NBER WORKING PAPER SERIES

CONSENSUS AND DISAGREEMENT:
INFORMATION AGGREGATION UNDER (NOT SO) NAIVE LEARNING

Abhijit Banerjee
Olivier Compte

Working Paper 29897
http://www.nber.org/papers/w29897

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2022

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w29897.ack

Consensus and Disagreement: Information Aggregation under (not so) Naive Learning
Abhijit Banerjee and Olivier Compte
NBER Working Paper No. 29897
April 2022
JEL No. D0

ABSTRACT

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

Abhijit Banerjee
Department of Economics, E52-540
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
banerjee@mit.edu

Olivier Compte
Paris School of Economics
48 Boulevard Jourdan
6th floor, office 54
Paris 75014
FRANCE
olivier.compte@gmail.com

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

Abhijit Banerjee and Olivier Compte

January 2022

Abstract

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

1 Introduction

As of May 2020, 41% of US Republicans were not planning to get vaccinated against Covid-19, as compared to 4% of Democrats.[1] We saw similar divergences in mask-wearing, social distancing etc, which protect against the disease. Since Covid-19 is a life-threatening ailment that had already taken more than 3.5 million lives so far world-wide, it is hard to think of these as being just empty gestures. There seems to be rather, a different reading of the facts on the ground; for example, in a Pew Research Center poll,[2] Re-

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

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

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

publicans were much more likely to say that Covid-19 is not a major threat to the health of the US population (53% compared to 15% of Democrats).

Given the increasing dominance of social media as a source of information it is natural to ask whether this divergence could be a result of the nature of social learning on networks. Indeed this is the argument that, for example, is made by Sunstein (2017) to explain the increasingly partisan nature of US politics. It turns out however that making this case theoretically is not as easy as one might imagine. Models of Bayesian social learning such as Acemoglu et al. (2011) propose relatively weak conditions on signals and network structure under which information is perfectly aggregated as the network grows to be very large. More recent work, in which agents repeatedly communicate (unlike in Acemoglu et al. (2011) where they communicate only once) include Mossel et al. 2015 who derive necessary conditions on the network structure under which Bayesian learning yields consensus and perfect information aggregation.[3] The general sense from this literature is that convergence to a consensus is likely even when the network exhibits a substantial degree of homophily (Republicans mostly talk to other Republicans) as long as everyone is ultimately connected.

This *Bayesian route* however requires that agents make correct inferences based on an understanding of all the possible ways information can transit through the network, which, at least for large networks, strains credibility.[4]

The alternative way to model learning on networks is to take a *non-Bayesian route*, which avoids these very demanding assumptions about information processing by postulating a simple rule that individuals use to aggregate own and neighbors' opinions. In recent years the economics literature has tended to favor the DeGroot (DG) rule, where agents update their current opinion by linearly averaging it with their neighbors' most recent opinions. As observed by DeMarzo et al. (2003), who brought it into the economics literature, the rule builds in a strong tendency towards consensus in any connected network, even when there is high degree of homophily and people put high weight on people like them.[5] Faced with this force towards consensus, Friedkin and Johnsen (1990) came up with a learning rule which

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

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

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

is similar to DG, but allows each individual to keep putting some weight on their own initial opinion.[6] This rule, for obvious reasons, does not lead to a consensus.

The first question we set out to answer here is whether the type of rule proposed by Friedkin-Johnsen would be chosen by individuals from a class of simple rules that include DG. To study this question we start from a broader class rules in the spirit of Friedkin-Johnsen (FJ) which can formally be written as

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

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

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

is the average report received by i from his neighbors (denoted by N_i) plus any error or noise in the transmission (or reception) of that signal.[7,8] When the weight m_i is 0, individual i is using a DG rule.[9]

Within this limited class of "natural" rules, we allow agents full discretion in the choice of rules and assume that each individual non-cooperatively selects the rule that best aggregates information (for her) in the long-run.[10] This is in the spirit of the approach advocated in Compte and Postlewaite (2018) to model mildly sophisticated agents.[11]

Our Result 1 says that each individual decision-maker will choose DG (and only DG) in the Nash equilibrium of the rule choice game if and only if

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

[7]The error term is an important ingredient of our analysis. We shall assume that ε_i^t has two components, a persistent one, drawn at the start of the process, and an idiosyncratic (time-independent) one.

[8]In Expression (1), all neighbors contribute symmetrically. In our formal exposition we will actually allow for non-symmetric weights but assume that relative weights across neighbors are fixed, not subject to optimization. This is to reduce the dimensionality of the rule-choice problem.

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

[10]That is, we assume that initial signals are correlated with some underlying state of the world, and that each i selects the rule for which the long-run opinion y_i is on average closest to the underlying state, assuming a quadratic loss function.

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

there is no noise in the transmission (or reception) of the signal. Otherwise a rule where $m_i > 0$ will be chosen and there will be no consensus even in the long run.

The fact that DG rules are an equilibrium absent any noise follows from the fact that if one person chooses $m_i > 0$ while the others don't, then everyone's opinion will converge to that one person's signal, because it is the only one that gets fed back in every period. This is in no one's interest, including person i, since this means that all the other signals are lost.

Moreover, DG (i.e. $m_i = 0$ for all i) is the only equilibrium absent noise. The reason for this is easy to see in the two-person case: through her decision rule, player i controls how the two initial signals are averaged in the steady state, and in the absence of noise, she can thus guarantee herself the efficient outcome. Since $m \neq 0$ leads to different averaging of initial signals hence inefficient information aggregation for at least one player, only DG can be an equilibrium. In fact this also shows that in the no-noise case the equilibrium choice of decision rules (i.e. the equilibrium vector of γ values chosen by everyone) is the one that generates optimal information aggregation. This result complements Golub-Jackson (2010) who show that when everyone does DG, information aggregation is almost perfect for large networks (under certain weak conditions) but generally imperfect in finite networks.[12]

The argument for the "only if" part starts from the observation that DG has the undesirable property that the variance of every decision-maker's belief grows without bound in the presence of any error or noise in communication (Proposition 1). Essentially 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, which themselves contain agent i's initial signal. This 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 error in every round. Given that the initial signals enter only at the beginning and the noise keeps coming in every period, it is no wonder errors come to dominate.

This intuition suggests that allowing each person's initial signal to come in with some weight every period would provide a countervailing force, which is what FJ allows: If $m_i > 0$ every decision-maker's average belief as well as

[12] The difference is that we assume endogenous weights γ_i.

its variance converges (Proposition 2). This explains why in the presence of noise, DG rules are actually dominated by rules with $m_i > 0$ and therefore only such rules will be chosen in Nash Equilibrium (Proposition 6)

Some disagreement therefore is to be expected in any Nash Equilibrium of the rule choice game. However there are two sources of divergence–a mechanical source arising from the presence of noise or errors but also the additional divergence that comes from always putting non-zero weight on one's initial signal (which is a choice, but one resulting from the noise). The next question is which of these is the main source of divergence. To simplify the analysis we start by assuming that the noise in transmission is fully persistent. This, we show, allows us to restrict the choice of rules to the sub-class of FJ rules where $\gamma = 1$. The only choice is now between rules which differ in m_i.

In this setting, our Result 2 shows that at least when the variance ω of persistent noise is close enough to zero, the second, non-mechanical, source dominates: in equilibrium, the weights m are comparable to $\omega^{1/3}$. The reason is that the errors that one makes are fed back into the network, and in a connected network, eventually bounce back: errors cumulate as a result of these echo effects and the long run variance is of the order of ω/m^2). Moreover, as a result, there is an incentive to significantly increase m.

We then ask whether there is too little or too much disagreement in any equilibrium relative to the social optimum. Result 3 shows that the equilibrium values of m_i are always lower than the socially optimal values. The reason is that there is no consensus in the long-run. When player i sets m_i optimally to minimize the variance of his long-run opinion y_i, he is not simultaneously minimizing the variance of y_j: while player j's long-run opinion is partially influenced by y_i, it is also "directly" influenced by the initial opinions of players other than i, and player j would prefer a marginal increase in m_i – this would generate a second-order increase in the variance of y_i, but a first-order reduction in the correlation between y_i and x_{-i}, which would reduce the variance of y_j.

Note that the discrepancy between social and private optima arises here only because of transmission noise and the lack of consensus that this generates. In the absence of noise, we already saw that DG rules will be chosen in equilibrium and will deliver both consensus and perfect information aggregation.

The failure of perfect information aggregation in our setting should be contrasted with Vives (1993,1997) who analyzes Bayesian social learning in a setting similar to ours (with agents receiving a noisy signal about current average choice in the population) and yet obtains long-run convergence to the truth. The reason is that agents perfectly process signals (as Bayesians are supposed to), without the kind of (processing) errors that we introduce in our model.

It should also be contrasted with Banerjee (1992) or Bhikchandani et al. (1992), where information aggregation fails due to coarse decisions (only two actions) and the network structure assumed.[13] Our model allows for arbitrary networks and the decisions are on a continuum. Moreover we allow agents to vary their decision rules continuously which guarantees perfect information aggregation in the absence of noise.

Finally, the discrepancy between social and private incentives takes the form of underweighting the private seed, which may be reminiscent of Vives (1997). The underweighting of private signals in Vives (1997) arises because agents fail to take into account a positive informational externality on learning: a stronger early reliance on private signals would speed up learning and benefit all. In our case, a higher reliance on private seeds (compared to equilibrium weights) improves welfare because this limits the correlation between information sources.[14]

The result that m_i is too low might suggest that there is always too little disagreement in equilibrium. This is true for two-person networks, but not in general. To see this consider a network where there are two dense clusters connected by one link (say). Such a network structure is not too dissimilar, for example, to the networks of Republicans and Democrats in the US, who mostly communicate with each other (Cox et al. (2020)). In this case, we show by example that there is a natural reason why lower m_i may be associated with a high degree of consensus within each cluster but extreme polarization across the groups, reminiscent of the situation of the Republicans and Democrats in the US. The general point, captured by Result 4 is that social efficiency requires the dispersion of opinions within and between subgroups to be of comparable magnitudes.

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

The rest of the paper is devoted to showing that these insights are robust. We first return to rule selection in the case where there are idiosyncratic shocks in information transmission in addition to permanent shocks. In this setting the speed of updating, γ_i, also plays a role. Slowing down updating

[13] Mossel et al. (2015) shows that with few states of the world, coarse communication can be circumvented for a large class of large networks, with consensus and almost perfect learning obtained.

[14] Averaging more strongly correlated sources hurts welfare.

by setting γ_i close to zero allows the agent to minimize the changes in opinions that result from these shocks, which average out over time. This is what Result 5 shows.

Since noise in transmission is central to the case we make for choosing $m_i > 0$, in the penultimate section of the paper we examine the robustness of our results to other ways of modeling the friction in the transmission of information. We start by examining the implications of agents adding a slant to the opinions they share–in other words adding noise that is biased in some direction. Recent results from survey experiments suggest this is a real problem–people on social media are more likely to pass on messages that they believe to be false than those that they believe to be true (Pennycook et al. (2021) and Vosoughi (2021)). We show that this does not produce any essential changes in our analysis, though there is a further shift towards reliance on one's own initial signal. A similar observation obtains when preferences are heterogenous and players have biased perceptions of others' preferences.

We next turn to the possibility of coarse communication–say each party only reports their current best guess about which of two actions is preferable, assuming that each has many neighbors. In this setting, the class of potentially "natural" rules include the infection models, studied in Jackson (2008) among (many) others, and the related class of models studied by Ellison and Fudenberg (1993, 1995), in which agents may rely on the popularity of a particular action among neighbors. We show that systematic errors in interpreting guesses by neighbors makes the long-run outcome from a DG-like rule entirely insensitive to the actual state of the world (Frick et al. (2019) report a related result for a proto-Bayesian rule), but this is not true for FJ-type rules.

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

We next allow for the possibility that not everyone speaks 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 under FJ rules, expected opinions are always independent of the communication protocol, whether or not there are errors in communication. However the long-run opinions under DG remain weighted averages of initial opinions, so the variance of long-run opinions induced by variations in protocols under DG remains bounded, unlike where there are transmission errors.

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

To end the introduction we briefly discuss the related literature. Our paper is related to and inspired by the recent upsurge of interest in the social learning with less than fully Bayesian agents. 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).

There are also a set of important recent papers that provide some justification for the DG rule in the absence of noise. Molavi et al. (2018) provide an axiomatic justification for DG-like (e.g. Log-linear learning) rules.[15] Dasaratha et al. 2020 argue that in their set-up the Bayesian rule is DG-like. Finally, Levy and Razin (2015) develop the Bayesian Peer Influence Paradigm to capture the idea of an almost Bayesian aggregation rule.

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.[16] 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

[15] In the work already mentioned by DeMarzo et al. (2003) and Golub and Jackson (2010), DG is justified by arguing that it coincides with the Bayesian rule in the static case.

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

recent opinion y_i^{t-1} and some summary perception z_i^{t-1} of his neighbors' opinions. Formally, this can be written as

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

where

$$z_i^t = A_i.y^t + \varepsilon_i^t \tag{2}$$

where y^t is the vector of all opinions at t, A_i is a row vector whose jth element A_{ij} is such that $\sum_j A_{ij} = 1$ and ε_i^t represents an error 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 $j \in N_i$. This average is then modified by 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 adjusted 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.[17,18]

To avoid technical difficulties once we give agents discretion in choosing their updating rule, we set $\underline{\gamma} > 0$ arbitrarily small and restrict attention to FJ rules where $\gamma_i \geq \underline{\gamma}$. We also assume that the matrix A of the A_i's is *connected* in the sense that for some positive integer k, the k^{th} power of A only has strictly positive elements, i.e., $A_{ij}^k > 0$ 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:

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

One can think of SFJ as a process that works like FJ, except that agents do not attempt to smooth out variations in their own opinion. In the absence of

[17]In 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 aiming to reduce the variance of his or her opinion. Of course, over time, a Bayesian would typically not keep that parameter constant, as the relative informative content of their own current opinion and that of others will in general not be constant.

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

idiosyncratic shocks on the perception of the opinions of others (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 which 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 shall discuss in Section 7.6 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 get to observe that change of opinion before they adjust their opinion in the following period. We will discuss what happens if we relax this assumption in Section 7.3.

2.2 Errors in opinion sharing

The term ε_i^t is an important ingredient of our model, meant to capture some imperfection in transmission.[19] It defines a distortion in what each individual "hears" that aggregates all the different sources of errors. Distortions may result from each individual being imprecise in expressing his or her opinion, or from an error in hearing or interpretation.

We assume that the error term 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.[20] The term ν_i^t is an *idiosyncratic* component drawn independently across agents and time.

[19] There has been several recent attempts to introduce noisy or biased transmission in networks. In Jackson et al. (2019), information is coarse (0 or 1), and noise can either induce a mutation of the signal (from 0 to 1 or 1 to 0) or a break in the chain of transmission (information 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 a systematic error in interpreting these actions because they have an erroneous model of the preferences of others. See Section 7.5.

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

We interpret ξ_i as a systematic bias that slants how opinions of others are *perceived*. For convenience, we assume that all error terms are homogenous across players and unbiased (that is, $E\xi_i = E\nu_i^t = 0$).[21] We let $\varpi_i = var(\xi_i)$ and $\varpi_0 = var(\nu_i^t)$ and assume that:

$$\varpi_i = \varpi$$

In most of the analysis, we think of the biases ξ_i as being drawn independently across players. We shall also discuss cases where these biases are positively correlated, reflecting a situation where a group of players is subject to a political bias.[22]

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, the initial opinions 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. We also have in mind that individuals choose a single rule to apply across different problems. This is why we focus on their ex ante performance.[23]

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

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

[22] Positively correlated errors may also arise when an agent has a bias that slants how her opinions are *expressed* to (all) her neighbors.

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

each other for all i and t and are also independent of θ. We assume that noise terms δ_i are unbiased, with variance $\sigma_i^2 > 0$. For convenience, except where we need to assume otherwise to make a specific point, we set $\sigma_i = 1$ for all i, but we do not actually need this assumption.

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

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

Define $\delta = (\delta_1, ..., \delta_n)$, $\xi = (\xi_1, ..., \xi_n)$ and $\nu_s = (\nu_{1s}, ..., \nu_{ns})$ 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 $\{d_{is}^t\}_{s=1}^t$.[24] It follows that

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

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

2.4 Methodological assumptions

The loss L_i depends on the profile of updating rules (m, γ), and our main methodological assumptions are that (i) there is a force towards the use of higher performing rules (e.g., justified by evolution or reinforcement learning), and (ii) in this quest for higher performing rules, each individual considers (or get feedback about) only a limited set of rules (e.g., the FJ class where each m_i belong to $[0, 1]$).

Formally, our analysis boils down to examining a rule choice game where, given the rules adopted by others, each agent aims at minimizing L_i (using the instrument m_i or γ_i available to her): the object of interest is the Nash

[24] θ enters additively in all opinions, and θ could thus be normalized to 0.

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

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

equilibrium of this rule choice game. Since L_i is an expectation across various realizations of initial signals and noise in transmission, we think of the person choosing one rule, parameterized by m_i and γ_i, to apply in many different life situations. These parameters are meant to capture some general features of opinion formation: specifically the *persistence* of initial opinions, and *speed of adjustment* of the current opinion.[26]

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

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

3 Some properties of the long run outcome of learning

It is well-known that in the DG case without noise ($m_i = 0$ for all i) learning converges to steady state values of y_i for all i and that these are all equal. The next subsections show that when there is noise there is still convergence under FJ as long as at least one person has $m_i > 0$, but not under DG. We then explore what determines the variance of the limit opinion in the case where such a limit opinion exists. In particular what part of it comes from the "signal"– the original seeds – and what part from the noise that gets added along the way?

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

3.1 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, for almost all realizations of the persistent component, y_i^t must diverge in expectation over time for all i.

To show this we fix x and ξ and define $\overline{y}_i^t = Ey_i^t$ and $V_i^t = var(y_i^t)$. We 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, $\lim \left|\overline{y}_i^t\right| = \infty$ for all i* and *x.*

For example, the proposition shows that an error ξ_1 in a single agent's perception may be enough to drive up the opinions of all: if $\xi_1 > 0$, say, the error creates a discrepancy with other's opinions, and each time others' opinions catch up, agent 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, there is a strictly positive vector of weights $\pi \in \Delta_n$ such that $\pi.B = \pi$,[27] so

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

Without noise, the limit weighted opinion $\pi.y$ coincides with the weighted initial opinion $\pi.x$. This explains why in the absence of noise the influence of initial opinions never dissipates (and also why all initial opinions matter – as $\pi >> 0$): 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, $\pi.y$ 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 $\pi.\overline{y}$ then diverge.

[27] This is because when $\gamma_i > 0$ for all i, $B = (B_i)_i$ is an irreducible probability matrix.

3.2 Anchored dynamics under FJ.

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

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

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

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

implying that $\overline{y}_1^t$ 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.[28]

The general argument for convergence runs as follows.[29] For any fixed x, ξ, the expected opinion evolves according to

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

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

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

[29] The argument follows Friedkin and Jensen (1999), extended to our setting.

[30] The 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 $m_{i_0} > 0$.

and the limit expected opinion $\overline{y}$ is the unique solution of

$$\overline{y} = X + (I - M)A\overline{y} \tag{3}$$

and this limit is independent of Γ. Next, defining $\eta^t = y^t - \overline{y}^t$ and $w^t_{ij} = E\eta^t_i\eta^t_j$, we have

$$\eta^t_i = (1 - m_i)\gamma_i\nu^t_i + B_i\eta^{t-1}$$

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

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

where $\overline{B}_{ij}$ is the row vector $(\overline{B}_{ij,hk})_{hk}$ with $\overline{B}_{ij,hk} = B_{ih}B_{jk}$ and Λ is the column vector with $\Lambda_{ii} = (1 - m_i)^2\gamma_i^2\varpi_0$. Proving convergence to the solution of

$$w = \Lambda + \overline{B}w \tag{4}$$

is also standard.[31]

Letting $\overline{L}_i \equiv var(\overline{y}_i)$, one immediate corollary of Equations (3) and (4) is that the loss L_i can be decomposed into

$$L_i = \overline{L}_i + V_i$$

where the variance $\overline{L}_i$ of the expected opinion $\overline{y}_i$ does not depend on γ (nor on ϖ_0) and the loss term V_i is proportional to the idiosyncratic component ϖ_0. When $\varpi_0 = 0$, $L_i = \overline{L}_i$, so this implies that 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 The dominance of noise under low m.

Although convergence is guaranteed when at least one player does not use DG, there is no discontinuity at the limit where *all* m_i get small: long-run opinions then become highly sensitive to the permanent component of the noise, and the variance induced by the idiosyncratic errors becomes very high. We have:

Proposition 3: Let $\overline{m} = \max m_i$. Then $\overline{L}_i \geq \frac{\varpi}{n}\frac{(1-\overline{m})^2}{\overline{m}^2}$ and $V_i \geq \frac{\varpi_0}{2n}\frac{\underline{\gamma}^2(1-\overline{m})^2}{\overline{m}}$.

[31] This is by the same logic as Footnote 30. 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 $\overline{B}^K$ has the contraction property.

The proof is in the Appendix. The lower bound on V_i is obtained as a simple extension of the proof of Proposition 1. We provide here the key step enabling us to obtain the lower bound on $\overline{L}_i$, as it highlights interesting properties of the FJ process.

The lower bound on $\overline{L}_i$ is obtained by showing that for given x, ξ, long-run expected opinions are weighted average of *modified initial opinions*, defined, whenever $m_i > 0$, as

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

When $m_i > 0$ for all i, one can write (using previous notations) $X = M\widetilde{x}$ and (3)) implies that each $\overline{y}_i$ is an average over modified initial opinions:

$$\overline{y} = M\widetilde{x} + (I - M)AM\widetilde{x} + ((I - M)A)^2 M\widetilde{x} + ... = P\widetilde{x} \tag{5}$$

where $P = (I - (I - M)A)^{-1}M$ is a probability matrix (see Lemma 4 in appendix). Intuitively, in each period, $\widetilde{x}_i$ can be thought of as the effective seed for individual i, and all long-run opinions are averages over effective seeds. For a fixed x_i, the variance of each $\widetilde{x}_i$ induced by the persistent component is bounded below by $\frac{\varpi(1-\overline{m})^2}{\overline{m}^2}$, so we obtain the desired lower bound.

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

$$\widetilde{\widetilde{x}}_i = \widetilde{x}_i + (1 - m_i)G_i\xi^0/m_i$$

where ξ^0 is the vector of perception errors of agents using DG, and G_i is a positive vector that only depends on the structure of the network and captures the influence of agents in N^0 on i (see Lemma 5). Our conclusion regarding $\overline{L}_i$ 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:

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

which further implies:

$$y_1^t = \sum_0^{K-1} \rho^k (m_1\widetilde{x}_1 + (1-m_1)m_2\widetilde{x}_2 + (1-m_1)\nu_1^{t-2k} + \rho\nu_2^{t-2k-1}) + \rho^K y_1^{t-2K} \quad (6)$$

which in turn gives us (7) and (8) below for the limits $\overline{y}_1$ and V_1:

$$\overline{y}_1 = p_1\widetilde{x}_1 + (1-p_1)\widetilde{x}_2 \text{ with } p_1 = m_1/(m_1 + (1-m_1)m_2). \quad (7)$$

$$V_1 = \varpi_0 \frac{(1-m_1)^2 + \rho^2}{1-\rho^2} \quad (8)$$

This example confirms that x_1 does not influence $\overline{y}_1$ 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 Understanding the difference between DG and FJ

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

The channel through which each player influences long-run opinions also differs substantially. In the absence of noise, and for a given network structure, relative influence in DG depends on relative speed of adjustment γ. More precisely, let $\rho \in \Delta_n$ be the vector such that $\rho.A = \rho$. When the γ_i's are identical across players, long-run opinions all converge to $\rho.x$, so ρ_i defines i's influence as determined by the network structure. When the γ_i's differ, long-run opinions all converge to $\pi.x$ where $\pi \in \Delta_n$ and

$$\pi_i/\pi_k = (\rho_i/\gamma_i)/(\rho_k/\gamma_k), \quad (9)$$

which explains how both the network and speeds of adjustment determine influence.[32]

[32] To see why (9) holds, observe that, up to a multiplicative constant, π is the unique solution of $\pi.B = \pi$ where $B = I - \Gamma + \Gamma A$. Now observe that $\pi = \rho\Gamma^{-1}$ is one such solution. Note that (9) also explains how speeds of adjustment γ can be set to induce the efficient weighting of signals π^* (which minimizes the variance of $\pi.x$).

In contrast, under FJ, only the m_i's (and the structure of the network) affect the expected long-run opinions $\overline{y}$. The speeds of adjustment γ have no effect on expected long-run opinion, they only affect the variance induced by idiosyncratic noise.

Regarding influence under FJ, it can be shown that at the limit where all m_i's are very small, all long-run expected opinions are close to another and close to $p.\widetilde{x}$ where all p_i are proportional to $m_i \rho_i$, that is:

$$p_i/p_k = m_i \rho_i/(m_k \rho_k),$$

Thus, close to the limit, m_i plays the same role as $1/\gamma_i$ does in DG and consensus obtains. As the m_i's go up however, consensus disappears: players "agree to disagree".

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

(c) **On the source of change in opinion**: One way to assess the difference between DG and 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:

$$Y_i^t = \gamma_i D_i^{t-1} \quad \text{(DG)}$$
$$Y_i^t = (1 - m_i) Z_i^t \quad \text{(SFJ)}$$

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 (unwanted) change of opinion will trigger a sequence of further changes – it will be partially incorporated in player 2's opinion, and therefore come back to player 1 again. This is what we call an *echo effect*. But, when $m_i > 0$ for at least one player, the echo effect will be smaller than the initial impact and will get even smaller over time. Hence over all it won't blow up. If all m_i are small however, the echo effects are not dampened enough, and the consequence is a high sensitivity of the final opinion to the magnitude of the errors.

3.5 How large are the echo effects?

In this subsection we develop the idea of echo effects more formally. The next proposition derives an expression for the cumulated error resulting from the echo effect and highlights the impact of the choice of low m_j by others on i's payoff's and strategic possibilities. This will be key in understanding the equilibrium impact of noise.

We consider an agent i who sets $m_i > 0$, which ensures that long-run expected opinions $\overline{y}$ are well-defined.

$$\overline{y}_i = m_i x_i + (1-m_i)\xi_i + (1-m_i)\widehat{y}_i \text{ with } \widehat{y}_i \equiv \sum_{k\neq i} A_{ik}\overline{y}_k \tag{10}$$

Equation (10) describes how i's expected opinion builds on the opinion $\widehat{y}_i$ of a (fictitious) *composite neighbor* who aggregates the opinions $\overline{y}_k$. Letting $\widetilde{A}^i_{kj} = \frac{A_{kj}}{1-A_{ki}}$, we rewrite (10) to describe how each opinion $\overline{y}_k$ builds on $\overline{y}_i$:

$$\overline{y}_k = m_k x_k + (1-m_k)\xi_k + (1-m_k)A_{ki}\overline{y}_i + (1-m_k)(1-A_{ki})\sum_{j\neq k,i} \widetilde{A}^i_{kj}\overline{y}_j \tag{11}$$

So in effect, in incorporating the opinion $\widehat{y}_i$, player i is partially incorporating her own opinion $\overline{y}_i$. In other words, the opinions that i gets from others are partially echoes of her own opinion. The following Proposition measures this echo effect for i.

Proposition 4: *Let M^i (resp. α^i) be the diagonal $N-1$ matrix for which $M^i_{kk} = m_k$ for $k \neq i$ (resp. $\alpha^i_{kk} = A_{ki}$). Define the matrix $Q^i =$*

$(I-(I-M^i)(I-\alpha^i)\widetilde{A}^i)^{-1}$ *and vector* R^i *such that* $R^i_j = \sum_k A_{ik}Q^i_{kj}$*. We have* $r_i \equiv \sum_{j\neq i} R^i_j m_j \leq 1$ *and*

$$\overline{y}_i = p_i x_i + (1-p_i)(\widehat{x}_i + \widehat{\xi}_i) \text{ where} \tag{12}$$
$$p_i = \frac{m_i}{m_i + (1-m_i)r_i},\ \widehat{x}_i = \frac{\sum_{j\neq i} R^i_j m_j x_j}{\sum_{j\neq i} R^i_j m_j}$$
$$\text{and } \widehat{\xi}_i = \frac{1}{r_i}(\xi_i + \sum_{j\neq i} R^i_j \xi_j (1-m_j))$$

In the long-run, player i's expected opinion is thus an average between own seed x_i and some *composite seed* $\widehat{x}_i$ (an average over the others' seeds) plus an error term $\widehat{\xi}_i$. This error term $\widehat{\xi}_i$ captures the cumulated error that i faces because of echo effects. At the limit where m_{-i} tends to 0, r_i tends to 0 and $\overline{y}_i$ tends to $x_i + \frac{1-m_i}{m_i}(\xi_i + \sum_{j\neq i} R^i_j \xi_j)$,[33] with echo effects thus rising without bound when m_i gets small as well.[34]

The next section builds upon this proposition and the previous propositions to characterize the equilibrium of the rule choice game. We are interested in the rule that people will choose in equilibrium both in terms of how much divergence in opinions we would expect to observe and how it relates to the efficient choice of rules.

4 Choosing the rule

4.1 When there is no noise

A direct implication of Proposition 4 is that in the absence of noise, the equilibrium must be DG and that in equilibrium, information aggregation must be perfect. Recall that in this case DG does converge to a stationary outcome. Formally, define π^* as the vector of weights on seeds that achieve perfect information aggregation, i.e., $\pi^* = \arg\min_\pi var(\sum_k \pi_k x_k)$, and let $v^* = var(\pi^*.x)$. We have

[33] In the limit where all m_j tend to 0, y_i remains well-defined because for a given network (characterized by A), Q^i and R^i are uniformly bounded (with a well-defined limit when all m_j tends to 0). In this case $(1-p_i)\widehat{x}_i$ tends to 0, and as expected, player i is not influenced by others' initial opinions.

[34] Note that at the other limit where $m_j = 1$ for all $j \neq i$, $R^i_j = A_{ij}$, $r_i = 1$ and $\overline{y}_i = m_i x_i + (1-m_i)(\widehat{x}_i + \xi_i)$. Other players do not incorporate i's opinions, and player i is thus not subject to echo effects.

Proposition 5: *In the absence of transmission errors, the equilibrium must be DG. In addition, in equilibrium,* $y_i = \pi^*.x$.

In other words, as long as there is no noise, we get perfect agreement in opinions in equilibrium and perfect information aggregation. As mentioned in introduction, the main difference with De Marzo et al. (2003) and Golub and Jackson (2010) is that we allow for endogenous weights γ_i, and for any connected network, this is enough to obtain efficiency in equilibrium.[35]

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

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

4.2 Rule choice when there is noise

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

Proposition 6: *Let* $\underline{m} = \varpi/(1+\varpi)$. *Any* (m_i, γ_i) *with* $m_i < \underline{m}$ *is dominated by* $(\underline{m}, \gamma_i)$*, from the individual and social point of view.*

Regarding the choice of the individually optimal rule, Proposition 6 builds on two ideas. First, if all other players use DG, then for agent i,

[35] As observed in the seminal paper by De Marzo et al. (2003), for generic networks, for any finite n, DG rules will not implement the efficient weighting of seeds π^*, 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)).

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

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

This sequence of propositions are summarized as

Result 1. *As long there is no noise in communication the unique equilibrium choice of rules is DG and this delivers perfect information aggregation. When there is any noise, DG rules are strictly dominated by rules that put a minimum amount of weight on the initial signal of the decision-maker in each period and will never be chosen.*

The next Proposition provides further characterization of the privately optimal choice of m_i. To simplify exposition, we focus on the case where the idiosyncratic component is absent ($\varpi_0 = 0$), so the outcome does not depend on γ. For the purpose of explaining how m_i is affected by the quality of initial signals and transmission errors, we allow both to vary across people. Recall σ_k^2 is the variance of k's initial opinion and ϖ_k the variance of $k's$ persistent component. Also let $W_k = \sigma_k^2 + \varpi_k(\frac{1-m_k}{m_k})^2$. The following Proposition is an immediate Corollary of Proposition 4:

Proposition 7. *Player i's optimal choice m_i satisfies:*

$$\frac{m_i}{1-m_i} = \frac{\varpi_i + \sum_{k \neq i}(R_k^i m_j)^2 W_k}{\sigma_i^2 r_i}$$

where R_k^i and r_i, defined in Proposition 4, 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 (σ_i^2) goes down or that of anyone else (σ_k^2) goes up. It also shifts up when the variance of the error term 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.

4.3 How big is the divergence in opinions?

Result 1 has the obvious implication that full consensus is never going to be an equilibrium when there are persistent errors–there are in fact two sources of deviation, the error itself (which mechanically prevents consensus) and the extra weight m_i on one's initial signal (which fuels further divergence unless m_i is small.)

Result 2 below shows that because of echo effects, the optimal weight put on one's own seed tends to be relatively large, comparable to $\varpi^{1/3}$.[36] As a result when ϖ is small, the extra weight on one's own seeds becomes the preponderant source of dispersion. Assuming no idiosyncratic noise ($\varpi_0 = 0$), we have:

Result 2: *For any given finite network and any $\varpi > 0$ small, all m_i and $p_i - \pi_i^*$ are positive and comparable to $\varpi^{1/3}$.*

One immediate implication of Result 2 is that all error terms $\widehat{\xi}_i$ have variance comparable to $\varpi^{1/3}$, thus directly driving the variance of opinions above the efficient level v^* by a term at least comparable to $\varpi^{1/3}$. There is another source of inefficiency in equilibrium, the fact that seeds are not efficiently weighted. But that inefficiency is comparable to $\varpi^{2/3}$: a socially optimal choice of weights m_i would trade-off more inefficient weighting (larger m) against decreasing the variance of the echo effect.

One consequence of this observation is that even between two neighbors, the dispersion of opinion is at least comparable to $m^2 \simeq \varpi^{2/3}$.

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

[36] When we mention that m is comparable to say $g(\varpi)$, we mean $m = O(g(\varpi))$, which means that $|m/g(\varpi)|$ has a finite limit when ϖ tends to 0.

[37] The proof also shows that there is no slack: all m_i and losses (away from efficient information aggregation) are comparable to $\varpi^{1/3}$ in equilibrium. The weights m_i cannot be larger for the same reason that the equilibrium without error terms must be DG: each player sets the weighting p_i of own seed x_i optimally, and this creates a force towards

Note that Result 2 focuses on the case where variances are small. When the m_i's rise, the relative weights on seeds eventually depart from efficient weighting sufficiently that this fuels a further rise in m. The examples in Section 5 will illustrate this.

4.4 Privately versus socially optimal choices

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

Result 3. *At any Nash equilibrium, a marginal increase of m_i by any player i would increase aggregate social welfare.*

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

$$\overline{y}_j = (1-\mu_{ji})C^{ji}\widetilde{x}_{-i} + \mu_{ji}\overline{y}_i \tag{13}$$

where C^{ji} is a probability vector and $\mu_{ji} \in (0,1)$,[39] with μ_{ji} and C^{ji} both independent of m_i.[40]

The expression above highlights that when player i chooses m_i optimally (for him) to minimize the variance of $\overline{y}_i$, he may not be minimizing the variance of $\overline{y}_j$. There is no consensus, thus incentives may not be aligned.

To check incentives and the direction of departure that improves welfare, we use (13) to separate the loss L_j into three terms:

$$L_j = (1-\mu_{ji})^2 var(C^{ji}\widetilde{x}_{-i}) + \mu_{ji}L_i + 2(1-\mu_{ji})\mu_{ji}Cov(C^{ji}\widetilde{x}_{-i}, \overline{y}_i). \tag{14}$$

When m_i is raised above $i's$ private optimum, there is no effect on the first term. There is a second-order effect on the second term (because we start at

optimal information aggregation

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

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

[40]The expression follows from solving the system of equations (11) for all $\overline{y}_k$ with $k \neq i$ (see Lemma 6 in Appendix). The coefficients of this system do not involve m_i, so C^{ji} and μ_{ji} are independent of m_i. Intuitively, μ_{ji} characterizes the influence of $\overline{y}_i$ on j's opinion (through paths that reach j from i), and $(1-\mu_{ji})C_k^{ji}$ characterizes the direct influence of seed $\widetilde{x}_k$ on $\overline{y}_j$ (through paths that reach j from k *without going through player i*). These contributions are independent of m_i because information does not transit through i along these paths.

$i's$ private optimum). The last term is what creates a discrepancy between private and social incentives.

This last term depends on the correlation between seeds other than i ($\widetilde{x}_{-i}$) and the opinion of i ($\overline{y}_i$), and how these separately contribute to y_j (given C^{ji} and μ_{ji}). When all $m_k \in (0,1)$, this term is positive even without persistent errors because in a connected network all x_k covary with $\overline{y}_i$.[41] The effect is amplified with persistent errors (because $\widetilde{x}_k = x_k + \frac{1-m_k}{m_k}\xi_k$, and ξ_k and $\overline{y}_i$ covary), and even more so when persistent errors are positively correlated.[42]

Increasing m_i has no effect on μ_{ji} and C^{ji}, but when m_i increases, the influence of each $k \neq i$ on i's opinion is reduced, and the correlation between $\overline{y}_i$ and $\widetilde{x}_k$ is also reduced. Overall, starting at a Nash equilibrium, L_j goes down when m_i is raised.

5 Equilibrium, efficiency and polarization in simple networks.

In this section we explore the quantitative significance of our results through a set of simple examples. We are particularly interested in whether a small amount of noise can lead to large distortions of information aggregation, how much of extra information loss comes from non-cooperative behavior, and the connection between individuals over-weighting their own signal and polarization at the population level. We focus on the case where there are only persistent errors, unless mentioned otherwise.

We start with the example of a large directed circle, where the intuition for Result 3 would suggest that there would be no divergence between private and social incentives: for a fixed m and a long enough loop, an individual does not need to worry about his opinion traveling back to him. This is confirmed by our analysis. Nevertheless there is very substantial information loss even in the presence of small amounts of noise.

5.1 Large circle case.

Social optimum Consider a large circle where information transmission is directed and one-sided: player i communicates to player $i+1$, who com-

[41] A single loop is enough. When there is a path from i to j, $\mu_{ji} \in (0,1)$, and when there is a path from j to i, $C_j^{ji} > 0$ and x_j covaries with $\overline{y}_i$.

[42] Although we do not prove it, we expect the terms $Cov(x_k, \overline{y}_i)$ to be preponderant in equilibrium, relative to $\frac{1-m_k}{m_k}Cov(\xi_k, \overline{y}_i)$, hence also expect Result 3 to hold even with negatively correlated errors.

municates to $i+2$, and so on.[43] Long-run opinions satisfy

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

Hence if player i chooses m_i and all other players choose m, we have

$$y_i = m_i\widetilde{x}_i + (1-m_i)(Z + (1-m)^{n-1}y_i) \text{ where } Z = m\sum_{k=0}^{n-2}(1-m)^k\widetilde{x}_{i-1-k}. \quad (15)$$

One can use this expression to derive y_i and its variance $J(m)$ when all choose the same m and n is set arbitrarily large.[44] When transmission errors ξ_i are independent,

$$J(m) = \frac{m}{2-m}(1+\mathcal{X}(m)), \quad (16)$$

where $\mathcal{X}(m) = \varpi\frac{(1-m)^2}{m^2}$ represents the effect of cumulated errors. Expression (16) captures the trade-off between improving information aggregation (which calls for reducing all m_i close to 0) and reducing the amplification of communication errors (which calls for increasing all m_i). It is minimized at m^{**} for which $J'(m^{**}) = 0$, that is, m^* solving $\frac{m^2}{1-m} = \varpi$.[45]

Note that for small variance, $m^{**} \simeq \varpi^{1/2}$, which seems to contradict Result 2. This is because we consider here the limit where the network becomes large: with independent errors and for this specific network, transmission errors tend to compensate one another to some extent.[46,47] With correlated errors, no such compensation occurs and $m^{**} \simeq (4\varpi)^{1/3}$.[48]

Private and social incentives. We first check directly that private and social incentives coincide when persistent errors are independent across players. One may use (15) to write:

$$L_i = (m_i)^2(1+\mathcal{X}(m_i)) + (1-m_i)^2J(m) = (m_i)^2 + (1-m_i)^2(\varpi + J(m)) \quad (17)$$

[43] Player $n+1$ coincides with player 1.

[44] Since $\sigma_i^2 = 1$ for all i, the social optimum is symmetric (See Appendix B)

[45] It is easy to check that at the social optimum, $J(m^{**}) = m^{**}$.

[46] When m is positive, the influence of neighbors living many steps away becomes very small. But when m is small, many neighbors remain influential, and with comparable influence.

[47] Nevertheless, for most pairs i, j in the network, each i and j build their opinions on the opinions of disconnected sets of players, so the dispersion of opinion is significant, comparable to $\varpi^{1/2}$.

[48] With perfectly correlated errors, Expression (16) becomes $J(m) = \frac{m}{2-m} + \mathcal{X}(m)$.

At a symmetric Nash equilibrium m^*, private incentives require $m_i = (1 - m_i)(\varpi + J(m^*))$ with $m_i = m^*$. Since $J(m) = m^2 + (1-m)^2(\varpi + J(m))$, this implies $J'(m^*) = 0$, so m^* is also a social optimum.

To connect this result with the intuition provided earlier on the source of discrepancy between private and social incentives, consider $j = i + k$, that is, j is k communication steps away from i. We have

$$y_j = m \sum_{s=0}^{k-1} (1-m)^s \widetilde{x}_{j-s} + \mu_{ji} y_i \text{ where } \mu_{ji} = (1-m)^k$$

As explained earlier, the magnitude of the terms $\mu_{ji} Cov(\widetilde{x}_{j-s}, \overline{y}_i)$ is key. When n is large, either k is large and $\mu_{ji} = (1-m)^k$ is negligible, or $n - k$ is large and then $Cov(\widetilde{x}_{j-s}, y_i)$ is small (because $j - s$ is at least $n - k$ communication steps away from i). So as the circle gets very large, private and social incentives coincide.

Note that this alignment of social and private incentives fails with correlated errors. When errors are perfectly correlated, $\widetilde{x}_i$ and Z are correlated, and this adds a term $2(1-m_i)^2 \varpi (1-m)/m$ to the expression of L_i in (17).[49] Similarly, even when j is many communication steps away from i, the covariance $Cov(\widetilde{x}_j, \overline{y}_i)$ is positive. Computations show that in equilibrium, $m^* \simeq (2\varpi)^{1/3} < m^{**}$.

Information aggregation and welfare loss. Despite the convergence in private and social incentives, the loss in welfare is significant relative to the benchmark case where players would observe (with transmission errors) the initial opinion of each of the other players in the network and perfectly aggregate these signals.[50] Under this benchmark and with transmission errors independent across players, each player's loss would be close to 0 when the circle grows large.[51]

Under FJ however, the weight m^* and the loss $J(m^*)$ remain bounded away from 0 as the circle grows large, and $N = 1/J(m^*)$ measures 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 $\varpi = 0.05$, $m^* = 0.2$, $J(m^*) = 0.2$, and a player's long-run information

[49] With $J(m)$ thus now equal to $m^2 + (1-m)^2(\varpi + J(m) + 2\varpi(1-m)/m))$.

[50] For 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.

[51] In the benchmark, 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 weighting of these opinions would lead to an opinion with variance $(\sum 1/v_k)^{-1}$, which goes to 0 with k.

is comparable to her having received only five independent signals (hence only four additional signals), out of the infinite pool that is available when the circle is arbitrarily large. We draw N as a function of ϖ.

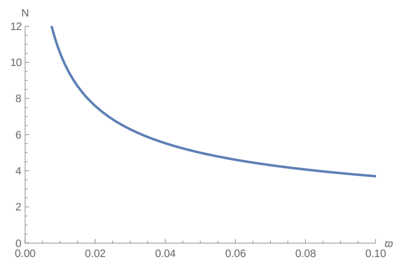

Figure 1: Information aggregation: large circle

5.2 Two-player case.

We next move to another canonical example, where there are just two players. Here the simple loop creates powerful echo effects (which favor a higher value of m) and strong correlations between one's own seed and the opinion of the other, which tend to drive a divergence between the equilibrium and optimal choices of m's.

Social optimum. Recall that with two players (see (7))

$$y_i = p_i \widetilde{x}_i + (1 - p_i)\widetilde{x}_j \text{ with } p_i = m_i/(m_i + (1 - m_i)m_j) \qquad (18)$$

With independent transmission errors this 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 and $\mathcal{X}(m)$ is as before. The total social loss is $L = L_1+L_2$.

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 ($m = m_1 = m_2$), $p_i = \frac{1}{2-m}$ and the social loss is:

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

As in the long circle case, the expression highlights a trade-off between decreasing m for information aggregation purposes ($I(p)$ is minimized at

$p = 1/2$),[52] and increasing m to limit the effect of cumulated communication errors (when $\varpi > 0$ and m is small, communication errors are hugely amplified).

Welfare is maximized for an m^{**} that optimally trades off these two effects and the socially efficient weight m^{**} (which minimizes L) can be significantly different from 0 even when ϖ is small. Specifically, for $\varpi = 0.0001$, $m^{**} = 0.13$ and for $\varpi = 0.001$, $m^{**} = 0.21$. Furthermore, for ϖ small, $m^{**} \simeq (4\varpi)^{1/4}$.[53]

Nash Equilibrium. We now assume that individuals choose their rules non-cooperatively. Equation (18) can be rewritten as

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

This means that in effect, i uses m_i (and hence p_i) to optimally weight two independent signals: her own seed x_i and a noisy correlate of the other player's seed x_j, where the noise term $\widehat{\xi_i}$ is independent of m_i.[54] Hence she should set

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

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

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

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

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

When ϖ is small, we have $m^* \simeq (2\varpi)^{1/3}$ when ϖ is small.

Figure 2 plots the best responses for $\varpi = 0.01$.

[52] This is under our assumption that seeds are equally informative ($\sigma_i^2 = 1$ for all i).

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

[54] Note that the choice of m_i affects $\widehat{\xi}_j$, with a low m_i amplifying the noise term $\widehat{\xi}_j$, which is another reason why the social optimum requires higher m.

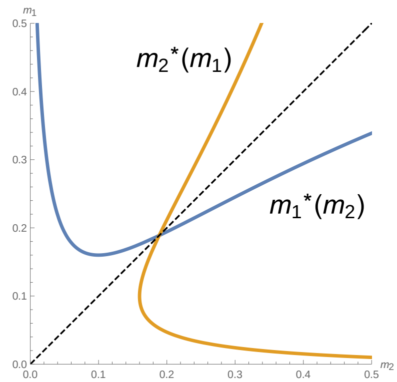

Figure 2: Best responses, $\varpi = 0.01$

Since $m^* \simeq (2\varpi)^{1/3}$ when ϖ is small while $m^{**} \simeq (4\varpi)^{1/4}$, the ratio of m^{**} to m^* explodes when ϖ is small.

5.3 The star network

We consider a network consisting of K peripheral players, each of whom has a seed x_i, and a central player 0 who has no seed but aggregates the opinions of all the peripheral players with some error (ξ_0). We assume that the central player uses a DG rule since he has no seed, and his long-run opinion is $y_0 = \frac{1}{K}\sum_k y_k + \xi_0$. Simple computations (see Appendix) show that long-run opinions have a simple expression, similar to the one we derived for the two-player case, which in turn yields simple closed form solution for the Nash outcome and social optimum.

When all players other than i use m, Proposition 4 gives:

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

with $p_i = \frac{m_i(1+(K-1)m)}{m_i+(K-1)m}$, a composite seed $\widehat{x}_i = \frac{1}{K-1}\sum_{k\neq i} x_k$ (the average seed of agents other than i) and a cumulated error term

$$\widehat{\xi}_i = \frac{K}{m(K-1)}(\xi_0 + \frac{1}{K}\sum_k \xi_k) + \xi_i - \frac{1}{K-1}\sum_{k\neq i}\xi_k,$$

Social optimum. When all players use the same m, the loss for a player is

$$L = p^2 + (1-p)^2\frac{1}{K-1} + (1-p)^2 E\widehat{\xi}_i^2 \tag{19}$$

where $p = \frac{1}{K} + \frac{K-1}{K}m$. With independent errors,

$$E\widehat{\xi}_i^2 = \frac{K}{m^2(K-1)}(\varpi_0 + \frac{\varpi}{K}(1 + (1-m)^2)) \tag{20}$$

and (19) illustrates again the tradeoff between decreasing m for information aggregation purposes (the first terms are minimized for $p = 1/K$) and raising m to limit the effect of cumulated communication errors. When ϖ_0 and ϖ are small, $m^{**} \simeq (2(1-\frac{1}{K})(\varpi_0 + \frac{2}{K}\varpi))^{1/4}$.[55]

Nash Equilibrium. Regarding private incentives, $\widehat{\xi}_i$ does not depend on m_i, which implies that m_i is optimally set so that

$$\frac{p_i^*}{1-p_i^*} = \frac{1}{K-1} + E\widehat{\xi}_i^2$$

condition which can be used to solve for the symmetric equilibrium weight m^*:

$$\frac{m^*}{1-m^*} = (1-\frac{1}{K})E\widehat{\xi}_i^2$$

When errors are independent, from (20), this gives an equilibrium weight $m^* \simeq (\varpi_0 + \frac{2}{K}\varpi))^{1/3}$, which is significantly below the social optimum.

Note that in contrast with the large circle case, m^* remains at least comparable to $\varpi_0^{1/3}$ even when K grows large and the transmission errors are independent across players. With independence, the increase in the size of the star helps because the persistent component of peripheral players are averaged out, so m^* is then essentially driven by the variance of ξ_0. But since the central player's opinion affects all peripheral players, the cumulated error still grows like ϖ_0/m^2 with m, so m^* remains comparable to $\varpi^{1/3}$.

5.4 Implications for the divergence of opinions

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

[55] With correlated errors, $E\widehat{\xi}_i^2 = \frac{K}{K-1}(\varpi_0 + \varpi)$ and one obtains $m^{**} \simeq (2(1-\frac{1}{K})(\varpi_0 + \varpi))^{1/4}$.

In the presence of noise, two things may happen. A divergence of long-run opinions $\overline{y}$ away from y^*, which means a *divergence of opinions across networks* (fueled by within-networks echo effects), and a *dispersion of opinions within networks* (resulting from the use of FJ rules with non-zero weights m). This section argues that there is a connection between consensus within subgroups (low dispersion) and polarization (high divergence across subgroups).

To fix ideas, we consider below the case of two large disconnected star networks modeled as above.[56] This description generally fits the maps of social networks in the US population with the two stars representing Democrats and Republicans (Cox et al. 2020). Assuming that in each star network all peripheral players use the same weight m, we have

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

Hence the dispersion of opinions between two peripheral players is

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

For m not too small, the dispersion d increases with m. Across the networks however, opinions are independent (conditional on θ) and the dispersion of opinion D between two individuals belonging to different networks is equal to twice the variance, which yields, assuming independent errors ξ_i,

$$D = 2L = 2Ey_i^2 = 2(m^2 + \frac{\varpi_0}{m^2}) \simeq d + \frac{4\varpi_0}{d}$$

In other words, the social optimum requires a high enough dispersion of opinion within subgroups (with m set so that $d = 2\varpi_0^{1/2}$ and m comparable to $\varpi_0^{1/4}$) and $D \simeq 2d$.

Result 4: *Social efficiency requires that the dispersion of opinions within and across subgroups be of comparable magnitude. Too much consensus within subgroups (low d) favors polarization across subgroups (high D).*

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

[56]Result 4 below would also hold if the set of cross-star links were a vanishingly small proportion of the total number of links.

6 Choosing among a richer class of rules

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. The next Proposition examines the effect of γ on the variance V_i induced by the idiosyncratic errors, as well as incentives for an individual to choose a low γ_i:

Result 5: *Fix $\underline{m}$. We have:*
(i)There exists c such that for any $\gamma > 0$ and $m \geq \underline{m}$, $V_i \leq c \max \gamma_j$.
(ii) For any $\gamma_{-i} > 0$, there exists c such that for all $m \geq \underline{m}$, $V_i \leq c\gamma_i$.

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

While the incentive is clear, a technical issue potentially arises if players wish to set γ_i arbitrarily small, as $\gamma_i > 0$ is an open interval.[57] We address this issue in the Appendix by showing that when all γ's are restricted to be above some lower bound $\underline{\gamma}$, any player i can secure a loss V_i no larger than $1/ \mid \log \underline{\gamma} \mid$ by choosing $\gamma_i = \underline{\gamma}$.[58] So if $\underline{\gamma}$ is small, V_i must be small in equilibrium. This also implies that investigating the properties of the game without idiosyncratic noise is a good enough approximation when $\underline{\gamma}$ is small.[59]

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

[57]Note that the limit opinion-formation process where γ_i tends to 0 is *not* the process where $\gamma_i = 0$ (under which no change in opinion would occur).

[58]$1/ \mid \log \underline{\gamma} \mid$ is small number when $\underline{\gamma}$ is small.

[59]In particular, if (m^*, γ^*) is an equilibrium of the game, then m^* is an $\varepsilon-$equilibrium of the game with no idiosyncratic noise, with ε comparable to $1/ \mid \log \underline{\gamma} \mid$.

7 Extensions and interpretations

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

7.1 Biased persistent errors

We have so far assumed that the persistent error is drawn from a distribution that is mean zero. One can however imagine settings where it is more reasonable to assume that the persistent error is biased, centered on ξ_i^0 for player i, for example because some individuals are systematically 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:

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

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, theses shifts imply as before (see (5)) that $\overline{y}_i$ is a linear combination P_i of the modified opinions $\widetilde{x}_j$ where

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

If c_i can be adjusted optimally, each i can set c_i to fully offset the systematic bias in transmission and this turns out to be optimal.[60] If c_i cannot be adjusted (e.g., $c_i = 0$ for all i) however, then the bias ξ_i^0 adds a fixed

[60] Letting 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 $L_i = L_i^0 + (P_i C)^2$ where C_i is a choice variable for i, 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.

amount to all the cumulated error terms $\widehat{\xi}_j$ defined in Proposition 4, thus increasing the variance of these error terms and generating for each player further incentives to increase m.

These effects are obviously reinforced if all players in the network have biases that have the same sign, while otherwise, and depending on the structure of the network and the distribution of biases within the network, they could partially offset one another.

7.2 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:

$$\theta_i = \theta + b_i \tag{21}$$

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

$$x_i = \theta_i + \delta_i \tag{22}$$

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 errors) 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 further 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 nature and extent of heterogeneity, and players are unable to tune c_i to each realization of the heterogeneity. Then the problem is formally equivalent to one where preferences are homogenous and a persistent transmission term $A_i\beta_i$ is added.[61]

The general take-away should be that there are many potential sources of errors 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.

7.3 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.[62] 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.[63] Imperfect communication is modeled as before, through the addition of an error term ε_i^t that slants what i hears. Together these give us

$$z_{i,j}^t = y_j^{t-1} + \varepsilon_i^t \text{ if } j \in N_i^t$$
$$z_{i,j}^t = z_{i,j}^{t-1} \text{ if } j \in N_i \backslash N_i^t$$

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.[64] We continue to assume FJ updating.

As before, for fixed x, ξ, we define $\overline{y}_i^t = Ey_i^t$. We have:

[61] One 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.

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

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

[64] We 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).

Proposition 8: *Assume at least one player, say i_0, updates according to FJ with $m_{i_0} > 0$. Then for any fixed x, ξ, $\overline{y}^t$ converges and the limit vector of expected opinions $\overline{y}$ is independent of the protocol.*[65]

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

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

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

So, the process evolves as if weights where $\gamma_1' = 1-(1-\gamma_1)^3 > \gamma_1$ and $\gamma_2' = \gamma_2(1-\gamma_1)^2 < \gamma_2$. This means that with DG rules, changes in the frequencies with which players communicate amount to changes in the values of γ_i (when you hear less often from others, your opinion changes more slowly, effectively reducing γ_i). And when communication is noiseless, these changes modify long-run opinions: if γ_i goes down, long-run opinions get closer to i's opinions (see Section 3.4 Equation (9)).

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

7.4 Uncertainty over the precision of initial signals.

We examine here another variation of the model, assuming that the precision of initial signals is a random variable. We will argue that in the absence of transmission errors, this type of shock does not affect the performance of DG and therefore, unlike where there is say noise in transmission, there is no incentive for players to use the instrument m_i.

[65] So long as the condition in footnote 63 holds.

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

To see why, recall that under DG, the consensual long-run opinion is a weighted average of initial opinions, with weights proportional to ρ_i/γ_i (see (9)). So if the μ_i's are proportional to ρ_i, the weights become proportional to ρ_i/γ_i, hence proportional to $1/\sigma_i^2$, implying that perfect aggregation obtains for each vector of realization $(\sigma_1, ..., \sigma_n)$.

7.5 Coarse communication

In the social learning literature, it is common to focus on choice problems where there are two possible actions, and the information being aggregated is which of the two is being recommended by others. Coarse communication is potentially a source of herding, but when agents have many neighbors, the fraction of players choosing a given action may become an accurate signal of the underlying state. We explain below how our model can accommodate an economic environment of this kind, and we use this to relate our findings to Ellison and Fudenberg (1993,1995) and Frick et al. (2019).

Preferences are heterogenous as in Section 7.2, with $\theta_i = \theta + b_i$ characterizing i's value from choosing 1 over 0, so the optimal action a_i^* is 1 when $\theta_i > 0$, 0 otherwise.[66] Agent i knows b_i but does not know θ perfectly. He has an initial opinion $x_i = \theta + \delta_i$ and aggregates opinions of others to sharpen his assessment of θ. Assume the b_i's are drawn from identical distribution g (and cumulative denoted G) with full support on $\mathcal{R}$.

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

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

Long-run opinions clearly depend on the inference rule assumed, but there is a natural candidate. Let $h(y) \equiv G(-y) = \Pr(y + b_i < 0)$ be the

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

fraction of agents that choose $a = 0$ when their opinions are all equal to y.[67] The inverse function $\phi \equiv h^{-1}$ is a natural candidate for the inference function ψ_i: if others' opinions are correct and equal to θ, a fraction $f \simeq h(\theta)$ choose $a = 0$ and $h^{-1}(f)$ is a good proxy for θ. Of course this assumes that agents know the distribution over preferences. In the spirit of our previous analysis, let's assume that

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

where ξ_i is a persistent error in interpreting f.[68,69] To fix ideas, we assume correlated errors ($\xi_i = \xi$ for all i) with variance ϖ.

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

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

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

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

[67] $h(y) \equiv \Pr(y + b_i < 0) = G(-y)$

[68] Following Frick et al. (2018), ξ_i could stem from an erroneous prior $g_i \neq g$, with agents using the inference function $\psi_i = h_i^{-1}$ where $h_i(\theta) = G_i(-\theta)$. The difference $\xi(f) = \psi(f) - \phi(f)$ is a systematic (and possibly non-uniform) error in making inferences. With preferences centered on $\bar{b}$, and agent having an erroneously translated prior centered on $\bar{b}_i$, the error is uniform and equal to $\xi_i = \bar{b}_i - \bar{b}$.

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

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 changes in others' opinions (rather than discrepancies between others' and own opinions), they end up being more robust, avoiding this particular form of fragility.

7.6 Non-stationary weights.

The updating processes that we consider have stationary weights. Agents do not attempt to exploit the possibility that early reports possibly reveal more information than 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.

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

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

8 Concluding remarks

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

One premise of our model is that everyone has a well-defined initial signal.[70]. 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

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

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 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 *changes* in others' opinions, rather than because of a difference between one's and others' opinions. In particular, players' opinions may differ in the long run. One can clearly envisage applying a similar idea to beliefs about the state of world. With two states for example, one could let $y_i^t = \ln p_i^t/(1-p_i^t)$ measure the belief of i over the underlying state,[71] $x_i \equiv y_i^0$ the initial belief, and assume an updating process for y_i in the spirit of SFJ rules:

$$y_i^t = m_i x_i + (1-m_i) z_i^{t-1} \text{ with } z_i^{t-1} = A y^{t-1} + \varepsilon_i^t$$

These rules keep beliefs anchored on initial beliefs x_i, with own beliefs only varying in response to variations in others' opinions (i.e., $y_i^t - y_i^{t-1} = (1-m_i)(z_i^{t-1} - z_i^{t-2})$). This contrasts with Bayes-inspired rules, which would incorporate initial beliefs only once and then adjust beliefs as a response to the spread *(used as a new informative signal)* between other's opinions (such as z_i^{t-1}) and own opinion (y_i^{t-1}), an adjustment process conducive to consensus. By allowing diverse long-run beliefs and keeping agents from over-reacting to (*the signal embodied in*) variations in other's beliefs (i.e., with m_i not too low), some FJ-like updating rules turn out to be less sensitive to transmission noise or biases (hence eventually better information aggregators) than Bayes-inspired rules.

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

References

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

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

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

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

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

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

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.

Cox, Daniel A., Ryan Streeter, Samuel J. Abrams and Jacqueline Clemence (2020), "Socially distant: How our divided social networks explain our politics," Survey Center on American Life.

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

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

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

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

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

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

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

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

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

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

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.

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

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

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

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

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

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

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

Sunstein, C. R. (2017) *#Republic: Divided Democracy in the Age of Social Media*. Princeton, NJ: Princeton University Press.

Vosoughi S, Roy D, Aral S. (2018, Mar 9) "The spread of true and false news online." *Science*. 359(6380):1146-1151.

Appendix A

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

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

and, whenever $m_i > 0$, we let $\widetilde{x}_i = x_i + \xi_i(1 - m_i)/m_i$ denote the modified initial opinion, and $\widetilde{x} = (\widetilde{x}_i)_i$ the vector.

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

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

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

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

$$y^t = \Gamma(X + (I - M)\nu) + By^{t-1} \tag{23}$$

$$\overline{y}^t = \Gamma X + B\overline{y}^{t-1}, \tag{24}$$

from which we obtain:

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

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

$$w^t = \Lambda + \overline{B}w^{t-1} \tag{25}$$

We relegate to Appendix B the proof that the matrices $H \equiv \sum_{k \geq 0} B^k$ and $\overline{H} \equiv \sum_{k \geq 0} \overline{B}^k$ are well-defined (see Lemma 1 and 2) or equivalently that the inverse $(I - B)^{-1}$ and $(I - \overline{B})^{-1}$ are well-defined, which implies that $\overline{y}^t$ and w^t have well-defined limits

$$\overline{y} = H\Gamma X \text{ and } w = \overline{H}\Lambda. \tag{26}$$

The limit $\overline{y}$ must solve $\overline{y} = \Gamma X + B\overline{y}$, which yields

$$\overline{y} = X + (I - M)A\overline{y}. \tag{27}$$

The long-run expected opinion is thus independent of Γ (Proposition 2).[72]

We also relegate to Appendix B the proof that long-run opinions are weighted average of suitably modified initial opinions. Specifically, we will show**:**

Lemma 4. *Assume $m_i > 0$ for all i. Then for each i, there exists* $P_i \in \Delta_n$ such that for all x, ξ, $\overline{y}_i = P_i.\widetilde{x}$.

We will also show how Lemma 4 extends to cases where a strict subset $N^0 \subsetneq N$ of agents follow a DG rule ($m_i = 0$). Denote by ξ^0 the vector of persistent errors of players. We have

Lemma 5. *Assume* $m_i > 0$ for all $i \notin N^0$. *There exist G and Q (defined independently of m) and a probability matrix P such that $\overline{y} = P\widetilde{\widetilde{x}} + Q\xi^0$ and $\widetilde{\widetilde{x}}_i = \widetilde{x}_i + (1-m_i)G_i\xi^0/m_i$ for each $i \notin N^0$.*

We now turn to the proof of our main Propositions.

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:[73]

$$p.y^t = q.y^{t-1} + \sum_i p_i \gamma_i \varepsilon_i^t. \tag{28}$$

Define $\underline{V}^t = \min_{p \in \Delta_n} var(p.y^t)$. We have $V_i^t \geq \underline{V}^t$ and since $\gamma_i \geq \underline{\gamma}$ for all i, Equality (28) 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 $\overline{y}^t = B\overline{y}^{t-1} + \Gamma$, which implies:

$$\overline{y}^t = \sum_{0 \leq k < t} B^k \Gamma + B^t x$$

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 $\pi.\xi \neq 0$, say $\pi.\xi > 0$. 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 $\overline{y}^t$.■

Proof of Proposition 3. The lower bound on $\overline{L}_i$ follows immediately from Lemma 4 and 5. We focus here on loss V_i induced by the idiosyncratic shocks. Recall

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

[72] For the purpose of computing $\overline{y}_i$ and its variance $\underline{L}_i$, we can thus set $\Gamma = I$ (i.e., $\gamma_i = 1$ for all i, the SFJ rule) and $B = (I-M)A$.

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

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

$$p.\eta^t = q.\eta^{t-1} + \sum_i \gamma_i(1-m_i)p_i\nu_i^t \text{ and } \sum_i q_i \geq 1 - \underline{m} \tag{29}$$

Define $\underline{V}^t = \min_{p\in\Delta_n} var(p.\eta^t)$. Note that $V_i^t \geq \underline{V}^t$. Since $var(q.\eta^{t-1}) \geq (1-\underline{m})^2\underline{V}^{t-1}$, Equality (29) implies $\underline{V}^t \geq (1-\underline{m})^2\underline{V}^{t-1} + \frac{1}{n}\underline{\gamma}^2(1-\underline{m})^2\varpi_0$, which yields the desired lower bound.■

Proof of Proposition 4: Assume $m >> 0$ so $\widetilde{x}_j$ is well-defined for all j.[74] For $j \neq i$ let $X_j = m_j\widetilde{x}_j + (1-m_j)A_{ji}y_i$ and $c_j^i = m_j + (1-m_j)A_{ji}$. (11) can be written in matrix form to obtain, by definition of Q^i, $\overline{y}_{-i} = Q^iX$. Note that if $\widetilde{x}_j = 1$ for all j and $\overline{y}_i = 1$, then $\overline{y}_k = 1$ for all k, so $\sum_{j\neq i} Q_{kj}^i c_j^i = 1$ for all k, which implies

$$\sum_{j\neq i} Q_{kj}^i(1-m_j)A_{ji} = 1 - \sum_{j\neq i} Q_{kj}^i m_j, \tag{30}$$

and, since Q^i is a positive matrix,[75] $\sum_{j\neq i} Q_{kj}^i m_j \leq 1$, so $\sum_{j\neq i} R_j^i m_j \leq 1$. (30) further implies

$$\overline{y}_k = \sum_{j\neq i} Q_{kj}^i m_j\widetilde{x}_j + (1 - \sum_{j\neq i} Q_{kj}^i m_j)\overline{y}_i, \tag{31}$$

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

$$\overline{y}_i = m_i\widetilde{x}_i + (1-m_i)(\sum_{j\neq i} R_j^i m_j\widetilde{x}_j + \overline{y}_i(1 - \sum_{j\neq i} R_j^i m_j) \tag{32}$$

which, since $m_j\widetilde{x}_j = m_jx_j + (1-m_i)\xi_j$ and $r_i = \sum_{j\neq i} R_j^i m_j$, gives the desired expressions (12) for $\overline{y}_i$, $\widehat{x}_i$, p_i and $\widehat{\xi}_i$.■

Proof of Proposition 5: There are two parts in this proof. We first prove that the $m_i's$ cannot be positive. Next we show that the equilibrium outcome must be efficient. Recall $\pi^* = \arg\min_\pi v(\sum_k \pi_k x_k)$ is the efficient weighting of seeds and $v^* \equiv v(\pi^*.x)$.

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

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

Assume by contradiction that $m_j > 0$. Then (12) implies that $m_i > 0$ for all i, so $m >> 0$. Next, from (32) we obtain $\widehat{y}_i = r_i \widehat{x}_i + (1 - r_i) y_i$, hence substituting y_i,

$$\widehat{y}_i = (1 - r_i) p_i x_i + (1 - (1 - r_i) p_i) \widehat{x}_i. \tag{33}$$

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

$$v(y_i) < v(\widehat{y}_i) \le \max_{j \neq i} v(y_j), \tag{34}$$

where the second inequality follows from $\widehat{y}_i$ being an average of the y_j's. Since (34) cannot be true for all i, we get a contradiction. The equilibrium must thus be DG.

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

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

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

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

which implies

$$\pi_i = \pi_i^* + \frac{(1 - p_i)(1 - \pi_i^*)}{\sigma_i^2} (v(\widehat{x}_i) - v_i^*) \tag{35}$$

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

Before showing Proposition 6, we start with two intermediate results that we also use to prove Result 3:

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

$$\overline{y}_j = (1 - \mu_{ji}) C^{ji} \widetilde{x}_{-i} + \mu_{ji} \overline{y}_i \tag{36}$$

Proof: This immediately follows from Expression (30) in the proof of Proposition 4.■

Lemma 7: if $\frac{\partial \overline{L}_i}{\partial m_i} \leq 0$, *then* $\frac{\partial \overline{L}_j}{\partial m_i} < 0$ *for all* j.

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

$$\frac{\partial \overline{L}_j}{\partial m_i} = (\mu_{ji})^2 \frac{\partial \overline{L}_i}{\partial m_i} + \mu_{ji}(1-\mu_{ji}) \sum_{k \neq i} C_k^{ji} \frac{\partial Cov(\widetilde{x}_k \overline{y}_i)}{\partial m_i}$$

We substitute $\overline{y}_i = p_i x_i + (1-p_i)(\widehat{x}_i + \widehat{\xi}_i)$ (see (12)). Since $\widetilde{x}_k$ and x_i are independent, and since $\widehat{x}_i, \widetilde{x}_k$ and $\widehat{\xi}_i$ do not depend on m_i, we get

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

The terms $\frac{\partial p_i}{\partial m_i}$ and $Cov(x_k \widehat{x}_i)$ are positive, and so are the terms $Cov(\widetilde{x}_k \widehat{\xi}_i)$ when persistent errors are independent or positively correlated. The sum on the right side is thus positive (and the effect is amplified with errors), which proves Lemma 7.[76]■

Proof of Proposition 6. Let $\underline{m} = \varpi/(1+\varpi)$. We show that DG and all strategies $m_i < \underline{m}$ are dominated by $\underline{m}$.

Assume first that all other players use DG. Then, if player i uses DG as well, L_i^t diverges and by Proposition 4, for any $m_i > 0$, $\overline{y}_i = \widehat{x}_i = x_i + (1-m_i)(\xi_i + R^i \xi_{-i})/m_i$. The variance of $\overline{y}_i$ thus decreases strictly with m_i.

Now assume that at least one player j chooses $m_j > 0$. Then $\overline{L}_i = p_i^2 + (1-p_i)^2 v(\widehat{x}_i + \widehat{\xi}_i)$. Whether persistent errors are independent or fully correlated, the variance of $\widehat{\xi}_i$ is at least equal to ϖ/r_i^2, which implies that $\overline{L}_i$ strictly decreases for all p_i such that $\frac{p_i}{1-p_i} < \varpi/r_i^2$, hence also for any m_i such that $\frac{m_i}{1-m_i} < \varpi/r_i$, and from Lemma 7, we conclude that $\overline{L}_j$ increases as well (on this range of m_i).

We now examine the effect of m_i on the vector of covariances w where $w_{jk} = \lim E(y_j^t - \overline{y}_j^t)(y_k^t - \overline{y}_k^t)$. Recall $w = \Lambda + \overline{B}w$. Since Λ and $\overline{B}$ are non-increasing in m_i and Λ_{ii} is strictly decreasing in m_i, w_{ii} strictly decreases with m_i, and w is non-increasing in m_i. Combining all steps, over the range

[76]With negatively correlated errors the terms $Cov(\widetilde{x}_k \widehat{\xi}_i) = \frac{1-m_k}{m_k} E\xi_k \widehat{\xi}_i$ could be negative. However, we suspect that in equilibrium, with the m_k's set optimally, the terms $Cov(x_k \widehat{x}_i)$ remain the preponderant ones – with small errors for example, $Cov(\widetilde{x}_k \widehat{\xi}_i)$ is comparable to $\varpi/m^2 \simeq \varpi^{1/3}$, hence we expect that Result 3 continues to hold.

$m_i < \underline{m}$, $L_i = \overline{L}_i + w_{ii}$ strictly decreases with m_i, and $\sum_k L_k$ also strictly decreases with m_i.■

Proof of Proposition 7. When there are no idiosyncratic errors, $y_i = \overline{y}_i$ and since x_i and $\widehat{x}_i + \widehat{\xi}_i$ are independent variables, unaffected by m_i, player i optimally sets p_i such that $\frac{p_i}{1-p_i} = \frac{v(\widehat{x}_i+\widehat{\xi}_i)}{v(x_i)}$, from which we derive the desired expression for m_i.■

Proof of Result 2:
Step 1: lowerbounds on $\overline{m}^i \equiv \max_{j\neq i} m_j$.
With transmission errors, optimal weighting of x_i and $\widehat{x}_i$ implies

$$\frac{p_i}{1-p_i} = \frac{v(\widehat{x}_i) + v(\widehat{\xi}_i)}{\sigma_i^2} \tag{37}$$

and (35) becomes

$$p_i = \pi_i^* + \frac{(1-p_i)(1-\pi_i^*)}{\sigma_i^2}(v(\widehat{x}_i) - v_i^* + v(\widehat{\xi}_i)) \tag{38}$$

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

While (38) implies a lower bound on p_i, as (35) did, there is a major difference here with the no noise case where DG is used by all: p_i is the weight that i puts on own seed, but since there is no consensus, the sum $\sum_i p_i$ is not constrained to be below 1. Nevertheless, when all m are small, players are close to consensus, and $\sum_i p_i$ is close to 1, and this allows us to bound $v(\widehat{\xi}_i)$ (and the difference $v(\widehat{x}_i) - v_i^*$), as we now explain.

From Proposition 4, each opinion y_i may be written as $y_i = P^i x + (1 - P_i^i)\widehat{\xi}_i$, where P^i is a weighting vector (such that $P_i^i = p_i$). (31) implies that when all m are small, the vectors P^i must be close to one another: seeds must be weighted in almost the same way, and differences in opinions are mostly driven by the terms $\widehat{\xi}_i$. Specifically, let $\overline{m}^i = \max_{j\neq i} m_j$. (31) implies that for all $k \neq i$,

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

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

$$1 - p_i = \sum_{k\neq i} P_k^i \geq \sum_{k\neq i} p_k - Kc\overline{m}^i \geq 1 - \pi_i^* - Kc\overline{m}^i \tag{39}$$

which, combined with (38) yields

$$\overline{m}^i \geq d(v(\widehat{x}_i) - v_i^* + \frac{\varpi}{(\overline{m}^i)^2}). \tag{40}$$

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

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

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

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

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

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

It only remains to check that m_{i_0} cannot be large. From (39), $p_{i_0} \leq \pi_{i_0}^* + O(\varpi^{1/3})$, and since $p_{i_0} \geq \frac{1}{1 + r_{i_0}/m_{i_0}}$, we conclude that all m_i (and thus $\overline{m}^i$) are comparable to $\varpi^{1/3}$, which further implies that all variances $v(\widehat{\xi}_i)$ are comparable to $\varpi^{1/3}$.

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

Proof of Result 3: this follows from Lemma 7 since at equilibrium $\frac{\partial \overline{L}_i}{\partial m_i} = 0$. ■

Appendix B (for on-line publication)

We first prove that the matrices $H \equiv \sum_{k\geq 0} B^k$ and $\overline{H} \equiv \sum_{k\geq 0} \overline{B}^k$ are well-defined (Lemma 1 and 2), and obtain Proposition 2 as a Corollary. Next we prove that long-run expected opinions are weighted average of suitably-defined modified initial opinions (Lemma 3 to 5).

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, $Y^t = X^0 + CY^{t-1}$ converges to HX^0.*

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

Proof of Proposition 2: We iteratively substitute in (24) to get:

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

where $X^0 = D\Gamma X$ with $D \equiv I + B + ... + B^{K-1}$, and $C = B^K$. By Lemma 2, Lemma 1 applies to C, so convergence of $\overline{y}^t$ to $\overline{y}$ is ensured, and $I - B$ has an inverse, which we denote H. We have $\overline{y} = H\Gamma X$, hence the conclusion that $\overline{y}$ does not depend on x_i when $m_i = 0$ (since X does not depend on x_i when $m_i = 0$). From (27), we also conclude that $\overline{y}$ is not affected by γ.

Regarding the covariance vector, we iteratively substitute in (25) to get

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

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

We now report standard results (Lemma 3 and Corollary 3 below) enabling us to show that long-opinions are weighted average of suitably defined modified opinions (Lemma 4 and 5). 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 $A^0 1_{N^0} + A^1.1_{N^1} = 1_{N^0}$. Then $P = (I - A^0)^{-1}A^1$ is a $N^0 \times N^1$ probability matrix, i.e., $P1_{N_1} = 1_{N_0}$.*

We apply Lemma 3 to the case where $A^1 = M$ and $A^0 = B = (I - M)A$. By construction $A^0 1_N + A^1.1_N = 1_N$ 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.*

Lemma 4. *Assume $m_i > 0$ for all i. Then for each i, there exists* $P_i \in \Delta_n$ such that for all x, ξ, $\overline{y}_i = P_i.\widetilde{x}$.

Proof of Lemma 4: When $m_i > 0$ for all i, the condition of Proposition 2 applies. Let $H = (I - B)^{-1}$ and $P = HM$. (23) can be rewritten as:

$$\overline{y} = M\widetilde{x} + B\overline{y}$$

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

Lemma 4 can be generalized to the case where a strict subset $N^0 \subsetneq N$ 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 $\overline{y}^0$ and $\overline{y}^1$, and the vectors of persistent errors ξ^0 and ξ^1. We have:

Lemma 5. *Assume $m_i > 0$ for all $i \notin N^0$. There exist G and Q (defined independently of m) and a probability matrix P such that $\overline{y} = P\widetilde{\widetilde{x}} + Q\xi^0$ and $\widetilde{\widetilde{x}}_i = \widetilde{x}_i + (1 - m_i)G_i\xi^0/m_i$ for each $i \notin N^0$.*

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

Before turning to the proof of Lemma 2, we define *sequences*, *paths* and probabilities over paths associated with a probability matrix $A = (A_{ij})_{ij}$. For any sequence $q = (i_1, ..., i_K)$, we let $\pi^A(q) \equiv \prod_{k=1}^{K-1} A_{i_k,i_{k+1}}$, and for any set of sequences Q, we abuse notations and let $\pi^A(Q) = \sum_{q \in Q} \pi^A(q)$. We define a *path* as a sequence q for which $\pi^A(q) > 0$.

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

$$A^K_{ij} \equiv \pi^A(Q^K_{i,j}) \text{ and } \sum_{j \in N} A^K_{ij} = \pi^A(Q^K_i) = 1 \tag{41}$$

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

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

Proof of Lemma 2: We consider A connected, that is, such that $A^k_{ij} > 0$ for all i, j, and consider $K \geq 2k$. 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, $\pi^B(q) \leq (1 - m_{i_0})\pi^A(q)$.[77] This implies

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

where the last inequality follows from (41) and Q_i^{K,i_0} non empty for $K \geq 2k$. This implies that C satisfies the condition of Lemma 1, hence $I - C$ has an inverse. Let $D \equiv I + B + ... + B^{K-1}$ and $H = (I - C)^{-1}D$. We have

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

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

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

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

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

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

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

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

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

Proof of Lemma 5: Let $\widetilde{x}^1$ denote the vector of modified initial opinions of players in N^1, and M^1 the restriction of M to N^1. We have:

$$\overline{y}^0 = A^{00}\overline{y}^0 + A^{01}\overline{y}^1 + \xi^0 \quad (42)$$

$$\overline{y}^1 = M^1\widetilde{x}^1 + (I^1 - M^1)(A^{10}\overline{y}^0 + A^{11}\overline{y}^1) \quad (43)$$

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:

$$\overline{y}^0 = P^0\overline{y}^1 + H^0\xi^0 \quad (44)$$

where $P^0 \equiv H^0A^{01}$ is a probability matrix (by Lemma 3 and because $A^{01}.1_{N^1} + A^{00}.1_{N^0} = 1_{N^0}$).[78]

Substituting $\overline{y}^0$ in (43), and letting $G = A^{10}H^0$ and $\widehat{x}_i = \widetilde{x}_i + (1 - m_i)G_i\xi^0/m_i$, we get

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

Since P^0 is a probability matrix, so is $\widehat{A}$, and $C^1 = (I^1 - M^1)\widehat{A}$ 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 $\overline{y}^1 = P^1\widehat{x}$ where $P^1 = H^1M^1$. Again, P^1 is a probability matrix because $\widehat{A}$ is a probability matrix and because $P^1 = M^1 + (I^1 - M^1)\widehat{A}P^1$. Substituting $\overline{y}^1$ in (44) we finally get $\overline{y}^0 = P^0P^1\widehat{x} + H^0\xi^0$ and $\overline{y}^1 = P^1\widehat{x}$, which concludes the proof.■

Proof of Result 5. *In addition to item (i) and (ii), we shall prove the following statement:* (iii) *If the lower bound $\underline{\gamma}$ on the choice set is sufficiently low and $\gamma_i = \underline{\gamma}$, $V_i \leq 1/\mid \log \underline{\gamma} \mid$ for all $m \geq \underline{m}$ and γ within the choice set.*

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

$$w_{ij} = \sum_{h,k} B_{ih}B_{jk}w_{hk} + \Lambda_{ij} \quad (45)$$

where $\Lambda_{ij} = 0$ if $i \neq j$ and $\Lambda_{ii} = (1 - m_i)^2(\gamma_i)^2\varpi_0$, and $B_{ii} = 1 - \gamma_i$, $B_{ij} = \gamma_iA_{ij}(1 - m_i)$.

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

$$w_{ij} \leq c\overline{\gamma} \quad (46)$$

[78] Indeed, for any $i \in N^0$, $\sum_{j \in N^0} A^{00}_{ij} + \sum_{j \in N^1} A^{01}_{ij} = \sum_{j \in N} A_{ij} = 1$

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

$$\widehat{w} \leq (1 - \underline{m})C\widehat{w} + \Gamma \tag{47}$$

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

Finally, to prove (iii), we consider two cases. Either $\overline{\gamma}$ is "small" and (46) applies, or we can separate individuals into a subgroup J where all have a small γ_j, and the rest of them with significantly larger γ_j. In the latter case, we redefine $\widehat{w} = (w_{jk})_{j\in J,k}$ as the vector of covariances involving some $j \in J$, and obtain inequality (47) with $\Gamma_{jk} \leq dp_{jk}$ for $k \notin J$ and $\Gamma_{jk} \leq d\gamma_j$ for $k \in J$, for some d. By definition of J, all γ_j and p_{jk} are small, and all Γ_{jk} are thus small, which will prove (iii). Details are below.

Step 1 (item (i)) $w_{ij} \leq c\overline{\gamma}$ with $c = \varpi_0/\underline{m}$.

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

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

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

$$w_{ii} \leq \gamma_i p_{ii}\overline{v} + (1 - \underline{m})\sum_k A_{ik}w_{ik} \tag{48}$$
$$w_{ij} \leq \gamma_j p_{ij}\overline{v} + (1 - \underline{m})(p_{ij}\sum_k A_{ik}w_{kj} + p_{ji}\sum_k A_{jk}w_{ik}) \tag{49}$$

These inequalities are obtained by solving for w_{ij} in equation (45), 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 (48-49).

Step 3 (item (ii)). It is immediate from (48-49) that (47) holds with $C_{jk} \equiv A_{jk}$ and $\Gamma_j = p_{ij}\gamma_j\overline{v} + p_{ij}c\overline{\gamma} \leq p_{ij}\overline{\gamma}(\overline{v}+c) \leq d\gamma_i$ for all j, for some d, which permits to conclude that $\widehat{w} \leq d\gamma_i/\underline{m}$.

Step 4 (item (iii)). Let $\varepsilon = \frac{1}{K|Log\underline{\gamma}|}$ with $K = 5\varpi_0/\underline{m}^2$ and set $\gamma_i = \underline{\gamma}$. Let us reorder individuals by increasing order of γ_j. Consider first the case where $\gamma_{j+1} \leq \gamma_j/\varepsilon$ for all $j = 1, ..., N-1$. Then $\overline{\gamma} < \underline{\gamma}/\varepsilon^{N-1}$, and for $\underline{\gamma}$ small enough, $\underline{\gamma}/\varepsilon^{N-1} < \varepsilon$, so $V_i \leq c\varepsilon < 1/\mid Log\underline{\gamma}\mid$.

Otherwise, there exists j_0 such that $\gamma_j \leq \underline{\gamma}/\varepsilon^{j_0-1}$ for all $j \in J$, and $\gamma_k > \gamma_j/\varepsilon$ for all $k \notin J$ and $j \in J$. It is immediate from (48-49) that (47) holds with Γ such that, for any $j \in J$,

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

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

Large Circle case. We check that the social optimum is symmetric: if $\underline{L}$ is the minimum loss that a player experiences at the social optimum. Define $\phi(L) = \min_m (m)^2 + (1-m)^2(\varpi + L)$. We have $L_i \geq \phi(\underline{L})$ for all i, so $\underline{L} \geq \phi(\underline{L})$, which implies $\underline{L} \geq \lim_n \phi^n(0)$. Since $\phi(L) = (\varpi+L)/(1+\varpi+L)$, ϕ is a contraction and since $\phi(J(m^{**})) = J(m^{**})$, $\lim_n \phi^n(0) = J(m^{**})$.

Proof of Proposition 8:

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

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

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

Coarse communication:

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

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

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

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

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

Assume now that player chooses m_i while others choose m. For player i, the estimation error is $\Delta_i \equiv m_i\delta_i + (1 - m_i)(\widehat{\xi} + \xi) \simeq m_i\delta_i + (1 - m_i)\frac{\xi}{m}$. Assuming that θ is drawn from a flat distribution with large support, the expected loss $L_i(\Delta)$ from estimating θ with an error Δ_i is quadratic in Δ_i and independent of b_i,[79] so $L(\Delta)$ is proportional to the variance of the error that i makes. To minimize the variance of Δ_i, player i sets $m_i = \frac{\varpi}{m^2}$, so in equilibrium $m^* = \varpi^{1/3}$.

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

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