Consensus and Disagreement:
Information Aggregation under (not so) Naive Learning*

Abhijit Banerjee and Olivier Compte

April 2023

Abstract

We explore a model of non-Bayesian information aggregation in networks. Agents non-cooperatively choose among Friedkin-Johnsen type aggregation rules to maximize payoffs. The DeGroot rule is chosen in equilibrium if and only if there is noiseless information transmission... leading to consensus. With noisy transmission, while some disagreement is inevitable, the optimal choice of rule blows up disagreement: even with little noise, individuals place substantial weight on their own initial opinion in every period, which inflates the disagreement. We use this framework to think about equilibrium versus socially efficient choice of rules and its connection to polarization of opinions across groups.

1 Introduction

As of May 2020, 41% of US Republicans were not planning to get vaccinated against Covid-19, as compared to 4% of Democrats.[1] We saw similar divergences in mask-wearing, social distancing etc, which protect against the disease. Since Covid-19 is a life-threatening ailment that had already taken more than 3.5 million lives so far world-wide, it is hard to think of these as being just empty gestures or entirely reflective of different preferences, though there is surely some of that. There seems to be rather, a different reading of the facts on the ground; for example, in a Pew Research Center poll,[2] Republicans were much more likely to say that Covid-19 is not a major threat to the health of the US population (53% compared to 15% of Democrats). This goes with a general deepening in the political divide between Democrats and Republicans in recent years.[3]

The source of this shift is a subject of much discussion: one potential source of change is the massive growth in the use of the internet. However the evidence from the careful work by Gentzkow and Shapiro (2011) suggests that online news consumption is not more segregated by political leanings than other sources of information that already existed, contrary to the concerns expressed for example by Sunstein (2001).[4] The most segregated sources of

*This paper was previously entitled "Information Aggregation under (not so) Naive Learning".

[1] https://www.pbs.org/newshour/health/as-more-americans-get-vaccinated-41-of-republicans-still-refuse-covid-19-shots

[2] https://www.pewresearch.org/fact-tank/2020/07/22/republicans-remain-far-less-likely-than-democrats-to-view-covid-19-as-a-major-threat-to-public-health.

[3] Pew Center (2014) documents such a shift on a 10-point scale of political values for the period 1994-2014)

[4] Though Guess (2021) suggests that the segregation in news consumption has been increasing in recent years.

information, according to Gentzkow and Shapiro (2011) seem to be social networks (voluntary associations, work, neighborhoods, family, “people you trust”, etc), which were of course always there. However there is evidence that online networks such as Facebook are substantially more segregated than other social networks and as a result, news that comes from being shared on Facebook tends to be more segregated than news from other media sources (Bakshy et al. (2015)).[5] It is true that social media are still a relatively small, though growing, part of news consumption, but the volume of ”information” that can be quickly shared on Facebook may be larger than other more traditional sources. Moreover while information was always shared through social connections, the evidence of growing affective polarization along political lines, especially in the US (Boxell et al. (2022)), raises the concern that the actual exchange of sensitive information in the social network is increasingly confined to those on the same side.

We feel therefore that it is worth exploring theoretically when and why social learning on networks can lead to large and persistent disagreements. As a starting point, we note that models of Bayesian social learning such as Acemoglu et al. (2011) propose relatively weak conditions on signals and network structure under which information is perfectly aggregated as the network grows to be very large. More recent work, in which agents repeatedly communicate (unlike in Acemoglu et al. (2011) where they communicate only once) include Mossel et al. 2015 who derive necessary conditions on the network structure under which Bayesian learning yields consensus and perfect information aggregation.[6] The general sense from this literature is that convergence to a consensus is likely even when the network exhibits a substantial degree of homophily (Republicans mostly talk to other Republicans) as long as everyone is ultimately connected.

This *Bayesian route* however requires that agents make correct inferences based on an understanding of all the possible ways information can transit through the network, which, at least for large networks, strains credibility.[7]

The alternative way to model learning on networks is to take a *non-Bayesian route*, which avoids these very demanding assumptions about information processing by postulating a simple rule that individuals use to aggregate own and neighbors' opinions. In recent years the economics literature has tended to favor the DeGroot (DG) rule, where agents update their current opinion by linearly averaging it with their neighbors' most recent opinions. As observed by DeMarzo et al. (2003), who brought it into the economics literature, the rule builds in a strong tendency towards consensus in any connected network, even when there is high degree of homophily and people put high weight on people like them, though convergence between those far from each other in the network can be very slow.[8] Faced with this force towards consensus, Friedkin and Johnsen (1990) came up with a learning rule which is similar to DG, but allows each individual to keep putting some weight on their own initial opinion.[9] This

[5]The Facebook news feed turns out to be even more seggregated (Levy (2021)).

[6]They build on Rosenberg et al. (2009) and the literature on “Agreeing to Disagree”that goes back to Aumann (1976).

[7]A Bayesian needs to think through all possible sequences of signals that could be received as a function of the underlying state and all the possible pathways through which each observed sequence of signals could have reached them. As discussed in Alatas et al. (2016, page 1681), there is obviously an extremely large number of such pathways.

[8]Moreover as shown by Golub-Jackson (2010), DG has the striking property that, under some restrictions on network structure and weights on neighbors, learning converges to perfect information aggregation in large networks.

[9]Friedkin and Johnsen (1999, page 3) write, referring to the work of DeGroot and other precursors: “These initial formulations described the formation of group consensus, but did not provide an adequate account of

rule, for obvious reasons, does not lead to a consensus.

The first question we set out to answer here is which type of rule, i.e., Friedkin-Johnsen (FJ) or DG would be favored by individuals given a choice. In other words are there good reasons to prefer rules where individuals anchor themselves to their initial beliefs even while updating their opinions based on what they are hearing from others?

To study this question, we start from a broad class rules in the spirit of Friedkin-Johnsen (FJ), which includes DG and can formally be written as

$$y_i^t = (1-\gamma_i)y_i^{t-1} + \gamma_i(m_i x_i + (1-m_i)z_i^{t-1}) \tag{FJ}$$

where y_i^t is i's belief in period t, x_i is the initial signal that i received, correlated with some underlying state of the world, and

$$z_i^t = \sum_{j \in N_i} A_{ij} y_j^t + \varepsilon_i^t \tag{1}$$

is the weighted average of reports received by i from his neighbors (denoted N_i),[10] plus any processing or transmission error. This error term is an important ingredient of our analysis. We assume that ε_i^t has two components, a persistent one, drawn at the start of the process, and an idiosyncratic one, drawn at each date, though, to simplify the exposition, much of the paper focuses on persistent errors. When the weight m_i is 0, individual i is using a DG rule.[11]

Within this limited class of "natural" rules, parameterized by γ_i and m_i,[12] we allow agents full discretion in the choice of rules and assume that each individual non-cooperatively selects m_i and γ_i to ensure that the long-run opinion y_i is on average closest to the underlying state. This is in the spirit of the approach advocated in Compte and Postlewaite (2018) to model mildly sophisticated agents.[13]

Our results highlight the major role of errors in shaping equilibrium choices and outcomes. Result 1 says that *absent errors*, each individual decision-maker will choose DG ($m_i = 0$) in the Nash equilibrium of the rule-choice game, hence there will be consensus. Moreover, we show that each individual will choose γ_i in such a way that information is efficiently aggregated. This result thus complements Golub-Jackson (2010) who show that when everyone does DG (but do not choose their γ_i), information aggregation in large networks is almost perfect under certain weak conditions, but generally imperfect in finite networks.

In contrast, Result 2 shows that in the presence of any error in transmission, each decision-maker must choose $m_i > 0$ in equilibrium, so there will be no consensus even in the long run. The reason is that when all the m_i are small (a fortiori when everyone uses DG) the errors tend to cumulate, with the result that long-run opinions explode. Intuitively, a positive error by i pushes up i's opinion, which raises the opinions of others', fueling a further rise in i's opinion, and so on –We call these *echo effects*. Raising m_i allows individuals to limit this cumulation of errors, at the cost of potentially putting too much weight on their own seeds. Moreover there is no way to use γ_i to mitigate this problem: in fact as long as there is no *idiosyncratic* error

settled patterns of disagreement".

[10]The matrix $A = (A_{ij})_{ij}$ defines the weight A_{ij} that i puts on j's opinion, with $A_{ij} > 0$ if and only if $j \in N_i$, and $\sum_j A_{ij} = 1$.

[11]Throughout our analysis, we assume that all γ_i are strictly positive.

[12]We assume that the weights A_{ij} are fixed, not subject to optimization.

[13]The limitation to a specific class of rules is key. Otherwise the individually optimal way to process signals among all possible signal processing rules would be the Bayesian rule.

and $m_i > 0$ for a least one player, γ_i's play no role: long-run opinions are *fully* determined by the m_i's. Later in the paper we show that γ_i does play an important role in controlling the effects of idiosyncratic errors, but that does not change the need to set $m_i > 0$.

It should be clear that in any Nash Equilibrium of the rule choice game, there are two sources of divergence of opinions–the errors themselves, but also the additional divergence that comes from always putting non-zero weight on one's initial signal (which is a choice, but one resulting from the presence of errors). The next question is which is the main source of divergence.

Result 3 shows that at least when the variance ϖ of persistent error is close enough to zero, the second, non-mechanical, source dominates: specifically we show that in equilibrium, the weights m are comparable to $\varpi^{1/3}$. A rough intuition goes as follows: from the perspective of player i, when other players use $m_j \simeq m$, the cumulated error he faces has a long-run variance of the order of ϖ/m^2. i will want to set m_i to counterbalance this, which means at the order of ϖ/m^2. Therefore in equilibrium, $m \simeq o(\varpi/m^2)$.

We then ask whether there is too little or too much disagreement in any equilibrium relative to the social optimum. Result 4 shows that the equilibrium values of m_i are always lower than the socially optimal values. One reason is that in setting m_i optimally, player i does not take into account the fact that lowering m_i raises the cumulated error faced by j. But this is not the only reason. A higher value of m_i reduces the influence of the transmission error, but it also reduces the weight on the opinions of others, which, especially in the long run, enables i to aggregate signals from all over the network and therefore provides very valuable information not contained i's own signals. This is the tradeoff that i takes account of in equilibrium. What he does not take into account is the fact that when m_i goes up, y_i reflects more the information contained in i's signal as against what i learnt from everyone else (which in the long run is very close to what i's neighbors **too** learnt from everyone else) and this is valuable for aggregate welfare. Technically, raising m_i diminishes the correlation between y_i and others' signals, and this enhances others' welfare.

We turn next to comparisons of the efficiency of information aggregation on specific simple networks – the complete network, the directed circle and the star network). A central aspect of our analysis is the characterization of cumulated errors that each individual faces, and how then each player mitigates the consequence of these errors by controlling the weight p_i of her own seed x_i in his or her own long-run opinion, through the choice of m_i. The network structure (through the weight matrix A) matters in two ways: it affects the variance $\widehat{\varpi}_i$ of cumulated errors (for given m), and it affects the elasticity $h_i = \partial p_i / \partial m_i$. Higher variance and higher elasticity both lead to worse information aggregation (higher elasticity hurts because it reduces the incentive to raise m_i, pushing it further away from the social optimum). We find that the star network performs worse than the two others, essentially because the central player propagates correlated errors to all peripheral players, thus raising cumulated errors. The comparison between complete and directed circle networks depend on the size of the network, small size giving an advantage to the directed circle (because of lower elasticity – due to less pronounced echo effects), while large size giving an advantage to the full network because of more effective averaging of the signals.

We next use our example of the star network to address the issue of polarization. The result that m_i is too low might suggests that there is always too little disagreement in equilibrium. This is true for two-person networks, but not in general. To see this consider a network where there are two dense clusters (modeled as stars) connected by one link (say). Such a network structure is not too dissimilar, for example, to the networks of Republicans and Democrats in the US, who mostly communicate with each other (Cox et al. (2020)). In this case, we show by example that there is a natural reason why lower m_i may be associated with a high degree of consensus within each cluster but extreme polarization across the groups, reminiscent of the situation of the Republicans and Democrats in the US. The general point, captured by Result 5, is that social efficiency requires the dispersion of opinions *within* and *between* subgroups to have same orders of magnitude.

Our very simple model therefore tells a useful story why disagreements are necessary, but also about why there could be wide divisions in opinions and when such disagreements are costly.

The rest of the paper is devoted to showing that our basic insights are robust. We first return to rule selection in the case where there are idiosyncratic shocks in information transmission in addition to permanent shocks. In this setting, the speed of updating, γ_i, also plays a role. Slowing down updating by setting γ_i close to zero allows the agent to minimize the changes in opinions that result from these shocks, which is an advantage because the shocks average out over time. This is what Result 6 shows.

Since errors in transmission are central to the case we make for choosing $m_i > 0$, in the penultimate section of the paper we examine the robustness of our results to other ways of modeling frictions. We start by examining the implications of agents adding a slant to the opinions they share–in other words adding errors that are biased in some direction. Recent results from a survey experiment suggest this is a real problem–people on social media are more likely to pass on messages that are more concordant with their political opinions, somewhat irrespective of the accuracy of the message (Pennycook et al. (2021). We note that biased errors do not produce any essential changes in our analysis, though there is a further shift towards reliance on one's own initial signal (higher m).

We next turn to the possibility of coarse communication–say each party only reports their current best guess about which of two actions is preferable. In this setting, the class of potentially "natural" rules include the infection models, studied in Jackson (2008) among (many) others, and the related class of models studied by Ellison and Fudenberg (1993, 1995), in which agents may rely on the popularity of a particular action among neighbors. We show that systematic errors in interpreting actions by neighbors makes the long-run outcome from a DG-like rule entirely insensitive to the actual state of the world, but this is not true for FJ-type rules. We use this framework to discuss the connection between the errors we introduce and mis-specifications in Bayesian models (as in Frick et al (2020) and Bohren and Hauser (2021) and the related (non-)robustness of long-run beliefs.

To end Section 7 we highlight some examples where our findings *are* qualitatively altered. We have so far assumed that agents know the precision of everyone's initial signals. We now explore the possibility that uncertainty about the precision of everyone else's signal is the only source of friction in communication. We find that, in the absence of transmission errors, this *does not* undermine the performance of DG-type rules. As a matter of fact, in a set-up where each participant only knows the precision of their own initial signal, perfect information aggregation can be achieved under DG, by choosing γ_i that is suitably scaled to the precision. This observation delineates the key role played by transmission shocks in our analysis, as opposed to other sources of shocks.

We next allow for the possibility that a friction comes from variations in who speaks when. We show that under FJ rules long-run opinions are independent of the communication protocol.

In contrast, we show by example that the outcome with DG rules is sensitive to who speaks when. So even in the absence of noise, under protocol uncertainty, the performance of DG rules would be impaired (though to a lesser extent than that induced by cumulated errors– long-run opinions would *not* blow up, but remain weighted averages of initial opinions).

We conclude with a discussion of non-stationary rules and when and why they may not always be appropriate.

1.1 Related Literature

Our paper contributes to the large literature on learning in social network (see the excellent review by Golub and Sadler (2016)). We study non-Bayesian learning on general networks with continuous choices and general networks. Within Bayesian social learning, Vives (1993,1997) studies a setting similar to ours (with agents receiving a noisy signal) and, unlike us, obtains long-run convergence to the truth. The reason is that with continuous choice sets Bayesian agents are able to perfectly extract the information content of the noisy signals. When the choice set is coarser, aggregation can fail even with Bayesian agents, as shown by Banerjee (1992) or Bhikchandani et al. (1992).[14]

In Vives (1997), like in this paper, agents underweight their private seed: in his set up a stronger reliance on private signals *in the initial phase* would speed up learning and benefit all.[15] In our case, the weight cannot be altered over time: however a higher reliance on private seeds compared to equilibrium weights improves welfare because this limits both the correlation between information sources and cumulated errors.

Our paper is also related to and inspired by the recent upsurge of interest in the social learning with "almost" Bayesian agents. Sethi and Yildiz (2012, 2016, 2019) allow for heterogenous and *unobservable* priors about the state, and since players exchange beliefs (but not priors), there can be long-run disagreement. However the divergence cannot exceed the spread in initial biases because agents interpret correctly the reports of others based on the known distribution of priors. In contrast, Eyster Rabin (2010), Frick and al. (2020), Bohren and Hauser (2021) and Gentzkow et al. (2021), among others, introduce mis-specifications that lead agents to *incorrectly interpret* reports or actions of others. In Eyster Rabin, the errors are assumed to be significant enough to generate incorrect long-run beliefs for many signal realizations. By contrast, Frick and al. (2020) show that even small systematic mis-specifications can lead to interpretation errors that accumulate over time, though a restriction to a small number states and common priors can prevent this drift as shown in Bohren and Hauser (2021) (See Section 7.5 for an extended discussion of the connection between these two papers and ours). Finally, in Gentzkow et al. (2021), uncertain precision of signals and mis-specifications lead players to overestimate the precision of signals received by others who are similarly biased,

Other papers directly modify the updating rule itself. Jadbabie et al. (2012) introduce rules that combine Bayesian updating of own signals with a DG-like averaging over neighbors' beliefs, while Levy and Razin (2018) consider a rule which involves cumulating log likelihood-ratios, which they justify, like DG, on the ground that it mimics what a subjective Bayesian (with an erroneous model of the world) would do (see also Dasaratha et al. 2020). Finally

[14] Mossel et al. (2015) shows that this result also depends on the network structure and that for a large class of large networks, consensus and almost perfect learning is possible even with coarse communication.

[15] In the context of non-Bayesian learning, Mueller-Frank and Neri (2021) argue in related terms in favor of non-stationary rules that aggregate information in a sufficiently dense part of the network, before other agents get contaminated.

Molavi et al. (2018) provides axiomatic justification(s) (motivated by imperfect recall) for DG style linear aggregation (and averaging) of log belief-ratios.[16]

By contrast we take an evolutionary approach to rule selection, assuming selection within a *restricted* family of plausible stationary rules. There is of course a vast literature on the evolutionary selection of general behavioral rules, going back to Axelrod (1984). Fudenberg and Levine (1998) provide an excellent introduction to the selection of strategies in game theoretic settings. Our focus is on selecting rules for aggregating information in potentially large and complex network settings.

2 Basic Model

2.1 Transmission on the network

We consider a finite network with n agents, assume noisy transmission/reception of information and define a simple class of rules that players may use to update their opinions.

Formally, each agent i in the network has an *initial opinion* x_i and, at date t, *an opinion* y_i^t that can both be represented as real numbers.[17] Taking as given the matrix A characterizing the weights A_{ij} that i puts on j's opinion, we consider the class of updating rules (FJ) parameterized by the weights m_i and γ_i and specified in the introduction.

When $m_i = 0$, the rule corresponds to the well-studied DeGroot rule (DG). When $m_i > 0$, then in each period the rule mixes decision-maker's own initial opinion x_i with DG. This perpetual use of the initial opinion in the updating process gives FJ a non-Bayesian flavor, since for a Bayesian, their prior (i.e., the seed) is already integrated into y_i^{t-1} and therefore there is no reason to go back to it.[18]

To avoid technical difficulties once we give agents discretion in choosing their updating rule, we set $\underline{\gamma} > 0$ arbitrarily small and restrict attention to FJ rules where $\gamma_i \geq \underline{\gamma}$. We also assume that the matrix A is *connected* in the sense that for some positive integer k, the k^{th} power of A only has strictly positive elements, i.e., $A^k_{ij} > 0$ for all i, j. In other words everyone is within a finite number of steps of the rest.

Note that all the rules considered here are stationary, in the sense that the weighting parameters m_i and γ_i do not vary over time.[19] We see these as plausible ways in which boundedly rational agents might incorporate others' opinions into their current opinion. We recognize that with enough knowledge of the structure of the network and the process by which new information gets incorporated, adjusting the weights over time may make sense and return to this possibility in Section 7.6.

We also impose the assumption that everyone operates on the same time schedule: periods are defined so that everyone changes their opinion once every period and everyone else get to observe that change of opinion before they adjust their opinion in the following period. We will discuss what happens if we relax this assumption in Section 7.2.

[16] Attempts to provide axiomatic foundations of the DG rule in the statistics literature go back to Genest and Zidek (1986).

[17] This opinion can be interpreted as a point-belief about some underlying state, which will eventually be used to undertake an action.

[18] In fact, as mentioned already, the one obvious attraction of *DG* is its quasi-Bayesian flavor. Note that although formally the expression (FJ) encompasses the DG rule, we shall refer to FJ as a rule for which $m_i > 0$.

[19] In this sense even DG is only quasi-Bayesian, since for Bayesian the weight on new reports goes down over time.

2.2 Errors in opinion sharing

The term ε_i^t is an important ingredient of our model, meant to capture some imperfection in transmission.[20] It represents a distortion in what each individual "hears" that aggregates all the different sources of errors. Until Section 6, we assume that the error term is persistent, realized at the start of the process and applying for the duration of the updating process.[21] We shall denote by ξ_i this persistent error, so

$$\varepsilon_i^t \equiv \xi_i$$

In Section 6, we extend the model and incorporate idiosyncratic errors:

$$\varepsilon_i^t = \xi_i + \nu_i^t,$$

where ν_i^t are i.i.d. across time and agents.

We interpret ξ_i as a systematic bias that slants how opinions of others are *processed* by i. Biases ξ_i may be drawn independently across players, but we shall also discuss cases where they are positively correlated, such as when a group of friends share a political bias. Also note that although errors are indexed by i, our formulation can accommodate biases that result from both "hearing" errors and "sending" errors.[22]

For convenience, we assume that all error terms are unbiased (that is $E\xi_i = 0$ and $E\nu_i^t = 0$)[23] and homogenous across players, so we let

$$\varpi = \varpi_i = var(\xi_i)$$

2.3 The objective function

There is an underlying state θ, and agents want their decision to be as close as possible to that underlying state, where the decision is normalized to be the same as the agent's long-run opinion. In other words, we visualize a process where agents exchange opinions a large number of times before the decision needs to be taken.

Given this private objective, we explore each agent's incentives to choose his updating rule within the class of FJ rules to maximize the above objective on average across many different realizations of the underlying state of the world, the initial opinions and the transmission errors. We have in mind the idea that individuals choose a single rule to apply to many different problems. This is why we focus on their ex ante performance.[24] The set of possible updating rules is extraordinary vast, so the limitation to FJ rules is of course a restriction. Our motivation is to examine the incentives of *mildly* sophisticated agents who have some limited discretion over how they update opinions.

[20] There has been several recent attempts to introduce noisy or biased transmission in networks. In Jackson et al. (2019), information is coarse (0 or 1), and noise can either induce a mutation of the signal (from 0 to 1 or 1 to 0) or a break in the chain of transmission (information does not get communicated to the network neighbor).

[21] One interpretation is that each information aggregation problem is characterized by the realization of an initial opinion vector x and persistent bias vector ξ, and that agents face a distribution over problems.

[22] For example, if there were both "hearing" errors labelled ξ_i^h and "sending" errors labelled ξ_i^s, one could define $\xi_i = \xi_i^h + \sum_j A_{ij}\xi_j^s$ as the resulting processing error. Sending errors naturally generate correlations across the ξ_i's, and a profile of errors that depend on the network structure A. We shall discuss this in Section 7.5.

[23] We shall come back to the case where $E\xi_i \neq 0$ in Section 7.

[24] That is, on average over states, initial opinions and transmission errors.

Formally, we assume that the initial signals are given by

$$x_i = \theta + \delta_i$$

where the θ are drawn from some distribution $G(\theta)$ with mean zero and finite variance, δ_i, ξ_i and ν_{it} are random variables that are independent of each other for all i and t and are also independent of θ. We assume that noise terms δ_i are unbiased, with variance $\sigma_i^2 > 0$. For convenience, except where we need to assume otherwise to make a specific point, we set $\sigma_i = 1$ for all i, but we do not actually need this assumption.

For any t, each profile of updating rules (m, γ) generates at any date t, a distribution over date t opinions. We now define the expected loss (where the expectation is taken across realizations of θ, δ_i, and ε_i^t for all i and t):

$$L_i^t = E(y_i^t - \theta)^2$$

We then define the limit loss $L_i = \lim_{t \nearrow \infty} L_i^t$.[25]

2.4 Methodological assumptions

The loss L_i depends on the profile of updating rules (m, γ), and our main methodological assumptions are that (i) there is a force towards the use of higher performing rules (e.g., justified by evolution or reinforcement learning), and (ii) in this quest for higher performing rules, each individual considers (and gets feedback about) only a limited set of rules (i.e., the FJ class).

Formally, our analysis boils down to examining a rule-choice game where, given the rules adopted by others, each agent aims at minimizing L_i (using the instruments m_i and γ_i available to her): the object of interest is the Nash equilibrium of this rule choice game. Since L_i is an expectation across various realizations of initial signals and noise in transmission, we think of the person choosing one rule, parameterized by m_i and γ_i, to apply in many different life situations. These parameters are meant to capture some general features of opinion formation: specifically the *persistence* of initial opinions, and *speed of adjustment* of the current opinion.[26]

It is precisely this fact that rules apply across many different problems, and that a limited set of rules are considered, that makes our third route cognitively less demanding than the Bayesian route. While we agree that choosing m_i and γ_i optimally is a difficult problem which in principle requires knowledge of the structure of the model, there is no reason why the standard justification of Nash Equilibrium as a resting point of an (un-modeled) learning/evolutionary process would not apply here. Moreover, one of our most important results is that DG rules, and indeed all rules that put too little weight (m_i) on initial opinions, are dominated when there is noise in transmission, suggesting a strong force away from DG even if agents find it difficult to find the exact optimal value of m_i.

In the next Section we start by exploring the long-run properties of different learning rules within the DG and FJ class, with and without errors. Then we turn to the optimal choice of learning rules.

[25]Alternatively, one could define $L_i = \lim_{h \searrow 0}(1-h)\sum h^{t-1} L_i^t$, assuming that the agent makes a decision at a random date far away in the future and that his preference over decisions is $u_i(a_i, \theta) = -(a_i - \theta)^2$.

L_i is well-defined for any vector m, γ so long as $m \neq 0$. As it will turn out, for $m = 0$, L_i is infinite. Note that each player can secure $L_i \leq var(\delta_i) = \sigma_i^2 = 1$ by ignoring everyone else's opinions ($m_i = 1$).

[26]Our view is that these features probably do adjust to the broad economic environment agents face, but for each opinion-formation problem within a certain context, the actual sequence of opinions is mechanically generated given these features.

3 Some properties of the long-run opinions

We start by studying the properties of long-run opinions under DG and FJ with and without errors. In particular, we shall show that in the presence of errors there is convergence under FJ as long as at least one person i_0 has $m_{i_0} > 0$, but not under DG. We then explore what determines the variance of the limit opinion in the case where such a limit opinion exists. In particular what part of it comes from the "signal"– the original seeds – and what part from the noise that gets added along the way?

3.1 DG without errors.

It is well-known that in the DG case without errors ($m_i = 0$ for all i) learning converges to consensus and steady state values of y_i for all i. Define Γ as the diagonal matrix such that $\Gamma_{ii} = \gamma_i$. In matrix form, the dynamic of the vector of opinions $y^t = (y_i^t)_i$ under DG without noise can be expressed as

$$y^t = B_0 y^{t-1} \text{ where } B_0 = I - \Gamma + \Gamma A, \tag{2}$$

implying that

$$y^t = (B_0)^t x \tag{3}$$

where x is the vector of initial opinions. Let Δ_n be the set of vectors of non-negative weights $p = \{p_i\}_i$ with $\sum p_i = 1$. Because the network is connected, A is a irreducible stochastic matrix,[27] so there is a (unique) strictly positive vector of weights $\rho \in \Delta_n$ such that $\rho A = \rho$. When $\gamma_i > 0$ for all i, B_0 is also an irreducible stochastic matrix, so there is a unique vector $\pi \in \Delta_n$ such that $\pi B_0 = \pi$, and we must have[28]

$$\frac{\pi_i}{\pi_j} \equiv \frac{\rho_i}{\rho_j}\frac{\gamma_j}{\gamma_i} \tag{4}$$

When t gets large, all rows of $(B_0)^t$ converge to π, so all opinions y_i^t converge to the same limit opinion $\pi.x$, i.e.,

$$y_i = \pi.x \text{ for all } i. \tag{5}$$

So although the direct contribution of i's initial signal to i's opinion vanishes, it surfaces back from the influence of neighbors' opinions (which increasingly incorporate i's initial signal), settling at a limit weight equal to π_i.

Using (4), one may rewrite (5) to highlight how the speed of adjustment γ_i affects player i's influence on long-run opinions. We have:

$$y_i = \pi_i x_i + (1 - \pi_i) q^i.x_{-i} \text{ where } \frac{\pi_i}{1 - \pi_i} = \frac{1}{\gamma_i}\frac{\rho_i}{\sum_{j \neq i} \rho_j/\gamma_j} \tag{6}$$

and where q^i is a probability vector in Δ_{n-1} that does not depend γ_i. In other words, the network structure determines ρ. Given ρ, player i can use γ_i to control her influence on the long run opinion, π_i, but she cannot control the weights q^i.

[27] This is because A^k only has strictly positive elements for some large k.

[28] This is because $\pi^0 \equiv \rho\Gamma^{-1}$ solves $\pi^0 B_0 = \pi^0 - \rho + \rho = \pi^0$. Thus, since π is unique, π must be proportional to π^0.

3.2 DG with errors: exploding dynamics.

We show below that if all agents follow a DG rule, then for almost all realization of ξ, the long-run opinions diverge.

Proposition 1. *Assume that* $m_i = 0$ *for all* i*. Then* for almost all *realizations of* ξ, $\lim \left|y_i^t\right| = \infty$ *for all* i and x.

This proposition shows, for one, that an error ξ_1 in a single agent's perception is enough to drive everyone's opinions arbitrarily far from the truth: if $\xi_1 > 0$, say, the error creates a discrepancy between 1's opinion and that of the others, but every time the others' opinions catch up with him, agent 1 further raises his opinion compared to others, prompting another round of catching up, and eventually all opinions blow up.

Proof: With errors, Equations 2 and 3 become $y^t = B_0 y^{t-1} + \Gamma\xi$ and

$$y^t = (B_0)^t x + \sum_{0 \le k < t} (B_0)^k \Gamma \xi$$

For k large enough, each row of $(B_0)^k$ is close to π, so y_i^t diverges for all i whenever $\pi\Gamma\xi \neq 0$.

3.3 Anchored dynamics under FJ.

Fixing again x and ξ, we now examine long-run dynamics under FJ.

Proposition 2. *Assume at least one player, say* i_0*, updates according to FJ (with* $m_{i_0} > 0$*). Then, for any fixed* x *and* ξ*,* y^t *converges, and the limit vector of opinions* y *does not depend on* γ *nor on the signal* x_i *of any individual with* $m_i = 0$**.**

Proposition 2 shows that to prevent all the opinions from drifting away, it is enough that there is one player who continues to put at least a minimum amount of weight on his own initial opinion in forming his opinion in every period. Proposition 2 also shows that when $m_i = 0$, the signal initially received by i has no influence on the players' long-run opinions. A detailed proof is in the Appendix.

The general argument for convergence runs as follows.[29] For any fixed x, ξ, y^t evolves according to

$$y^t = \Gamma X + B y^{t-1} \text{ with } B = I - \Gamma + \Gamma(I - M)A$$

where $X_i = m_i x_i + (1 - m_i)\xi_i$, Γ and M are diagonal matrices with $\Gamma_{ii} = \gamma_i$ and $M_{ii} = m_i$. When $m_{i_0} > 0$ for some i_0, proving convergence is standard.[30] The limit opinion y then solves

$$y_i = (1 - \gamma_i)y_i + \gamma_i(X_i + (1 - m_i)A_i y) \text{ for all } i$$

which implies that it is also the solution of

$$y = X + (I - M)Ay. \tag{7}$$

This limit is thus independent of Γ. Furthermore, letting $C = (I - M)A$, we obtain

$$y = DX \text{ where } D = \sum_{k \ge 0} C^k = (I - C)^{-1} \tag{8}$$

[29] The argument follows Friedkin and Jensen (1999).

[30] The key to convergence is whether $\sum_j B_{ij} < 1$ for all i (or whether this *contraction* property holds for some power of B). This property holds when A is connected and $m_{i_0} > 0$ for some i_0.

explaining why long-run opinions only involves the seeds x_i of players for whom $m_i > 0$, as for others, $X_i = \xi_i$.

3.4 The dominance of noise under low m.

Although convergence is guaranteed when at least one player does not use DG, there is no discontinuity at the limit where *all* m_i get small: long-run opinions then become highly sensitive to the persistent error ξ. We have:

Proposition 3: Let $\overline{m} = \max m_i$. Then $L_i \geq \frac{\varpi}{n}\frac{(1-\overline{m})^2}{\overline{m}^2}$.

The detailed proof is in the Appendix. The lower bound on L_i is obtained by showing that for given x, ξ, long-run expected opinions are a weighted average of *modified initial opinions*, defined, whenever $m_i > 0$, as

$$\widetilde{x}_i = x_i + (1 - m_i)\xi_i/m_i.$$

To fix ideas, assume $m_i > 0$ for all i.[31] Then one can write (using the previous notation) $X = M\widetilde{x}$ and obtain, using (7) and (8)

$$y = M\widetilde{x} + (I - M)Ay = P\widetilde{x} \tag{9}$$

where $P \equiv DM$ is a probability matrix.[32] Intuitively, $\widetilde{x}_i$ can be thought of as the effective seed for individual i, and long-run opinions are averages over effective seeds. Since the variance of each $\widetilde{x}_i$ is bounded below by $\frac{\varpi(1-\overline{m})^2}{\overline{m}^2}$, we obtain the desired lower bound.

The two-player case. The two-player case provides a useful illustration. With two players, assuming m_1 and m_2 strictly positive, long-run opinions solve

$$y_i = m_i\widetilde{x}_i + (1 - m_i)y_j = m_i\widetilde{x}_i + (1 - m_i)(m_j\widetilde{x}_j + (1 - m_j)y_i$$

which further implies

$$y_i = p_i\widetilde{x}_i + (1 - p_i)\widetilde{x}_j \text{ where } p_i = \frac{m_i}{m_i + (1 - m_i)m_j} \tag{10}$$

confirming that long-run opinions are weighted average of modified opinions. Furthermore

$$y_i = p_i x_i + (1 - p_i)(x_j + \widehat{\xi}_i) \text{ where } \widehat{\xi}_i = \frac{\xi_i + \xi_j}{m_j} - \xi_j. \tag{11}$$

The term $\widehat{\xi}_i$ can be interpreted as the *cumulated error* that player i faces, resulting from each player repeatedly processing the other's opinion with an error, while p_i characterizes how player i's own seed *influences* her long-run opinion. Since $p_i + p_j = \frac{m_1+m_2}{m_1+m_2-m_1m_2} > 1$, it must be that players differ in the weight they each put in the long-run on their seeds, so there is disagreement, and the magnitude of the disagreements rises with m.

In networks, *echo effects* arise because players incorporate opinions that they contributed to shape, and these echoes shape both *long-run influence* and *cumulated errors*: when m_i and m_i/m_j are both small, the influence of player i is large because although i puts a large weight on y_j, y_j has been mostly shaped by x_i; and when m_j is small, a single loop of communication generates a combined error of $\xi_i + \xi_j$, which is (partially – but almost entirely when m_j is small) added to all opinions and thus cumulates over time.

---

[31] The argument generalizes to the case where a subset $N^0 \subsetneq N$ of agents follows DG ($m_i = 0$). (see Appendix).

[32] This means that each line of P is a probability vector. P is the limit of P^t defined recursively by $P^{t+1} = M + (I - M)AP^t$ and $P^1 = I$. By induction, each P^t (and P) is a probability matrix.

3.5 Influence under FJ rules and cumulated errors.

Under DG rules and no errors, a player can control her influence by modifying γ_i. Under FJ rules, the long-run opinions do not depend on γ_i–instead, as the previous two-player example illustrates, the limit opinions depend on the vector of weights m. Here we characterize both influence and cumulated errors for more general networks.

When at least one player i_0 sets $m_{i_0} > 0$, long-run opinions converge and we have

$$y_i = m_i x_i + (1 - m_i)\xi_i + (1 - m_i)\widehat{y}_i \text{ with } \widehat{y}_i \equiv \sum_{k \neq i} A_{ik} y_k \tag{12}$$

Player i's opinion thus builds on the opinion $\widehat{y}_i$ of a (fictitious) *composite neighbor* who aggregates the opinions y_k, to which the error ξ_i is added. Letting $\widetilde{A}^i_{kj} = \frac{A_{kj}}{1-A_{ki}}$, we rewrite (12) to describe how each opinion y_k builds on y_i:

$$y_k = m_k x_k + (1 - m_k)\xi_k + (1 - m_k)A_{ki} y_i + (1 - m_k)(1 - A_{ki}) \sum_{j \neq k,i} \widetilde{A}^i_{kj} y_j \tag{13}$$

So in effect, in incorporating the composite opinion $\widehat{y}_i$, player i is (partially) incorporating her own opinion y_i: the opinions that i gets from others are partially echoes of her own opinion. So even if her per-period reliance on x_i is small (i.e. m_i small), her seed x_i may eventually have a large influence on long-run opinions. Another aspect in that in incorporating the composite opinion $\widehat{y}_i$, each player i is (partially) adding other players' error terms to her own, and any opinion that contributes to $\widehat{y}_i$ is itself subject to errors. Proposition 4 below characterizes both effects: long-run influence and cumulated errors.

Let M^i (resp. α^i) be the diagonal $N-1$ matrix for which $M^i_{kk} = m_k$ for $k \neq i$ (resp. $\alpha^i_{kk} = A_{ki}$) and define the matrix $Q^i = (I - (I - M^i)(I - \alpha^i)\widetilde{A}^i)^{-1}$ and vector R^i such that $R^i_j = \sum_k A_{ik} Q^i_{kj}$. Also let $h_i \equiv 1/\sum_{j \neq i} R^i_j m_j$. We have:

Proposition 4: *Assume player* $i_0 \neq i$ has $m_{i_0} > 0$. *Then* $h_i \geq 1$ *and*

$$y_i = p_i x_i + (1 - p_i)(\widehat{x}_i + \widehat{\xi}_i) \textit{ where } \widehat{x}_i = q^i.x_{-i}, \tag{14}$$

$$\frac{p_i}{1-p_i} = \frac{m_i h_i}{(1-m_i)},\ q^i_j = \frac{R^i_j m_j}{\sum_{j \neq i} R^i_j m_j} \text{ and}$$

$$\widehat{\xi}_i = h_i(\xi_i + \sum_{j \neq i} R^i_j \xi_j (1 - m_j))$$

Proposition 4 provides an analog of Expression 6 when at least one player uses an FJ rule. Without errors, player i's long-run opinion is an average between her own seed x_i and a *composite seed* $\widehat{x}_i$ (an average over the others' seeds). The weight p_i defines how player i's own seed influences her long-run opinion, and through the choice of m_i player i has full control over this weight. Player i however has no control over the composite seed $\widehat{x}_i$, as the vector of weights $q^i \in \Delta_{n-1}$ is fully determined by A and m_{-i}.

In the presence of errors, the weights p_i and q^i remain the same. The difference is that when attempting to incorporate the composite seeds, player i faces a cumulated error term $\widehat{\xi}_i$. This error term can be very large when all m_j are small.

Proposition 4 also confirms an insight suggested by Proposition 2: the seed x_j of any individual that sets $m_j = 0$ has no influence on long-run opinion (either own or others). Finally, to complete all cases, we have:

Proposition 5: *If $m_{-i} = 0$ and $m_i > 0$, then $y_i = x_i + \frac{1-m_i}{m_i}(\xi_i + \sum_{j \neq i} R^i_j \xi_j)$ where R is as defined in Proposition 4.*

Consistent with Proposition 3, echo effects rise without bound when m_i gets small. Proposition 4 and 5 imply that if all players but i use DG, all players opinion's will build on x_i only, however small m_i is. Mueller-Franck (2017) makes a similar observation in a model without errors (concluding to the lack of robustness of DG outcomes to small departures from DG).

Section 4 will build upon Propositions 4 and 5 to characterize the equilibrium of the rule choice game. We conclude this Section with further comments on DG and FJ rules.

3.6 Understanding the difference between DG and FJ

(a) **On anchoring, influence and consensus:** DG and FJ generate a very different dynamic of opinions. Permanently putting weight on one's initial opinion is equivalent to putting a weight on the opinion of an individual that never changes opinion: it anchors one's opinion, preventing too much drift. As a result, it also anchors the opinions of one's neighbors, hence, the opinions of everyone in the (connected) network.

The channel through which each player influences long-run opinions also differs substantially. In the absence of noise, and for a given network structure, relative influence in DG depends on relative speed of adjustment γ, with lower speed increasing influence (see (4)).

In contrast, under FJ, the speeds of adjustment γ have no effect on long-run opinions y. Only the m_i's (and the structure of the network) matter. These m_i's determine *player-specific* vectors of weights, but at the limit where all m_i's are very small, these vectors converge to one another (see Appendix), with the weight p_i on i's seed proportional to $m_i \rho_i$, that is:

$$\frac{p_i}{p_k} = \frac{m_i \rho_i}{m_k \rho_j} \tag{15}$$

This is an analog to (4) showing that close to the limit, m_i plays the same role as $1/\gamma_i$ does in DG and consensus obtains. As the m_i's go up however, consensus disappears: players "agree to disagree".

(b) **On the fragility of DG:**

There is something inherently fragile about the long-run evolution of opinions under DG. Since individuals don't put any weight on their own initial signal after the first period, the direct route for that signal to stay relevant is through the weight put on their own previous period's opinion. This source clearly has dwindling importance over time. This gets compensated by the growing weight on the indirect route–each individual i adjusts his or her opinion based on the opinions of their neighbors, and these are in turn influenced by i's past opinions and through those, by i's initial signal. In DG without transmission errors, the second force at least partly offsets the first one – but this is no longer true when there is any transmission error because of the cumulative effect of noise that comes with the feedback from others.

(c) **On the source of change in opinion**: One way to assess the difference between DG and FJ is to express them in terms of changes of opinions and opinion spreads. Defining the change of opinion $Y_i^t = y_i^t - y_i^{t-1}$, the neighbors' average opinion $\widehat{y}_i^t$ and the spread $D_i^t = \widehat{y}_i^t - y_i^t$ between others' and own opinions, and setting $\gamma_i = 1$ for all i for the FJ process, we have the following expressions:

$$Y_i^t = \gamma_i(D_i^{t-1} + \xi_i) \tag{DG}$$
$$Y_i^t = (1 - m_i)A_i Y^{t-1} \tag{FJ}$$

Under DG, one changes one's opinion whenever there is a (perceived) difference between that opinion and the opinions of one's neighbors: any difference generates an adjustment aimed at reducing it. In the absence of errors, this creates a force towards consensus, with D_i^t and Y_i^t eventually converging to 0. With errors however, this adjustment aimed at reducing the (perceived) spread actually keeps opinions moving:[33] errors are eventually incorporated into the opinions of all the players, and repeated errors tend to cumulate and generate a general drift in opinions. The force towards consensus is in this sense too strong.

By contrast, under FJ, players only incorporate *changes* in the opinions of others**.** So, in the case where the transmission error is fixed, ξ_1 will generate a *one time* change on 1's opinion, but it won't, by itself, generate any further changes for player 1. Of course, this initial (unwanted) change of opinion will trigger a sequence of further changes – it will be partially incorporated in player 2's opinion, and therefore come back to player 1 again. This is what we call an *echo effect*. But, when $m_i > 0$ for at least one player, the echo effect will be smaller than the initial impact and will get even smaller over time, and as result, opinions won't blow up: all Y_i^t's eventually converge to 0. Nevertheless, if all m_i are small, the echo effects are not dampened enough, and the consequence is a high sensitivity of the final opinion to the errors.

4 Choosing the rule

4.1 When there are no errors

We build upon Proposition 4 and 5 to characterize the equilibrium of the rule-choice game, starting with the case of no error. We show that the equilibrium must be DG and that in equilibrium, information aggregation must be perfect. Formally, define π^* as the vector of weights on seeds that achieve perfect information aggregation, i.e., $\pi^* = \arg\min_\pi var(\sum_k \pi_k x_k)$, and let $v^* = var(\pi^*.x)$. We have

Result 1: *In the absence of transmission errors, the equilibrium must be DG. In addition, in equilibrium,* $y_i = \pi^*.x$ *and* $L_i = v^*$.

In other words, as long as there is no noise, we get perfect agreement in opinions in equilibrium and perfect information aggregation. As mentioned in the introduction, the main difference with De Marzo et al. (2003) and Golub and Jackson (2010) is that we allow for endogenous weights γ_i. For any connected network, this is enough to obtain efficiency in equilibrium.

Intuitively, both y_i and the neighbor's composite limit opinion $\widehat{y}_i$ are weighted averages between x_i and the composite seed $\widehat{x}_i$, with different weights when players do not use DG rules. In equilibrium, i chooses optimally the weighting to reduce variance, so if the equilibrium is not DG, the variance $v(y_i)$ must be *strictly* smaller than the variance $v(\widehat{y}_i)$, which itself is no larger than the maximum variance $\max_k v(y_k)$. Since this cannot be true for all i, the equilibrium must be DG.

[33] Technically, opinions can never settle because this would require finding a vector y for which $D + \xi = 0$, hence $Ay - y + \xi = 0$ which is not possible unless $\rho.\xi = 0$.

Regarding efficiency, in a DG equilibrium, player i chooses the relative weight π_i on her own seed by modifying γ_i, and any departure from perfect information aggregation leads i to choose a relative weight π_i no smaller than π_i^*. In a DG equilibrium, π_i also characterizes the influence of x_i on the common long-run opinion (there is consensus), so the weights π_i must add up to 1. This can only happen if they coincide with the efficient weights π_i^*. Therefore there is a unique (and efficient) equilibrium outcome.

4.2 Rule choice when there is noise

We already saw that as soon as there is some noise, the outcome generated by any DG rule drifts very far from minimizing L_i. The loss grows without bound. Indeed from the point of view of the individual decision maker it would be better to ignore everyone else than to follow DG. In fact all strategies that put too little weight on their own seed (recall DG puts zero weight) are dominated from the point of view of the individual decision-maker, as well as being socially suboptimal.

Result 2: *Let $\underline{m} = \varpi/(1+\varpi)$. Any (m_i, γ_i) with $m_i < \underline{m}$ is dominated by $(\underline{m}, \gamma_i)$, from the individual and social point of view.*

Regarding the choice of the individually optimal rule, Result 2 builds on two ideas. First, if all other players use DG, then for agent i, any $m_i > 0$ is preferable to DG because everyone's opinion drifts off indefinitely if $m_i = 0$, as we saw above. Second, if some players use FJ (with $m_j > 0$), then initial opinions of these players x_j (plus any persistent noise in their reception of the signal) totally determines the long run outcome and the seeds of all the players that use *DG* do not get any weight – they end up as pure followers. This is not desirable for these DG players (and for the others) for the same reason why, in the absence of noise, each one wishes to let their own seed influence their long-run opinion. Hence the lower bound on m_i.

To see why this is also true of the socially optimal rule, i.e. the rule that minimizes $\sum_i L_i$, we observe that when $m_i = 0$, the only effect of information transmission by i to his neighbors is to introduce i's perception errors into the network. When i raises m_i above 0, he raises the quality of the information he transmits, while reducing the damaging echo effect that low m_i generates.

We now provide further characterization of the privately optimal choice of m_i, and its consequence for the loss L_i. Recall from Proposition 4 that $y_i = p_i x_i + (1-p_i)(\widehat{x}_i + \widehat{\xi}_i)$ where $\widehat{x}_i + \widehat{\xi}_i$ is a term that only *depends on the structure of the network and m_{-i}*, and which has variance

$$W_i \equiv var(\widehat{x}_i) + \widehat{\varpi}_i \text{ where } \widehat{\varpi}_i = E\widehat{\xi}_i^2 \tag{16}$$

Since player i fully controls p_i by adjusting m_i, we obtain:

Proposition 6: *For a given m_{-i}, the optimal choice of m_i satisfies $\frac{p_i}{1-p_i} = \frac{h_i m_i}{1-m_i} = \frac{W_i}{\sigma_i^2}$, resulting in a loss*

$$L_i = \sigma_i^2 p_i = \frac{W_i}{1 + W_i/\sigma_i^2}. \tag{17}$$

This Proposition implies that the best response is a continuous function (which we know maps into a compact set $[\underline{m}, 1]$), so existence of an equilibrium is guaranteed. It also implies that the loss L_i is fully determined by W_i.

In the absence of errors, efficient aggregation would obtain, resulting in the minimum feasible loss v^*, which satisfies

$$v^* = \sigma_i^2 \pi_i^* = \frac{W_i^*}{1 + W_i^*/\sigma_i^2} \tag{18}$$

where $W_i^* = \min_q var(q.x_{-i})$.[34] With transmission errors, W_i must be higher than W_i^*. Expression (16) highlights the two possible additional sources of losses that player i now faces: (i) the fact that seeds of others may not be efficiently aggregated (i.e. $var(\widehat{x}_i) > W_i^*$) and (ii) the presence of the cumulated error term $\widehat{\xi}_i$.

In next Section, we will see that when errors are small, the cumulated errors are the preponderant source of inefficiency. In Section 5, we will examine how the equilibrium W_i^* varies with network structure.

4.3 How big is the divergence in opinions?

Result 2 has the obvious implication that full consensus is never going to be an equilibrium when there are persistent errors–there are in fact two sources of deviation, the error itself (which mechanically prevents consensus) and the extra weight m_i on one's initial signal (which fuels further divergence.)

Result 3 below shows that because of cumulated errors, the optimal weight put on one's own seed tends to be relatively large, i.e. $O(\varpi^{1/3})$ (of the order of $\varpi^{1/3}$).[35] As a result when ϖ is small, the extra weight on one's own seeds becomes the preponderant source of dispersion. These extra weights also determine the equilibrium magnitude of $\widehat{\varpi}_i$ and L_i. We have:

Result 3: *For any given finite network and any $\varpi > 0$ small, all m_i, $p_i - \pi_i^*$, $\widehat{\varpi}_i$ and $L_i - v_i^*$ are positive and $O(\varpi^{1/3})$ in equilibrium.*

Note that in addition to cumulated errors, there is another source of inefficiency in equilibrium, the fact that seeds are not efficiently weighted. But that inefficiency is $O(\varpi^{2/3})$:[36] a socially optimal choice of weights m_i would trade-off more inefficient weighting (larger m) against decreasing the variance of cumulated errors.

The intuition for Result 3 runs as follows. The error terms $\widehat{\varpi}$ are $O(\varpi/m^2)$. These errors terms degrade the quality of information that each i gets (raising W_i above W_i^*), which in turn implies a weighting p_i of i's seed larger than the efficient weighing π_i^*, with $p_i - \pi_i^*$ at least $O(\varpi/m^2)$ (by (17) and (18)). When $m > 0$, players end up weighing seeds differently, but when all m are small, the spread between the weights is also small and $O(m)$. So if p_k is the weight that k puts on x_k, the weight that i puts on x_k must be $p_k + O(m)$. Since the weights that i puts on all seeds must add to 1, the p_k's must add up to at most $1 + O(m)$. And since the sum $\sum_k (p_k - \pi_k^*)$ is at least $O(\varpi/m^2)$, m must be at least $O(\varpi/m^2)$ in equilibrium, which gives m at least $O(\varpi^{1/3})$.[37]

[34] This is because $v^* = \min_\pi \pi.x = \min_{\pi_i} var(\pi_i x_i + (1-\pi_i)W_i^*)$.

[35] When we say that $m = O(g(\varpi))$, we mean that $m/g(\varpi)$ has a finite limit when ϖ tends to 0.

[36] This is because for an inefficient weighting of seeds $q \neq \pi^*$, the loss is second order in the differences $q_i - \pi_i^*$: $L_i - v^* = \sum (q_i^2 - \pi_i^{*2})\sigma_i^2 = \sum (q_i - \pi_i^*)^2 \sigma_i^2 + 2\sum (q_i - \pi_i^*)\pi_i^* \sigma_i^2$, and the last term is 0 because $\sum (q_i - \pi_i^*) = 1$ and at the optimum $\pi_i^* \sigma_i^2 = \pi_{j.}^* \sigma_j^2$ for all i, j

[37] The proof also shows m_i cannot increase beyond $O(\varpi^{1/3})$ in equilibrium for the same reason that the equilibrium without error terms must be DG: each player sets the weighting p_i of own seed x_i optimally, and this creates a force towards optimal information aggregation.

Note that Result 3 focuses on the case where variances are small. When the m_i's rise, the relative weights on seeds eventually diverge sufficiently from efficient weighting that this fuels a further rise in W_i hence in m_i.

4.4 Privately versus socially optimal choices

We already showed that both private and social optima must deviate from DG when there is noise. The next result shows that there is a sense in which, in the presence of noise, the Nash Equilibrium is closer to DG than is desirable from the point of view of social welfare maximization.[38]

Result 4. *At any Nash equilibrium, a marginal increase of m_i by any player i would increase aggregate social welfare.*

To see why this result holds, assume $m_j \in (0, 1)$ and observe that player j's opinion can be expressed as an average between the (modified) seeds $\widetilde{x}_{-i}$ of players other than i and player i's opinion

$$y_j = (1 - \mu_{ji})C^{ji}\widetilde{x}_{-i} + \mu_{ji}y_i \tag{19}$$

where C^{ji} is a probability vector and $\mu_{ji} \in (0, 1)$,[39] with μ_{ji} and C^{ji} both independent of m_i.

The expression above highlights that when player i chooses m_i optimally (for him) to minimize the variance of y_i, there is no reason why he would be also minimizing the variance of y_j. Specifically we use use (19) to separate the loss L_j into three terms:

$$L_j = (1 - \mu_{ji})^2 var(C^{ji}\widetilde{x}_{-i}) + \mu_{ji}L_i + 2(1 - \mu_{ji})\mu_{ji}Cov(C^{ji}\widetilde{x}_{-i}, y_i). \tag{20}$$

When m_i is raised above $i's$ private optimum, there is no effect on the first term. There is a second-order effect on the second term (because we start at $i's$ private optimum). The last term is what creates a discrepancy between private and social incentives.

This last term depends on the covariance between seeds other than that of i ($\widetilde{x}_{-i}$) and the opinion of i (y_i). When m_i increases, the influence of each $k \neq i$ on i's opinion is reduced, and the correlation between y_i and x_k (and even more so with $\widetilde{x}_k$) is also reduced. Therefore, starting at a Nash equilibrium, L_j goes down when m_i is raised.

5 Equilibrium, efficiency and polarization in simple networks.

In this section we explore the quantitative significance of our results through a set of simple examples where, unless specified otherwise. We assume that seeds are equally informative ($\sigma_i^2 = 1$ for all i) and each player treats all his neighbors symmetrically ($A_{ij} = 1/|N_i|$). We are particularly interested in whether a small amount of noise can lead to large distortions of information aggregation, how much of extra information loss comes from non-cooperative behavior, and the connection between disagreements between individuals and polarization at the population level.

We start with the simple two-player network and then examine two classes of symmetric networks, the complete network and the directed circle, followed by the star network class

[38] The result shows that a marginal increase over equilibrium weights enhances welfare, but we do not have a full characterization of socially efficient weights.

[39] This assumes $m_j \in (0, 1)$. C_j^{ji} is positive because j is using her own seed.

of asymmetric networks. Different networks lead to differences in how seeds are aggregated, differences in how errors cumulate, as well as differences in incentives to raise m, resulting in differences in the efficiency of information aggregation across networks, which we will try to highlight.

5.1 Two-player case.

Social optimum. Assuming independent errors, we obtain from (10)

$$L_1 = I(p_1) + (p_1)^2\mathcal{X}(m_1) + (1-p_1)^2\mathcal{X}(m_2)$$

where $I(p) = p^2 + (1-p)^2$ is the variance of long run opinion in the absence of transmission noise and $\mathcal{X}(m) = \varpi\frac{(1-m)^2}{m^2}$ represents the effect of cumulated noise. The total social loss is $L = L_1 + L_2$.

It is easy to check that, given the symmetry, minimizing the social loss requires setting identical values for m_1 and m_2. When both players use the same rule ($m = m_1 = m_2$), $p_i = \frac{1}{2-m}$ and the social loss is:

$$L = 2I(\frac{1}{2-m})(1+\mathcal{X}(m))$$

The expression highlights a trade-off between decreasing m for information aggregation purposes ($I(p)$ is minimized at $p = 1/2$), and increasing m to limit the effect of cumulated communication errors (when $\varpi > 0$ and m is small, communication errors are hugely amplified).

Welfare is maximized for an m^{**} that optimally trades off these two effects and the socially efficient weight m^{**} (which minimizes L) can be significantly different from 0 even when ϖ is small. Specifically, for $\varpi = 0.0001$, $m^{**} = 0.13$ and for $\varpi = 0.001$, $m^{**} = 0.21$. Furthermore, for ϖ small, $m^{**} \simeq (4\varpi)^{1/4}$.[40]

Nash Equilibrium. We now assume that individuals choose their rules non-cooperatively. Applying Proposition 5, we obtain $p_i/(1-p_i) = 1 + \widehat{\varpi}_i$, so

$$p_i = \frac{1+\widehat{\varpi}_i}{2+\widehat{\varpi}_i} \text{ where } \widehat{\varpi}_i = E\widehat{\xi}_i^2$$

which gives the best response for i, as a function of m_j:

$$m_i = \frac{m_j(1+\widehat{\varpi}_i)}{1+m_j(1+\widehat{\varpi}_i)}$$

Figure 1 plots the best responses for $\varpi = 0.01$.

In the absence of noise, $\widehat{\varpi}_i = 0$, and player 1 should set m_1 so that $p_1 = 1/2$ (for information aggregation purposes), which requires $m_1 < m_2$, which explains why there is no equilibrium with positive m (this is the force towards DG). With noise, the variance $\widehat{\varpi}_i$ explodes when m_j gets small, reflecting the cumulation of errors when m_j is low. This provides i with incentives to raise p_i (hence m_i) which in turn puts a lower bound equilibrium weights: in equilibrium, $m_1^* = m_2^* = m^*$ and m^* is a solution to

$$m^* = \frac{\widehat{\varpi}^*}{1+\widehat{\varpi}^*} \text{ with } \widehat{\varpi}^* = \varpi\frac{1+(1-m^*)^2}{m^{*2}}.$$

[40]This is because for ϖ small $L \simeq 1 + m^2/4 + \varpi/m^2$.

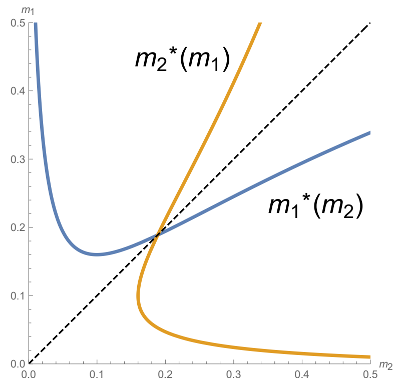

Figure 1: Best responses, $\varpi = 0.01$

When ϖ is small, we have $m^* \simeq (2\varpi)^{1/3}$. Since $m^{**} \simeq (4\varpi)^{1/4}$, the ratio of m^{**} to m^* become arbitrarily large when ϖ is small.

5.2 Larger networks

Before focusing on specific networks of size $n > 2$, we make preliminary observations that explain how one derives an equilibrium. Recall from Proposition 5 that player i's incentives yields

$$\frac{m_i}{1 - m_i} = W_i / h_i \tag{21}$$

where $W_i = var(\widehat{x}_i) + \widehat{\varpi}_i$. Both h_i and W_i depend only m_{-i} and the structure of the network, and the equilibrium values m_i^* are obtained by simultaneously solving these equations.

The terms h_i and W_i differ across networks and affect relative performance. For example, for a given W_i, the term h_i characterizes the degree to which a rise in m_i translates into a rise in i's influence on long-run opinions ($\frac{p_i}{1-p_i} = \frac{m_i h_i}{1-m_i}$). For a given target p_i, a higher h_i translates into a lower m_i, which raises the magnitude of echo effects $\widehat{\varpi}_j$ and hurts welfare.

To facilitate network comparisons, we continue to assume seeds of identical precision ($\sigma_i^2 = 1$), so that $\pi_i^* = 1/n$ and $W_i^* = 1/(n-1)$.[41] It will be useful to define $\widehat{\Delta}_i \equiv W_i - W_i^*$ (which characterizes how well information is aggregated by neighbors) and $\rho_i \equiv \frac{p_i}{1-p_i} - \frac{1}{n-1}$ (which characterizes the degree to which player i use its seed above the efficient information-aggregation level). The equilibrium condition

$$\rho_i = \widehat{\Delta}_i$$

determines m_i as a function of m_{-i}, while the socially efficient level m^{**} is determined by the minimization of $\Sigma_i L_i$ where $L_i = p_i^2 + (1 - p_i)^2 (W_i^* + \widehat{\Delta}_i)$.

In what follows, we derive explicit formula for h_i, $\widehat{x}_i$, $\widehat{\xi}_i$ for the networks we consider, a superscript c denoting the complete network, d the directed circle and s the star network. We also solve for equilibrium and the social optimum for each network, and discuss their relative performance, focussing on the limit cases where ϖ is small (for a fixed n) and where for a fixed ϖ small, n gets large.

[41] Given that all seeds have same precision, efficient information aggregation would require $p_i = 1/n$ hence $\rho_i = 0$.

5.2.1 The complete network.

In the complete network, all the players are connected to each other. We denote by $\overline{x}_{-i}$ (respectively $\overline{\xi}_{-i}$) the average seed (respectively error) of players other than i. We analyze a situation where all players but i use the same weight m and look for a symmetric equilibrium. We have:

Lemma 1: *For the complete network,* $h_i^c = 1 + \frac{1-m}{m(n-1)}$, $\widehat{x}_i^c = \overline{x}_{-i}$, *and* $\widehat{\xi}_i^c = h_i^c \xi_i + \overline{\xi}_{-i}\frac{1-m}{m}$.

This implies that $W_i^c = \frac{1}{n-1} + \widehat{\varpi}_i^c$, so for a given m, the losses L_i are solely generated by the cumulated error term $\widehat{\varpi}_i^c$. When everyone uses the same m, the loss is

$$L_i = p^2 + (1-p)^2(\frac{1}{n-1} + \widehat{\varpi}_i^c) \text{ where } \frac{p}{1-p} = \frac{1}{n-1} + \frac{m}{1-m}, \tag{22}$$

and it is minimized for some m_c^{**} which optimally solves a tradeoff between optimal information aggregation (achieved for $p = 1/n$) and the reduction of cumulated errors. In equilibrium, each i takes $\widehat{\varpi}_i^c$ as given while setting p_i (hence m_i) optimally, which explains why the equilibrium m_c^* is inefficiently low.

We examine below two interesting *limit cases*. Fixing n, we examine the limit case where ϖ gets small. Then fixing ϖ small, we examine the case where n gets large. We have:

Proposition 7: *Fix* n. *With independent errors,* $\widehat{\Delta}_c^* \simeq m_c^* \simeq \varpi^{1/3}(\frac{n}{(n-1)^2})^{1/3}$, *while with perfectly correlated errors* $\widehat{\Delta}_c^* \simeq m_c^* \simeq \varpi^{1/3}(\frac{n}{n-1})^{2/3}$. *Now fix* ϖ *small. For large* n *we have:* $\widehat{\Delta}_c^* \simeq m_c^* \simeq (\frac{\varpi}{n})^{1/3}$ *for independent errors, and* $\widehat{\Delta}_c^* \simeq m_c^* \simeq \varpi^{1/3}$ *for perfectly correlated errors.*

In the case where the errors are independent, the complete network thus has the desirable property of reducing the variance of the mean error that i incorporates in processing his neighbors' opinions (goes down to 0 when n is large). On the other hand, with correlated errors, the efficiency loss rises quickly with ϖ even at the limit where n is large. In fact the equilibrium and social losses under FJ are much larger, for one, than in the case where each player just aggregates the initial opinions of others with an error ξ, where the loss would be ϖ. For $\varpi = 0.01$, we find $m_c^* = 0.2$ and $m_c^{**} = 0.29$, with equilibrium and social losses respectively 0.2 and 0.14, as if each player had perfectly aggregated only 4 (respectively 6) additional signals.[42]

5.2.2 The directed circle

Consider a circle with n players where information transmission is directed and one-sided: player i communicates to player $i-1$, who communicates to $i-2$, and so on. Player 0 is player n. Long-run opinions satisfy

$$y_i = m_i \widetilde{x}_i + (1-m_i) y_{i+1}.$$

Assuming all other players choose m, and repeatedly substituting y_{i+k} for $k = 1, ..., n-1$, we obtain below explicit expressions for h_i^d, $\widehat{x}_i^d$ and the cumulated error $\widehat{\xi}_i^d$. Specifically, for any z_{-i}, let $\psi(z_{-i}) = \sum_{k=1}^{n-1}(1-m)^{k-1} z_{i+k} / \sum_{k=1}^{n-1}(1-m)^{k-1}$ be the weighted average over z_{i+k}'s where the weight of the $k-$step neighbor is diminished by a factor $(1-m)^{k-1}$. We have:

[42] To compute exactly the equilibrium and social losses at the large limit n, we use (22) to get $L_i = m_i^2 + (1-m_i)^2\varpi/m^2$

Lemma 2: *For the directed circle,* $h_i^d = \frac{1}{1-(1-m)^{n-1}} < h_i^c$ *and* $\widehat{x}_i^d = \psi(x_{-i})$ *and* $\widehat{\xi}_i^d = h_i^d \xi_i + \psi(\xi_{-i})\frac{1-m}{m}$.

There are two notable differences with the complete network: (i) the architecture of the network itself implies a stronger incentive to raise m_i (because $h_i^d < h_i^c$); (ii) A *less efficient* information aggregation by neighbors (because $\psi(x_{-i}) \neq \overline{x}_{-i}$), and a *less efficient* averaging of errors when errors not perfectly correlated (because then $\psi(\xi_{-i}) \neq \overline{\xi}_{-i}$).

Regarding relative performance of the network, item (i) works in favor of the directed circle because it raises the equilibrium weights m^* and therefore lowers cumulated errors in equilibrium. Item (ii) works against the directed circle. Our two limit cases illustrate that each effect may dominate the other:

Proposition 8: *Fix* $n > 2$. *Let* $\mu = \frac{n}{2(n-1)} < 1$. *Then for fixed* n and *small* ϖ, $\widehat{\Delta}_i^c \simeq \widehat{\Delta}_i^d$ *and in equilibrium,* $m_d^* \simeq \mu^{-1/3} m_c^*$ *and* $\widehat{\Delta}_d^* \simeq \mu^{2/3} \widehat{\Delta}_c^*$. *For a fixed small* ϖ, *taking the large* n *limit,* $\widehat{\Delta}_i^d > \widehat{\Delta}_i^c$ *and* $h_i^d \simeq h_i^c \simeq 1$, *and in equilibrium,* $m_d^* \simeq \widehat{\Delta}_d^* > m_c^* \simeq \widehat{\Delta}_c^*$, with $\widehat{\Delta}_d^* \simeq \varpi^{1/2}$ for independent errors, *and* $\widehat{\Delta}_d^* \simeq (2\varpi)^{1/3}$ *for correlated errors.*

Intuitively, for a fixed n and ϖ small enough, m is small and the inefficiencies (ii) are of order 2 in m (see Appendix)[43], explaining why $\widehat{\Delta}_i^c \simeq \widehat{\Delta}_i^d$. So the only first order effect comes from a stronger incentive to raise m, which reduces $\widehat{\Delta}_d^*$ and L_d^*. At the other limit, incentives to raise m given $\widehat{\Delta}$ are identical across networks (because $h_i^d \simeq h_i^c$), and the complete network dominates because of better aggregation properties.

5.2.3 The star network

We consider a network consisting of $n-1$ peripheral players labelled $k = 1, ...n-1$ and a central player, labelled 0, who aggregates the opinions of the peripheral players. All players have a seed of same precision and are subject to a processing error ξ_i, with same variance ϖ for peripheral players, and variance $\varpi_0 \leq \varpi$ for the central player.

We look for an equilibrium where all peripheral players use the same weight m_s^* and the central player uses m_0^*. We start with the central player:

Lemma 3: *In the star network, for the central player* $h_0^s = \frac{1}{m}$, $\widehat{x}_0^s = \overline{x}$ *and* $\widehat{\xi}_0^s = \frac{\xi_0 + (1-m)\overline{\xi}}{m}$.

This implies that $W_0^s = \frac{1}{n-1} + \widehat{\varpi}_0^s$, so for a given m, the loss L_0 is solely generated by the cumulated error term. The choice of m_0 is given by:

$$\frac{m_0}{(1-m_0)m} = \frac{1}{n-1} + \widehat{\varpi}_0^s \tag{23}$$

implying that for a fixed n and for small $\widehat{\varpi}_0^s$, $m_0 \simeq \frac{m}{n-1}$. The reason is that the central player's opinion influences (many) peripheral players, so for information aggregation purposes, the central player should compensate for that influence by setting a smaller m_0 compared to m. Furthermore, at the large n limit, $m_0 \simeq m\widehat{\varpi}_0^s$, so when $\widehat{\varpi}_0^s$ is small (which will be true in equilibrium when ϖ is small), his behavior becomes close to that of a DG player.

We now consider a peripheral player i using m_i while all other peripheral players use m, deriving h_i, z_i and $\widehat{\xi}_i$:

[43] This is because for small m, the average $\psi(x_{-i})$ has close to the efficient weights.

Lemma 4: *Let* $\rho_0 = \frac{m_0}{(1-m_0)m} - \frac{1}{n-1}$ and $q_0 = \frac{1/(n-1)+\rho_0}{1+\rho_0}$. *In the star network, for a peripheral player,* $\widehat{x}_i^s = q_0 x_0 + (1-q_0)\overline{x}_{-i}$, $h_i^s = 1 + \frac{1}{m(n-1)(1+\rho_0)}$ *and* $\widehat{\xi}_i^s = \xi_i(1 + \frac{1}{(n-1)(1+\rho_0)}) + \frac{1}{(1+\rho_0)}\widehat{\xi}_0$.

The Lemma tells us that when $\rho_0 > 0$, the aggregation of seeds is distorted (i.e, $v(\widehat{x}_i^s) > \frac{1}{n-1}$) and potentially, $h_i^s < h_i^c$. However in equilibrium, the incentive condition (23) of the central player implies $\rho_0 = \widehat{\varpi}_0^s$, so when $\widehat{\varpi}_0^s$ is small (which again will be true in equilibrium when ϖ is small), the distortion is negligible and $h_i^s \simeq h_i^c$. This implies that for small ϖ, losses essentially come from the cumulated error terms (i.e., $\widehat{\varpi}_i^s \equiv var(\widehat{\xi}_i^s)$) and that, for given cumulated errors, incentives to raise m_i for peripherical players are similar to that of the complete network (because $h_i^s \simeq h_i^c$). Computing $\widehat{\varpi}_i^s$ for small ϖ gives us:

Proposition 9: *Fix n. For small ϖ, with independent errors,* $m_s^* \simeq \widehat{\Delta}_s^* \simeq (\varpi_0 + \frac{\varpi}{n-1})^{1/3}$ *and with correlated errors,* $m_s^* = \widehat{\Delta}_s^* \simeq (var(\xi_0 + \xi))^{1/3}$. *If in addition* $\varpi_0 = \varpi$, then whether errors are independent or correlated, $\widehat{\Delta}_s^* > \widehat{\Delta}_c^* > \widehat{\Delta}_d$. *Furthermore, for fixed small ϖ, at the large n limit,* $\widehat{\Delta}_s^* > \widehat{\Delta}_d^* > \widehat{\Delta}_c^*$.

Details are in the Appendix. As explained above, for small ϖ, the main difference with the complete network comes from the magnitude of the cumulated error term $\widehat{\varpi}_i^s$, which is higher for the star network because the central player's errors contaminate all the other players in a systematic (i.e., correlated) way, unless ϖ_0 is significantly smaller than ϖ. Also observe that even when the star network gets large (which, unlike the directed network, leads to a close-to-efficient aggregation of seeds), the central player's errors lead to large cumulated errors that end up impairing performance. Unless ϖ_0 is significantly smaller than ϖ, the star network thus results in lower performance than the directed circle under both limit cases.

5.3 Implications for the divergence of opinions

In the absence of noise, and if players use DG with appropriate weights γ, long-run opinions converge to a consensus $y^* = \pi^*.x$ which efficiently aggregates seeds. In a large network, this opinion y^* will essentially coincide with the underlying state θ ($y^* \simeq \theta$), which implies that if we consider two such identical networks, there will be *consensus within* each network and *consensus across* networks.

In the presence of noise, two things may happen. A divergence of long-run opinions y away from y^*, which means a *divergence of average opinions between the networks*, as well as some *dispersion of opinions within networks.* This section argues that there is a connection between consensus within subgroups (low dispersion) and polarization (high divergence across subgroups).

To fix ideas, we consider below the case of two large disconnected star networks modeled as above.[44] This description generally fits the maps of social networks in the US population with the two stars representing Democrats and Republicans (Cox et al. 2020). We assume that in each star network all peripheral players use the same weight m and that central players behave as DG players, just aggregating peripheral players' opinions.[45] We are interested in the effect

[44]Result 5 below would also hold if the set of cross-star links were a vanishingly small proportion of the total number of links.

[45]In the Appendix, we consider the case where central players benevolently choose m_0 to minimize the losses of the peripheral players, given m.

on m on the distribution of opinions within the star and across stars. We have

$$y_i = mx_i + (1-m)(y_0 + \xi_i)$$

The dispersion of opinions between two peripheral players *within* a given star is

$$d \equiv E(y_i - y_j)^2 = 2(m)^2 + 2(1-m)^2\varpi$$

The *average* opinion of peripherical players is $\overline{y} = m\overline{x} + (1-m)(y_0 + \overline{\xi})$, and for a large network, with independent errors, only y_0 contributes to the variance of $\overline{y}$. *Across* the networks, average opinions are independent (conditional on θ) and the dispersion of opinion D between average opinions is thus:

$$D = 2v(\overline{y})$$

The following result establishes a relationship between d and D:

Result 5: *Fix ϖ_0 small and assume independent errors. At the social optimum m^{**}, $D \simeq d$ and for any $m \leq m^{**}$, $D \simeq \frac{4\varpi_0}{d}$.*

Proof: When the central player is DG, $y_0 = \overline{y} + \xi_0$, so for a large network and independent errors this immediately gives $\overline{y} = \frac{(1-m)\xi_0}{m}$, hence $D \simeq \frac{2\varpi_0}{m^2} \simeq \frac{4\varpi_0}{d}$ for small m. Writing $y_i = (y_i - \overline{y}) + \overline{y}$, we obtain $v(y_i) = \frac{1}{2}(d + D)$. Since $D \simeq \frac{4\varpi_0}{d}$, the loss $v(y_i)$ is minimized for $D \simeq d \simeq 2\varpi_0^{1/2}$ (hence $m^{**} \simeq \varpi_0^{1/4}$).■

Result 5 says that the social optimum is achieved for $D \simeq d$ and it establishes a relationship between consensus within each group (small d) and polarization across groups (high D): as m decreases below m^{**}, within-group consensus goes up but so does polarization across the groups.

Our equilibrium analysis provides one possible reason for m being too low, but there may be others. For example, imagine that for some issues, the errors ξ_i are correlated across network members (calling for higher m), while for other issues, the errors are independent (calling for lower m). If agents are unable to adjust m to the type of problem they face, the weights m will be inefficiently low for the problems where there are correlated errors, thus fostering too much consensus and polarization for these problems.

5.4 An alternative modeling of errors

To conclude this Section, we briefly comment on our modeling of errors. Given the way we index errors, it is natural to interpret ξ_i as a persistent error that i makes in *processing* or hearing others' opinions. We discuss below an alternative model where i does not make processing errors but makes a persistent error ζ_i^e in *expressing her opinion.* In this case Equation (1) becomes

$$z_i^t = A_i(y^t + \zeta^e)$$

so in effect, i is subject to an error $\xi_i \equiv A_i\zeta^e$. Our analysis thus extends to this alternative modelling with ξ_i appropriately re-defined. With perfectly correlated errors, this alternative modeling yields $\xi_i = \zeta_i^e$, so the analysis is unchanged. With independent errors, the errors ξ_i (hence the cumulated errors $\widehat{\xi_i}$) now potentially depends on the network structure. We re-examine our three network examples in light of this alternative modeling. Specifically, we compare the cumulated errors terms when agents are subject to processing errors ($\xi_i = \xi_i^p$)

(respectively expressing errors $\xi_i = A_i \zeta^e$), and denote by $\widehat{\varpi}_i^p$ and $\widehat{\varpi}_i^e$ the respective variances, assuming that all errors ξ_i^p and ζ_i^e are independent and homogenous. We let $\varpi = var\xi_i^p = var\zeta_i^e$. We have

Proposition 11: *For the directed circle, $\widehat{\varpi}_i^p = \widehat{\varpi}_i^e$. For the complete network, $\widehat{\varpi}_i^p = \widehat{\varpi}_i^e + \frac{n-2}{n-1}\varpi$. For the star network, $\widehat{\varpi}_0^p = \widehat{\varpi}_0^e + \frac{(2-m)(n-1)}{mn}\varpi$.*

The main insight of this Proposition is that although the magnitude of the one-shot error ξ_i that a player faces may differ substantially depending on whether we consider processing or expressing errors, the cumulated error terms do not differ much in the sense that terms of order ϖ/m^2 remain the same.[46] The consequence is that, while processing errors generate slightly larger cumulated errors than communication errors, the effect is negligible for small ϖ, and at least for the specific networks considered above, equilibrium analysis is then unchanged (see Appendix).

6 Idiosyncratic errors

We now introduce idiosyncratic errors and assume that

$$\varepsilon_i^t = \xi_i + \nu_i^t$$

where ν_i^t are i.i.d. across individuals and time.[47] We further assume $E\nu_i^t = 0$ and let $\varpi^0 = var(\nu_i^t)$. We wish to characterize the (additional) loss generated by these idiosyncratic errors, and examine the consequence regarding incentives.

In the absence of idiosyncratic elements, the speeds of adjustment γ_i plays no role when $m_{i_0} > 0$ for i_0. The main insight of this Section is that idiosyncratic errors induce temporary variations in opinions which are potentially costly, and players have incentives to reduce these variations by decreasing γ_i. Furthermore, when all players choose an arbitrarily small γ_i, long-run opinions essentially coincide with the ones obtained in the absence of idiosyncratic errors.

Formally, for any fixed m, x and ξ, we define the expected opinion vector $\overline{y}^t = Ey^t$ where the expectation is taken over all ν_i^s for $s \leq t$. We also let $\eta_t = y^t - \overline{y}^t$ and $V^t = var(\eta_t)$. Furthermore, we let y^0 denote the long-run opinion that would obtain *in the absence of idiosyncratic errors*, and $L_i^0 = var(y^0)$ the associated loss of player i computed over realizations of x and ξ. The next Proposition (proved in Appendix B) provides the analog of Propositions 1 to 3 for the idiosyncratic noise case:

Proposition 12: *If $m_i = 0$ for all i, V^t increases without bound.* If *$m_{i_0} > 0$ for some i_0, $\overline{y}^t$ and V^t both have well-defined limits $\overline{y}$ and V. Besides, $\overline{y} = y^0$, V is independent of x and ξ, and $L_i = L_i^0 + V$. Furthermore, if $m_i \leq \overline{m}$* and $\gamma_i \geq \underline{\gamma}$ for all i, $V_i \geq \frac{\varpi^0}{2n}\frac{\underline{\gamma}^2(1-\overline{m})^2}{\overline{m}}$.

For given m, $\gamma > 0$, expected long-run opinions eventually coincide with y_0, but long-run opinions are subject to temporary changes resulting from idiosyncratic communication errors. Proposition 12 shows that, for given γ, these temporary changes are significant and costly and when all m are small.

[46] For example, in a star network, $\xi_0 = \overline{\zeta}^e$ so $var\xi_0 = \varpi/n$ for expressing errors, and $var\xi_0 = \varpi_0 = \varpi$ for processing errors. Nevertheless, the cumulated errors are respectively $(\overline{\zeta}^e + (1-m)\zeta_0^e)/m$ and $(\xi_0 + (1-m)\overline{\xi})/m$.

[47] Implicitly, we think of ν_i^t as an error in interpreting the opinions expressed by others. Alternatively, one could consider errors in expressing one's opinion.

However choosing a lower γ_i slows down the adjustment of one's opinion. Result 6 below shows that for small enough γ_i, long-run opinions becomes essentially unaffected by temporary shocks in perceptions or temporary variations in others' opinions.

Result 6: *Fix $\underline{m}$. We have:*
(i) There exists c such that for any $\gamma > 0$ and $m \geq \underline{m}$, $V_i \leq c \max \gamma_j$.
(ii) For any $\gamma_{-i} > 0$, there exists c such that for all $m \geq \underline{m}$, $V_i \leq c\gamma_i$.

The proof is in Appendix B. Item (i) shows that when all γ_i are small, all V_i are small. Item (ii) shows that by choosing γ_i very small, a player can get rid of the additional variance induced by the idiosyncratic noise.

Note that the incentive to set γ_i arbitrarily small obviously depends on the assumption that players only care about long-run opinions. If players also cared about opinions at shorter horizons, then they would have incentives to increase γ_i to more quickly absorb information from the opinions of others: the trade-off is between increasing the rate of convergence (which is desirable when the relevant horizon is shorter) and increasing the variance induced by idiosyncratic noise (which is not desirable).

7 Extensions and interpretations

In this section we discuss extensions of and possible variations upon our base model, with the view to understand why different rules lead to different degrees of information aggregation in different settings.

7.1 Biased persistent errors

We have so far assumed that the persistent error is drawn from a distribution that is mean zero. One can however imagine settings where it is reasonable to assume that the persistent error is biased, centered on ξ_i^0 for player i. This could be because some individuals are systematically biased in what they report or process (for whatever reason), or because others erroneously believe that they are and wrongly correct for it. Another reason could be that preferences are heterogenous, say each person cares about $\theta_i = \theta + b_i$, observes $x_i = \theta_i + \delta_i$, but has an imprecise and potentially biased estimate of the vector of preference spreads $\beta_i = (b_j - b_i)_j$.

In either case, adding systematic biases ξ_i^0 can only raise the terms $\widehat{\varpi}_i = E\widehat{\xi_i}$, thus providing additional incentives to raise m_i and increasing the losses L_i.

7.2 Other communication protocols

We have followed the standard approach to modeling communication in this literature, with each player communicating with all his neighbors at every date.[48] We now consider an extension where each round of communication is one-sided and, at any date t, each agent i only hears from a subset $N_i^t \subset N_i$ of his neighbors but there exists K such that each player hears from all his neighbors at least once every K periods.[49] Imperfect communication is modeled

[48]Banerjee et al. (2019) introduce the idea of a Generalized DeGroot model where not everyone starts with a signal and therefore does not participate in the communication till they get a signal. They show that this partially weakens the "wisdom of crowds".

[49]That is, for all $t : \cup_{s=1,..,K} N_i^{t+s-1} = N_i$.

as before, through the addition of an error term ξ_i that slants what i hears. Together these give us

$$z^t_{i,j} = y^{t-1}_j + \xi_i \text{ if } j \in N^t_i$$
$$z^t_{i,j} = z^{t-1}_{i,j} \text{ if } j \in N_i \backslash N^t_i$$

where $z^t_{i,j}$ is i's current perception of j's opinion, based on the last time he has heard from j. Player i uses these perceptions to construct an average over neighbor's opinions

$$z^t_i = A_i Z^t_i$$

where $Z^t_i = (z^t_{i,j})_j$ is the vector of i's perceptions and $A_i = (A_{ij})_j$ defines as before how i averages others' opinions. We continue to assume FJ updating. We have:

Proposition 13: *Assume at least one player, say i_0, updates according to FJ with $m_{i_0} > 0$. Then for any fixed x, ξ, y^t converges and the limit vector of expected opinions y is independent of the protocol.*[50]

Intuitively, convergence obtains for standard reasons, and at the limit, since expected opinions do not change, the timing with which one hears others does not matter (see Appendix).

This robustness contrasts with what happens when players use DG rules. For example, consider two agents using DG rules and assume that agent 1 updates every period, while agent 2 updates every three other periods. At dates t where 2 updates, we have:

$$\begin{aligned} y^t_1 &= (1-\gamma_1)^3 y^{t-3}_1 + (1-(1-\gamma_1)^3) y^{t-3}_2 \\ y^t_2 &= (1-\gamma_2) y^{t-3}_2 + \gamma_2 y^{t-1}_1 \\ &= (1-\gamma_2) y^{t-3}_2 + \gamma_2((1-\gamma_1)^2 y^{t-3}_1 + (1-(1-\gamma_1)^2) y^{t-3}_2) \\ &= (1-\gamma_2(1-\gamma_1)^2) y^{t-3}_2 + \gamma_2(1-\gamma_1)^2 y^{t-3}_1 \end{aligned}$$

So, the process evolves as if weights were $\gamma'_1 = 1-(1-\gamma_1)^3 > \gamma_1$ and $\gamma'_2 = \gamma_2(1-\gamma_1)^2 < \gamma_2$. This means that with DG rules, changes in the frequencies with which players communicate amount to changes in the values of γ_i (when you hear less often from others, your opinion changes more slowly, effectively reducing γ_i). And even when communication is noiseless, these changes modify long-run opinions: if γ_i goes down, long-run opinions get closer to i's opinions (see Section 3.1 Equation (4)).

Thus, even in the absence of transmission errors, variations in the communication protocol induce additional variation in long-run opinions which can be mitigated by the use of FJ rules by all players. That said, in the absence of transmission errors, long-run opinions under DG remain averages over initial opinions, so the fragility is not as severe as the one already highlighted: the variance induced by variations in the protocol remains bounded even when $m_i = 0$.

7.3 Uncertainty over the precision of initial signals.

We examine here another variation of the model, assuming that the precision of initial signals is a random variable and that players are able to ajust the speed γ_i as a (linear) function of σ^2_i.

[50] So long as the condition in footnote 49 holds.

We **argue below** that in the absence of processing errors, this type of shock does not affect the performance of DG and therefore, unlike where there are errors, there is no incentive for players to use the instrument m_i.

Formally, assume that each the speed of adjustment γ_i as a linear function of the variance of signal, that is, $\gamma_i = \mu_i \sigma_i^2$. Then for well-suited coefficients $\mu^* = (\mu_i^*)_i$ information aggregation is perfect, which further implies that this particular μ^* is also a Nash Equilibrium of the game where each chooses μ_i.

To see why, recall that under DG, the consensual long-run opinion is a weighted average of initial opinions, with weights proportional to ρ_i/γ_i (see (4)). So if the μ_i's are proportional to ρ_i, the weights become proportional to ρ_i/γ_i, hence proportional to $1/\sigma_i^2$, implying that perfect aggregation obtains for each vector of realization $(\sigma_1, ..., \sigma_n)$.

7.4 Coarse communication

In the social learning literature, it is common to focus on choice problems where there are two possible actions, and the information being aggregated is which of the two is being recommended by others. Coarse communication is potentially a source of herding, but when agents have many neighbors, the fraction of players choosing a given action may become an accurate signal of the underlying state. We explain below how our model can accommodate an economic environment of this kind, and we use this to relate our findings to Ellison and Fudenberg (1993,1995) and Frick et al. (2020), as well as Bohren and Hauser (2021).

Assume heterogenous preferences with $\theta_i = \theta + b_i$ characterizing i's value from choosing 1 over 0, so the optimal action a_i^* is 1 when $\theta_i > 0$, 0 otherwise.[51] Agent i knows b_i but does not know θ perfectly. He has an initial opinion $x_i = \theta + \delta_i$ and aggregates opinions of others to sharpen his assessment of θ. Assume the b_i's are drawn from identical distribution g (and cumulative denoted G) with full support on $\mathcal{R}$.

We define, as before, y_i^t as agent i's opinion (about θ) at date t and we assume that an agent with current opinion y_i^t reports $a_i^t = 1$ if $y_i^t + b_i > 0$ and $a_i^t = 0$ otherwise. Each agent i observes the fraction f_i^t of neighbors that choose action 0, which she can use to make an inference $\psi_i(f_i^t)$ about θ, and update her opinion using an FJ-like rule:

$$y_i^{t+1} = (1 - \gamma_i) y_i^t + \gamma_i (m_i x_i^t + (1 - m_i) \psi_i(f_i^t))$$

Long-run opinions clearly depend on the inference rule assumed, but there is a natural candidate for ψ_i, the function $\phi \equiv h^{-1}$, where $h(y) \equiv G(-y) = \Pr(y + b_i < 0)$ is the fraction of agents that choose $a = 0$ when their opinions are all equal to y. If others have opinions that are correct and equal to θ, a fraction $f \simeq h(\theta)$ choose $a = 0$ and $h^{-1}(f)$ is a good proxy for θ. Of course this assumes that agents know the distribution over preferences. In the spirit of our previous analysis, let's assume that

$$\psi_i(f) = \phi(f) + \xi_i$$

[51] Thus for i with preference parameter b_i, choosing 0 when $\theta + b_i > 0$ costs $\theta + b_i$. When agents choose between products 1 or 0, θ represents a relative quality dimension affecting all preferences, as in Ellison and Fudenberg (1993).

where ξ_i is a persistent error in interpreting f.[52,53] To fix ideas, we assume correlated errors ($\xi_i = \xi$ for all i) with variance ϖ.

Within this extension, we may ask about the fragility of long-run opinions when m is small, as well as equilibrium and socially efficient weights (details are provided in the Appendix).

DG-like rules ($m = 0$) generate long-run opinions unanimously in favor of $a = 1$ if $\xi > 0$, $a = 0$ if $\xi < 0$, *independently of the underlying state and the initial signals received.*

Under FJ with small m, long-run opinions remain anchored on initial opinions, but long run opinions drift away from θ and converge to $\theta + \frac{(1-m)\xi}{m}$. The trade-off is thus similar to the one in our basic model. Raising m reduces fragility with respect to transmission noise, dampening the echo term $\frac{(1-m)\xi}{m}$. And agents continue to diagree even in the long-run. The consequence regarding social incentives and private incentives is as before, with m^* and m^{**} respectively comparable to $\varpi^{1/3}$ and $\varpi^{1/4}$: agents do not incorporate the damaging echo effect that an m_i set too low produces in their choice of m_i .

7.5 A connection to misspecified Bayesian models.

How does the results in the previous sub-section relate to the results from Bayesian models where agents have misspecified priors (and in particular Frick et al. (2020) and Bohren and Hauser (2021))? Consider a social learning environment related to **these** Bayesian models where players move in sequence and observe all previous choices. Preferences and signals are as defined above. Assume the true state is θ_0. Under Bayesian learning, if beliefs get highly concentrated on some θ, then private signals do not affect decisions much and the fraction f of people that choose $a = 0$ are approximately those for which $\theta + b < 0$ so $f \simeq G(-\theta)$. If agents have an erroneous prior about the distribution of b's and believe its cumulative is shifted by ξ (say, $\widehat{G}(b) \equiv G(b-\xi)$) then agents are expecting a fraction close to $\widehat{f} = \widehat{G}(-\theta) = G(-\theta-\xi)$, so if $\xi > 0$, $\widehat{f} < f$. When the subjective prior over states has full support, this should inevitably lead agents to believe that the state is lower than θ (to justify the higher-than-expected f observed) and so on...which explains the fragility result obtained in Frick and al. (2020)

Let us now introduce, as in Bohren and Hauser, a fraction q of autarkic players that only base their choice on their private signal x_i (thus ignoring the social information). Define $G^0(\theta)$ as the fraction of autarkic types that choose $a = 0$ when the state is θ, and to fix ideas, further assume that non-autarkic types have correct priors about G^0. When beliefs of non-autarkic types are concentrated on θ and the true state is θ_0, the fraction f becomes

$$f = qG^0(-\theta_0) + (1-q)G(-\theta)$$

while a fraction

$$\widehat{f} = qG^0(-\theta) + (1-q)G(-\theta-\xi)$$

[52]As in Frick et al. (2020), ξ_i could arise from an erroneous prior $g_i \neq g$, with agents using the inference function $\psi_i = h_i^{-1}$ where $h_i(\theta) = G_i(-\theta)$. The difference $\xi_i(f) \equiv \psi(f) - \phi(f)$ is an error in making inferences. With preferences centered on $\overline{b}$, and agent having an erroneously translated prior centered on $\overline{b}_i$, the error is independent of f and equal to $\xi_i \equiv \overline{b}_i - \overline{b}$.

[53]Ellisson and Fudenberg (1993, Section 1) examines social learning assuming $b_i = 0$ for all and $\psi_i(f) = f - 1/2$: choices are tilted in favor of the more popular one. EF find that small enough ms generate perfect learning in the long-run. A key aspect of the inference rule $\psi_i(f)$ is that it correctly maps the sign of $f - 1/2$ to the sign of θ, which, given homogeneity, is the only thing that agents care about. (Note that in EF, agents receive many signals x_i about the state, but, given their assumptions, their model is equivalent to the one proposed here where agents just receive one signal at the start).

would be expected. The observed f will meet expectations when

$$G^0(-\theta) - G^0(-\theta_0) = \frac{1-q}{q}(G(-\theta) - G(-\theta - \xi))$$

which implies a discrepancy $\Delta = \theta_0 - \theta$ comparable to $\frac{\xi}{q}$, which thus blows up when q is small.

To relate this to our paper, observe that a measure q of autarkic types generates an overall inefficiency comparable to q (because they are not using information so each experiences a loss comparable to 1), while when ξ is a random variable with variance ϖ, the loss induced by the discrepancy Δ is quadratic in Δ, so comparable to $\frac{\varpi}{q^2}$, which in turn implies that to implement the social optimum (to minimize the overall loss), q should be comparable to $\varpi^{1/3}$.

Autarkic types thus play a role similar to our weights m_i, helping the anchoring the beliefs of social types.[54] In our setup, the analog of social and autarkic types would be to assume that agents are either DG ($m_i = 0$) or use $m_i = 1$. In contrast, we have assumed that some intermediate m_i is feasible for each agent.

Another difference is that we focus on the optimal choices of m_i from the social or private points of view. In looking for a Nash equilibrium, we decentralize the choice of m_i and endogenize the weight each puts on social versus private information.[55]

The lesson we draw from this discussion is that both DG and Bayesian updating are sensitive to transmission or specification errors for a similar reason: they both incorporate a force towards consensus, but consensus is not feasible (given the errors), and beliefs are thus pushed to the boundaries of the feasible set of states. FJ-like rules, to the extent that they allow for sufficiently diverse opinions or beliefs, end up being more robust.

7.6 Non-stationary weights.

The updating processes that we consider have stationary weights. Agents do not attempt to exploit the possibility that early reports possibly reveal more information than later reports: later reports from neighbors may incorporate information that one has oneself transmitted to the network, and therefore should have lesser impact on own opinion.

As a matter of fact, with two players, one could imagine a process in which (i) player 1 combines the first report he gets with own opinion, yielding $y_1 = m_1 x_1 + (1 - m_1)(x_2 + \varepsilon)$, and then ignores any further reports from player 2; and (ii) player 2 follows DG. With m_1 set appropriately, such a process would permit player 1 to almost perfectly aggregate information and player 2 to benefit from that information aggregation performed by player 1.

There are however important issues with such time-dependent processes. In particular, it is not obvious how one extends these to larger networks since they require that each person knows his or her role in the network. They are also sensitive to the timing with which information gets transmitted or heard. With some randomness in the process of transmission, it could for example be that the first report y_2 that player 1 hears already incorporates player 1's own signal (because after a while y_2 starts being a mixture between x_2 and x_1), and as a result,

[54] Note that unlike Bohren and Hauser, we find here that the fraction q needs to be large enough. This is because, unlike BH who assume few states and correct priors over states, we assumed here that subjective priors on θ have full support.

[55] A similar decentralization exercise (endogenizing q) could be done in the BH environment with agents choosing ex ante whether to be autarkic or social, with the consequence that in equilibrium, they would have to be indifferent between the two roles, hence incur a significant loss (equal to that of the autarkic type).

player 1 should put more weight on the opinions of others. But of course, in events where $y_2 = x_2$, this increase in weight makes information aggregation worse.

To illustrate this strategic difficulty in a simple model with noisy transmission, assume that time is continuous, communication is one-sided (either 1->2 or 2->1), with each player getting opportunities to communicate at random dates. The processes generating such opportunities are assumed to be two independent Poisson process with (identical) parameter λ. Also assume that a report, once sent, gets to the other with probability p. Consider the time-dependent rule where each person communicates own current opinion, and their current opinion coincides with their initial opinion if one has not received any report ($y_i = x_i$), and otherwise coincides with $y_i = m_i x_i + (1 - m_i) z_i^f$ where z_i^f is the perception of the first report received. Even if perceptions are almost correct (i.e. perceptions almost coincide with the other's current opinion), the noise induced by the communication channel generates uncertainty about who updates first, contributing to variance in the final opinion for all m_i. For example, in events where player 1 already sent a report and receives one from player 2, it matters whether player 2 received the report that 1 sent and incorporated it into her opinion, or whether player 2 failed to receive the report, in which case what player 1 gets is player 2's initial opinion.

In contrast, the time-independent FJ is not sensitive to that noise and achieves reasonably good information aggregation for many values of $m = m_1 = m_2$. FJ rules conveniently address a key issue in networks: whether what I hear already incorporates some of what I said.

8 Concluding remarks

We end the paper with a discussion of issues that we have not dealt with, and which may provide fruitful directions for future research.

One premise of our model is that everyone has a well-defined initial signal. However the analysis here would be essentially unchanged if some players did not have an initial opinion to feed the network and were thus setting $m_i = 0$ for the entire process. FJ would aggregate the initial opinions of those who have one.

In real life many of our opinions come from others and in ways that we are not necessarily aware of, and the existence of a well-defined "initial opinion" could be legitimately challenged. In other words, people may have a choice over the particular opinion they want to hold on to and refer back to (in other words, the one that gets the weight m_i).

To see why this might matter, consider a variation of our model where some players (N^{dg}) have initial opinions but use DG rule (or set m_i very low), while other agents (N^{fj}) have no initial opinions (or very unreliable ones). In this environment, there is a risk that the initial opinions of the DG players eventually disappear from the system, and soon are overwhelmed by noise in transmission. The other (non-DG) players could provide the system with the necessary memory, using the initial communication phase to gradually build up an "initial opinion" based on the reports of their more knowledgeable DG neighbors, and then seed in perpetually that "initial opinion" into the network. In other words, in an environment where information is heterogeneous and weights m_i are set sub-optimally by some, there could be a value for some agent in adopting a more sophisticated strategy in which the "initial opinion" is updated for some period of time before it becomes anchored. In other words, it may be optimal for some of the less informed to listen and not speak for a while as they build up their own "initial opinions" before joining the public conversation.

Another important assumption of our model is that the underlying state θ is fixed. In particular, there would be no reason to keep on seeding in the initial opinions if the underlying state drifts. However it may still be useful to use a FJ-type rule where the private seed is periodically updated by each player to reflect the private signals about θ that each one accumulates.

Finally, our approach evaluates rules based on their fitness value. With a continuum of states and opinions modeled as point-beliefs, averaging opinions naturally has some fitness value. When there are few states and opinions take the form of probabilistic beliefs, averaging beliefs or log-beliefs will generally have poor (if not negative) fitness value (see for example Sobel (2014)). In this context, a promising FJ-like rule would consist in linearly aggregating the *initial change* in one's own log-belief (induced by one's initial signal) with the perceived *change* in a composite neighbor's log-beliefs: such a rule accommodates the intuition that *belief changes* potentially reveal information, and through appropriate weighting of one's own versus other's changes, it enables each player to deal with situations where initial belief updates are driven by interpretation errors (one then needs to filter out interpretation errors and averaging is good in these cases) and situations where independent information needs to be aggregated (adding changes in log-beliefs across all players would be called for). Furthermore, as in this paper, it allows beliefs to differ and the anchoring on one's own initial information (i.e., the initial change in one's own log-belief) can limit the damaging effects of cumulated processing errors.

References

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

Alatas, Vivi, Abhijit Banerjee, Arun G. Chandrasekhar, Rema Hanna, and Benjamin A. Olken (2016), "Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia." *The American Economic Review* 106, no. 7: 1663-704.

Aumann, Robert (1976), "Agreeing to disagree," *The Annals of Statistics*, 4(6):1236–1239.

Bala, Venkatesh and Sanjeev Goyal (1998), "Learning from Neighbours," *The Review of Economic Studies*, 65 (3), 595–621.

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

Banerjee, Abhijit, Emily Breza, Arun G. Chandrasekhar and Markus Mobius, 2019. "Naive Learning with Uninformed Agents," NBER Working Papers 25497.

Bakshy, Eytan and Solomon Messing and Lada A. Adamic (2015), "Exposure to ideologically diverse news and opinion on Facebook," *Science*, 348, 6239, pp1130-1132.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy,* 100 (5), 992-1026.

Bohren, J.A. and Hauser, D.N. (2021), "Learning With Heterogeneous Misspecified Models: Characterization and Robustness," *Econometrica*, 89: 3025-3077.

Boxell, Levi, Jacob Conway, James N. Druckman and Matthew Gentzkow (2022), "Affective Polarization Did Not Increase During the COVID-19 Pandemic", *Quarterly Journal of Political*

Science, Vol. 17: No. 4, pp 491-512.

Compte, Olivier and Andy Postlewaite (2018). Ignorance and Uncertainty (Econometric Society Monographs). Cambridge: Cambridge University Press.

Cox, Daniel A., Ryan Streeter, Samuel J. Abrams and Jacqueline Clemence (2020), "Socially distant: How our divided social networks explain our politics," Survey Center on American Life. https://www.americansurveycenter.org/research/socially-distant-how-our-divided-social-networks-explain-our-politics/

Degroot, Morris H. (1974), "Reaching a Consensus," *Journal of the American Statistical Association*, 69 (345), 118–121.

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

Dasaratha, Krishna, Benjamin Golub and Nir Hak (2020),"Learning from Neighbors About a Changing State," SSRN working paper n° 3097505.

Ellison, Glenn, and Drew Fudenberg (1993), "Rules of Thumb for Social Learning," *Journal of Political Economy* 101 (4), 1993, 612–43.

Ellison, Glenn, and Drew Fudenberg (1995), "Word-of-Mouth Communication and Social Learning," *The Quarterly Journal of Economics* 110 (1), 93–125.

Eyster, Erik, and Matthew Rabin (2010), "Naive Herding in Rich-Information Settings." *American Economic Journal: Microeconomics*, 2 (4), 221-43.

Friedkin, Noah E. and Eugene C. Johnsen (1990), "Social influence and opinions," *The Journal of Mathematical Sociology*, 15 (3-4), 193-206.

Friedkin, Noah E. and Eugene C. Johnsen (1999), "Social Influence Networks and Opinion Change," *Advances in Group Processes* 16: 1-29.(1999)

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

Fudenberg, Drew and David K. Levine (1998), *The Theory of Learning in Games*, The MIT Press.

Gentzkow, Matthew, Michael Wong and Allen T. Zhang (2018), "Ideological Bias and Trust in Information Sources," Working paper.

Genest, Christian and James V. Zidek (1986), "Combining Probability Distributions: A Critique and an Annotated Bibliography," *Statistical Science*, 1, 114-135.

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

Jackson, Matthew O., Suraj Malladi, and David McAdams (2019) "Learning through the Grapevine: the Impact of Message Mutation, Transmission Failure, and Deliberate Bias" Working paper

Jadbabaie, Ali, Pooya Molavi, Alvaro Sandroni, Alireza Tahbaz-Salehi (2012) "Non-Bayesian social learning," *Games and Economic Behavior*, 76 (1), 210-225.

Levy, Gilat, and Ronny Razin (2015) "Correlation Neglect, Voting Behavior, and Information Aggregation." *The American Economic Review* 105 (4), 1634–645.

Molavi, P., Tahbaz-Salehi, A. and Jadbabaie, A. (2018), A Theory of Non-Bayesian Social Learning. *Econometrica*, 86: 445-490.

Mossel, Elchanan, Allan Sly, and Omer Tamuz (2015), ”Strategic Learning and the Topology of Social Networks,” *Econometrica* 83, no. 5: 1755-794.

Pennycook, G., Epstein, Z., Mosleh, M. et al. Shifting attention to accuracy can reduce misinformation online. *Nature* 592, 590–595 (2021).

Roberts, Ben, and Dirk P. Kroese (2007), “Estimating the Number of s-t Paths in a Graph,” *Journal of Graph Algorithms and Applications,* 11 (1): 195–214.

Rosenberg, D., E. Solan, and N. Vieille (2009), ”Informational Externalities and Emergence of Consensus,” *Games and Economic Behavior*, 66 (2), 979–994.

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

Sethi, Rajiv, and Muhamet Yildiz (2016). ”Communication with unknown perspectives,” *Econometrica* 84 (6), 2029–69.

Sethi, Rajiv and Yildiz, Muhamet (2019), ”Culture and Communication,” SSRN working paper n° 3263560.

Sobel, Joel (2014), ”On the relationship between individual and group decisions,” *Theoretical Economics*, 9 (1): 163-185.

Sunstein, Cass. R. (2001), *Echo Chambers: Bush V. Gore, Impeachment, and Beyond,* Princeton, NJ: Princeton University Press.

Appendix A

Notations. Define M and Γ as the $N \times N$ diagonal matrices where $M_{ii} = m_i$ and $\Gamma_{ii} = \gamma_i$. For any fixed vectors of signals x and systematic bias ξ, we let

$$X = Mx + (I - M)\xi$$

and, whenever $m_i > 0$, we let $\widetilde{x}_i = x_i + \xi_i(1 - m_i)/m_i$ denote the modified initial opinion, and $\widetilde{x} = (\widetilde{x}_i)_i$ the vector. Next define the matrix $B = I - \Gamma + \Gamma(I - M)A$.

We shall say that P is a *probability matrix* if and only if $\sum_j P_{ij} = 1$ for all i. Note that A is a probability matrix and throughout, we assume that the power matrix A^k only has strictly positive elements for some k. Finally, we refer to $v(y)$ as the variance of y.

In the main text, we show that when $m_i > 0$ for all i, long-run opinions are weighted averages of modified opinions $\widetilde{x}$. Lemma 5 below (proved in Appendix B) generalizes this observation. Define $N^0 \subsetneq N$ as the set of n_0 agents following DG ($m_i = 0$). Denote by ξ^0 the vector of errors of these players. We have:

Lemma 5. *Assume $n_0 < n$. Then $y = P\widetilde{x} + Q\xi^0$ where P is a $(n, n - n_0)$-probability matrix.*

Proposition 3 is then obtained as an immediate corrolary of Lemma 5.

Proof of Proposition 3: From Lemma 5, $L_i = var(y) \geq \frac{1}{n} \min var(\widetilde{x}_i) \geq \frac{(1-m)^2\varpi}{n\ m^2}$.■

We now turn to the proof of our main Propositions.

Proof of Proposition 4: Assume $m >> 0$ so $\widetilde{x}_j$ is well-defined for all j.[56] For $j \neq i$ let $X_j = m_j\widetilde{x}_j + (1 - m_j)A_{ji}y_i$ and $c^i_j = m_j + (1 - m_j)A_{ji}$. (13) can be written in matrix form to obtain, by definition of Q^i, $y_{-i} = Q^iX$. Note that if $\widetilde{x}_j = 1$ for all j and $y_i = 1$, then $y_k = 1$ for all k, so $\sum_{j\neq i} Q^i_{kj}c^i_j = 1$ for all k, which implies

$$\sum_{j\neq i} Q^i_{kj}(1 - m_j)A_{ji} = 1 - \sum_{j\neq i} Q^i_{kj}m_j, \tag{24}$$

and, since Q^i is a positive matrix,[57] $\sum_{j\neq i} Q^i_{kj}m_j \leq 1$, so $\sum_{j\neq i} R^i_j m_j \leq 1$. (24) further implies

$$y_k = \sum_{j\neq i} Q^i_{kj}m_j\widetilde{x}_j + (1 - \sum_{j\neq i} Q^i_{kj}m_j)y_i, \tag{25}$$

thus characterizing the influence of y_i on k's opinion. In particular, the smaller $\sum_{j\neq i} Q^i_{kj}m_j$ the larger the influence of i on k. Averaging over all neighbors of i, and taking into account the weight A_{ik} that i puts on k, we obtain:

$$y_i = m_i\widetilde{x}_i + (1 - m_i)(\sum_{j\neq i} R^i_j m_j\widetilde{x}_j + y_i(1 - \sum_{j\neq i} R^i_j m_j) \tag{26}$$

which, since $m_j\widetilde{x}_j = m_jx_j + (1 - m_i)\xi_j$ and $h_i = 1/\sum_{j\neq i} R^i_j m_j$, gives the desired Expressions (14) for y_i, $\widehat{x}_i$, p_i and $\widehat{\xi}_i$.■

[56] Cases where some or all m_j are 0 can be derived by taking limits as Q^i remains well-defined.

[57] $Q^i = \sum_{n\geq 0}((I - M^i)(I - \alpha^i)\widetilde{A}^i)^n$ so Q^i is non-negative. If in addition, $m_{-i} << 1$, and since A is connected, then $Q^i >> 0$.

Proof of Proposition 5: Assume $m_i > 0$ and apply Proposition 4, taking the limit where all m_j tend to 0. For A given, Q^i and R^i are uniformly bounded (with a well-defined limit when all m_j tends to 0), and $(1-p_i)\widehat{x}_i$ tends to 0, which concludes the proof.■

Proof of Result 1: There are two parts in this proof. We first prove that the $m_i' s$ cannot be positive. Next we show that the equilibrium outcome must be efficient. Recall $\pi^* = \arg\min_\pi v(\sum_k \pi_k x_k)$ is the efficient weighting of seeds and $v^* \equiv v(\pi^*.x)$. Also let $r_i = 1/h_i$.

Assume by contradiction that $m_j > 0$. Then (14) implies that $m_i > 0$ for all i, so $m >> 0$. Next, from (26), and letting $r_i = 1/h_i$, we obtain $\widehat{y}_i = r_i\widehat{x}_i + (1-r_i)y_i$, hence substituting y_i,

$$\widehat{y}_i = (1-r_i)p_i x_i + (1-(1-r_i)p_i)\widehat{x}_i. \tag{27}$$

So both $\widehat{y}_i$ and y_i are weighted average between x_i and $\widehat{x}_i$, and since $m >> 0$, $r_i \in (0,1)$, the weights are different. Since i optimally weighs x_i and $\widehat{x}_i$ (using p_i on x_i), the weight $(1-r_i)p_i$ is suboptimal so

$$v(y_i) < v(\widehat{y}_i) \leq \max_{j\neq i} v(y_j), \tag{28}$$

where the second inequality follows from $\widehat{y}_i$ being an average of the y_j's. Since (28) cannot be true for all i, we get a contradiction. The equilibrium must thus be DG.

Consider now a DG equilibrium. Call $\pi = (\pi_i)_i$ the weights on seeds induced by γ and A, $\widehat{\pi}^i$ the relative weights on $k \neq i$, and $\widehat{x}_i = \widehat{\pi}^i.x_{-i}$. We have $y_i = \pi_i x_i + (1-\pi_i)\widehat{x}_i$, and modifying γ_i allows the agent to modify π_i without affecting $\widehat{x}_i$ (player i increases π_i by decreasing γ_i). Therefore the optimal choice π_i satisfies

$$\frac{\pi_i}{1-\pi_i} = \frac{v(\widehat{x}_i)}{\sigma_i^2}$$

Let $W_i^* = \min_q v(q.x_{-i})$. Since optimal weighting of all seeds requires optimal weighting on seeds other than i, we have:

$$\frac{\pi_i^*}{1-\pi_i^*} = \frac{W_i^*}{\sigma_i^2}$$

which implies

$$\pi_i = \pi_i^* + \frac{(1-p_i)(1-\pi_i^*)}{\sigma_i^2}(v(\widehat{x}_i) - W_i^*) \tag{29}$$

Since all π_i (and π_i^*) add up to one, one must have $v(\widehat{x}_i) - W_i^* \leq 0$, hence information aggregation is perfect.■

Before showing Result 2, we start with two Lemma that we also use to prove Result 3:

Lemma 6: *For each $j \neq i$, there exists μ_{ji} and a probability vector $C^{ji} \in \Delta_{N-1}$, each independent of m_i, such that*

$$y_j = (1-\mu_{ji})C^{ji}\widetilde{x}_{-i} + \mu_{ji}y_i \tag{30}$$

Proof: This immediately follows from Expression (24) in the proof of Proposition 4.■

Lemma 7: if $\frac{\partial L_i}{\partial m_i} \leq 0$, *then* $\frac{\partial L_j}{\partial m_i} < 0$ *for all* j.

Proof: Since μ_{ji} and C^{ji} are independent of m_i, we obtain:

$$\frac{\partial L_j}{\partial m_i} = (\mu_{ji})^2 \frac{\partial L_i}{\partial m_i} + \mu_{ji}(1-\mu_{ji}) \sum_{k \neq i} C_k^{ji} \frac{\partial Cov(\widetilde{x}_k y_i)}{\partial m_i}$$

We substitute $y_i = p_i x_i + (1-p_i)(\widehat{x}_i + \widehat{\xi}_i)$ (see (14)). Since $\widetilde{x}_k$ and x_i are independent, and since $\widehat{x}_i, \widetilde{x}_k$ and $\widehat{\xi}_i$ do not depend on m_i, we get

$$\frac{\partial L_j}{\partial m_i} = (\mu_{ji})^2 \frac{\partial L_i}{\partial m_i} - \mu_{ji}(1-\mu_{ji}) \frac{\partial p_i}{\partial m_i} \sum_{k \neq i} C_k^{ji} Cov(x_k \widehat{x}_i + \widetilde{x}_k \widehat{\xi}_i)$$

The terms $\frac{\partial p_i}{\partial m_i}$ and $Cov(x_k \widehat{x}_i)$ are positive, and so are the terms $Cov(\widetilde{x}_k \widehat{\xi}_i)$ when persistent errors are independent or positively correlated. The sum on the right side is thus positive (and the effect is amplified with errors), which proves Lemma 7.■

Proof of Result 2: Let $\underline{m} = \varpi/(1+\varpi)$. We show that DG and all strategies $m_i < \underline{m}$ are dominated by $\underline{m}$.

Assume first that all other players use DG. Then by Proposition 5, L_i decreases strictly with m_i. Now assume that at least one player j chooses $m_j > 0$. Then $L_i = p_i^2 + (1-p_i)^2 v(\widehat{x}_i + \widehat{\xi}_i)$. Whether persistent errors are independent or fully correlated, the variance of $\widehat{\xi}_i$ is at least equal to $h_i^2 \varpi$, which implies that L_i strictly decreases with p_i when $\frac{p_i}{1-p_i} < h_i^2 \varpi$, hence also with m_i when $\frac{m_i}{1-m_i} < h_i \varpi$, and from Lemma 7, we conclude that L_j increases as well (on this range of m_i).■

Proof of Proposition 6. Player i optimally sets p_i such that $\frac{p_i}{1-p_i} = \frac{v(\widehat{x}_i + \widehat{\xi}_i)}{v(x_i)} = \frac{W_i}{\sigma_i^2}$. Substituting p_i, we get the desired expression for L_i.■

Proof of Result 3:

Step 1: lowerbounds on $\overline{m}^i \equiv \max_{j \neq i} m_j$.

With transmission errors, optimal weighting of x_i and $\widehat{x}_i$ implies

$$\frac{p_i}{1-p_i} = \frac{v(\widehat{x}_i) + v(\widehat{\xi}_i)}{\sigma_i^2} \tag{31}$$

and (29) becomes

$$p_i = \pi_i^* + \frac{(1-p_i)(1-\pi_i^*)}{\sigma_i^2}(v(\widehat{x}_i) - v_i^* + v(\widehat{\xi}_i)) \tag{32}$$

The weight p_i is thus necessarily above the efficient level π_i^*, and there are now two motives for doing that: inefficient aggregation by others, and the cumulated error term $\widehat{\xi}_i$.

While (32) implies a lower bound on p_i, as (29) did, there is a major difference here with the no noise case where DG is used by all: p_i is the weight that i puts on own seed, but since there is no consensus, the sum $\sum_i p_i$ is not constrained to be below 1. Nevertheless, when all m are small, $\sum_i p_i = 1 + O(m)$ is close to 1, and this allows us to bound $v(\widehat{\xi}_i)$ (and the difference $v(\widehat{x}_i) - v_i^*$), as we now explain.

From Proposition 4, each opinion y_i may be written as $y_i = P^i x + (1-P_i^i)\widehat{\xi}_i$, where P^i is a weighting vector (such that $P_i^i = p_i$). (25) implies that when all m are small, the vectors P^i must be close to one another: seeds must be weighted in almost the same way, and differences in opinions are mostly driven by the terms $\widehat{\xi}_i$. Specifically, let $\overline{m}^i = \max_{j \neq i} m_j$. (25) implies that for all $k \neq i$,

$$p_k = P_k^k \leq P_k^i + c\overline{m}^i$$

for some constant c independent of m and k. Since $P_{kk} = p_k \geq \pi_k^*$, adding these inequalities yield

$$1 - p_i = \sum_{k \neq i} P_k^i \geq \sum_{k \neq i} p_k - Kc\overline{m}^i \geq 1 - \pi_i^* - Kc\overline{m}^i \tag{33}$$

which, combined with (32) yields, for some constant d,

$$\overline{m}^i \geq d(v(\widehat{x}_i) - v_i^* + \frac{\varpi}{(\overline{m}^i)^2}). \tag{34}$$

Since $var(\widehat{x}_i) - v_i^* \geq 0$, this implies $\overline{m}^i \geq (d\varpi)^{1/3}$, which further implies that the variance $v(\widehat{\xi}_i)$ is at most comparable to $\varpi^{1/3}$.

Step 2: upperbounds on $\overline{m}^i$. Let $r_i = \sum_{j \neq i} R_j m_j$ and $\widehat{y}_i = \sum_{k \neq i} A_{ik} y_k$. With transmission errors, we obtain:

$$\widehat{y}_i = (1 - r_i)p_i x_i + (1 - (1 - r_i)p_i)(\widehat{x}_i + \widehat{\xi}_i) + \overline{\xi}_i$$

where $\overline{\xi}_i = -p\xi_i + (1 - p_i) \sum_{j \neq i} R_j(1 - m_j)\xi_j$. Since p_i is set optimally by i, we have:

$$v(\widehat{y}_i) - v(y_i) \geq (r_i p_i)^2(\sigma_i^2 + v(\widehat{x}_i) + v(\widehat{\xi}_i)) - E\overline{\xi}_i - (1 - p_i)E\overline{\xi}_i\widehat{\xi}_i \geq cr_i^2 - \frac{d\varpi}{r_i}$$

for some constant c and d (independent of ϖ and m). Since $v(\widehat{y}_i) \leq \max v(y_k)$, the right-hand side cannot be positive for all i, so $r_{i_0} \leq (d\varpi/c)^{1/3}$ for some i_0. From step 1, we conclude that $\overline{m}^{i_0}$ and all m_j with $j \neq i_0$ are $O(\varpi^{1/3})$, and that m_{i_0} is thus *at least* $O(\varpi^{1/3})$.

It only remains to check that m_{i_0} cannot be large. From (33), $p_{i_0} \leq \pi_{i_0}^* + O(\varpi^{1/3})$, and since $p_{i_0} \geq \frac{1}{1 + r_{i_0}/m_{i_0}}$, we conclude that all m_i (and thus $\overline{m}^i$) are $O(\varpi^{1/3})$, which further implies that all variances $v(\widehat{\xi}_i)$ are $O(\varpi^{1/3})$.

These variances imply that $Ey_i^2 - v^*$ is at least $O(\varpi^{1/3})$. Ey_i^2 also rises because of inefficient weighting of seeds, but the loss is of the order of $(p_i - \pi_i^*)^2$, that is, $O(\varpi^{2/3})$, a significantly lower loss.■

Proof of Result 4: this follows from Lemma 7 since at equilibrium $\frac{\partial L_i}{\partial m_i} = 0$. ■

We turn to network comparisons.

Proof of Lemma 1 to 4: For each network, we write the equations determining long-run opinions. Through appropriate subsitutions, we derive these opinions as a function of seeds and errors.

For the complete network, to determine $h_i^c, \widehat{x}_i$ and $\widehat{\xi}_i^c$ we use

$$y_i = m_i\widetilde{x}_i + (1 - m_i)\overline{y}_{-i} \text{ and}$$

$$\overline{y}_{-i} = m\overline{\widetilde{x}}_{-i} + (1 - m)(\frac{1}{n-1}y_i + \frac{n-2}{n-1}\overline{y}_{-i})$$

where $\overline{\widetilde{x}}_{-i}$ (and $\overline{y}_{-i}$) refer to the mean modified seed (and opinion) of all players but i. Note $\overline{\widetilde{x}}_{-i} = \overline{x}_{-i} + \frac{1-m}{m}\overline{\xi}_{-i}$.

For the directed circle, we use $y_1 = m_1\widetilde{x}_1 + (1-m_1)y_2$ and repeatedly substitute $y_i = m\widetilde{x}_i + (1-m_i)y_{i+1}$ to obtain:

$$y_1 = m_1\widetilde{x}_1 + (1-m_1)(\sum_{k=0}^{N-2}(1-m)^k m\widetilde{x}_{k+2} + (1-m)^{N-1}y_1$$

To prove that $h_i^d < h_i^c$ for all $m \in (0,1)$, observe that the inequality holds for m close to 0 and that for any $m \in (0,1)$ that would satisfy $h_i^d(m) = h_i^c(m)$, we would have $\frac{\partial h_i^c}{\partial m}(m) > \frac{\partial h_i^d}{\partial m}(m)$, in contradiction with $h_i^d < h_i^c$ for m close to 0.

For the star network, we first determine h_0, $\widehat{x}_0$, and $\widehat{\xi}_0$ using

$$y_0 = m_0\widetilde{x}_0 + (1-m_0)\overline{y} \text{ and } \overline{y} = m\overline{\overline{x}} + (1-m)y_0$$

where $\overline{\overline{x}}$ (and $\overline{y}$) refer to the mean modified seed (and opinion) of *peripherical* players. Next, to determine h_i, $\widehat{x}_i$, and $\widehat{\xi}_i$, we use

$$\begin{aligned} y_0 &= m_0\widetilde{x}_0 + (1-m_0)(\frac{1}{n-1}y_i + (1-\frac{1}{n-1})\overline{y}_{-i}) \\ y_i &= m_i\widetilde{x}_i + (1-m_i)y_0 \\ \overline{y}_{-i} &= m\overline{\overline{x}}_{-i} + (1-m)y_0 \end{aligned}$$

where $\overline{\overline{x}}_{-i}$ (and $\overline{y}_{-i}$) refer to the mean modified seed (and opinion) of all peripherical players but i.■

Before turning to the proof of the Propositions, we prove:

Lemma 8: *For a fixed* n and *small* m, and for i.i.d random variables, $v(\psi(x_{-i})) = v(x_i)(\frac{1}{n-1} + cm^2)$ where $c = \frac{n(n-2)}{12(n-1)}$. *At the large* n *limit,* $v(\psi(x_{-i})) = v(x_i)\frac{m}{2}$.

Lemma 9: *For fixed* n*, small* m*, and* $m_i = m$, $h_i^c\frac{m}{1-m} = \frac{1}{n-1} + m$ *and* $h_i^d\frac{m}{1-m} = \frac{1}{n-1} + \mu m$ *where* $\mu = \frac{n}{2(n-1)}$

Lemma 8 implies that for fixed n, suboptimal weighing of independent seeds and errors in the directed circle are $O(m^2)$, while at the large n limit, the inefficiency is $O(m)$. Since $\mu < 1$ when $n > 2$, Lemma 9 will imply that in equilibrium, incentives to raise m are stronger in the directed circle.

Proof of Lemma 8 and 9: Let $r = 1-(1-m)^{n-1}$. We have:

$$v(\psi(x_{-i}))/v(x_i) = \frac{\sum_{k=0}^{n-2}(1-m)^{2k}}{(\sum_{k=0}^{n-2}(1-m)^k)^2} = \frac{(1-(1-m)^{2(n-1)})m^2}{m(2-m)r^2} = \frac{(2-r)m}{r(2-m)}$$

At the large n limit, $r = 1$, hence the desired result. For fixed n, compute $\Delta = v(\psi(x_{-i}))/v(x_i) - \frac{1}{n-1}$ considering terms of order up to 2 in m. We have $r = (n-1)m(1-\ell m)$ where $l = \frac{n-2}{2}(1 - \frac{(n-3)}{3}m)$, from which we obtain

$$(n-1)\Delta \simeq \frac{1-\frac{r}{2}}{(1-\frac{m}{2})(1-\ell m)} - 1 \simeq (\ell + \frac{1}{2})m - \frac{\ell m^2}{2} - \frac{r}{2} \simeq m^2\frac{n(n-2)}{12}$$

Regarding Lemma 9, the first statement is immediate. Regarding the directed circle, $h_i^d = 1/r$, so we have

$$(n-1)h_i^d\frac{m}{1-m} - 1 \simeq \frac{1}{(1-\ell m)(1-m)} - 1 \simeq (1+\ell)m \simeq \frac{n}{2}m■$$

Proof of Propositions 7 to 9.

(i) *For the complete network*, at $m_i = m$, $\frac{p_i}{1-p_i} - \frac{1}{n-1} = \frac{m}{1-m}$, so the equilibrium condition gives, omitting terms of higher order in m, $m = v(\widehat{\xi}_i^c)$. For independent errors, $v(\widehat{\xi}_i^c) \simeq \frac{1}{m^2}(\frac{1}{(n-1)^2} + \frac{1}{n-1})$, while for correlated errors, $v(\widehat{\xi}_i^c) = \frac{1}{m^2}(\frac{1}{n-1}+1)^2$, from which the expressions for m_c^* follow, as well as for the large n limit.

(ii) *For the directed network*, the equilibrium condition now gives, by Lemma 8 and 9, $\mu m \simeq v(\widehat{\xi}_i^d)$ for fixed n, and $m \simeq \frac{m}{2} + v(\widehat{\xi}_i^d)$ for the large n limit (because $h_i^d = 1$ and $v(\psi(x_{-i})) = \frac{m}{2}$). For fixed n, $v(\widehat{\xi}_i^d) \simeq v(\widehat{\xi}_i^c)$ (by Lemma 8 for independent errors, and because $\psi(\xi_{-i})) = \overline{\xi}_{-i}$ for correlated errors). It follows that $m_d^* = \mu^{-1/3} m_c^*$ and $\Delta_d^* = \mu m_d^* = \mu^{2/3}\Delta_c^*$ in both cases. At the large n limit, Lemma 8 implies $v(\psi(\xi_{-i})) = \frac{m}{2}\varpi$ for independent errors, so $m_d^* \simeq \Delta_d^* \simeq \varpi^{1/2}$. For correlated errors, $v(\widehat{\xi}_i^d) = v(\widehat{\xi}_i^c)$, so the equilibrium condition gives $m_d^* \simeq 2v(\widehat{\xi}_d^*) \simeq (2\varpi)^{1/3}$ and $L_d^* = \frac{m_d^*}{2} + v(\widehat{\xi}_d^*) \simeq m_d^* \simeq 2^{1/3} m_c^*$.

Note that at the large n limit, in contrast to full network where inefficiencies are solely driven by cumulated errors, the cumulated errors and the poor averaging of seeds *equally contribute* to the overall loss.

(iii) *For the star network*, the equilibrium condition for the central player gives $\rho_0 = \widehat{\varpi}_0^s$, and for a peripheral player it gives, for small m,

$$\frac{h_i^s m}{1-m} - \frac{1}{n-1} \simeq v(\widehat{x}_i^s) - \frac{1}{n-1} + v(\widehat{\xi}_i^s) \tag{35}$$

Omitting terms of order 2 in ρ_0 or m, we have $\frac{h_i^s m}{1-m} - \frac{1}{n-1} \simeq m + \frac{m-\rho_0}{n-1}$, $v(\widehat{x}_i^s) - \frac{1}{n-1} \simeq 0$, and $v(\widehat{\xi}_i^s) \simeq v(\widehat{\xi}_0^s) = \rho_0$, so (35) implies

$$m \simeq \rho_0 = \widehat{\varpi}_0^s. \tag{36}$$

For low ϖ, we thus have $m \simeq O(\varpi^{1/3})$, justifying the omission of terms of higher order. (36) also implies that incentives are approximately the same as in the complete network. So the inefficiency is entirely driven by cumulated errors: for independent errors, $\widehat{\varpi}_0^s \simeq (\varpi_0 + \frac{\varpi}{n-1})/m^2$, which yields $m_s^* \simeq (\varpi_0 + \frac{\varpi}{n-1})^{1/3} \simeq \widehat{\varpi}_s^*$, and for correlated errors, $m_s^* \simeq (\varpi_0 + \varpi)^{1/3} \simeq \widehat{\varpi}_s^*$. ■

Proof of Expression (15). Call p_j^i the weight that i puts on j and $\overline{R}_i$ the limit of R_i when m_{-i} tends to 0. It follows from Proposition 4 when all m are small, $(p_j^i/m_j)/(p_i/m_i) \simeq \overline{R}_{ij}$. To compute $\overline{R}_{ij}$, consider the case where $m_i = m$ for all i. Then $y = m\widetilde{x} + (1-m)A\widetilde{x} = \sum m(1-m)^k A^k \widetilde{x}$. Since all lines of A^k are close to ρ when k is large enough, $y_i \simeq \rho\widetilde{x}$ for all i, so $\overline{R}_{ij} = \rho_j/\rho_i$.

Proof of (generalized) Result 5: Rather than assuming that the central player is DG, we consider here a central player who uses her seed x_0 optimally to minimize the loss $v(\overline{y})$, given m. We have $\overline{y} = (1-m)y_0$ and $y_0 = m_0 x_0 + (1-m_0)(\overline{y} + \xi_0)$. This gives $\overline{y} = (1-m)(p_0 x_0 + (1-p_0)\frac{\xi_0}{m})$ where the central player controls p_0. The variance $v(\overline{y})$ is minimized for $\frac{p_0}{1-p_0} = \frac{\varpi_0}{m^2}$, and we get $v(\overline{y}) = (1-m)^2 \frac{\varpi_0/m^2}{1+\varpi_0/m^2}$. So long as $m >> (\varpi_0)^{1/2}$, we obtain $D \simeq \frac{4\varpi_0}{d}$ as for the DG case. Note that when $m \leq O(\varpi_0)^{1/2}$, cumulated errors are potentially huge and the (benevolent) central player mitigates them by choosing a large m_0: since she is benevolent, the loss cannot exceed 1 (the variance of her own seed).■

Appendix B (for on-line publication)

We first prove that the matrix $H \equiv \sum_{k\geq 0} B^k$ is well-defined (Lemma 1 and 2), and obtain Proposition 2 as a Corollary.

Lemma B1: *Consider any non-negative matrix $C = (c_{ij})_{ij}$ such that $\mu = \min_i(1 - \sum_j c_{ij}) > 0$. Then $I - C$ has an inverse $H \equiv \sum_{k\geq 0} C^k$, and for any X^0 and Y^0, $Y^t = X^0 + CY^{t-1}$ converges to HX^0.*

Lemma B2: *If $m_{i_0} > 0$, then for K large enough, $C = B^K$ satisfies the condition of Lemma 1, and $I - B$ has an inverse.*

Proof of Proposition 2: We just need to check that y^t converges. We iteratively substitute in (39) to get:

$$y^t = X^0 + Cy^{t-K}$$

where $X^0 = D\Gamma X$ with $D \equiv I + B + ... + B^{K-1}$, and $C = B^K$. By Lemma 2, Lemma 1 applies to C, so convergence of y^t is ensured.■

Proof of Lemma B1: Consider the matrix $H^t = (h^t_{ij})_{ij}$ defined recursively by $H^0 = I$ and $H^t = I + CH^{t-1}$. Let $z^t = \max_{ij} |h^t_{ij} - h^{t-1}_{ij}|$. We have $z^t \leq (1-\mu)z^{t-1}$, implying that H^t has a well-defined limit H, which satisfies $H \equiv \sum_{k\geq 0} C^k$. By construction, $(I-C)H = H(I-C) = I$, so $H = (I-C)^{-1}$. Similarly, defining $z^t = \max_i \left|Y^t_i - Y^{t-1}_i\right|$, we obtain that Y^t has a limit Y which satisfies $(I - C)Y = X^0$, implying $Y = HX^0$.■

Before turning to the proof of Lemma B2, we define *sequences*, *paths* and probabilities over paths associated with a probability matrix $A = (A_{ij})_{ij}$. For any sequence $q = (i_1, ..., i_K)$, we let $\pi^A(q) \equiv \prod_{k=1}^{K-1} A_{i_k, i_{k+1}}$, and for any set of sequences Q, we abuse notations and let $\pi^A(Q) = \sum_{q\in Q} \pi^A(q)$. We define a *path* as a sequence q for which $\pi^A(q) > 0$.

Denote by $Q^K_{i,j}$ the set of paths of length K from i to j, and Q^K_i the set of paths of length K that start from i. $Q^K_i = \cup_j Q^K_{i,j}$ and by construction, for any i, j

$$A^K_{ij} \equiv \pi^A(Q^K_{i,j}) \text{ and } \sum_{j\in N} A^K_{ij} = \pi^A(Q^K_i) = 1 \tag{37}$$

where A^K is the K^{th} power of matrix A.

Proof of Lemma B2: We consider A connected, that is, such that $A^k_{ij} > 0$ for all i, j, and consider $K \geq 2k$. Call $Q^{K,i_0}_i \subset Q^K_i$ the set of paths of length K that start from i (to some j) and go through i_0. For any such path, $\pi^B(q) \leq (1 - \underline{\gamma} m_{i_0})\pi^A(q)$. This implies

$$\sum_j C_{ij} \equiv \pi^B(Q^K_i) \leq (1 - \underline{\gamma} m_{i_0})\pi^A(Q^{K,i_0}_i) + \pi^A(Q^K_i \backslash Q^{K,i_0}_i) < 1$$

where the last inequality follows from (37) and Q^{K,i_0}_i non empty for $K \geq 2k$. This implies that C satisfies the condition of Lemma 1, hence $I - C$ has an inverse. Let $D \equiv I + B + ... + B^{K-1}$ and $H = (I - C)^{-1}D$. We have

$$\sum_{k\geq 0} B^k = \sum_{k\geq 0} C^k D = H,$$

so $H(I-B) = (I-B)H = I$ and $I - B$ also has an inverse.

Proof of Lemma 5: Using the recursive equation $y = X + (I-M)Ay$, and $X_i = m_i\widetilde{x}_i$ for $i \notin N^0$ and $X_i = \xi_i^0$ for $i \in N^0$, we define recursively the $(n, n-n_0)$ and (n, n_0) matrices P^t and Q^t as follows: for $i \notin N^0$, we let $P_i^t = m_i + (1-m_i)A_iP^{t-1}$ and $Q_i^t = (1-m_i)A_iQ^{t-1}$, and for $i \in N^0$, $P_i^t = A_iP^{t-1}$ and $Q_i^t = I + A_iQ^{t-1}$. Also we let $P_{ii}^1 = 1$ for $i \notin N^0$, and all other P_{ij}^1 and all Q_{ij}^1 equal to 0. By construction, $y = P\widetilde{x} + Q\xi_0$ where P and Q are the limit of P^t and Q^t respectively. Besides, by induction on t, each P^t is a probability matrix, hence so is the limit P.

Proof of Proposition 11: Lemma 1 to 4 provide cumulated error terms for processing errors. We use these Lemma to derive the cumulated error terms for expressing errors, using $\xi_i \equiv A_i\zeta^e$. For the directed circle, $\xi_i = \zeta_{i+1}^e$, so we immediately obtain $\widehat{\varpi}_i^p = \widehat{\varpi}_i^e$. For the full network, $\xi_i = \overline{\zeta}_{-i}^e$, so $\overline{\xi}_{-i} = \frac{1}{n-1}\zeta_i^e + \overline{\zeta}_{-i}^e(1-\frac{1}{n-1})$, which further implies $\widehat{\xi}_i^e = \frac{1-m}{m(n-1)}\zeta_i^e + \frac{1}{m}\overline{\zeta}_{-i}^e$, hence the desired comparison. For the star network, $\xi_i = \zeta_0^e$ and $\xi_0 = \overline{\zeta}^e$, so $\widehat{\xi}_0^e = \frac{1-m}{m}\zeta_0^e + \frac{1}{m}\overline{\zeta}^e$, hence the desired comparison. Note that, for the cumulated errors faced by peripherical players, one can compute $\widehat{\varpi}_i^p$ and $\widehat{\varpi}_i^e$ for fixed ρ_0. In equilibrium, for small ϖ, omitting terms of higher orders, one can check that in equilibrium, $\widehat{\varpi}^p - \widehat{\varpi}^e \simeq (1-\frac{1}{n^2})\varpi/m^*$ with $m^* \simeq ((1+1/n)\varpi)^{1/3}$.

The case with noise.

For any fixed (x,ξ), we define the expected opinion at t, $\overline{y}_i^t = Ey_i^t$ and the vector of expected opinions $\overline{y}^t = (\overline{y}_i^t)_i$. We further define $\eta^t = y^t - \overline{y}^t$, $w_{ij}^t = E\eta_i^t\eta_j^t$ and the vector of covariances $w^t = (w_{ij}^t)_{ij}$.

We define the N^2 vector Λ with $\Lambda_{ij} = 0$ if $i \neq j$, $\Lambda_{ii} = (\gamma_i(1-m_i))^2\varpi^0$ and $\overline{B}$ the $(N^2 \times N^2)$ matrix where $\overline{B}_{ij}$ is the row vector $(\overline{B}_{ij,hk})_{hk}$ with $\overline{B}_{ij,hk} = B_{ih}B_{jk}$.

For any fixed (x,ξ), we define the expected opinion at t, $\overline{y}_i^t = Ey_i^t$ and the vector of expected opinions $\overline{y}^t = (\overline{y}_i^t)_i$. We further define $\eta^t = y^t - \overline{y}^t$, $w_{ij}^t = E\eta_i^t\eta_j^t$ and the vector of covariances $w^t = (w_{ij}^t)_{ij}$.

The evolution of opinions and expected opinions (given x, ξ) follows

$$y^t = \Gamma(X + (I-M)\nu^t) + By^{t-1} \tag{38}$$

$$\overline{y}^t = \Gamma X + B\overline{y}^{t-1}, \tag{39}$$

from which we obtain:

$$\eta^t = \Gamma(I-M)\nu^t + B\eta^{t-1}$$

Since the ν_i^t are independent random variables, the evolution of the vector of covariances follows:

$$w^t = \Lambda + \overline{B}w^{t-1} \tag{40}$$

The evolution of $\overline{y}^t$ coincides with the case where there is no noise. Lemma B3 below extends Lemma B2, showing that $\overline{H} \equiv \sum_{k\geq 0}\overline{B}^k$ (or the inverse $(I-\overline{B})^{-1}$) are well-defined, which implies that w^t has a well-defined limit

$$w = \overline{H}\Lambda, \tag{41}$$

Lemma B3: *For K large enough, $\overline{B}^K$ satisfies the condition of Lemma 1.*

Proof of Lemma B3. We extend the notion of sequences and paths to pairs $ij \in N^2$ (rather than individuals). For any sequence of pairs $\overline{q} = (i_1 j_1, ..., i_K j_K)$ (or equivalently, any pair of sequences $\overline{q} = (q^1, q^2) = ((i_1, ..., i_K), (j_1, ..., j_K)))$ and any matrix $A = (A_{ij})_{ij}$, and we let $\overline{\pi}^A(\overline{q}) = \pi^A(q^1)\pi^A(q^2)$. We define a path $\overline{q}$ as a sequence such that $\overline{\pi}^A(\overline{q}) > 0$.

We apply the argument of Lemma B2 to paths $\overline{q}$ of pairs rather than paths q of individuals. Let $\overline{C} = \overline{B}^K$. Call $\overline{Q}_{ij}^K$ the set of paths $\overline{q} = (q^1, q^2)$ of length K that start from ij (to some hk), $\overline{Q}_i^{K,i_0}$ those for which q^1 goes through i_0. We have

$$\sum_{hk} \overline{C}_{ij,hk} \equiv \overline{\pi}^B(\overline{Q}_{ij}^K) \leq (1 - \underline{\gamma} m_{i_0})\overline{\pi}^A(\overline{Q}_i^{K,i_0}) + \overline{\pi}^A(\overline{Q}_i^K \backslash \overline{Q}_i^{K,i_0}) < 1$$

hence $\overline{C}$ satisfies the condition of Lemma 1, $I - \overline{C}$ has an inverse, and so does $I - \overline{B}$.■

Proof of Proposition 12.

(i) Let $\overline{C} = \overline{B}^K$ and $\overline{D} = I + \overline{B} + ... + \overline{B}^{K-1}$. Repeated substitutions in (40) yield

$$w^t = \Lambda^0 + \overline{C} w^{t-K}$$

where $\Lambda^0 = \overline{D}\Lambda$. By Lemma 2b, Lemma 1 applies to $\overline{C}$, so convergence of w^t to w is ensured.

(ii) We bound the loss V_i induced by the idiosyncratic errors. Recall

$$\eta_i^t = \gamma_i(1 - m_i)\nu_i^t + (1 - \gamma_i)\eta_i^{t-1} + \gamma_i(1 - m_i)A_i\eta^{t-1}$$

This implies that for any $p \in \Delta_n$, there exists $q \in \Delta_n$ such that:

$$p.\eta^t = q.\eta^{t-1} + \sum_i \gamma_i(1 - m_i)p_i\nu_i^t \text{ and } \sum_i q_i \geq 1 - \underline{m} \tag{42}$$

Define $\underline{V}^t = \min_{p\in\Delta_n} var(p.\eta^t)$. Note that $V_i^t \geq \underline{V}^t$. Since $var(q.\eta^{t-1}) \geq (1 - \underline{m})^2\underline{V}^{t-1}$, Equality (42) implies $\underline{V}^t \geq (1 - \underline{m})^2\underline{V}^{t-1} + \frac{1}{n}\underline{\gamma}^2(1 - \underline{m})^2\varpi^0$, which yields the desired lower bound at the limit.

(iii) We now re-examine Result 2. We consider the effect of m_i on the vector of covariances w where $w_{jk} = \lim E(y_j^t - \overline{y}_j^t)(y_k^t - \overline{y}_k^t)$. Recall $w = \Lambda + \overline{B}w$. Since Λ and $\overline{B}$ are non-increasing in m_i and Λ_{ii} is strictly decreasing in m_i, w_{ii} strictly decreases with m_i, and w is non-increasing in m_i. Combining all steps, over the range $m_i < \underline{m}$, $L_i = \overline{L}_i + w_{ii}$ strictly decreases with m_i, and $\sum_k L_k$ also strictly decreases with m_i.■

Proof of Result 6. *In addition to item (i) and (ii), we shall prove the following statement:* (iii) *If the lower bound $\underline{\gamma}$ on the choice set is sufficiently low and $\gamma_i = \underline{\gamma}$, $V_i \leq 1/ \mid \log \underline{\gamma} \mid$ for all $m \geq \underline{m}$ and γ within the choice set.*

Let $\overline{\gamma} = \max \gamma_i$ and recall:

$$w_{ij} = \sum_{h,k} B_{ih}B_{jk}w_{hk} + \Lambda_{ij} \tag{43}$$

where $\Lambda_{ij} = 0$ if $i \neq j$ and $\Lambda_{ii} = (1 - m_i)^2(\gamma_i)^2\varpi^0$, and $B_{ii} = 1 - \gamma_i$, $B_{ij} = \gamma_i A_{ij}(1 - m_i)$.

The proof starts by proving item (i), that is, computing a uniform upper bound on all w_{ij} of the form (see step 1)

$$w_{ij} \leq c\overline{\gamma} \tag{44}$$

To prove (ii), we define $\widehat{w} = (w_{ij})_j$ as the vector of covariances involving i, and show that there exists a matrix C for which $\sum_k C_{jk} \leq 1$ for all j and such that

$$\widehat{w} \leq (1-\underline{m})C\widehat{w} + \Gamma \tag{45}$$

where $\Gamma_j \leq dp_{ij}$ for some d, with $p_{ij} = \gamma_i/(\gamma_i+\gamma_j)$. This in turn implies that $\max_j w_{ij} \leq \max_j \Gamma_i/\underline{m}$, which will prove (ii) (see step 3).

Finally, to prove (iii), we consider two cases. Either $\overline{\gamma}$ is "small" and (44) applies, or we can separate individuals into a subgroup J where all have a small γ_j, and the rest of them with significantly larger γ_j. In the latter case, we redefine $\widehat{w} = (w_{jk})_{j\in J,k}$ as the vector of covariances involving some $j \in J$, and obtain inequality (45) with $\Gamma_{jk} \leq dp_{jk}$ for $k \notin J$ and $\Gamma_{jk} \leq d\gamma_j$ for $k \in J$, for some d. By definition of J, all γ_j and p_{jk} are small, and all Γ_{jk} are thus small, which will prove (iii). Details are below.

Step 1 (item (i)) $w_{ij} \leq c\overline{\gamma}$ with $c = \varpi^0/\underline{m}$.

Let $\overline{V} = \max_i w_{ii}$ and $\overline{w} = \max_{i,j\neq i} w_{ij}$ and $\overline{w} = \max w_i$. For all $j \neq i$, w_{ij} is a weighted average between all $w_{h,k}$ and 0, so $w_{ij} < \max(\overline{w}, \overline{V})$, hence $\overline{w} < \max(\overline{w}, \overline{V})$, which thus implies $\overline{w} \leq \overline{V}$. Consider i that achieves $\overline{V}$. Since $\sum_{h,k} B_{ih}B_{ik} = (1-\gamma_i m_i)^2$, we have:

$$\overline{V} = w_{ii} \leq (1-\gamma_i m_i)^2\overline{V} + \gamma_i^2(1-m_i)^2\varpi^0 \text{ hence}$$
$$\overline{V} \leq \frac{\gamma_i(1-m_i)^2}{m_i}\varpi^0 \leq \frac{\varpi^0\overline{\gamma}}{\underline{m}}$$

Step 2. Let $p_{ij} = \gamma_i/(\gamma_i+\gamma_j)$ and $\overline{v} = 2(c\overline{\gamma}+\omega_0)$. We have:

$$w_{ii} \leq \gamma_i p_{ii}\overline{v} + (1-\underline{m})\sum_k A_{ik}w_{ik} \tag{46}$$
$$w_{ij} \leq \gamma_j p_{ij}\overline{v} + (1-\underline{m})(p_{ij}\sum_k A_{ik}w_{kj} + p_{ji}\sum_k A_{jk}w_{ik}) \tag{47}$$

These inequalities are obtained by solving for w_{ij} in equation (43), that is, we write

$$(1-B_{ii}B_{jj})w_{ij} = \Gamma_{ij} + \sum_{k\neq i} B_{ii}B_{jk}w_{ik} + \sum_{k\neq i} B_{jj}B_{ik}w_{kj} + \sum_{k\neq i,h\neq j} B_{jk}B_{ih}w_{kj}.$$

Observing that $2B_{ii}B_{ik}/(1-B_{ii}B_{jj}) \leq (1-m_i)A_{jk}$, $B_{ii}B_{jk}/(1-B_{ii}B_{jj}) \leq (1-m_j)p_{ji}A_{jk}$, and $B_{jk}B_{ih}/(1-B_{ii}B_{jj}) \leq 2\gamma_j p_{ij}A_{jk}A_{ih}$ and $\Gamma_{ii}/(1-B_{ii}B_{jj}) \leq \gamma_i\omega^0$ yields (46-47).

Step 3 (item (ii)). It is immediate from (46-47) that (45) holds with $C_{jk} \equiv A_{jk}$ and $\Gamma_j = p_{ij}\gamma_j\overline{v} + p_{ij}c\overline{\gamma} \leq p_{ij}\overline{\gamma}(\overline{v}+c) \leq d\gamma_i$ for all j, for some d, which permits to conclude that $\widehat{w} \leq d\gamma_i/\underline{m}$.

Step 4 (item (iii)). Let $\varepsilon = \frac{1}{K|Log\underline{\gamma}|}$ with $K = 5\varpi^0/\underline{m}^2$ and set $\gamma_i = \underline{\gamma}$. Let us reorder individuals by increasing order of γ_j. Consider first the case where $\gamma_{j+1} \leq \gamma_j/\varepsilon$ for all $j = 1, ..., N-1$. Then $\overline{\gamma} < \underline{\gamma}/\varepsilon^{N-1}$, and for $\underline{\gamma}$ small enough, $\underline{\gamma}/\varepsilon^{N-1} < \varepsilon$, so $V_i \leq c\varepsilon < 1/\mid Log\underline{\gamma}\mid$.

Otherwise, there exists j_0 such that $\gamma_j \leq \underline{\gamma}/\varepsilon^{j_0-1}$ for all $j \in J$, and $\gamma_k > \gamma_j/\varepsilon$ for all $k \notin J$ and $j \in J$. It is immediate from (46-47) that (45) holds with Γ such that, for any $j \in J$,

$$\Gamma_{jk} = \gamma_j\overline{v} \text{ if } k \in J \text{ and}$$
$$\Gamma_{jk} = \gamma_j\overline{v} + p_{jk}\sum_{h\notin J} A_{jh}w_{hk} \text{ if } k \notin J$$

By definition of J, for all $j \in J$, $\gamma_j \leq \underline{\gamma}/\varepsilon^{N-1} < \varepsilon$ and for all $k \notin J$, $p_{jk} \leq \varepsilon$, which further that all Γ_{jk} are bounded by $\varepsilon(\overline{v} + c) \leq \underline{m}/ \mid Log\underline{\gamma} \mid$, which concludes the proof.■

Proof of Proposition 13:

For fixed x, ξ, let $Y_i^t = (y_i^{t-k})_{k=0,..,K}$ denote the column vector of i's past recent opinions, and $Y^t = (Y_i^t)_i$. One can write $Y^t = X + BY^{t-1}$. Y^t converges for standard reasons, to some uniquely defined Y. Consider now the vector y solution to

$$y_i = m_i x_i + (1 - m_i)A_i(y + \xi_i)$$

and let $Y_i = (y_i, ..., y_i)$ and $Y = (Y_i)_i$. By construction, under this profile of opinions, it does not matter when i heard from j because opinions do not change. Y thus solves $Y = X + BY$ and it coincides with Y. The limit expected opinion vector under FJ is thus independent of the communication protocol.

Coarse communication:

Recall f is the fraction of agents choosing $a = 0$, and call $y = \phi(f)$ the associated "population opinion". We now consider two cases:

Case 1: $m = 0$. Set $\xi > 0$ and assume $f > 0$. Each makes an inference z_i at least equal to $y + \xi$ regarding neighbors' opinions, so eventually, under DG, each player of type b_i may only report 0 if $b_i + y + \xi < 0$. Under the large number approximation, a fraction at most equal to $f' = h(y + \xi) < f$ reports 0, hence the fraction of agents reporting 0 eventually vanishes.

Case 2: m small. When $m > 0$, agents with signal x_i believe the state is $mx_i+(1-m)(y+\xi)$, which generates, under the large number approximation, a fraction $f = Eh(mx_i+(1-m)(y+\xi))$ choosing $a = 0$. The long-run opinion y thus solves

$$y = h^{-1}(Eh(m(\theta + \delta_i) + (1 - m)(y + \xi)))$$

Call $\widehat{\xi} = y - \theta$ the resulting population estimation error. When m is small, h is locally linear, so, since $E\delta_i = 0$, $y \simeq h^{-1}h(m(\theta + \frac{1-m}{m}\xi) + (1 - m)y)$, which implies $\widehat{\xi} \simeq \frac{1-m}{m}\xi$.

Assume now that player chooses m_i while others choose m. For player i, the estimation error is $\Delta_i \equiv m_i\delta_i + (1 - m_i)(\widehat{\xi} + \xi) \simeq m_i\delta_i + (1 - m_i)\frac{\xi}{m}$. Assuming that θ is drawn from a flat distribution with large support, the expected loss $L_i(\Delta)$ from estimating θ with an error Δ_i is quadratic in Δ_i and independent of b_i,[58] so $L(\Delta)$ is proportional to the variance of the error that i makes. To minimize the variance of Δ_i, player i sets $m_i = \frac{\varpi}{m^2}$, so in equilibrium $m^* = \varpi^{1/3}$.

Regarding the social optimum, when all choose m, the estimation error is $m\delta_i + (1 - m)\frac{\xi}{m}$. For ϖ small, the variance of this error is minimized for $m \simeq (2\varpi^{1/4})$.

[58] When $\Delta_i > 0$, $L(\Delta_i) = \int_{-\Delta_i - b_i}^{-b_i} -(\theta + b_i)d\theta = \frac{\Delta^2}{2}$.