriors under ignorance for 1, 2, …, n under the assumptions that $\gamma_0 = \gamma > 0$ *but σ_0 is still played*. Denote these by $\overline{\theta}^d(1, \gamma), \ldots, \overline{\theta}^d(n, \gamma)$. It can be shown directly that $\overline{\theta}^d(1, \gamma) < \ldots < \overline{\theta}^d(n, \gamma)$ for all $\gamma \in (0, 1)$, and moreover that $\frac{\partial \overline{\theta}^d(1,\gamma)}{\partial \gamma}|_{\gamma=0} < \ldots < \frac{\partial \overline{\theta}^d(n,\gamma)}{\partial \gamma}|_{\gamma=0}$. (This follows from the fact that σ_0 is segmented for each θ.)

Next, it is straightforward to show that the equilibrium σ_0 is strict for a generic set of values of $\theta \in [-1, 1]$. Indeed, σ_0 restricted to a certain realization of $\theta > 0$ is guaranteed to be strict if, for all i and all $A \subseteq N$, the expression $-b_A + b_i + \frac{\theta}{2}$ is nonzero—a finite number of linear equations in θ, in all of which θ has nonzero coefficient. Hence σ_0 is strict when restricted to all but finitely many values of θ.

[17] Note that, for θ close enough to zero, the network will no longer be well-connected, even if i is linked to $i+1$ for all i, since the differences $\frac{b_{i+1}-b_i}{\theta}$ will grow bigger than $\frac{1}{2}$. However, by arguments analogous to the proof of Proposition 2, the equilibrium will still be natural; when the adjusted bias gaps between consecutive agents grow bigger than $\frac{1}{2}$, this will simply lead to one-person segments.

As a result, there is a function $m(\gamma)$, with $m(\gamma) \xrightarrow[\gamma\to 0]{} 0$, such that if we assume any mean posterior belief paths under ignorance bounded by $\pm\frac{\gamma}{2-\gamma}$, σ_0 is still a strict equilibrium for all values of θ except for a set of measure up to $m(\gamma)$. Therefore, $\frac{\partial\overline{\theta}^d(i,\gamma)}{\partial\gamma}|_{\gamma=0} = \frac{\partial\overline{\theta}(i,\gamma)}{\partial\gamma}|_{\gamma=0}$. This implies that $\overline{\theta}(1,\gamma_0) < \ldots < \overline{\theta}(n,\gamma_0)$ for small $\gamma_0 > 0$. It also implies that, for all $\gamma_0 \in (0,\gamma)$, the equilibrium is unique and natural for a set of values of θ of measure at least $2 - m(\gamma)$. This proves (i).

For (ii), we will show that there is a strict equilibrium with full learning when $\theta = 0$. By continuity of the utility functions in θ and the backward induction arguments employed in Proposition 2, this will imply that there is full learning for all θ in a neighborhood of 0, as we want.

Consider first a modified game in which the mean posterior belief paths under ignorance are identically equal, by assumption, to their long-run values in the real equilibrium. That is, assume that the players' actions under ignorance are as if $\overline{\theta}(i,t) = \overline{\theta}(i)$ for all i, t. This game has an equilibrium which is generically unique, pure and Markovian, by the same arguments from Proposition 2. Suppose that, in this equilibrium, full learning fails to obtain with positive probability when $\theta = 0$. Again, by the same arguments from Proposition 2, this means that a terminal state with partial learning is reached with probability bounded away from zero, $P \geq P_0 > 0$.

Now take $T \subset N$ with the following property: beginning in state $(0, T, \mathbf{0})$, full learning fails to obtain with positive probability, but for any $T' \supset T$, full learning obtains w.p. 1 beginning in state $(0, T', \mathbf{0})$. By construction, the set of informed players in the long run, beginning in state $(0, T, \mathbf{0})$, must be either T or N with probabilities adding up to 1, as anyone else being informed triggers full learning.

As in Proposition 2, we can reduce a player i's incentives to inform one of her neighbors j to the effect of such a message on the set of informed players in the terminal state, plus second-order terms. Thus, if there is an informed player $i \in T$ that would strictly prefer all of $N - T$ being informed vs. not, and who is linked to an uninformed player j, i would always inform j at the first opportunity, leading to full learning—a contradiction. Thus, no such i can exist.

Let i_0 be the lowest informed player who is adjacent to an uninformed player (i.e., either $i_0 - 1$ or $i_0 + 1$ is uninformed), and let i_1 be the highest. Then it must be that both prefer to leave all of $N - T$ uninformed rather than inform them. Suppose that $b_{N-T} \in [b_{i_0}, b_{i_1}]$. Informing the members of $N - T$ has two effects: it may increase or decrease the average action chosen by them, depending on whether $\sum_{j\in N-T} \frac{\overline{\theta}(j)}{|N-T|}$ is smaller or greater than 0, and it will decrease the variance of their actions—something desired by both i_0 and i_1—since $\overline{\theta}(j)$ is increasing in j. Then, if $\sum_{j \in N-T} \frac{\overline{\theta}(j)}{|N-T|} < 0$, i_1 would strictly prefer to inform the group, while in the opposite case i_0 would strictly prefer to inform, and in case of equality, they both would, a contradiction in any case.

Then we must have $b_{i_0} \leq b_{i_1} < b_{N-T}$ or $b_{N-T} < b_{i_0} \leq b_{i_1}$. WLOG assume the former. For i_1 to prefer not to inform $N-T$, we must have $\sum_{j \in N-T} \frac{\overline{\theta}(j)}{|N-T|} < 0$. This implies that $b_{i_1} > b_N$, as otherwise all the players missing from N in $N-T$ would be players with below-average bias, hence with negative mean posterior under ignorance (and note that the population average of the mean posteriors must be zero, by the symmetry assumption).

Divide the set of players into five groups: $D = \{i_1\}$; her opposite $B = \{n+1-i_1\}$ (where, since $b_{i_1} > b_N$, we must have $i_1 > n+1-i_1$); $E = \{i_1+1, \ldots, n\}$; their opposites $A = \{1, \ldots, n-i_1\}$; and the players in the middle, $C = \{n+2-i_1, \ldots, i_1-1\}$. i_1 strictly prefers to inform each player j with $b_j < b_N$ (since $b_j < b_{i_1}$, and $b_j < b_N$ implies $\overline{\theta}(j) < 0$), as well as each player j with $b_j > b_{i_1}$ (since $b_j > b_{i_1} > b_N$ means that $\overline{\theta}(j) > 0$). In particular, i_1 strictly prefers to inform any players in $A \cup E$ that are still uninformed.

As for players in $B \cup C$, note the following. By construction $E \subseteq N-T$, while in general only some subset of A is included in $N-T$. Hence, by the symmetry assumption, $\sum_{j \in (A \cup E) \cap (N-T)} \overline{\theta}(j) \geq 0$, which implies $\sum_{j \in (B \cup C) \cap (N-T)} \overline{\theta}(j) < 0$. Then i_1 also strictly prefers informing $(B \cup C) \cap (N-T)$ versus not, since informing them will increase their average action (note $b_j < b_{i_1}$ for all $j \in B \cup C$) and reduce the variance of their actions. Thus i_1 strictly prefers to inform a neighbor, thus guaranteeing full learning, versus not, and full learning must obtain with probability 1.

□

Proof of Proposition 4. Assume WLOG that $\theta = 1$. The proof is by induction on k.

For any player i with $k = 1$, if i is informed by her upstream neighbor, each of her other (downstream) neighbors j has no one else to inform. If $b_j < b_i + \frac{1}{2}$, then, since $\overline{\theta}(j,t) \in [-\frac{\gamma_0}{2-\gamma_0}, \frac{\gamma_0}{2-\gamma_0}]$, for $\gamma_0 > 0$ low enough, i strictly prefers for j to be informed at all times, so i informs j at the first opportunity. If $b_j > b_i + \frac{1}{2}$, i strictly prefers for j to be uninformed forever, and never informs j. The case $b_j = b_i + \frac{1}{2}$ is non-generic.

Now, suppose the result is true up to $k-1$. Let i be a player at distance $\overline{l} - k$ from i_0, and let j be a downstream neighbor. By the inductive hypothesis, if j is informed, then every member of A is informed, and no one else (downstream of j).

Moreover, because all players who inform others do so at the first opportunity, and the network is finite, the number of periods until any member of A is informed is uniformly bounded in expectation. In addition, if i informs j at time t, the expected delay (counting from t) until any other agent in A is informed is not a function of t (again, because messages are sent at the first opportunity). Thus i's marginal payoff from informing j at time t, relative to doing so at time $t+1$, can be written as

$$\sum_{j' \in A} \sum_{s=0}^{\infty} \delta^s (1-\delta) \left[-(b_{j'} - b_i)^2 + (b_{j'} + \overline{\theta}(j', t+s) - b_i - 1)^2 \right] Q(j', s),$$

where $Q(j', s)$ is the probability that it would take s periods for j' to be informed, after i informs j. As $\delta \to 1$ and $\gamma_0 \to 0$, this can be written as

$$\begin{aligned} (1-\delta) \sum_{j' \in A} [1 - 2(b_{j'} - b_i) + \mathcal{O}(1-\delta) + \mathcal{O}(\gamma_0)] = \\ =(1-\delta)|A| [1 - 2(b_A - b_i)] + \mathcal{O}(1-\delta) + \mathcal{O}(\gamma_0). \end{aligned}$$

This follows from the fact that $\overline{\theta}(j', t+s) \in [-\frac{\gamma_0}{2-\gamma_0}, \frac{\gamma_0}{2-\gamma_0}]$ for all j', s, and from the fact that the expected delay until each j' is informed is finite.

In the generic case where $b_A \neq b_i + \frac{1}{2}$, if $b_A < b_i + \frac{1}{2}$, then this expression is positive for all δ close enough to 1 and γ_0 low enough. Hence i prefers informing at t over informing at $t+1$ for all t. Therefore, i prefers informing immediately to waiting until any future opportunity (or never informing). If $b_A > b_i + \frac{1}{2}$, then informing at t is worse than informing at $t+1$ for all t, and by the same logic it is optimal to never inform j.

□

Proof of Remark 4. We will show a stronger statement: that, for each i, j and t, $E\left((a_{jt} - \theta - b_i)^2\right) \leq E\left((\tilde{a}_{jt} - \theta - b_i)^2\right)$. (Here a_{jt}, $\tilde{a}_{jt}$ denotes j's action at time t under m and $\tilde{m}$, respectively, and E is an expectation given only the prior.)

Clearly this is true for $i = j$, as j can make better decisions with strictly more information. In addition,

$$\begin{aligned} E\left((a_{jt} - \theta - b_i)^2\right) - E\left((a_{jt} - \theta - b_j)^2\right) &= E\left(a_{jt} - \theta\right) 2(b_j - b_i) + b_i^2 - b_j^2 = \\ &= 2b_j(b_j - b_i) + b_i^2 - b_j^2, \end{aligned}$$

where the last equality follows from the fact that $E(a_{jt}) = E(\theta) + b_j$ for any information structure, by Remark 2 and the law of iterated expectations. Because the right-hand side is independent of the message strategy profile,

$$E\left((a_{jt} - \theta - b_i)^2\right) - E\left((\tilde{a}_{jt} - \theta - b_i)^2\right) = E\left((a_{jt} - \theta - b_j)^2\right) - E\left((\tilde{a}_{jt} - \theta - b_j)^2\right) \leq 0.$$

□

Proposition 6. *Let $G(n,p)$ be a large Erdõs-Rényi random graph, taking the population $N = \{1, \ldots, n\}$ and the distribution of biases $(b_1, \ldots, b_n)$ as fixed. Let $K = \frac{2(b_n - b_1)(2-\gamma_0)}{1-\gamma_0}$ and $m = \frac{n}{K}$. Suppose that $(b_1, \ldots, b_n)$ are roughly uniformly distributed, i.e., $(1-\epsilon)\frac{b_j - b_i}{b_n - b_1} \leq \frac{|\{i, \ldots, j\}|}{n} \leq (1+\epsilon)\frac{b_j - b_i}{b_n - b_1}$ whenever $b_j - b_i \geq \frac{1-\gamma_0}{2(2-\gamma_0)}$. Let $np = \frac{(\ln m + c)K}{1-\epsilon}$. Then (N, G) is well-connected with probability $1 - 3(K-1)\frac{e^{-c}}{1-\epsilon}$ or greater for n large enough.*

Proof of Proposition 6. Partition the interval $[b_1, b_n]$ into subintervals of size $Z = \frac{(1-\gamma_0)}{2(2-\gamma_0)}$. In other words, write $[b_1, b_n] = I_1 \cup \ldots \cup I_{\overline{m}}$, where $I_l = [b_1 + (l-1)Z, b_1 + lZ)$ for $l < \overline{m}$ and $I_{\overline{m}} = [b_1 + (\overline{m} - 1)Z, b_n]$.

The proof is in two steps. First we argue that, to guarantee well-connectedness, it is sufficient that, for each l s.t. $1 < l < m$,

(i) each interval I_l is connected;

(ii) each $i \in I_{l-1}$ has a link to some agent in I_l, and each $i \in I_{l+1}$ has a link to some agent in I_l.

The proof is as follows. Let $(a, b) \subseteq [b_1, b_n]$ be an interval such that $b - a = \frac{1-\gamma_0}{2-\gamma_0}$. By construction, (a, b) contains at least one subinterval I_l (for $1 < l < \overline{m}$). Let i be an agent with $b_i \in (a, b)$. Then we must have $i \in I_{l-1}$, I_l or I_{l+1}. If $i \in I_l$, i is connected to the rest of I_l by (i). If $i \in I_{l-1}$, she is connected to someone in I_l by (ii). Thus, all agents in (a, b) are connected through I_l.

The second step is to calculate the probability $P(n)$ that a random network on n agents (satisfying the conditions imposed on the distribution of biases, as well as p) satisfies the conditions (i) and (ii). We can provide a lower bound for $P(n)$:

$$P(n) \geq \prod_{l=2}^{\overline{m}-1} P_l(n) - \sum_{l=2}^{\overline{m}-1} \left[1 - \left(1 - (1-p)^{|I_l|}\right)^{|I_{l-1}| + |I_{l+1}|}\right],$$

where $P_l(n)$ is the probability that I_l is connected, $\left(1-(1-p)^{|I_l|}\right)^{|I_{l-1}|+|I_{l+1}|}$ is the probability that all members of I_{l-1} and I_{l+1} have a link to I_l, and we are using the fact that $P(A \cap B_1 \cap \ldots \cap B_k) \geq P(A) - \sum_{i=1}^{k} P(B_i)$.

By our assumptions, $\frac{|I_l|}{n} \geq \frac{(1-\epsilon)Z}{b_n - b_1} = \frac{1-\epsilon}{K}$, i.e., $|I_l| \geq (1-\epsilon)m =: m'$. On the other hand, $p = \frac{\ln m + c}{(1-\epsilon)m} = \frac{\ln m' + c + \ln(1-\epsilon)}{m'}$.

It is known that, in an Erdõs-Rényi random graph (a, q) with link probability $q = \frac{\ln a + c}{a}$, the probability of the graph being connected converges to $e^{-e^{-c}}$ as $a \to \infty$ (Erdos and Rényi, 1960). Therefore, for any sequence of n-random networks satisfying the conditions on the bias vectors, $\liminf_{n\to\infty} P_l(n) \geq e^{-e^{-c-\ln(1-\epsilon)}}$.

In addition, $(1-p)^a \to e^{-ap}$ as $a \to \infty$ if ap converges to a constant. Hence $(1-p)^{|I_l|} \approx e^{-|I_l|p}$, and $\left(1-(1-p)^{|I_l|}\right)^{|I_{l'}|} \approx e^{-e^{-|I_l|p}|I_{l'}|}$. Since $(1-\epsilon)m \leq |I_l| \leq (1+\epsilon)m$, we have

$$e^{-e^{-|I_l|p}|I_{l'}|} \geq e^{-e^{-(1-\epsilon)m\frac{\ln m + c}{(1-\epsilon)m}}(1+\epsilon)m} = e^{-e^{-c}(1+\epsilon)}$$

$$\sum_{l=2}^{\overline{m}-1}\left[1-\left(1-(1-p)^{|I_l|}\right)^{|I_{l-1}|+|I_{l+1}|}\right] \approx \sum_{l=2}^{\overline{m}-1}\left[1-e^{-e^{-|I_l|p}(|I_{l-1}|+|I_{l+1}|)}\right] \leq$$
$$\leq (\overline{m}-2)\left(1-e^{-2e^{-c}(1+\epsilon)}\right) \leq (\overline{m}-2)2e^{-c}(1+\epsilon).$$

Thus

$$\begin{aligned}\liminf_{n\to\infty} P(n) \geq & e^{-e^{-c-\ln(1-\epsilon)}(\overline{m}-2)} - (\overline{m}-2)2e^{-c}(1+\epsilon) \geq \\ \geq & 1 - e^{-c-\ln(1-\epsilon)}(\overline{m}-2) - (\overline{m}-2)2e^{-c}(1+\epsilon) \geq 1 - 3(\overline{m}-2)\frac{e^{-c}}{1-\epsilon}.\end{aligned}$$

Finally, by construction, $\overline{m} = \lceil K \rceil$ so $\overline{m} \leq K+1$. Then

$$\liminf_{n\to\infty} P(n) \geq 1 - 3(K-1)\frac{e^{-c}}{1-\epsilon}.$$

□

If we take $n = 320.000.000$, γ_0 close to zero, $(b_1, b_n) = (-2, 2)$ and the distribution of biases close to uniform (that is, ϵ close to zero), then $K = 16$, $m = 20.000.000$, $\ln(m) \approx 16.81$, and $\liminf P(n) \geq 1 - 45e^{-c}$. Then, if we take $c = 8.42$, at least 99% of the resulting networks are well-connected. This choice of c implies an average degree $np = 404$.

If the network exhibits homophily, the condition that a player's average global degree be $(\ln m + c)K$ is no longer necessary. Consider, for instance, a model in which a link between each pair of players i and j is formed with probability p_1 if $|b_j - b_i| \leq \frac{1-\gamma_0}{2-\gamma_0}$, and with probability p_2 otherwise. Such a model generates networks with ideological homophily if $p_1 > p_2$.[18]

Note that whether conditions (i) and (ii) are satisfied by a network depends exclusively on links of the first type, as both conditions involve links within an interval I_l or between consecutive intervals, and each of these intervals has length $\frac{1-\gamma_0}{2(2-\gamma_0)}$. Then, reproducing the logic of Proposition 6, (N, G) is well-connected with probability $1 - 3(K-1)\frac{e^{-c}}{1-\epsilon}$ or greater for large n, if $np_1 = \frac{(\ln m + c)K}{1-\epsilon}$. This implies an average degree of at most $\frac{4+\epsilon}{K}\frac{(\ln m+c)K}{1-\epsilon} + \left(1 - \frac{4-\epsilon}{K}\right) np_2$, since the fraction of players at distance $\frac{1-\gamma_0}{2-\gamma_0}$ or less from a player i is at most approximately $\frac{4}{K}$ (if i is not at an extreme of the distribution). For instance, in our example, taking $p_1 = \frac{1}{n}\frac{(\ln m+c)K}{1-\epsilon}$ and $p_2 = 0$, we obtain networks with a probability of 99% or higher of being well-connected, but average degree of at most $\frac{4}{16}404 + \frac{12}{16}0 = 101$.

Proof of Proposition 5. Given the natural segments, construct a strategy profile where, when $\theta = 1$, each $i \in \overline{M}_l$ always messages every member of $\overline{M}_l$ and lower segments, but no one in higher segments; and analogously for $\theta = -1$. By construction, this strategy profile is natural. For it to not be an equilibrium, there would have to be a strictly profitable deviation for some player. But the only deviations which change the long-run payoffs are those that send messages to the wrong segments (e.g., when $\theta = 1$, $i \in \overline{M}_l$ would have to message someone in $\overline{M}_{l'}$, with $l' > l$), which are strictly unprofitable by construction. □

[18] This is a particular kind of Social Distance Attachment model (Boguná et al., 2004).