

Information Aggregation under (not so) Naive Learning

Abhijit Banerjee and Olivier Compte

September 30th 2020

Abstract

We explore a model of non-Bayesian information aggregation in networks. Agents non-cooperatively choose an aggregation rule from the Friedkin-Johnsen (FJ) class to maximize private payoffs in the presence of noise in information transmission. We characterize rules that get chosen. The well-known DeGroot (DG) rule, nested in FJ, is never chosen – all near optimal rules have individuals putting enough weight on their own initial opinion in every period unlike in DG. This precludes full consensus even in the long-run but ensures better information aggregation. This trade-off extends to a wider range of environments and a broader class of rules.

1 Introduction

Living in a world dominated by Facebook, Twitter and their ilk, it is hard to avoid wondering about the quality of information aggregation on networks. We constantly get information passed onto us by others and in turn pass it on to our network neighbors. What are properties of such a process? How well is information aggregated?

The literature that theoretically explores these questions typically takes one of two routes: a Bayesian route, in which agents make perfect inferences as if they knew the process that generates signals through the network; and a non-Bayesian route, which is in part reacting to the very demanding assumptions about information processing that the Bayesian approach requires. In this alternative route, one postulates a simple rule that individuals use to aggregate own and neighbors opinions. One such rule is the DeGroot (DG) rule, which updates one's current opinion by computing an average between one's and neighbors' most recent opinions. This updating is justified by

arguments that this is like Bayesian in the sense that it approximates the Bayesian decision rule in certain very specific settings.

This paper proposes a third route. It considers a *class* of simple aggregation rules and postulates that, within this class, each individual selects its favorite aggregation rule, based on its instrumental value.¹ The class we consider is the class of Friedkin-Johnsen (FJ) models (Friedkin and Johnsen (1990)). It nests DG, but it allows each individual to keep putting some weight on their own initial opinion. Our motivation is the following. Rather than emphasizing some similarity to Bayesian decision rules, we justify the use of rules by individuals by assuming that individuals will/should favor rules with higher instrumental value. To this end, we introduce some discretion in the choice of rule and assume that each individual selects one within a class of “natural” rules according to how well it performs (for her) information aggregation in the long-run.²

We explore this third route in a standard setup (i.e., each individual initially gets a noisy signal correlated with some underlying state of the world) in which information transmission is assumed to be noisy. The class of Friedkin-Johnsen (FJ) models can formally be written as

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

where y_i^t is i 's belief in period t , x_i is the initial signal that i received and

$$z_i^t = \frac{1}{|N_i|} \sum_{j \in N_i} y_j^t + \varepsilon_i^t.$$

is the average report received by i from his neighbors (denoted by N_i) plus any noise in the transmission (or reception) of that signal. When the weight m_i is 0, individual i is using a DG rule.³ One key assumption is that there is noise in communication of signals (or alternately, there are biases in the reports individuals get from others). Also, as should be evident from the above expression, we impose the restriction that the decision-maker treats signals from everyone else symmetrically. This is to reduce the dimensionality of the rule choice problem:⁴ only m_i (and possibly γ_i) will be subject

¹This is in the spirit of the approach advocated in Compte and Postlewaite (2018) to model mildly sophisticated agents.

²The limitation to a class of rules is key. With no limitation, the optimal rule would be the Bayesian rule, as the Bayesian rule is the individually optimal way to process signals among all possible signal processing rules.

³Throughout our analysis, we shall assume that all γ_i are bounded away from 0.

⁴Symmetry is for convenience. We will actually allow for non-symmetric weights. The main simplification is that we assume that relative weights across neighbors are fixed, not subject to optimization.

to choice.

We then assume that the initial signals x_i are correlated with some underlying given state of the world θ and that the individuals have a utility function which is a decreasing function of distance between the state θ and their long-run belief about θ , formed after the signal exchange process has had a long enough run. Given these preferences and the rules chosen by the other participants, they each choose the rule within the permitted set of rules (e.g., the FJ class where each m_i belong to $[0, 1]$) that maximizes this utility function subject to the uncertainty about the state of the world and the noise in communication: we are interested in the Nash equilibrium of this rule choice game.

Our main results and the logic behind them are as follows. We start by observing that in this environment, DG has the undesirable property that the variance of every decision-maker's belief grows without bound in the presence of any noise in communication (Proposition 1).

This points to a specific sense in which DG is fragile: when $\gamma_i > 0$, agent i puts less than one hundred percent weight on his most recent belief, so the weight that agent i directly puts on his own initial signal is going to 0. Without noise in transmission, this is at least partly offset by the weight agent i puts on reports from others that themselves contain agent i 's initial signal, which is why the influence of initial signals on current beliefs does not dissipate. In other words, the agent holds on to his own signal only through the feed-back from others. The problem is that when transmission is noisy, you only get the feedback at the cost of some extra noise in every round. Given that the initial signals enter only at the beginning, the noise keeps coming in every period, it is no wonder the noise comes to dominate. This contrasts with DG in a noiseless environment, which has been shown by Golub-Jackson (2010) to have the attractive property that, under some restrictions on network structure and weights on neighbors, learning converges to perfect information aggregation (indeed within the FJ class DG is the only rule with this property).

This intuition suggests that a countervailing force would be to allow the initial signal to come in with some weight every period, which is what FJ does. We next show that with FJ rules, every decision-maker's average belief as well as its variance converges (Proposition 2). To see exactly why, note that the prime mover of beliefs is a change in the beliefs of an agent's neighbors, who in turn change their beliefs because of changes in the beliefs of their own neighbors. Potentially, a single belief change by one agent may thus induce a sequence of changes in the beliefs of others, hence echo effects

that could swamp all players' beliefs. However, with (FJ) rules where agent's put some weight on their own initial signals, these echo effects are on average dampened at each round and this dampening guarantees convergence.

In fact we show that convergence in FG requires only one agent to put positive weight on his initial signal. And we also show that if the weight on initial signals is small for *all* players, the variance of the long term outcome will tend to be very large.

These observations suggest that the presence of noise tends to favor the use of FG type rules with *significant* weight on the own signal. In fact Proposition 3 shows that rules where m_i is too low are dominated. One step of the argument is obvious: essentially, if every other agent chooses DG, then choosing FG is the only way to stop the variance from blowing up. Another step is that if some players choose FG, then those who stick with DG become followers: their initial opinions disappear from current opinions in the long-run, because they face players that constantly feeds in their own initial signals. Long-run opinions are a weighted average of the signals of those who are putting positive weight on their signals. Starting at the point where an individual is putting zero weight on his own signal, the effect of increasing that weight slightly always reduces the ex ante variance of his final opinion (because the contribution of an additional independent signal is always reducing ex ante variance).

In fact we can go a step further. Restricting attention to the simpler case where noise in transmission is modeled as a systematic bias drawn at the start of the interaction, long-run outcomes are independent of γ_i and we can characterize the equilibrium where all the players are non-cooperatively choosing the weight m_i to put on their own initial signal (Proposition 4). Moreover the equilibrium always involves putting too little weight on one's own signal relative to the social optimum (Proposition 5). In equilibrium each player is trying too hard to free ride on the collective wisdom of the others, not fully taking into account the fact that setting m_i too low increases the correlation across opinions.⁵

We then return to the question of rule choice when idiosyncratic noise in transmission is also present. In the presence of idiosyncratic noise, FJ rules with large γ_i generates a lot of bouncing around because i is reacting every period to the latest reports from others and each of those comes with

⁵Player i takes into account the correlation in opinions when setting own m_i optimally, but she does not fully take into account the consequence for player j , in particular, the positive correlation between the sources (including j 's own signal) that directly contribute to j 's opinion independently of i 's mediation, and player i 's opinion which also contributes to j 's opinion and partially contains these same sources.

a different piece of noise in every period. i can limit the churn by putting some weight on past beliefs (i.e., a smaller γ_i). Said differently, the problem with FJ rules with large γ_i is that it is missing the stabilizing benefit of updating the belief slowly (a property that DG also offers). This is precisely what Proposition 6 addresses.

In the penultimate section of the paper we examine other potential sources of noise. In particular we have so far assumed that the noise comes from errors in transmission. Our first exercises show that whether the noise is deliberately slanted in a particular direction or results from heterogeneous preferences (or biased perceptions of others' preferences), nothing essential changes other than a further shift towards reliance on one's own initial signal. The next sub-section shows that another key difference between FJ and DG rules comes from the way they deal with uncertainty over the exact communication protocol – for example not everyone may speak in every period. We show by example that the outcome from using DG rules is sensitive to who speaks when, even in the absence of noise, whereas the outcomes from FJ rules are always independent of the communication protocol, whether or not there is noise in communication.

We next turn to the possibility that communication may be coarse in the sense that each party only reports their current best guess about which of two actions is preferable (or equivalently, which among two states of the world is the one we are in). In this setting, a systematic error in interpreting guesses by neighbors makes the long-run outcome from a DG-type rule entirely insensitive to the actual state of the world (Frick et al. (2019) report a related result for Bayesian-type rule), but this is not true for FJ-type rules.

In the next sub-section, we move away from the linear aggregation rule assumption. We introduce a class of non-linear aggregation rules that remain in the spirit of DG rules, and show that the non-linearity may actually exacerbate the long-run drifting of beliefs, suggesting that linear rules are the best case scenario for DeGroot rules.

The final sub-section discusses non-stationary rules. It is easy to see that the way to avoid double-counting the same signal is to stop listening at some point, though different people in the network may need to stop at different times, depending on how the network and communication are structured. This is easy to implement in the case of communication between two people, but in large networks it is much less clear – one reason to continue to listen to one's neighbor is that they may be a conduit for new information from farther out in the network. Even in simple networks however, one issue is in determining the degree to which one's information is already incorporated in one's neighbors opinion. When messages get lost with positive probabil-

ity, the later issue becomes central, and non-stationary rules with strategic listening will typically fail to address it, while FJ-type rules will be immune to these errors.

This paper is related to and inspired by the recent upsurge of interest in the social learning in environments that are not fully Bayesian. Eyster and Rabin (2010), Sethi and Yildiz (2012, 2016, 2019), Jadbabie et al. (2012) and Gentzkow et al. (2018), among others, explore the implications of applying Bayes rule when the underlying information structure is misspecified, as does the previously mentioned paper by Frick et al. (2019).

To end this section we provide a brief review of the closely related literature on Bayesian and non-Bayesian learning rules.

Acemoglu et al. (2011) study of social learning on general networks within the Bayesian paradigm.⁶ They provide conditions under which information is perfectly aggregated as the network grows very large. It is not clear however how realistic the assumption of full Bayesian decision-making is in their setting: a Bayesian needs to think all possible sequences of signals that could be received as a function of the underlying state and of all the possible pathways through which each observed sequence of signals could have reached them, and there is a very large number of such pathways.⁷

For this reason, the literature has mostly taken a different route. The idea is to assume that individuals only get to observe a low-dimensional summary of the information their network neighbors have—usually the neighbors’ summary belief about the state of the world or their actions. Then the models postulate a simple rule that individuals use to aggregate these signals they get from their neighbors, while not worrying about the signals that their neighbors in turn received (which is typically different from what a Bayesian would do) and study the consequences, at the level of the network as a whole, of individual using these kinds of simple rules rather than Bayes rule.

As mentioned, one well-known such rule is the one proposed by DeGroot (1974) and brought into economics literature by DeMarzo et al. (2003), which the literature calls the DeGroot rule (henceforth DG). The literature motivates the use of DG by underlying its close connection with Bayes rule

⁶Bayesian models of social learning go back to Banerjee 1992, and Bhikchandani, Hirshleifer and Welch 1992 but the network structure they study is extremely particular.

⁷Alatas et al. (2016) remark ”To give a sense of scale to this computation, note that enumerating all such paths is # P-complete and a random graph with n nodes and edges with probability p_n has an expected number of paths between nodes 1 and n given by $(n - 2)!p_n^{n-1}e(1 + o(1))$, which is potentially an enormous number (Roberts and Kroese (2007)).”

in simple environments. For example, DeMarzo et al. (2003) and Golub and Jackson (2010) justify DG by arguing that it coincides with the Bayesian rule in the static case (i.e., when information is transmitted just once). Bala and Goyal (1998) in their seminal paper on learning on networks describe their agents as "bounded Bayesians". More recently Levy and Razin (2015) have introduced the idea of Bayesian Peer Influence Paradigm to capture the idea of an almost Bayesian aggregation rule. Likewise, in a setting where the state of the world that is being learnt changes, Alatas et al. (2016) propose an aggregation rule that they justify by suggesting that it mimics full Bayesianism for certain very specific network structures. A different approach is taken by Molavi et al. (2018), who provide an axiomatic justification for DG/DG-like (e.g. Log-linear learning) rules.

In settings where the set of choices are discrete, and only the choices (and not the underlying beliefs) are observed by network neighbors, the class of "natural" rules are somewhat different: for example there are the infection models, studied in Jackson (2008) among (many) others or the related class of models studied by Ellison and Fudenberg (1993, 1995). We discuss the applicability of our results to these and other settings in the concluding section.

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, at any date t , each agent i in the network has an opinion that can be represented as a real number.⁸ We consider a class of updating rules due to Friedkin and Johnsen (1990) (henceforth FJ), in which player i 's current opinion y_i^t is a convex combination of his initial opinion x_i , his most recent opinion y_i^{t-1} and some summary perception z_i^{t-1} of his neighbors' opinions. Formally, this can be written as

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

where

$$z_i^t = A_i \cdot y^t + \varepsilon_i^t \quad (1)$$

⁸This opinion can be interpreted as a point-belief about some underlying state, which will eventually be used to undertake an action.

where y^t is the vector of all opinions at t , A_i is a row vector whose j th element A_{ij} is such that $\sum_j A_{ij} = 1$ and ε_i^t represents noise in transmission or reception. z_i^t is meant to be some average of the opinions of i 's neighbors (denoted N_i), so the presumption is that $A_{ij} > 0$ for $j \in N_i$, plus some noise in transmission or reception.

When $m_i = 0$, the rule corresponds to the well-studied DeGroot rule (DG). When $m_i > 0$, the updating process works like DG, but the perception of other's opinions is corrected using the decision-maker's own initial opinion as a perpetual seed. 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.^{9,10}

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 of the A_i 's is connected in the sense that for some positive integer k , $A_{ij}^k > 0$ for all i, j . In other words everyone is within a finite number of steps of the rest.

Finally, before proceeding, it is useful to define a simplified version of the rule FJ , where $\gamma_i = 1$. We refer to it at SFJ :

$$y_i^t = m_i x_i + (1 - m_i) z_i^{t-1} \quad (\text{SFJ})$$

One can think of SFJ as a process that works as FJ, except that agents do not attempt to smooth out variations in their own opinion. In the absence of idiosyncratic shocks on the perception of other's opinions (see details below), SFJ and FJ will generate identical long-run opinions.

Note that all the rules considered are stationary, in the sense that the weighting parameters m_i and γ_i do not vary over time. We are interested in these rules not only because they have been studied in the literature, but also because we see them as plausible ways by agents might incorporate others' opinions into their current opinion. Of course, with some knowledge of the structure of the network, and the process by which information gets incorporated, an agent might want to adjust the weights over time. We

⁹In fact, as mentioned already, the one obvious attraction of DG has been its quasi-Bayesian flavor. If y_i^{t-1} is viewed as a summary statistic of past signals, and z_i^{t-1} as a new signal, then the linear weight γ_i can be seen as the optimal weighting strategy of a Bayesian that would aim to reduce the variance of one's opinion. Of course, over time, a Bayesian would typically not keep that parameter constant, as the relative informative content of recent own and recent other's opinion will in general not be constant.

¹⁰Note that although the expression (FJ) encompasses the DG rule, we shall refer to FJ as a rule for which $m_i > 0$.

shall discuss in Section 6.6 the risks that such elaborate adjustments be misguided, in particular when there is randomness over the dates at which communication takes place.

We have also imposed the assumption that everyone operates on the same time schedule: periods are so defined that everyone changes their opinion once every period and everyone else gets 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 6.3.

2.2 Noise in opinion sharing

The noise term ε_i^t is an important ingredient of our model, meant to capture some imperfection in transmission.¹¹ We assume that there is a single piece of noise in what each individual “hears” that aggregates all the different sources of noise. There may be noise that results from each individual being imprecise in expressing his or her opinion, or from an error in hearing or interpretation.

We assume that the noise has two components:

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

The term ξ_i is a *persistent* component realized at the start of the process, that applies for the duration of the updating process.¹² The term ν_i^t is an *idiosyncratic* component drawn independently across agents and time. We interpret ξ_i as a systematic bias that slants how opinions of others are perceived. For convenience, we assume that all noise terms are homogenous across players and unbiased (that is, $E\xi_i = E\nu_i^t = 0$).¹³ We let $\varpi_i = \text{var}(\xi_i)$ and $\varpi_0 = \text{var}(\nu_i^t)$ and assume that:

$$\varpi_i = \varpi > 0$$

¹¹There has been several recent attempts to introduce noisy 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 is not communicated to the next neighbor). In Frick et al. (2019), agents communicate through a choice of action $a \in \{0, 1\}$ correlated with an unknown underlying state, and they make errors in interpreting these actions because they have an erroneous model of the preferences of others.

¹²One 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.

¹³The assumption $E\nu_i^t = 0$ is without loss of generality. We shall come back to the case where $E\xi_i \neq 0$ in the Discussion Section.

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 his objective on average across realizations of the underlying state of the world and the transmission errors. 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. In particular we have in mind examining whether there are forces away from DG rules, and whether private and social incentives differ.

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 σ_i^2 . For convenience, we assume that $\sigma_i = 1$ 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 , ξ_i and η_{it} , for all i and t):

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

Define $\delta = (\delta_1, \dots, \delta_n)$, $\xi = (\xi_1, \dots, \xi_n)$ and $\nu_s = (\nu_{1s}, \dots, \nu_{ns})$ for all s . Now given the set of updating rules that we consider, it will become evident that

$$y_i^t = b_i^t \delta + c_i^t \xi + \sum_{s=1}^t d_{is}^t \nu_s + \theta$$

for some non-negative vectors b_i^t , c_i^t and $\{d_{is}^t\}_{s=1}^t$.¹⁴ It follows that

$$L_i^t = E[b_i^t \delta + c_i^t \xi + \sum_{s=1}^t d_{is}^t \nu_s]^2$$

¹⁴This is because θ enters additively in all opinions.

We define the limit loss $L_i = \lim_{t \nearrow \infty} L_i^t$.¹⁵ We assume that each agent i aims at minimizing L_i . Now whenever L_i is finite, we can write it as

$$L_i = L_i^0 + V_i \text{ where } L_i^0 \equiv E(b_i\delta + c_i\xi)^2.$$

The term L_i^0 results from variations in initial opinions and the persistent component, while the term V_i results from the idiosyncratic components only. Note that the distribution over θ plays no role, so θ can be normalized to 0.

In the next Section we start by exploring the long-run properties of different learning rules within the FJ class. Then we turn to the optimal choice of learning rules.

3 Properties of learning rules

We are interested in long-run opinions: whether they converge to some limit opinion and if they do, what determines the variance of the limit opinion. 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 Exploding dynamics under DG

Our first result shows that if all agents follow a DG rule, as long as there is any idiosyncratic component in the noise, the variance of long-run opinions diverges. Moreover if there is a permanent component to the noise and the permanent component is either positive for everyone and strictly positive for some, or negative for everyone and strictly negative for some, then y_i^t must diverge in expectation over time for all i .

To show this we fix x and ξ and define $\bar{y}_i^t = Ey_i^t$ and $V_i^t = \text{var}(y_i^t)$. We then have:

Proposition 1: *Assume that $m_i = 0$ for all i . (i) If $\varpi_0 > 0$, then for all i and any fixed x, ξ , $\lim_t V_i^t = \infty$. (ii) For almost all realizations of the persistent components ξ , $\lim |\bar{y}_i^t| = \infty$ for all i and x .*

¹⁵Alternatively, 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 large date 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 $L_i \leq \text{var}(\delta_i) = \sigma_i^2 = 1$ by ignoring everyone else’s opinions ($m_i = 1$).

For example, the proposition shows that a bias in a single player's perception ξ_1 may be enough to drive up the opinions of all: if $\xi_1 > 0$, say, the bias creates a discrepancy with other's opinions, and each time others' opinions catch up, player 1 further raises his opinion compared to others, prompting another round of catching up, and eventually all opinions blow up.

We present here some intuition that explains why DG works well without noise and becomes fragile as soon as there is some noise. Let y^t denote the vector of opinions at t . Let Δ_n be the set of vectors of non-negative weights $p = \{p_i\}_i$ with $\sum p_i = 1$. For any i , we have $y_i^t = B_i y^{t-1} + \gamma_i \varepsilon_i^t$ with $B_i \in \Delta_n$. Because the network is connected, $B = (B_i)_i$ is an irreducible probability matrix with *strictly positive* eigenvector, implying that there is a vector of weights $\pi \in \Delta_n$ such that

$$\pi \cdot y^t = \pi \cdot y^{t-1} + \sum_i \pi_i \gamma_i \varepsilon_i^t.$$

Without noise, the limit weighted opinion $\pi \cdot y$ coincides with the weighted initial opinion $\pi \cdot x$. This explains why in the absence of noise the influence of initial opinions never dissipates (and also why all initial opinions matter – as $\pi \gg 0$): the direct contribution of i 's initial signal to i 's opinion vanishes, but 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 .

With noise however, $\pi \cdot y$ is a random walk, explaining why the influence of initial opinions vanishes and why the variance diverge. Besides, the random walk has a drift when $\sum_i \pi_i \gamma_i \xi_i \neq 0$, explaining why $\pi \cdot \bar{y}$ then diverge.

3.2 Anchored dynamics under FJ.

Fixing again x and ξ , we now examine long-run dynamics under FJ. Define $\bar{y}^t = (\bar{y}_i^t)_i$ and $V^t = (V_i^t)_i$ as the vector of expected opinions and variances.

Proposition 2. *Assume at least one player updates according to FJ with $m_{i_0} > 0$. Then, for any fixed x, ξ , \bar{y}^t and V^t converge. Besides, the limit variance V does not vary with x and ξ , and if $m_i = 0$, the limit vector of expected opinions \bar{y} does not vary with x_i .*

Proposition 2 shows that to avoid that all opinions drift, 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 players' long-run opinions. A detailed proof is in the Appendix.

Before providing some intuition for the proof, let us consider a two-player example where we set $m_2 = 1$ and $m_1 = 0$. Then player 2 always keeps the same opinion ($y_2^t = x_2$ for all t) and

$$\begin{aligned}\bar{y}_1^t &= \gamma_1(x_2 + \xi_1) + (1 - \gamma_1)\bar{y}_1^{t-1} + \\ &= (x_2 + \xi_1)\gamma_1(1 + (1 - \gamma_1) + \dots + (1 - \gamma_1)^{k-1}) + (1 - \gamma_1)^k \bar{y}_1^{t-k}\end{aligned}$$

implying that \bar{y}_1^t converges to $x_2 + \xi_1$ as t grows large, independently of player 1's initial opinion. Player 2 serves as an anchor that prevents agent 1's opinion from drifting. Long-run opinions however only incorporate player 2's initial opinion.¹⁶

The general argument for convergence runs as follows. For any fixed x, ξ , the expected opinion evolves according to $\bar{y}^t = X + B\bar{y}^{t-1}$, where $X_i = m_i x_i + (1 - m_i)\xi_i$, and where the i^{th} row vector of B is $B_i = (1 - m_i)A_i$. When $m_{i_0} > 0$ for some i_0 , proving convergence is standard:¹⁷ the limit expected opinion \bar{y} is the unique solution of

$$\bar{y} = X + B\bar{y} \tag{2}$$

Next, defining $\eta^t = y^t - \bar{y}^t$ and $w_{ij}^t = E\eta_i^t \eta_j^t$, we have

$$\eta_i^t = (1 - m_i)v_i^t + B_i \eta^{t-1}$$

implying an expression for the evolution of the covariance vector $w^t = (w_{ij}^t)$ of the form

$$w^t = \Lambda + \bar{B}w^{t-1},$$

where \bar{B}_{ij} is the row vector $(\bar{B}_{ij,hk})_{hk}$ with $\bar{B}_{ij,hk} = B_{ih}B_{jk}$. Proving convergence to the solution of

$$w = \Lambda + \bar{B}w \tag{3}$$

¹⁶More generally, up to noise terms, long-run opinions are determined by the opinions of agents for which $m_i > 0$.

¹⁷The key to convergence is whether $\sum_j B_{ij} < 1$ for all i . When this is the case, we say that B has the *contraction property*. When $m_i > 0$ for all i , this property trivially holds: $\sum_j B_{ij} = (1 - \gamma_i) + \gamma_i(1 - m_i)\sum_j A_{ij} < 1$ for all i . When $m_i > 0$ for only some players, we use the fact that the network is connected to conclude that for some large enough K , $C = B^K$ has the contraction property: with large enough K , then for any i , there are paths of length K that go through i_0 for which $m_{i_0} > 0$.

is also standard.¹⁸

One immediate corollary of Equations (2) and (3) is that the loss terms L_i^0 do not depend on γ or on the magnitude of the idiosyncratic component ϖ_0 , while the loss terms V_i are proportional to the idiosyncratic component ϖ_0 (and equal to 0 when $\varpi_0 = 0$). This also implies that when $\varpi_0 = 0$, the parameters γ have no effect on L_i , and we can focus on the weights m and the analysis of the rule *SFJ*.

3.3 Fragility under low m .

Although convergence is guaranteed when at least one player does not use DG, long-run opinions may nevertheless be highly sensitive to the permanent component of the noise, and the variance V may be very high. The next Corollary makes this idea precise.

Whenever $m_i > 0$, we define the *modified initial opinion*

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

The following Corollary shows that long-run expected opinions are weighted averages of modified initial opinions (the proof is in the appendix):

Corollary 1. *Assume $m_i > 0$ for all i . Then for each i , there exists $P_i \in \Delta_n$ such that for all $x, \xi, \bar{y}_i = P_i.\tilde{x}$.*

Intuitively, in each period, as a quick glance at equation FJ should make clear, \tilde{x}_i can be thought of as the effective seed for individual i so in the long-run, it is not surprising that there exist probability vectors $P_i \in \Delta_n$ such that if $\bar{y}_i = P_i.\tilde{x}$. An immediate consequence of Corollary 1 is that if all m_i are smaller than \underline{m} , then the variance of each \tilde{x}_i is at least $\varpi((1 - \underline{m})/\underline{m})^2$, hence the variance of \bar{y} is of the order of $\varpi/(n \underline{m}^2)$.

Corollary 1 can be generalized to the case where a subset N^0 of agents follows DG ($m_i = 0$). Then long-run opinions become linear combinations of modified opinions of the agents *not in* N^0 , and these modified opinions are

$$\hat{x}_i = \tilde{x}_i + (1 - m_i)R_i\xi^0/m_i$$

where R_i is a positive vector that only depends on the structure of the network which captures the influence of agents in N^0 on i (see Corollary

¹⁸This is by the same logic as Footnote 17. Among all the K -step paths that start in ij , there is at least one that goes through i_0k for some k implying that \bar{B}^K has the contraction property.

2 in the Appendix). Our conclusion regarding the sensitivity of long-run opinions to m when all the m_i are small extends to this case.

The two-player case. With two players, assuming m_1 and m_2 strictly positive and $\gamma_1 = \gamma_2 = 1$ (both use SFJ), the model can be solved by directly substituting y_2^{t-2} , then y_1^{t-2} , and so on. Letting $\rho = (1 - m_1)(1 - m_2)$, we have:

$$\begin{aligned} y_1^t &= m_1 \tilde{x}_1 + (1 - m_1) \nu_1^t + (1 - m_1) y_2^{t-1} \\ &= m_1 \tilde{x}_1 + (1 - m_1) m_2 \tilde{x}_2 + (1 - m_1) \nu_1^t + \rho \nu_2^{t-1} + \rho y_1^{t-2} \end{aligned}$$

which further implies:

$$y_1^t = \sum_0^{K-1} \rho^k (m_1 \tilde{x}_1 + (1 - m_1) m_2 \tilde{x}_2 + (1 - m_1) \nu_1^{t-2k} + \rho \nu_2^{t-2k-1}) + \rho^K y_1^{t-2K} \quad (4)$$

which in turn gives us (5) and (6) below for the limits \bar{y}_1 and V_1 :

$$\bar{y}_1 = p_1 \tilde{x}_1 + (1 - p_1) \tilde{x}_2 \text{ with } p_1 = m_1 / (m_1 + (1 - m_1) m_2). \quad (5)$$

$$V_1 = \varpi_0 \frac{(1 - m_1)^2 + \rho^2}{1 - \rho^2} \quad (6)$$

This example confirms that $p_1 = 0$ when $m_1 = 0$ and it illustrates that when both m_1 and m_2 get close to 0, $1 - \rho \simeq m_1 + m_2$, and the variance of opinion V_1 induced by the *idiosyncratic* noise gets arbitrarily high, approximately equal to $\varpi_0 / (m_1 + m_2)$.

3.4 Comments

(a) **On anchoring:** 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.

(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 loss due to the force—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 SFJ 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 change in perception of neighbors’ opinions $Z_i^t = z_i^t - z_i^{t-1}$, and the spread between own and neighbors’ opinions $D_i^t = z_i^t - y_i^t$, we have the following expressions:

$$Y_i^t = \gamma_i D_i^{t-1} \tag{DG}$$

$$Y_i^t = (1 - m_i) Z_i^t \tag{SFJ}$$

Under DG, one changes one’s opinion whenever there is a difference between that opinion and the opinions of one’s neighbors: any difference generates an adjustment, which is why the evolution is so sensitive to transmission errors. 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 too strong.

At the opposite extreme, under SFJ, players only incorporate *changes* in the opinions of others. So, in the case where the transmission error is always the same, 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 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. But, when $m_i > 0$ for at least one player, the knock-on effect will be smaller than the initial impact and will get even smaller over time. Hence over all it won’t blow up. If all m_i are small however, these indirect effects are not dampened enough, and the consequence is a high sensitivity of the final opinion to the magnitude of the errors.

(d) **On talking and listening.** We have so far introduced noise in the reception of opinions. Other sources of noise are also plausible: for example, not everyone needs to express their opinions to their neighbors every period,

potentially generating some randomness in the communication protocol. Or there may be randomness in whether expressed opinions are actually heard or processed. We will argue in Section 6.6 that these add to the fragility of DG but leave SFJ largely unaffected.

4 Choosing the rule

4.1 What are good rules?

Equipped with these insights about the properties of different rules, we now return to the main question of this paper: what rule will people choose and what rule should they choose (and are these the same)? Recall that we already specified the objective function of any individual i , which is to minimize the loss function.

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

To fix ideas it is worth starting with the case where there is no noise. In this case

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

As t becomes large, this system of equations must converge to limiting values of y_i , given by the system of equations

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

for all i . This implies that

$$y_i = \sum_j \beta_{ij} x_j = \sum_j \beta_{ij} (\theta + \delta_j) = \theta + \sum_j \beta_{ij} \delta_j$$

where $\beta_{ij} \geq 0$ and $\sum_j \beta_{ij} = 1$. Therefore

$$L_i = E\left(\sum_j \beta_{ij} \delta_j\right)^2$$

Under our assumption that the variance of the δ_j 's are identical, it is evident that across all possible probability vectors $(\beta_{ij})_j$, L_i is minimized by choosing $\beta_{ij} = \frac{1}{n}$ so that all the y_i would be equal to $\sum_j \frac{1}{n} x_j$. However β_{ij} are endogenously determined by the underlying rules that the players adopt, so there is no guarantee that these weights will be implemented as a result

of a rule that results from individual choice. In particular, if all the y_i are going to be the same then they must all satisfy

$$y_i = m_i x_i + (1 - m_i) y_i \text{ or equivalently, } m_i y_i = m_i x_i.$$

Since the realizations of x_i may differ, the equality of the y_i requires $m_i = 0$ for all i . There is thus no way to get to $\beta_{ij} = \frac{1}{n}$ unless $m_i = 0$ for all i . In other words, within the class of rules we consider, DG rules are the only ones that even offer the possibility of reaching the lowest feasible L_i .¹⁹

However as soon as there is some noise, we already saw that 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.

Proposition 3: *Let $\underline{m} = \varpi/(1 + \varpi)$. Any (m_i, γ_i) with $m_i < \underline{m}$ is dominated by $(\underline{m}, \gamma_i)$, individually and socially.*

Regarding the choice of the individually optimal rule, Proposition 3 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 the same reason why, in the absence of noise, the ideal rule puts strictly positive weight on all the seeds. 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.

¹⁹In fact as observed in the seminal paper by De Marzo et al. (2003), for generic networks, for any finite n , even DG rules will not implement these weights, though for large n the outcomes generated by DG rules will approximately minimize L_i , for a large class of networks (Golub and Jackson (2010)).

The next Proposition provides further characterization of the privately optimal choice of m_i . In particular it clarifies how the quality of initial signals and transmission noise affects it. In addition it tells us that for any fixed γ , an equilibrium where each player i chooses m_i optimally always exists. To simplify exposition, we focus on the case where the idiosyncratic component is null ($\varpi_0 = 0$), so the outcome does not depend on γ . However we allow here for heterogenous quality of signals and noise terms. We let $W_k = \sigma_k^2 + \varpi_k(\frac{1-m_k}{m_k})^2$, where σ_k^2 is the variance of k 's initial opinion and ϖ_k the variance of k 's persistent component. We have:

Proposition 4. *Player i 's optimal choice m_i satisfies:*

$$\frac{m_i}{1 - m_i} = \frac{\varpi_i + (1 - \lambda_i)^2 \sum_{k \neq i} (r_{ik})^2 W_k}{\sigma_i^2 (1 - \lambda_i)}$$

where λ_i and $r_i = (r_{ik})_{k \neq i} \in \Delta_{N-1}$ only depend on A and m_{-i}

We see from this that m_i 's response shifts up when the variance of his own signal (ϖ_i) or that of anyone else (ϖ_k) goes up. It further 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.

Finally, although we are unable to fully characterize the social optimum, the following Proposition clarifies the relationship between the private and social optima.

Proposition 5: *Assume $\frac{\partial L_i}{\partial m_i} = 0$. Then $\frac{\partial L_j}{\partial m_i} < 0$ for all j .*

This implies that at any Nash equilibrium, any player would increase aggregate social welfare by increasing m_i further. The next subsection fully studies a simpler environment (two players), where we explore the equilibrium, the social optimum and the relation between the two in greater detail.

We provide below some intuition for Proposition 5. A convenient way to express player j 's opinion is by seeing it as an average between sources different from i that are unmediated by i 's opinion, and player i 's opinion (see Lemma 4 in Appendix):

$$\bar{y}_j = (1 - \mu_{ji}) Q_j^i \tilde{x}_{-i} + \mu_{ji} \bar{y}_i$$

where $\mu_{ji} \in (0, 1)$ and Q_j^i is a probability vector. Since we consider the influence of \tilde{x}_{-i} through channels that are unmediated by i , μ_{ji} and Q_j^i are independent of m_i .

The expression above permits to separate the loss L_j into three terms: the variance of sources other than i (including j 's own modified opinion),

the variance i 's long-run opinion (i.e., L_i), and a last term that comes from the positive correlation between sources other i and the opinion of i . The correlation is positive because the network is connected: any report made by j incorporates \tilde{x}_j and eventually reaches i who also incorporates it; also, whenever there are paths that go from k to j without going through i and others that go from k to i without going through j , \tilde{x}_k contributes to j 's opinion through both the unmediated channel and i 's opinion, hence the correlation.

Now when m_i is raised away from the private optimum, there is no effect on the first term, there is a second order effect on the second term L_i (because we start at the private optimum), and there is a first-order reduction on the last term: when m_i increases, the influence of each $k \neq i$ on i 's opinion is reduced, and the correlation between \bar{y}_i and \tilde{x}_k is also reduced.

4.2 Simple networks.

We start with the two-player case with fully persistent noise.

4.2.1 Social optimum.

Recall from (5) that $y_i = p_i \tilde{x}_i + (1 - p_i) \tilde{x}_j$ with $p_i = m_i / (m_i + (1 - m_i) m_j)$, which yields

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

where $I(p) = p^2 + (1 - p)^2$ is the variance of long run opinion in the absence of transmission noise (minimized at $p = 1/2$),²⁰ and $\mathcal{X}(m) = \varpi \frac{(1-m)^2}{m^2}$ measures the variance of cumulated error term. The total social loss is $L_1 + L_2$. It is easy to check that minimizing the social loss requires setting identical values for m_1 and m_2 . When both players use the same rule ($m = m_1 = m_2$), the loss is:

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

Given that initial opinions are equally informative, optimal information aggregation in the absence of noise by both players (which amounts to minimizing $I(p)$) would require setting $p_i = p_j = 1/2$. This is not feasible, but if m is small enough, p_i and p_j are both close to $1/2$ and $I(p)$ is potentially

²⁰This is under our assumption that the variance σ_i^2 is equal to 1 for all i .

close to its minimum. When m goes up from close to zero, the contribution of initial opinions to long-run opinions become asymmetric ($p_i, p_j > 1/2$), and this pushes $I(p)$ up. Indeed as the expression for $I(\frac{1}{2-m})$ makes clear, the logic of standard information aggregation gives the players a joint incentive to reduce m . However as the second term in the expression for L , $(1 + \varpi \frac{(1-m)^2}{m^2})$, makes evident, there is also a cost to lowering m . 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 (for $\varpi = 0.0001$, $m^{**} = 0.13$ and for $\varpi = 0.001$, $m^{**} = 0.21$).²¹

Asymmetric noise. Social incentives to raise m away from zero are large even when only one player is subject to noise (say $\xi_2 \equiv 0$). Although player 1 is then the only one subject to errors, this affects the opinion of player 2, which in turns affects the opinion of player 1, and so on. These errors bounce back and cumulate, generating *for both players* a loss term of substantial magnitude if m_1 is small:

$$L_1 = I(p_1) + (p_1)^2 \mathcal{X}(m_1) \text{ and } L_2 = I(p_2) + (1 - p_2)^2 \mathcal{X}(m_1)$$

Once again the term $I(p_1) + I(p_2)$ calls for setting p_1 and p_2 close to $1/2$, hence m_1 and m_2 close to zero. The term $((p_1)^2 + (1 - p_2)^2) \mathcal{X}(m_1)$, on the other hand, calls for increasing m_1 as in the previous case.

4.2.2 Nash Equilibrium.

So far we have focused on the socially efficient choice of rules. It is perhaps more natural however to assume that individuals choose their rules non-cooperatively. To see what this does to the choice of rules, consider the asymmetric case where only player 1 is subject to noise.

For very low m_2 , player 1 should choose m_1 close to m_2 for information aggregation purposes, but this would generate very high cumulated error, and player 1 is better off ignoring player 2 (m_1 close 1). For higher m_2 , information aggregation is the main issue, and getting p_1 close to $1/2$ requires choosing $m_1 < m_2$. A similar best response curve obtains for player 2.²²

The point where curves cross defines the equilibrium weights (m_1^*, m_2^*) .

²¹For ϖ arbitrarily small, $2L \simeq 1 + m^2/4 + \varpi/m^2$, so $m^{**} \simeq (4\varpi)^{1/4}$.

²²The figure assumes $\varpi = 0.001$




Figure 1: Best responses, $\sigma^2 = 0.02$

4.2.3 Private versus social incentives.

Starting from low values of m , each player can reduce the impact of the \mathcal{X} term on private (and social) losses, by increasing their m_i . However increasing m_i also has the adverse effect of increasing p_i . This latter effect can be partially mitigated if the other player raises their m as well, but in a Nash Equilibrium each player ignores the other player's incentive to raise m and therefore raises their own m too little.

Figure 2 below plots equilibrium and socially optimal levels of (m_1, m_2) as a function of ϖ .

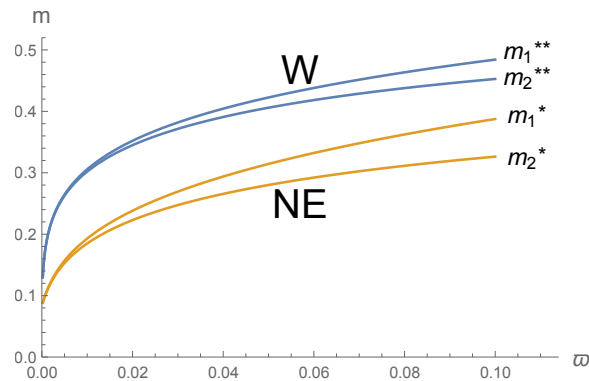


Figure 2: Equilibrium and socially optimal weights (m_1, m_2)

Finally, we observe that adding idiosyncratic noise to the persistent noise makes the incentive to increase m_i even stronger. With two players, we can use (4) to get:²³

$$L_1 = I(p_1) + ((p_1)^2 \mathcal{X}(m_1) + (1 - p_1)^2 \mathcal{X}(m_2)) \left(1 + \frac{1 - \rho \varpi_0}{1 + \rho \omega}\right)$$

With idiosyncratic transmission errors, the variance of the error term is amplified proportionally. Efficient and Equilibrium weights on one's own signal both go up.

4.3 An example where the Nash Equilibrium is efficient

As explained earlier, the reason why participants end up with too little weight on their own seed is that (i) there are echo effects in the network: reports that individuals send eventually (partially) come back to them and (ii) signals may reach an individual through two different channels. Both motives create correlations that are not minimized when each player only considers private incentives, and they would all benefit from raising m to reduce these correlations.

To illustrate this, we examine below a network where both effects above are suppressed: a long circle where information transmission is directed and one-sided. Player i communicates to player $i+1$, who communicates to $i+2$, and so on. Since the circle is long, there is (almost) no echo effect because reports sent by a player, say j , have almost no influence on the reports he receives. And given the one-sided nature of the network, signals reach j through a unique channel. We check below that in such a network there is no scope for collaborating in jointly raising m 's: the social and private outcome coincide.

Formally, if player i chooses m_i and all other players choose m , long-run opinions satisfy

$$y_i = m_i \tilde{x}_i + (1 - m_i) (Z + (1 - m)^{K-1} y_i) \text{ where } Z = m \sum_{k=0}^{K-2} (1 - m)^k \tilde{x}_{i-1-k}$$

One can use this expression to derive y_i and next $m_i^*(m)$ that minimizes L_i , the variance of y_i .

At the limit where K is arbitrarily large, the variance of Z is

$$J(m) = \frac{m}{2 - m} (1 + \mathcal{X}(m)),$$

²³Recall $\rho = 1 - (m_1 + (1 - m_1)m_2)$.

L_{i-1} (the variance of y_{i-1}) coincides with $J(m)$ and the correlation between \tilde{x}_i and $y_{i-1} = Z + (1 - m)^{K-1}y_i$ becomes negligible. As a consequence, the variance of y_i is

$$L_i = (m_i)^2(1 + \mathcal{X}(m_i)) + (1 - m_i)^2 L_{i-1} = (m_i)^2 + (1 - m_i)^2(\varpi + J(m)). \quad (7)$$

We can derive the optimal choice of m_i for player i ,²⁴ from which we obtain the equilibrium value m^* and equilibrium loss L^* , which satisfy:

$$m^* = \phi(m^*) \text{ where } \phi(m) = \frac{\varpi + m}{1 + \varpi + m} \text{ and } L^* = m^* = J(m^*)$$

It is straightforward to check that social incentives coincide with private incentives. Indeed, let \underline{L} be the minimum loss that a player can obtain at the social optimum. From (7), we get $\underline{L} \geq \min_{m_i} (m_i)^2 + (1 - m_i)^2(\varpi + \underline{L}) = \phi(\underline{L})$. It follows that $\underline{L} \geq \phi^{(n)}(0)$ for all n . Since $\lim_n \phi^{(n)}(0) = L^*$,²⁵ we conclude that $\underline{L} = L^*$.

4.4 Information aggregation

In the case where there is large circle, m^* is of the order of $\varpi^{1/2}$ when ϖ is very small, and $N = 1/m^* = 1/J(m^*)$ is a measure of the quality of the aggregation of information: it represents the number of signals that are eventually aggregated into the information of each player. For example, with $\varpi = 0.05$, $m^* = 0.2$, and a player's ultimate information is comparable to her receiving 5 independent signals, out of the infinite pool that is available when the circle is arbitrarily large. We draw N as a function of ϖ .

The loss in welfare is significant relative to a benchmark case where players observe (with transmission noise) the initial opinion of other players in the network.²⁶ Information aggregation would be perfect in the large circle case, in this benchmark.²⁷ On the other hand the quality of the aggregation remains bounded away from perfect aggregation in our case.

For the two-player example, a persistent noise of magnitude $\varpi = 0.05$ would yields a total variance $L = 1.025$ if each player could observe the

²⁴From the first-order condition $m_i = (\omega + J(m))/(1 + \omega + J(m))$. Substituting m_i , this yields a minimum loss equal to $L_i = m_i$.

²⁵This is because $L^* = \phi(L^*)$ and $|\phi(L) - \phi(L')| \leq \frac{|L-L'|}{1+\varpi}$.

²⁶For i , this corresponds to getting the opinion $z_j = x_j + \varepsilon_i$ if j is a neighbor, and $z_k = x_k + \varepsilon_i + \varepsilon_j$ if k is not a neighbor of i but a neighbor of j , and so on.

²⁷Even with transmission errors, the variance of the opinion of a neighbor at k steps is $v_k = 1 + k\varpi$. Since it grows linearly with k , the optimal weighing of these opinions would lead to an opinion with variance $(\sum 1/v_k)^{-1}$, which goes to 0 with k .




Figure 3: Information aggregation: large circle

other's initial opinion with noise. Under FJ, the minimum loss rises up to 1.164. The equilibrium loss is even higher, 1.187.²⁸ Figure 4 summarizes how welfare levels compare between the Nash equilibrium, the social optimum and the benchmark case as ϖ varies.

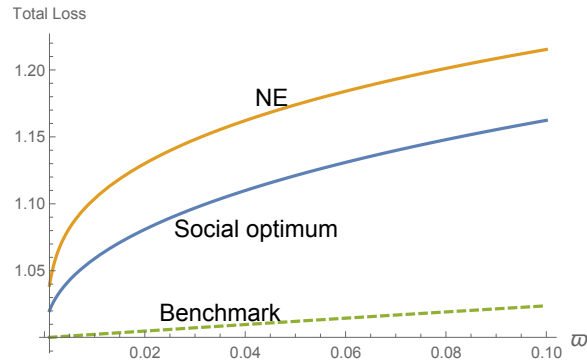


Figure 4: Information aggregation: two players

4.5 Dispersion of opinions

Even when transmission errors are small, equilibrium weights may be high, implying some significant dispersion in opinions in the long-run. That opinions do not converge to one another is built in the FJ updating rule. We

²⁸The spread between minimum loss and equilibrium loss is higher in the asymmetric case.

would be close to consensus if all weights on one's own initial opinion but as we saw, this is never optimal. Players will agree to disagree, unlike under standard Bayesian communication.

5 Choosing among a richer class of rules.

5.1 Choosing γ

In Section 4, we examined incentives to modify the weight m_i . We now turn to the other sets of weights, the γ_i . A potential issue with FJ where γ_i is large is that long-run opinions are sensitive to idiosyncratic noise in transmission, and more generally to temporary changes in other's opinions. Choosing a lower γ_i slows down these reactions, hence opinions are only mildly affected by temporary shocks on perception and temporary variations in others' opinions. Recall that $L_i = L_i^0 + V_i$ where $L_i^0 = E(\bar{y}_i^2)$ represents the loss induced by the error in the original signal and the persistent component noise, and where V_i is the variance induced by the idiosyncratic noise. L_i^0 does not depend on γ . The next Proposition examines the effect of γ on V_i as well as incentives for an individual to choose a low γ_i :

Proposition 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 $m \geq \underline{m}$, $V_i \leq c\gamma_i$.
- (iii) If the lower bound $\underline{\gamma}$ on the choice set is sufficiently low and $\gamma_i = \underline{\gamma}$, $V_i \leq 1/|\log \underline{\gamma}|$ for all $m \geq \underline{m}$ and γ within the choice set.

The proof is in the Appendix. 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. Finally item (iii) examines the case where all γ_i are restricted to be above some threshold $\underline{\gamma}$. It shows that by choosing $\gamma_i = \underline{\gamma}$, player i can ensure that the variance is a small²⁹ independently of the choice of other individuals. This further implies that if (m^*, γ^*) is an equilibrium of the game, then m^* is an ε -equilibrium of the game with no idiosyncratic noise, with ε comparable to $1/|\log \underline{\gamma}|$.

This result 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

²⁹ $1/|\log \underline{\gamma}|$ is small number when $\underline{\gamma}$ is small.

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).

6 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. The general point is that some long run dispersion of opinion remains part of the answer and indeed there are reasons to expect that adding the new elements exacerbates this property.

6.1 Biased persistent noise

We have so far assumed that the persistent noise is drawn from a distribution that is mean zero. One can however imagine settings where it is more reasonable to assume that the persistent noise is biased, centered on ξ_i^0 for player i , for example because some individuals are biased in what they report (for whatever reason). That could for example be because they are truly biased and therefore try to sway opinion in the direction of their bias, or because they believe that others are biased and try to correct for it. In any case it makes sense to consider a variant of the updating rule FJ in which the agent can shift the opinion he incorporates by a constant c_i , so as to try to undo the systematic biases in his perception or perception of others:

$$y_i^t = (1 - \gamma_i)y_i^{t-1} + \gamma_i(c_i + m_i x_i + (1 - m_i)z_i^{t-1}) \quad (\text{FJc})$$

Suppose that in all other respects, the model is as before. For any (m, γ, c) , this shift does not affect the variance of opinions resulting from idiosyncratic noise, but it shifts all long-run opinions. Regarding expected long-run opinions, these shifts imply as before that \bar{y}_i is a linear combination P_i of the modified opinions \tilde{x}_j where

$$\tilde{x}_j = x_j + \frac{(1 - m_j)\xi_j + c_j}{m_j}$$

The linear combination P_i is independent of c , so for any fixed (m, γ) , each i can in principle set c_i to fully offset the systematic bias in transmission and

this turns out to be optimal.³⁰ If c_i cannot be adjusted (e.g., $c_i = 0$ for all i), then the bias ξ_i^0 amounts to an increase in the variance of own error term ϖ_i , which, as Proposition 4 explains, generates further incentives to increase m_i : intuitively, when ξ_i^0 increases, opinions of others become a less accurate estimate of θ , and i prefers to put more weight on own.

6.2 Heterogenous preferences

Assume that preferences of player i are quadratic (i.e., $u(a, \theta_i) = -(a - \theta_i)^2$) but vary in their relation to the common component θ :

$$\theta_i = \theta + b_i \tag{8}$$

and that x_i is a noisy estimate of one's preferred point, that is,

$$x_i = \theta_i + \delta_i \tag{9}$$

Define $Y_i^t = y_i^t - b_i$, $X_i = x_i - b_i$ and $\beta_i = (b_j - b_i)_j$. The "debiased" opinions Y_i^t evolve according

$$Y_i^t = (1 - \gamma_i)Y_i^{t-1} + \gamma_i(c_i + m_i X_i + (1 - m_i)(z_i^{t-1} + A_i \beta_i))$$

and the problem becomes formally equivalent to the homogenous preference case with a persistent transmission term $A_i \beta_i$ added. If the biases b are fixed and if players can adjust c_i optimally, then like in the previous case, in equilibrium players can offset the bias by setting

$$c_i = -(1 - m_i)A_i \beta_i.$$

and the analysis is formally equivalent to the homogenous preference case, and the issue we raised (in particular, the fragility of long-run opinions to transmission noise) apply. In contrast, if players are unable to adjust c_i (e.g., $c_i = 0$), then the term $A_i \beta_i$ is akin to a systematic bias ξ_i^0 , which, as explained in the previous subsection generates incentives to increase m_i .

Finally, consider the intermediate case where players can adjust c_i , but biases are not fixed and players can only adjust c_i on average across realizations of the β 's. Said differently, across problems, there are variations in the heterogeneity, and players are unable to tune c_i to each realization of the heterogeneity. Then the problem is formally equivalent to one where

³⁰Letting L_i^0 denote the loss when all ξ_i are centered on 0, and $C_i = (1 - m_i)\xi_i^0 + c_i$. We have $L_i = L_i^0 + (P_i C)^2$, and L_i is minimized for $C_i = -\sum_{j \neq i} P_{i,j} C_j / P_{i,i}$. The equilibrium loss thus coincides with L_i^0 , and $c_i = -(1 - m_i)\xi_i$ for all i is an equilibrium.

preferences are homogenous and a persistent but noisy transmission term $A_i\beta_i$ is added.³¹

The general take-away should be that there are many potential sources of noise which will favor the choice of FJ over DG rules. To illustrate with one final example, assume that i misperceives other's preferences. He perceives $\widehat{\beta}_i$ instead of β_i and erroneously sets $c_i = -A_i\widehat{\beta}_i(1-m_i)$. Then the difference $(1-m_i)A_i(\beta_i - \widehat{\beta}_i)$ is akin to an additional (independent) source of persistent bias/noise in transmission.

6.3 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.³² 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.³³ Noisy communication is modeled as before, through the addition of a noise term ε_i^t that slants what i hears. Together these give us

$$\begin{aligned} z_{i,j}^t &= y_j^{t-1} + \varepsilon_i^t \text{ if } j \in N_i^t \\ z_{i,j}^t &= z_{i,j}^{t-1} \text{ if } j \in N_i \setminus N_i^t \end{aligned}$$

where $z_{i,j}^t$ 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

$$z_i^t = A_i Z_i^t$$

where $Z_i^t = (z_{i,j}^t)_j$ is the vector of i 's perceptions and A_i defines as before how i averages neighbors' opinions.³⁴ We continue to assume FJ updating.

For fixed x, ξ , define $\bar{y}_i^t = E y_i^t$, $Y_i^t = (\bar{y}_i^{t-k})_{k=0,\dots,K}$, the column vector of i 's past recent opinions, and $Y^t = (Y_i^t)_i$. One can write $Y^t = X + B Y^{t-1}$.

³¹One difference with the case examined in the basic model however is that the β_i 's are correlated: with two players, $A_i\beta_i = b_j - b_i = -A_j\beta_j$. Nevertheless, so long as there still exists a persistent noise term ξ_i (with all ξ_i drawn independently of the β_i 's), Proposition 3 applies, as for each realization of β , all $m_i < \underline{m}$ are dominated.

³²Banerjee 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".

³³That is, for all $t : \cup_{s=1,\dots,K} N_i^{t+s-1} = N_i$.

³⁴We abuse previous notations here, using the restriction of vector A_i to i 's neighbors (A_i was previously defined over all players, with weight 0 on non-neighbors).

Y^t converges for standard reasons, to some uniquely defined Y . Consider now the vector \bar{y} solution to

$$\bar{y}_i = m_i x_i + (1 - m_i) A_i (\bar{y} + \xi_i)$$

and let $\bar{Y}_i = (\bar{y}_i, \dots, \bar{y}_i)$ and $\bar{Y} = (\bar{Y}_i)_i$. By construction, under this profile of opinions, it does not matter when i heard from j because opinions do not change. \bar{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.³⁵

This robustness contrasts with what happens when players use DG rules. As we explain in Appendix with a simple example, changes in the protocol and in particular, 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). To assess the effect of changing weights γ , consider the two-player case *with noiseless communication*. Under DG, for i and $j \neq i$,

$$y_i^t = (1 - \gamma_i) y_i^{t-1} + \gamma_i y_j^{t-1},$$

so for any α, β such that $\beta\gamma_2 = \alpha\gamma_1$,

$$\begin{aligned} \alpha y_1^t + \beta y_2^t &= ((1 - \gamma_1)\alpha + \beta\gamma_2) y_1^{t-1} + ((1 - \gamma_2)\beta + \alpha\gamma_1) y_2^{t-1} \\ &= \alpha y_1^{t-1} + \beta y_2^{t-1} = \alpha x_1 + \beta x_2 \end{aligned}$$

Since $y_1^t - y_2^t$ converge to 0,³⁶ which implies a common long-run opinion

$$y = \frac{\alpha x_1 + \beta x_2}{\alpha + \beta} = \frac{\gamma_2 x_1 + \gamma_1 x_2}{\gamma_1 + \gamma_2}$$

Thus when γ_1/γ_2 rises, long run opinions get closer to player 2's initial opinion.

6.4 Coarse communication

In the social learning literature, it is common to focus on cases where the choice problem is about whether action 1 or action 0 should be taken, and the information being aggregated is which of the two is being recommended by

³⁵So long as the condition in footnote 33 holds.

³⁶The usual contraction argument works: $y_2^t - y_1^t = (1 - (\gamma_2 + \gamma_1))(y_2^{t-1} - y_1^{t-1})$ so $|y_2^t - y_1^t| \leq k |y_2^{t-1} - y_1^{t-1}|$ for $k < 1$

others. Frick et al. (2019) explore this class of models under the assumption that the agents may have erroneous priors, which is related to our emphasis on the role of errors in transmission in social learning.

We introduce this into our model by assuming that $\theta_i = \theta + b_i$ characterizes i 's preference (as in (8)) but that the optimal action a_i^* is 1 when $\theta_i > \theta^*$, 0 otherwise. Agent i knows the bias b_i and 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 F with full support on \mathcal{R} . Call $q = h(\theta)$ the fraction of agents that would choose $a = 0$ if their opinion was θ , and let $\phi(q) \equiv h^{-1}(q)$.³⁷

A player with current opinion y_i^t about θ reports $a_i^t = 1$ to neighbors if $y_i^t + b_i > \theta^*$ and $a_i^t = 0$ otherwise. From a vector of reports, he computes the fraction f_i^t that report 0, and uses this as an input to make an inference about others' opinions. A plausible rule is

$$z_i^t = \phi(f_i^{t-1}) + \xi_i$$

where ξ_i is a persistent bias in making inferences.³⁸ He next incorporates z_i^t using FJ to generate a new opinion y_i^{t+1} .³⁹

In the special case where the number of agents is large, each player hears from all other and all agents are subject to a perception bias $\xi_i = \xi > 0$, DG rules generate a dynamic that induces all players to report 1 independently of the state of the world.⁴⁰ Assume a fraction at most equal to $f > 0$ reports 0. Each makes an inference z_i at least equal to $\phi(f) + \xi$ regarding neighbors' opinions, so eventually, under DG, each player of type b_i may only report 0 if $b_i + \phi(f) + \xi < \theta^*$. Under the large number approximation, a fraction at

³⁷ $h(\theta) = \Pr(\theta + b_i < \theta^*) = F(\theta^* - \theta)$. Choosing an F that arises from a density with full and unbounded support ensures that h is strictly decreasing from \mathcal{R} to $(0, 1)$.

³⁸In the terminology of Frick et al. (2018), ξ_i could stem from an erroneous prior $\hat{F} \neq F$. Indeed, define $\hat{h}(\theta) = \hat{F}(\theta^* - \theta)$ and $\hat{\phi} = \hat{h}^{-1}$. If agents use $\hat{\phi}$ to make inferences, i.e., $z_i^t = \hat{\phi}(f_i^t)$, then the difference $\hat{\phi}(f_i^t) - \phi(f_i^t)$ is a systematic bias in making inferences.

³⁹Note that the loss function is no longer quadratic, but once one defines utilities $u(a, \theta, \theta_i)$, one can define loss functions, hence, further, express long-run expected losses as a function of the profile of updating rule. The optimal action given θ is $\sigma^*(\theta, \theta_i) = \arg \max u(a, \theta, \theta_i)$. Players form opinions y_i^t and report $a_i^t = \sigma^*(y_i^t, \theta_i)$. We assume that, eventually, if they take a decision at t , they mechanically use y_i^t and choose $a_i^t = \sigma^*(y_i^t, \theta_i)$. The loss associated with a decision taken at t is

$$L_i^t = E(u_i^*(\theta, \theta_i) - u(\sigma^*(y_i^t, \theta_i), \theta_i))$$

where $u_i^*(\theta, \theta_i) = u(\sigma^*(\theta, \theta_i), \theta_i)$, and agent i cares about minimizing the long-run expected loss $L_i = \lim_{t \nearrow \infty} L_i^t$.

⁴⁰The assumption can be weakened to heterogenous biases ξ_i at least equal to $\xi > 0$.

most equal to $f' = h(\phi(f) + \xi) < f$ report 0, hence this fraction of agents reporting 0 eventually vanishes.

In contrast, under FJ with m_i sufficiently large, long-run opinions remain anchored on initial opinions, and opinions remain bounded (and correlated with the underlying state θ). For example, when all signals x_i coincide (say, $x_i = x = \theta + \delta$), the long-run opinion must solve:

$$h\left(m\left(x + \frac{(1-m)\xi}{m}\right)\right) + (1-m)\phi(f) = f$$

hence

$$\phi(f) = x + \frac{(1-m)\xi}{m}.$$

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 $\frac{(1-m)\xi}{m}$. However, it creates heterogeneity in agents beliefs when initial opinions x_i differ.⁴¹

Said differently, with agents who constantly seed in their own initial opinion x_i , the drift in opinions remains bounded. Information aggregation is not perfect because of the positive weight m_i , opinions remain dispersed in the long run, but they remain correlated with the underlying state.

Frick et al. (2019) obtain a fragility result similar to the one obtained above under DG. They consider players who naively apply Bayesian updating to their erroneous priors. Like DG, Bayesian updating incorporates a strong forces towards consensus, which eventually makes both processes (DG and Bayesian updating) fragile to errors.

FJ processes can be seen as a potential fix to the fragility of DG or Bayesian processes: by allowing for heterogenous opinions or beliefs and by triggering updates based on variations in others opinions (rather than discrepancies between others' and own opinions), they end up being more robust, not subject to this particular form of fragility.

⁴¹With a large number of neighbors, and heterogenous initial opinions, the long-run fraction f must solve:

$$Eh\left(m\left(\theta + \delta_i + \frac{(1-m)\xi}{m}\right)\right) + (1-m)\phi(f) = f$$

If h is locally linear around, this yields $\phi(f) = \theta + \frac{(1-m)\xi}{m}$: aggregation of private signals allows a perfect inference θ , up to the persistent echo effect $\frac{(1-m)\xi}{m}$.

6.5 Non-linear aggregation rules

We consider next an extension to non-linear updating rules. In DG, opinions adjust as a function of the spread between of own and others' opinions, and the adjustment is linear in the spread. We examine a two-player example where the adjustment is linear for player 1, and non-linear for player 2.

Formally, denote by $\Delta_i^t = y_j^t - y_i^t$ the spread of opinion between j and i , and assume that

$$y_i^t = y_i^{t-1} + \gamma_i(\phi_i(\Delta_i^{t-1}) + \varepsilon_i^t)$$

where

$$\phi_1(\Delta) = \Delta \text{ and } \phi_2(\Delta) = \Delta - d\rho(\Delta)$$

with $\rho(\Delta) = 1 - \exp -\Delta^2$. In other words, player 1 adopts the standard linear DG rule, while player 2 adopts a rule in the spirit of DG but less sensitive to bigger variations in Δ than a linear rule (for small Δ , $\phi_2(\Delta) \simeq \Delta - d\Delta^2$).

Choose α and β such that $\alpha + \beta = 1$ and $\alpha\gamma_1 = \beta\gamma_2$. Next define

$$Y^t = \alpha y_1^t + \beta y_2^t \text{ and } \Delta_1^t = y_2^t - y_1^t.$$

Letting $\varepsilon^t = \alpha\gamma_1\varepsilon_1^t + \beta\gamma_2\varepsilon_2^t$, we have:

$$\begin{aligned} Y^t &= Y^{t-1} + \alpha\gamma_1\phi_1(\Delta_1^{t-1}) + \beta\gamma_2\phi_2(-\Delta_1^{t-1}) + \varepsilon^t \\ &= Y^{t-1} - d\beta\gamma_2\rho(\Delta_1^{t-1}) + \varepsilon^t \end{aligned}$$

In other words, when both players use the linear DG, Y^t is a random walk. When players do not both use the linear DG rule, and one player uses an adjustment that is more conservative for large spread, then Y^t is a random walk with a negative drift. The drift is determined by $\rho(\Delta_1^{t-1})$, so it is vanishing if Δ_1^t tends to 0, but, for any t , Δ^t is actually bounded away from 0 with positive probability,⁴² which implies that Y^t diverge.

6.6 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

⁴²Indeed, the evolution of Δ_1^t is determined by

$$\Delta_1^{t+1} = (1 - (\gamma_1 + \gamma_2))\Delta_1^t - d\beta\gamma_2\rho(\Delta_1^t) + \eta^t$$

where $\eta^t = \gamma_2\varepsilon_2^t - \gamma_1\varepsilon_1^t$. This implies that the spread Δ_1^t tends to revert to 0, but the noise term η^t keeps it up bounded away from 0 with positive probability. Hence the negative drift for Y^t .

more information than latter 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.

Such time-dependent processes however have important weaknesses. It is not obvious how one extends these to larger networks. 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, 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-dependant rule where each person communicates own current opinion, and 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, hence 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.

7 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.⁴³ 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 be overwhelmed by noise in transmission. The FJ players could provide the system with the necessary memory, using the initial communication phase to 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 temporarily updated until 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

⁴³As mentioned earlier, Banerjee et al. (2019) introduce the idea of a Generalized DeGroot model where not everyone starts with a signal

to use a FJ type rules where the private seed is periodically updated by each player to reflect the private signals about θ that each one accumulates.

Finally, one interesting property of FJ type rules that we already emphasized is that one's opinions vary as a result of *variations* in others' opinions vary, rather than because of a difference between one's and others' opinions. In particular, players' opinions may differ in the long run. One could imagine applying a similar idea to beliefs about the state of world. With two states for example, one could let $y_i = \ln p_i / (1 - p_i)$ measure the belief of i over the underlying state⁴⁴ and assume an updating process to y_i in the spirit of SFJ rule:

$$y_i^t - y_i^{t-1} = (1 - m_i)(z_i^{t-1} - z_i^{t-2})$$

where

$$z_i^t - z_i^{t-2} = \frac{1}{N_i} \sum_{j \in N_i} (y_j^{t-1} - y_j^{t-2}) + \varepsilon_i^t$$

measures the perceived variations in others' opinions. This updating process allows for diverse beliefs in the population, and also transmission of information when some signals in the network induce variations in beliefs. These updating rules (with m_i set appropriately) could turn out be more robust than Bayesian rules when the updating process is subject to noise or biases.

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 (2016): 1663-704.

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.

⁴⁴This is in the spirit of log-linear learning rules that use the logarithm of likelihood ratios, as in Molavi et al. (2018).

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

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.

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

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.

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), "Naïve Herding in Rich-Information Settings." *American Economic Journal: Microeconomics*, 2 (4), 221-43.

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

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

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

Golub, Benjamin, and Matthew O. Jackson (2010), "Naïve 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

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

Mira Frick, Ryota Iijima, and Yuhta Ishii (2019), Misinterpreting Others and the Fragility of Social Learning, Cowles foundation paper n° 2160.

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

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.

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.

Appendix.

Notations. Define M as the $N \times N$ diagonal matrix where $M_{ii} = m_i$ (and $M_{ij} = 0$ for $j \neq i$). 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 $\tilde{x}_i = x_i + \xi_i(1 - m_i)/m_i$ denote the modified initial opinion.

Next define $B_{ij} = (1 - m_i)A_{ij}$ and the $N \times N$ matrix $B = (I - M)A$. Also define the (N^2) vector Λ with $\Lambda_{ij} = 0$ if $i \neq j$, $\Lambda_{ii} = (1 - m_i)^2\varpi_0$ and \bar{B} the $(N^2 \times N^2)$ matrix where \bar{B}_{ij} is the row vector $(\bar{B}_{ij,hk})_{hk}$ with $\bar{B}_{ij,hk} = B_{ih}B_{jk}$.

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

Finally, 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.

Evolution of expected opinions and covariances. Under SFJ, the evolution of opinions and expected opinions (given x, ξ) follows

$$y^t = X + (I - M)\nu + By^{t-1} \quad (10)$$

$$\bar{y}^t = X + B\bar{y}^{t-1}, \quad (11)$$

from which we obtain:

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

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

$$w^t = \Lambda + \bar{B}w^{t-1} \quad (12)$$

In the general case (FJ rather than SFJ), the evolution is defined similarly, with $X_i = \gamma_i(m_ix_i + (1 - m_i)\xi_i)$ and $B_{ij} = (1 - \gamma_i)I_{ij} + \gamma_i(1 - m_i)A_{ij}$ and $\Lambda_{ii} = (\gamma_i(1 - m_i))^2\varpi_0$.

Paths. For any K , any K -sequence $q = (i_1, \dots, i_K)$ and any probability matrix $D = (D_{ij})_{ij}$, we let $\pi^D(q) \equiv \prod_{k=1}^{K-1} D_{i_k, i_{k+1}}$, and for any set of sequences Q , we abuse notations and let $\pi^D(Q) = \sum_{q \in Q} \pi^D(q)$. We define a K -*path* as a K -sequence q for which $\pi^D(q) > 0$. For any i, j is a K -neighbor of i if there exists a K -path ending in j , and we denote by \bar{N}_i^D the set of individuals that are K -neighbors of i for some K , under D .

Assumption 1: $\bar{N}_i^A = N$ for all i .

We denote by $Q_{i,j}^K$ the set of paths of length K from i to j , and Q_i^K the set of paths of length K that start from i . $Q_i^K = \cup_j Q_{i,j}^K$ and by construction, for any i, j

$$A_{ij}^K \equiv \pi^A(Q_{i,j}^K) \text{ and } \sum_{j \in N} A_{ij}^K = \pi^A(Q_i^K) = 1 \quad (13)$$

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

Proof of Proposition 1: Let y^t denote the vector of opinions at t . Let Δ_n be the set of vectors of non-negative weights $p = \{p_i\}_i$ with $\sum p_i = 1$. We have $y_i^t = B_i y^{t-1} + \gamma_i \varepsilon_i^t$ with $B_i \in \Delta_n$. So for any $p \in \Delta_n$, there exists $q \in \Delta_n$ such that:⁴⁵

$$p \cdot y^t = q \cdot y^{t-1} + \sum_i p_i \gamma_i \varepsilon_i^t. \quad (14)$$

Define $\underline{V}^t = \min_{p \in \Delta_n} \text{var}(p \cdot y^t)$. We have $V_i^t \geq \underline{V}^t$ and since $\gamma_i \geq \underline{\gamma}$ for all i , Equality (14) implies $\underline{V}^t \geq \underline{V}^{t-1} + \frac{1}{n} \underline{\gamma}^2 E(\varepsilon_i^t)^2$, hence the divergence.

Next let $\Gamma = (\gamma_i \xi_i)_i$. In matrix form, we have $\bar{y}^t = B \bar{y}^{t-1} + \Gamma$, which implies:

$$\bar{y}^t = \sum_k B^k \Gamma + B^t x$$

Since the network is connected, for some large enough k , B^k is a strictly positive probability matrix. Let π be the stationary distribution ($\pi B = \pi$). Consider a realization ξ such that $\pi \cdot \xi \neq 0$, say $\pi \cdot \xi > 0$. For k large enough, each row of B^k is close to π , implying that for k large enough, all $B^k \Gamma$ are positive and bounded away from 0, which proves the divergence of \bar{y}^t . ■

Before proving Proposition 2, we start with two standard results.

Lemma 1: 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 , $Y^t = X^0 + C Y^{t-1}$ converges to $H X^0$.

Lemma 2: Under Assumption 1, if $m_{i_0} > 0$, then for K large enough, $C = B^K$ and $\bar{C} = \bar{B}^K$ both satisfy the condition of Lemma 1, and $I - B$ and $I - \bar{B}$ have an inverse.

⁴⁵ $B_{ii} = 1 - \gamma_i$ and $B_{ij} = \gamma_i A_{ij}$. $q_i = \sum_j p_j B_{ji} = p_i (1 - \gamma_i) + \sum_j \gamma_j p_j A_{ji}$.

Proof of Proposition 2: We iteratively substitute in (11) to get:

$$\bar{y}^t = X^0 + C\bar{y}^{t-K}$$

where $X^0 = DX$ with $D \equiv I + B + \dots + B^{K-1}$, and $C = B^K$. By Lemma 2, Lemma 1 applies to C , so convergence of \bar{y}^t to \bar{y} is ensured, and $I - B$ has an inverse, which we denote H . We have $\bar{y} = HX$, hence the conclusion that \bar{y} does not depend on x_i when $m_i = 0$ (since X does not depend on x_i when $m_i = 0$).

Regarding the covariance vector, we iteratively substitute in (12) to get

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

where $\Lambda^0 = \bar{D}\Lambda$ with $\bar{D} = I + \bar{B} + \dots + \bar{B}^{K-1}$ and $\bar{C} = \bar{B}^K$. By Lemma 2, Lemma 1 applies to \bar{C} , so convergence of w^t to w is ensured, and $I - \bar{B}$ has an inverse which we denote \bar{H} . We have $w = \bar{H}\Lambda$, which is thus independent of initial opinions. ■

Before proving Corollary 1 and 2, we report another standard result. Let 1_N denote the column vector of dimension N for which all elements are equal to 1.

Lemma 3: *Let A^0 be a non-negative $N^0 \times N^0$ matrix and A^1 a non-negative $N^0 \times N^1$ matrix. Assume $I - A^0$ has an inverse and $A^0 1_{N^0} + A^1 \cdot 1_{N^1} = 1_{N^0}$. Then $P = (I - A^0)^{-1} A^1$ is a $N^0 \times N^1$ probability matrix, i.e., $P 1_{N^1} = 1_{N^0}$.*

We apply Lemma 3 to the case where $A^1 = M$ and $A^0 = B = (I - M)A$. By construction $A^0 1_N + A^1 \cdot 1_N = 1_N$ holds, which gives the following immediate corollary:

Corollary 3: *Assume $m_{i_0} > 0$ and let $P = (I - B)^{-1} M$. Then P is a probability matrix*

Proof of Corollary 1: When $m_i > 0$ for all i , the condition of Proposition 2 applies. Let $H = (I - B)^{-1}$ and $P = HM$. (10) can be rewritten as:

$$\bar{y} = M\tilde{x} + B\bar{y}$$

implying that $\bar{y} = P\tilde{x}$ with $P = (I - B)^{-1} M$, and P is a probability matrix by Corollary 3. ■

Proof of Proposition 3. 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 either use DG . Then, if player i uses DG as well, L_i^t diverges and by Corollary 2, for any $m_i > 0$, $y_i = \hat{x}_i = x_i + (1 - m_i)(\xi_i + R_i\xi_{-i})/m_i$. The variance of y_i thus decreases strictly with m_i .

Now assume that at least one player j chooses $m_j > 0$. Then, long-run expected opinions converge to \bar{y} . Now define $Y_j^i \equiv E\bar{y}_j X_i$ and the N vector of covariances $Y^i \equiv (Y_j^i)_j$. Also define $\bar{Y}_{jk} = E\bar{y}_j \bar{y}_k$ and $\bar{Y}_{jk} = (\bar{Y}_{jk})_{jk}$ as the N^2 vector of covariances. From (11) we have:

$$\begin{aligned} Y^i &= \Gamma^i + B^i Y^i \text{ and} \\ \bar{Y} &= \bar{\Gamma} + \bar{B} \bar{Y} \end{aligned}$$

where $\Gamma_j^i = EX_j X_i$ and $\Gamma^i = (\Gamma_j^i)_j$, and $\bar{\Gamma}_{jk} = \Gamma_j^k + (1 - m_k)A_k Y^j + (1 - m_j)A_j Y^k$ and $\bar{\Gamma} = (\bar{\Gamma}_{jk})_{jk}$.

Given our independence assumptions, Γ^i has just one positive element, $\Gamma_i^i = EX_i^2$ and for any $m_i < \underline{m}$, Γ_i^i is strictly decreasing in m_i . Next observe that B^i only has non-negative elements, and that B^i is non-increasing in m_i . So Y^i is strictly decreasing in m_i for all $m_i < \underline{m}$. Applying the same argument to the vector $Y^k \equiv EyX_k$, and since $\Gamma^k = EXX_k$ does not vary with m_i , we obtain that Y^j is non-increasing in m_i .

It follows that all terms $\bar{\Gamma}_{jk}$ for $k \neq i$ and $j \neq i$ are non-increasing in m_i and all terms $\bar{\Gamma}_{ij}$ are strictly decreasing on the range $m_i < \underline{m}$. Since \bar{B} is non-increasing in m_i , it follows again that for all $m_i < \underline{m}$, \bar{Y} is non-increasing in m_i and \bar{Y}_{ii} is strictly decreasing in m_i , and \bar{Y} is non-increasing in m_i .

We now examine the effect of m_i on the vector of covariances w where $w_{jk} = \lim E(y_j^t - \bar{y}_j^t)(y_k^t - \bar{y}_k^t)$. Recall $w = \Lambda + \bar{B}w$. Since Λ and \bar{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 $m_i < \underline{m}$, $L_i = \bar{Y}_{ii} + w_{ii}$ strictly decreases with m_i , and $\sum_k L_k$ also strictly decreases with m_i . ■

From Corollary 1, $\bar{y}_i = P_i \tilde{x}$ where P_i is the probability vector. We now characterize P_i and derive how each P_k varies with m_i . As a preliminary observation, we express \bar{y}_{-i} as an average of \tilde{x}_{-i} and \bar{y}_i .

Lemma 4: *For each $k \neq i$, there exists μ_{ji} and a probability vector $Q_j^i \in \Delta_{N-1}$, each independent of m_i , such that*

$$\bar{y}_j = (1 - \mu_{ji})Q_j^i \tilde{x}_{-i} + \mu_{ji} \bar{y}_i \tag{15}$$

The proof consists in using $\bar{y} = M\tilde{x} + B\bar{y}$ to solve \bar{y}_{-i} as an average over \tilde{x}_{-i} and \bar{y}_i . Since for each \bar{y}_j with $j \neq i$, all coefficients in B_j are independent of m_i , the result follows. Details are in Appendix B. Note that Lemma 4 immediately implies

$$\frac{\partial P_j}{\partial m_i} = \mu_{ji} \frac{\partial P_i}{\partial m_i} \text{ and } P_j = (1 - \mu_{ji})\bar{Q}_j^i + \mu_{ji}P_i \quad (16)$$

where $\bar{Q}^i \in \Delta_N$ with $\bar{Q}_{jk}^i = Q_{jk}^i$ and $\bar{Q}_{ji}^i = 0$.

In the expression $\bar{y}_i = m_i\tilde{x}_i + B_i\bar{y}$, we now substitute each \bar{y}_k obtained in Lemma 4. This permits us to get an expression of P_i that makes explicit the dependence on m_i . Specifically, we have (see Appendix B):

Lemma 5: *There exists λ_i and $r_i = (r_{ik})_k \in \Delta_{N-1}$ that only depend on A and m_{-i} such that*

$$P_{ii} = \frac{m_i}{1 - \lambda_i + m_i\lambda_i} \text{ and } P_{ik} = \frac{(1 - m_i)(1 - \lambda_i)r_{ik}}{1 - \lambda_i + m_i\lambda_i} \quad (17)$$

We now express the loss L_i^0 as a function of P_i . Define $W_k = \sigma_k^2 + \varpi_k(\frac{1-m_k}{m_k})^2$, where σ_k^2 is the variance of k 's initial opinion and ϖ_k the variance of k 's persistent component. We have:

$$L_i^0 = \sum_k (P_{ik})^2 W_k$$

Given the expression for P_i (see (17), optimization over m_i is simple, yielding Proposition 4 below. Next, Proposition 5 is also obtained as a simple corollary of (16) and (17):

Proposition 4: *Let $c_i = \sum_{k \neq i} (r_{ik})^2 W_k$. Player i 's optimal choice m_i is uniquely defined and satisfies:*

$$\frac{m_i}{1 - m_i} = \frac{\varpi_i + (1 - \lambda_i)^2 c_i}{\sigma_i^2 (1 - \lambda_i)}$$

Proof. Let $d_i = (\varpi_i + (1 - \lambda_i)^2 c_i) / (\sigma_i^2 (1 - \lambda_i))$. Rewrite L_i^0 as

$$L_i^0 = \sigma_i^2 \left(\frac{(m_i)^2 + (1 - m_i)^2 (1 - \lambda_i) d_i}{(m_i + (1 - m_i)(1 - \lambda_i))^2} \right)$$

Since d_i and λ_i do not depend on m_i , checking the first order condition yields $m_i / (1 - m_i) = d_i$, as desired. ■

Proposition 5. Assume $\frac{\partial L_i}{\partial m_i} \leq 0$. $\frac{\partial L_j}{\partial m_i} \leq 0$ for all j .

Proof. Using (16), and assuming no idiosyncratic noise, we rewrite

$$L_j = \sum_k ((1 - \mu_{ji})\bar{Q}_{jk}^i + \mu_{ji}P_{ik})^2 W_k$$

Since μ_{ji} , \bar{Q}_{jk}^i and W_k are independent of m_i for $k \neq i$, and since $\bar{Q}_{ji}^i = 0$, we obtain:

$$\frac{\partial L_j}{\partial m_i} = \mu_{ji}(1 - \mu_{ji}) \sum_{k \neq i} Q_{jk}^i \frac{\partial P_{ik}}{\partial m_i} W_k + (\mu_{ji})^2 \frac{\partial L_i}{\partial m_i}, \quad (18)$$

By Lemma 2, we further have $\frac{\partial P_{ik}}{\partial m_i} < 0$ for all $k \neq i$, which concludes the proof. ■

Comment: Long circle case (... $i - 1 \rightarrow i \rightarrow i + 1 \rightarrow i + 2 \dots$), that is $y_i = m_i x_i + (1 - m_i)y_{i-1}$. If the path from i to j is long, μ_{ji} is close to 0 (because i 's initial opinion has almost no influence on j), so $\frac{\partial L_j}{\partial m_i} \simeq 0$. Otherwise, j is a close neighbor and to the right of i . Then $Q_{jk}^i > 0$ for all $k \in \{i + 1, \dots, j\}$ and $Q_{jk}^i = 0$ otherwise, but for any $k \in \{i + 1, \dots, j\}$, $P_{ik} \simeq 0$ (because the path from any such k is long so the weight put by i on signals coming from these individuals is negligible), so from (18), $\frac{\partial L_j}{\partial m_i} \simeq 0$ whenever $\frac{\partial L_i}{\partial m_i} = 0$, i.e., private and social incentives coincide.

Example with modified protocol of communication. We illustrate below how changing protocol amounts to changing the weights γ_i . We consider two players and assume that player 1 updates every period, while player 2 updates every other three periods. Then, at dates t where 2 updates, we have:

$$\begin{aligned} y_1^t &= (1 - \gamma_1)^3 y_1^{t-3} + (1 - (1 - \gamma_1)^3) y_2^{t-3} \\ y_2^t &= (1 - \gamma_2) y_2^{t-3} + \gamma_2 y_1^{t-1} \\ &= (1 - \gamma_2) y_2^{t-3} + \gamma_2 ((1 - \gamma_1)^2 y_1^{t-3} + \gamma_2 (1 - (1 - \gamma_1)^2) y_2^{t-3}) \\ &= (1 - \gamma_2 (1 - \gamma_1)^2) y_2^{t-3} + \gamma_2 (1 - \gamma_1)^2 y_1^{t-3} \end{aligned}$$

So, the process evolves as if weights where $\gamma_1' = 1 - (1 - \gamma_1)^3 > \gamma_1$ and $\gamma_2' = \gamma_2 (1 - \gamma_1)^2 < \gamma_2$.

Appendix B (for on line publication)

Proof of Lemma 1: Consider the matrix $H^t = (h_{ij}^t)_{ij}$ defined recursively by $H^0 = I$ and $H^t = I + CH^{t-1}$. Let $z^t = \max_{ij} |h_{ij}^t - h_{ij}^{t-1}|$. 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 $z^t = \max_i |Y_i^t - Y_i^{t-1}|$, we obtain that Y^t has a limit Y which satisfies $(I - C)Y = X^0$, implying $Y = HX^0$. ■

Proof of Lemma 2: Call $Q_i^{K,i_0} \subset Q_i^K$ the set of paths of length K that start from i (to some j) and go through i_0 . For any such path, $\pi^B(q) \leq (1 - m_{i_0})\pi^A(q)$.⁴⁶ This implies

$$\sum_j C_{ij} \equiv \pi^B(Q_i^K) \leq (1 - m_{i_0})\pi^A(Q_i^{K,i_0}) + \pi^A(Q_i^K \setminus Q_i^{K,i_0}) < 1$$

where the last inequality follows from (13) and Q_i^{K,i_0} non empty for K large enough.⁴⁷

This implies that C satisfies the condition of Lemma 1, hence $I - C$ has an inverse. Let $D \equiv I + B + \dots + B^{K-1}$ 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.

Regarding \bar{C} , the argument is similar. We work on paths \bar{q} of pairs rather than paths q of individuals. Call \bar{Q}_{ij}^K the set of paths $\bar{q} = (q^1, q^2)$ of length K that start from ij (to some hk), \bar{Q}_i^{K,i_0} those for which q^1 goes through i_0 . We have

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

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

Proof of Lemma 3: Let $q = P1_{N^1} - 1_{N^0}$. $P = A^1 + A^0P$ so $P_{ij} = A_{ij}^1 + \sum_{k \in N^0} A_{ik}^0 P_{kj}$. Since $\sum_{j \in N^1} A_{ij}^1 = 1 - \sum_{j \in N^0} A_{ij}^0$ we have

$$q_i = \sum_{j \in N^1} P_{ij} - 1 = \sum_{k \in N^0, j \in N^1} A_{ik}^0 P_{kj} - \sum_{j \in N^0} A_{ij}^0 = A_i^0 q$$

⁴⁶In the general case (FJ rather than SFJ), $\pi^B(q) \leq (1 - m_{i_0}\gamma)\pi^A(q)$.

⁴⁷ N is finite, so K can be chosen large enough that Q_i^{K,i_0} is non-empty for all i .

implying that $q = A^0 q$, hence, since $I - A^0$ has an inverse, $q = 0$. ■

Proof of Lemma 4. Let $\widehat{m}_{ji} = 1 - (1 - m_j)(1 - a_{ji})$. Define M^i and \widehat{M}^i as $(N - 1) \times (N - 1)$ diagonal matrices where $\widehat{M}^i_{jj} = \widehat{m}_{ji}$ and $M^i_{jj} = m_j$. Let $\widehat{A}^i_{jk} = \frac{(1 - m_j)a_{jk}}{1 - \widehat{m}_{ji}}$ defined for all j, k different from i , and $g_{ji} = \frac{(1 - m_j)a_{ji}}{\widehat{m}_{ji}}$. \widehat{A}^i is a probability matrix. Also let $\widehat{X}_j = (1 - g_{ji})\tilde{x}_j + g_{ji}y_i$. By construction $y_j = \widehat{m}_{ji}\widehat{X}_j + (1 - \widehat{m}_{ji})\sum_{k \neq i} \widehat{A}^i_{jk}y_k$, which in matrix form gives

$$y_{-i} = \widehat{M}^i \widehat{X}_{-i} + \widehat{B}^i y_{-i}$$

where $\widehat{B}^i = (I - \widehat{M}^i)\widehat{A}^i$, which in turn yields $y_{-i} = R^i \widehat{X}_{-i}$ where $R^i \equiv (I - \widehat{B}^i)^{-1}\widehat{M}^i$ is a probability matrix (by Lemma 1). Letting $\mu_{ji} = \sum_{k \neq i} R^i_{jk}g_{ki}$ and $Q^i_{jk} = R^i_{jk}(1 - g_{ki})/(1 - \mu_{ji})$, we obtain the desired expression for y_j , and Q^i is by construction a probability matrix. Note that M^i , \widehat{M}^i and \widehat{A}^i depend on A and m_{-i} only, so the same is true for Q^i and μ_{ji} for all j . ■

Proof of Lemma 5. Using $\bar{y}_j = P_j \tilde{x}$ and (16) we get

$$\bar{y}_i = m_i \tilde{x}_i + (1 - m_i) \sum_{j \neq i} A_{ij} ((1 - \mu_{ji}) \bar{Q}_j^i + \mu_{ji} P_i) \tilde{x}$$

Letting $\lambda_i = \sum_j A_{ij} \mu_{ji}$, and $r_{ik} = \sum_{j \neq i} A_{ij} ((1 - \mu_{ji}) Q^i_{jk} / (1 - \lambda_i))$, and using $\bar{y}_i = \sum_k P_{ik} \tilde{x}_{ik}$, we get the desired expressions.⁴⁸ Since μ_{ji} and Q^i depend only on A and m_{-i} , the same is true for λ_i and r_{ik} . ■

Corollary 1 can be generalized to the case where a subset N^0 of agents has $m_i = 0$. Call N^1 the set of agents with $m_i > 0$, and accordingly define the vectors of expected long-run opinions \bar{y}^0 and \bar{y}^1 , and the vectors of persistent errors ξ^0 and ξ^1 . We have:

Corollary 2. Fix N^0 . There exists R and Q (defined independently of m) such that, for any m , there exists a probability matrix P such that $\bar{y} = P\hat{x} + Q\xi^0$ and $\hat{x}_i = \tilde{x}_i + (1 - m_i)R_i\xi^0/m_i$ for each $i \in N^1$.

Proof of Corollary 2: Let \tilde{x}^1 denote the vector of modified initial opinions of players in N^1 , and M^1 the restriction of M to N^1 . We have:

$$\bar{y}^0 = A^{00}\bar{y}^0 + A^{01}\bar{y}^1 + \xi^0 \quad (19)$$

$$\bar{y}^1 = M^1\tilde{x}^1 + (I^1 - M^1)(A^{10}\bar{y}^0 + A^{11}\bar{y}^1) \quad (20)$$

⁴⁸For example, focusing on the contribution of \tilde{x}_i , and since $\bar{Q}_{ji}^i = 0$, we get $P_{ii} = m_i + (1 - m_i)\sum_{j \neq i} A_{ij}\mu_{ji}P_{ii}$.

Under A , for K large enough, all agents in N^0 have a K -neighbor in N^1 , so $(A^{00})^K$ satisfies the condition of Lemma 1 and $I - A^{00}$ has an inverse, which we denote H^0 . We thus have:

$$\bar{y}^0 = P^0 \bar{y}^1 + H^0 \xi^0 \quad (21)$$

where $P^0 \equiv H^0 A^{01}$ is a probability matrix (by Lemma 3 and because $A^{01} \cdot 1_{N^1} + A^{00} \cdot 1_{N^0} = 1_{N^0}$).⁴⁹

Substituting \bar{y}^0 in (20), and letting $R = A^{10} H^0$ and $\hat{x}_i = \tilde{x}_i + (1 - m_i) R_i \xi^0 / m_i$, we get

$$\bar{y}^1 = M^1 \hat{x} + (I^1 - M^1) \hat{A} \bar{y}^1 \text{ where } \hat{A} \equiv A^{11} + A^{10} P^0$$

Since P^0 is a probability matrix, so is \hat{A} , and $C^1 = (I^1 - M^1) \hat{A}$ therefore satisfies the condition of Lemma 1 (as all $m_i > 0$ for $i \in N^1$). Letting $H^1 = (I^1 - C^1)^{-1}$, we get $\bar{y}^1 = P^1 \hat{x}$ where $P^1 = H^1 M^1$. Again, P^1 is a probability matrix because \hat{A} is a probability matrix and because $P^1 = M^1 + (I^1 - M^1) \hat{A} P^1$. Substituting \bar{y}^1 in (21) we finally get $\bar{y}^0 = P^0 P^1 \hat{x} + H^0 \xi^0$ and $\bar{y}^1 = P^1 \hat{x}$, which concludes the proof. ■

Proof of Proposition 6. Let $\bar{\gamma} = \max \gamma_i$ and recall:

$$w_{ij} = \sum_{h,k} B_{ih} B_{jk} w_{hk} + \Lambda_{ij} \quad (22)$$

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 upperbound on all w_{ij} of the form (see step 1)

$$w_{ij} \leq c \bar{\gamma} \quad (23)$$

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

$$\hat{w} \leq (1 - \underline{m}) C \hat{w} + \Gamma \quad (24)$$

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

Finally, to prove (iii), we consider two cases. Either $\bar{\gamma}$ is ‘‘small’’ and (23) applies, or we can separate individuals into a subgroup J where all

⁴⁹Indeed, for any $i \in N^0$, $\sum_{j \in N^0} A_{ij}^{00} + \sum_{j \in N^1} A_{ij}^{01} = \sum_{j \in N} A_{ij} = 1$

have a small γ_j , and the rest of them with significantly larger γ_j . In the later case, we redefine $\hat{w} = (w_{jk})_{j \in J, k}$ as the vector of co-variances involving some $j \in J$, and obtain inequality (24) 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)) $w_{ij} \leq c\bar{\gamma}$ with $c = \varpi_0/\underline{m}$.

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

$$\begin{aligned} \bar{V} = w_{ii} &\leq (1 - \gamma_i m_i)^2 \bar{V} + \gamma_i^2 (1 - m_i)^2 \varpi_0 \text{ hence} \\ \bar{V} &\leq \frac{\gamma_i (1 - m_i)^2}{m_i} \varpi_0 \leq \frac{\varpi_0 \bar{\gamma}}{\underline{m}} \end{aligned}$$

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

$$w_{ii} \leq \gamma_i p_{ii} \bar{v} + (1 - \underline{m}) \sum_k A_{ik} w_{ik} \quad (25)$$

$$w_{ij} \leq \gamma_j p_{ij} \bar{v} + (1 - \underline{m}) (p_{ij} \sum_k A_{ik} w_{kj} + p_{ji} \sum_k A_{jk} w_{ik}) \quad (26)$$

These inequalities are obtained by solving for w_{ij} in equation (22), 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 (25-26).

Step 3 (item (ii)). It is immediate from (25-26) that (24) holds with $C_{jk} \equiv A_{jk}$ and $\Gamma_j = p_{ij}\gamma_j\bar{v} + p_{ij}c\bar{\gamma} \leq p_{ij}\bar{\gamma}(\bar{v} + c) \leq d\gamma_i$ for all j , for some d , which permits to conclude that $\hat{w} \leq d\gamma_i/\underline{m}$.

Step 4 (item (iii)). Let $\varepsilon = \frac{1}{K|\text{Log}\underline{\gamma}|}$ 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 $j = 1, \dots, N-1$. Then $\bar{\gamma} < \underline{\gamma}/\varepsilon^{N-1}$, and for $\underline{\gamma}$ small enough, $\underline{\gamma}/\varepsilon^{N-1} < \varepsilon$, so $V_i \leq c\varepsilon < 1/|\text{Log}\underline{\gamma}|$.

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 (25-26) that (24)

holds with Γ such that, for any $j \in J$,

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

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 $\varepsilon(\bar{v} + c) \leq \underline{m}/|\text{Log}\underline{\gamma}|$, which concludes the proof. ■