Consensus and Disagreement: Information Aggregation under (Not So) Naive Learning

Abhijit Banerjee

Massachusetts Institute of Technology

Olivier Compte

Paris School of Economics and École des ponts ParisTech

We explore a model of non-Bayesian information aggregation in networks. Agents noncooperatively 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 amplifies the disagreement: even with little noise, individuals place substantial weight on their own initial opinion in every period, exacerbating 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.

I. Introduction

As of May 2020, 41% of US Republicans were not planning to get vaccinated against COVID-19, compared with 4% of Democrats.[1] We saw similar divergences in mask wearing, social distancing, and so on, which protect

We thank our editor Emir Kamenica and the referees for extremely useful comments and suggestions. This paper was previously entitled "Information Aggregation under (Not So) Naive Learning."

[1] This finding is according to a 2021 PBS NewsHour/NPR/Marist poll (https://www.pbs.org/newshour/health/as-more-americans-get-vaccinated-41-of-republicans-still-refuse-covid-19-shots).

Electronically published July 23, 2024

Journal of Political Economy, volume 132, number 8, August 2024.

https://doi.org/10.1086/729448

against the disease. Since COVID-19 is a life-threatening ailment that had already taken more than 3.5 million lives worldwide, it is hard to think of these as being simply empty gestures or entirely reflective of different preferences, though there is surely some of that. Rather, there seems to be a different reading of the facts on the ground; for example, in a Pew Research Center poll,[2] Republicans were much more likely to say that COVID-19 is not a major threat to the health of the US population (53% vs. 15% of Democrats). This goes with a general deepening in the political divide between Democrats and Republicans in recent years.[3]

The source of this shift is a subject of much discussion: one potential source of change is the massive growth in the use of the internet. However, the evidence from the careful work by Gentzkow and Shapiro (2011) suggests that online news consumption is not more segregated by political leanings than other sources of information that already existed, contrary to the concerns expressed by, for example, Sunstein (2001).[4] The most segregated sources of information, according to Gentzkow and Shapiro (2011), seem to be social networks (voluntary associations, work, neighborhoods, family, "people you trust," etc.), which were of course always there. However, there is evidence that online networks such as Facebook are substantially more segregated than other social networks and as a result, news that comes from being shared on Facebook tends to be more segregated than news from other media sources (Bakshy, Messing, and Adamic 2015).[5] It is true that social media are still a relatively small (though growing) part of news consumption, but the volume of "information" that can be quickly shared on Facebook may be much larger than other, more traditional sources. Moreover, while information was always shared through social connections, the evidence of growing affective polarization along political lines, especially in the United States (Boxell et al. 2022), raises the concern that the actual exchange of sensitive information in the social network is increasingly confined to those with similar views.

Given this evidence, we feel that it is worth exploring theoretically when and why social learning on networks can lead to large and persistent disagreements. As a starting point, we note that models of Bayesian social learning, such as Acemoglu et al. (2011), propose relatively weak conditions on signals and network structure under which information is perfectly aggregated as the network grows to be very large. More recent work, in

[2] See the 2020 poll at https://www.pewresearch.org/short-reads/2020/07/22/republicans-remain-far-less-likely-than-democrats-to-view-covid-19-as-a-major-threat-to-public-health.

[3] A 2014 Pew Research Center report (https://www.pewresearch.org/politics/2014/06/12/political-polarization-in-the-american-public) documents such a shift of political values for the period 1994–2014. See also Gentzkow (2016) and Bertrand and Kamenica (2023).

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

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

which agents repeatedly communicate (unlike in Acemoglu et al. [2011] where they communicate only once), includes Mossel, Sly, and Tamuz (2015), who derive necessary conditions on the network structure under which Bayesian learning yields consensus and perfect information aggregation.[6] The general sense from this literature is that convergence to a consensus is likely even when the network exhibits a substantial degree of homophily (Republicans talk mostly to other Republicans) as long as everyone is ultimately connected. This *Bayesian route*, however, requires that agents make correct inferences based on an understanding of all the possible ways information can transit through the network, which, at least for large networks, strains credibility.[7]

The alternative way to model learning on networks is to take a *non-Bayesian route*, which avoids these very demanding assumptions about information processing by postulating a simple rule that individuals use to aggregate own and neighbors' opinions. In recent years, the economics literature has tended to favor the DeGroot (DG) rule, where agents update their current opinion by linearly averaging it with their neighbors' most recent opinions. As observed by DeMarzo, Vayanos, and Zwiebel (2003), who brought it into the economics literature, the rule builds in a strong tendency toward consensus in any connected network, even when there is a high degree of homophily and people put high weight on people like them, though convergence between those far from each other in the network can be very slow.[8] Faced with this force toward consensus, Friedkin and Johnsen (1990) came up with a learning rule that is similar to DG but allows each individual to keep putting some weight on their own initial opinion.[9] For obvious reasons, this rule does not lead to a consensus.

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

[6] They build on Rosenberg, Solan, and Vieille (2009) and the literature on "agreeing to disagree" that goes back to Aumann (1976).

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

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

[9] Friedkin and Johnsen (1999, 3) write, referring to the work of DeGroot (1974) and other precursors: "These initial formulations described the formation of group consensus, but did not provide an adequate account of settled patterns of disagreement."

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

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

where y_i^t represents i's belief in period t; x_i represents the initial signal that i received, correlated with some underlying state of the world (we refer to x_i as i's initial opinion or *seed*); and

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

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

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

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

In contrast, result 2 shows that in the presence of any error in transmission, each decision-maker must choose $m_i > 0$ in equilibrium, so there will be no consensus even in the long run. The reason is that when all the m_i are small (a fortiori when everyone uses DG) the errors tend to cumulate, with the result that long-run opinions explode. Intuitively, a positive error by i pushes up i's opinion, which raises the opinions of others,

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

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

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

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

fueling a further rise in i's opinion, and so on—we call these *echo effects.* Raising m_i allows individuals to limit this cumulation of errors, at the cost of potentially putting too much weight on their own seeds. Moreover, there is no way to use γ_i to mitigate this problem; in fact, as long as there is no *idiosyncratic* error and $m_i > 0$ for at least one player, γ_i's play no role: long-run opinions are *fully* determined by the m_i's. Later in this paper, we show that γ_i does play an important role in controlling the effects of idiosyncratic errors, but that does not change the need to set $m_i > 0$.

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

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

We then compare the extent of disagreement in any equilibrium with the social optimum. Result 4 shows that there is too little—equilibrium values of m_i are always lower than the socially optimal values. One reason is that in setting m_i optimally, player i does not take into account the fact that lowering m_i raises the cumulated error faced by j. But this is not the only reason. In choosing m_i, i trades off the fact that a higher value of m_i reduces the influence of the transmission error with the fact that it reduces the weight on the opinions of others (which, especially in the long run, enables i to aggregate signals from all over the network and therefore provides very valuable information not contained in i's own signals). But he does take account of the fact that when m_i goes up, y_i better reflects the information contained in i's signal compared with what i learned from everyone else (which in the long run is very close to what i's neighbors *also* learned from everyone else), and this is valuable for aggregate welfare. Technically, raising m_i diminishes the correlation between y_i and others' signals, and this enhances the welfare of others.

Next we turn to comparisons of the efficiency of information aggregation on specific simple and oft-studied networks—the complete network, the directed circle, and the star network. At the heart of our analysis is the characterization of cumulated errors that each individual faces and how each player then mitigates the consequence of these errors by controlling the weight of her own seed x_i in her own long-run opinion, through the choice of m_i. We find that the star network performs worse than the two others, essentially because the central player propagates correlated errors to all peripheral players, thus raising cumulated errors.

In section V, we use our example of the star network to address the key issue of polarization. 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 (modeled as stars) connected by, say, one link. Such a network structure is not too dissimilar, for example, to the networks of Republicans and Democrats in the United States, who communicate mostly with each other (Cox et al. 2020). In this case, we show that lower m_i is associated with a high degree of consensus within each cluster but more extreme polarization across the groups, reminiscent of the situation of the Republicans and Democrats in the United States. The general point, captured by result 5, is that social efficiency requires the dispersion of opinions *within* and *between* subgroups to have the same orders of magnitude. Our very simple model therefore tells a useful story about why disagreements are necessary, but it also helps us understand why the resulting divergence of opinions can be surprisingly large and when they are likely to be costly.

The rest of the paper is devoted to two extensions.[14] In section VI, we allow for the possibility of idiosyncratic shocks in information transmission in addition to permanent shocks. In this setting, the speed of updating, γ_i, which plays no role in the previous analysis, also comes into play. Slowing down updating by setting γ_i close to zero allows the agent to minimize the changes in opinions that result from these shocks, which is an advantage because the shocks average out over time. This is what result 6 shows.

In section VII, we turn to the possibility of coarse communication—say each party reports only their current best guess about which of two actions is preferable. In this setting, the class of potentially "natural" rules includes 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 work with a version of this class of models where preferences are heterogeneous and each player has many neighbors. We show that systematic errors in interpreting actions by neighbors makes the long-run outcome from a DG-like rule entirely insensitive to the actual state of the world, but this is not true for FJ-type rules. We use this framework to discuss the connection between the errors we introduce and misspecifications in Bayesian models (as in Frick, Iijima, and Ishii 2020 and Bohren and Hauser 2021) and the related (non)robustness of long-run beliefs.

[14] Other extensions are examined in the working paper version (Banerjee and Compte 2023).

Related literature.—Our paper contributes to the large literature on learning in social network (see the excellent review by Golub and Sadler 2017). We study non-Bayesian learning on general networks with continuous choices and general networks. Within Bayesian social learning, Vives (1993, 1997) studies a setting similar to ours (with agents receiving a noisy signal) and, unlike us, obtains long-run convergence to the truth. The reason is that with continuous choice sets, Bayesian agents are able to perfectly extract the information content of the noisy signals. When the choice set is coarser, aggregation can fail even with Bayesian agents, as shown by Banerjee (1992) or Bikhchandani, Hirshleifer, and Welch (1992).[15]

In Vives (1997), like in this paper, agents underweight their private seed; in his setup, a stronger reliance on private signals *in the initial phase* would speed up learning and benefit all.[16] In our case, the weight cannot be altered over time; however, a higher reliance on private seeds compared with equilibrium weights improves welfare because this limits both the correlation between information sources and the cumulated errors.

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

Other papers directly modify the updating rule itself. Jadbabaie et al. (2012) introduce rules that combine Bayesian updating of own signals with

[15] Mossel, Sly, and Tamuz (2015) show that this result also depends on the network structure and that for a large class of large networks, consensus and almost perfect learning are possible even with coarse communication.

[16] In the context of non-Bayesian learning, Mueller-Frank and Neri (2021) argue in related terms in favor of nonstationary rules that first aggregate information in a sufficiently dense part of the network.

a DG-like averaging over neighbors' beliefs, while Levy and Razin (2015) consider a rule that involves cumulating log-likelihood ratios, which they justify, like DG, on the grounds that it mimics what a subjective Bayesian (with an erroneous model of the world) would do (see also Dasaratha, Golub, and Hak 2023, 11). Finally, Molavi, Tahbaz-Salehi, and Jadbabaie (2018) provide axiomatic justification(s) (motivated by imperfect recall) for DG-style linear aggregation (and averaging) of log belief ratios.[17]

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

II. Basic Model

A. *Transmission on the Network*

We consider a finite network with n agents, assume noisy transmission/reception of information, and define a simple class of rules that players may use to update their opinions. Formally, each agent i in the network has an *initial opinion* x_i and, at date t, *an opinion* y_i^t, where both can be represented as real numbers.[18] Taking as given the matrix A characterizing the weights A_{ij} that i puts on j's opinion, we consider the class of updating rules (FJ) parameterized by the weights m_i and γ_i and specified in the introduction. Along with expression (1) for transmission errors, the dynamic of opinions for player i is

$$y_i^t = (1-\gamma_i)y_i^{t-1} + \gamma_i\Big(m_i x_i + (1-m_i)\Big(\sum_{j\in N_i} A_{ij} y_j^{t-1} + \varepsilon_i^t\Big)\Big).$$

When $m_i = 0$, the rule corresponds to the well-studied DG rule. When $m_i > 0$, in each period the rule mixes the decision-maker's own initial opinion x_i with DG. This perpetual use of the initial opinion in the updating process gives FJ a non-Bayesian flavor, since for a Bayesian, their prior (i.e., the seed) is already integrated into y_i^{t-1} and therefore there is no reason to go back to it.[19]

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

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

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

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

Note that all the rules considered here are stationary, in the sense that the weighting parameters m_i and γ_i do not vary over time.[20] We see these as plausible ways in which boundedly rational agents might incorporate others' opinions into their current opinion. We recognize that with enough knowledge of the structure of the network and the process by which new information gets incorporated, adjusting the weights over time may make sense. In Banerjee and Compte (2023), the working paper version of this paper, we discuss this.

We also impose the assumption that everyone operates on the same time schedule: periods are defined so that everyone changes their opinion once every period and everyone else gets to observe that change of opinion before they adjust their opinion in the following period. We relax this assumption in Banerjee and Compte (2023).

B. *Errors in Opinion Sharing*

The term ε_i^t is an important ingredient of our model, meant to capture some imperfection in transmission.[21] It represents a distortion in what each individual "hears" that aggregates all the different sources of errors. Until section VI, we assume that the error term is persistent, realized at the start of the process and applying for the duration of the updating process.[22] We denote by ξ_i this persistent error, so

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

In section VI, we extend the model and incorporate idiosyncratic errors:

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

where ν_i^t is independent and identically distributed (i.i.d.) across time and agents.

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

[21] There have been several recent attempts to introduce noisy or biased transmission in networks. In Jackson, Malladi, and McAdams (2019), information is coarse (zero or one), and noise can either induce a mutation of the signal (from zero to one or from one to zero) or a break in the chain of transmission (information does not get communicated to the network neighbor).

[22] 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 *processed* by *i*. Biases ξ_i may be drawn independently across players, but we also discuss cases where they are positively correlated, such as when a group of friends share a political bias. Also note that although errors are indexed by *i*, our formulation can accommodate biases that result from both "hearing" errors and "sending" errors.[23]

For convenience, we assume that all error terms are unbiased (i.e., $E\xi_i = 0$ and $E\nu_i^t = 0$) and homogeneous across players, so we let

$$\varpi = \varpi_i = \mathrm{var}(\xi_i).$$

C. *The Objective Function*

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

Given this private objective, we explore each agent's incentives to choose his updating rule within the class of FJ rules to maximize the above objective on average across many different realizations of the underlying state of the world, the initial opinions, and the transmission errors. We have in mind the idea that individuals choose a single rule to apply to many different problems. This is why we focus on their ex ante performance.[24] The set of possible updating rules is extraordinarily 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.

Formally, we assume that the initial signals are given by

$$x_i = \theta + \delta_i,$$

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

[23] For example, if there were both "hearing" errors labeled ξ_i^h and "sending" errors labeled ξ_i^s, one could define $\xi_i = \xi_i^h + \Sigma_j A_{ij}\xi_j^s$ as the resulting processing error. Sending errors naturally generate correlations across the ξ_i's and a profile of errors that depend on the network structure *A*. This is further discussed in Banerjee and Compte (2023).

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

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

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

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

D. *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 toward the use of higher-performing rules (e.g., justified by evolution or reinforcement learning) and (ii) in this quest for higher-performing rules, each individual considers (and gets feedback about) only a limited set of rules (i.e., the FJ class).

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

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

[25] Alternatively, one could define $L_i = \lim_{h \searrow 0}(1 - h)\Sigma h^{t-1} L_i^t$, 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 $L_i \leq \text{var}(\delta_i) = \sigma_i^2 = 1$ by ignoring everyone else's opinions ($m_i = 1$).

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

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

III. Some Properties of the Long-Run Opinions

In this paper, we make a distinction between results, which are meant to be of substantive interest, and propositions, which are more technical and meant to explain and lead up to the results. This section reports a number of propositions that provide the bulwark for our main results in section IV. We start by studying the properties of long-run opinions under DG and FJ with and without errors. In particular, we show that in the presence of errors there is convergence under FJ as long as at least one person i_0 has $m_{i_0} > 0$ but not under DG. We then explore what determines the variance of the limit opinion in the case where such a limit opinion exists. In particular, what part of it comes from the "signal"—the original seeds—and what part comes from the noise that gets added along the way? We also explore the degree to which a player can influence long-run opinions through the choice of m_i and γ_i.

A. *DG without Errors*

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

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

implying that

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

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

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

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

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

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

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

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

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

PROPOSITION 0. When $m_i = 0$ for all i and in the absence of errors, long-run opinions all converge to the same limit opinion $\pi.x$ and

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

and where q^i is a probability vector in Δ_{n-1} that does not depend on γ_i.

In other words, the network structure determines ρ. Given ρ, player i can use γ_i to control her influence on the long-run opinion, π_i, but she cannot control the relative weights on the opinions of others, captured by q^i.

B. *DG with Errors: Exploding Dynamics*

Below we show that if all agents follow a DG rule, then for almost all realizations of ξ, the long-run opinions diverge.

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

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

Proof. With errors, equations (2) and (3) become $y^t = B_0 y^{t-1} + \Gamma\xi$ and

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

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

C. *Anchored Dynamics under FJ*

Again fixing x and ξ, we now examine long-run dynamics under FJ.

PROPOSITION 2. Assume that at least one player—say, i_0—updates according to FJ (with $m_{i_0} > 0$). Then, for any fixed x and ξ, y^t converges and the limit vector of opinions y does not depend on γ or on the signal x_i of any individual with $m_i = 0$.

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

When $m_{i_0} > 0$ for some i_0, proving convergence is standard.[29] The limit opinion y then solves

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

where $X_i = m_i x_i + (1 - m_i)\xi_i$, which implies that, in matrix form, it is also the solution of

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

where M is the diagonal matrix with $M_{ii} = m_i$. This expression implies that limit opinions are independent of the γ_i's. It also explains why long-run opinions involve only the seeds x_i of players for whom $m_i > 0$, since for the others, $X_i = \xi_i$.

D. *The Dominance of Noise under Low m*

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

PROPOSITION 3. Let $\bar{m} = \max m_i$. Then $L_i \geq (\varpi/n)[(1 - \bar{m})^2/\bar{m}^2]$.

The detailed proof is in appendix A2. The lower bound on L_i is obtained by showing that for given x, ξ, long-run expected opinions are a weighted average of *modified initial opinions*, defined, whenever $m_i > 0$, as

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

To fix ideas, assume that $m_i > 0$ for all i.[30] Then, using the previous notation, one can write $X = M\tilde{x}$ and, using (7), obtain

$$y = M\tilde{x} + (I - M)Ay \equiv P\tilde{x}, \tag{8}$$

where P is a probability matrix.[31] Intuitively, x_i can be thought of as the *seed* that individual i plants in her belief in every period and $\tilde{x}_i$ as the

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

[30] The argument generalizes to the case where a subset $N^0 \nsubseteq N$ of agents follows DG ($m_i = 0$). (See app. sec. A2.)

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

effective seed given processing errors. Long-run opinions are averages over effective seeds. Since the variance of each $\tilde{x}_i$ is bounded below by $\varpi(1-\bar{m})^2/\bar{m}^2$, we obtain the desired lower bound.

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

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

which further implies that

$$y_i = p_i\tilde{x}_i + (1-p_i)\tilde{x}_j, \text{ where } p_i = \frac{m_i}{m_i + (1-m_i)m_j}, \tag{9}$$

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

$$y_i = p_i x_i + (1-p_i)(x_j + \hat{\xi}_i), \text{ where } \hat{\xi}_i = \frac{\xi_i + \xi_j}{m_j} - \xi_j. \tag{10}$$

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

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

E. *Influence under FJ Rules and Cumulated Errors*

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

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

$$y_i = m_i x_i + (1 - m_i)\xi_i + (1 - m_i)\hat{y}_i, \text{ with } \hat{y}_i \equiv \sum_{k \neq i} A_{ik} y_k. \tag{11}$$

Player i's opinion thus builds on the opinion $\hat{y}_i$ of a (fictitious) *composite neighbor* who aggregates the opinions y_k, to which the error ξ_i is added. Letting $\tilde{A}^i_{kj} = A_{kj}/(1 - A_{ki})$, we rewrite (11) to describe how each opinion y_k builds on y_i:

$$y_k = m_k x_k + (1 - m_k)\xi_k + (1 - m_k)A_{ki} y_i + (1 - m_k)(1 - A_{ki})\sum_{j \neq k,i} \tilde{A}^i_{kj} y_j. \tag{12}$$

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

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

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

$$y_i = p_i x_i + (1 - p_i)(\hat{x}_i + \hat{\xi}_i), \text{ where } \hat{x}_i = q^i . x_{-i},$$

$$\frac{p_i}{1 - p_i} = \frac{m_i h_i}{(1 - m_i)}, \quad q^i_j = \frac{R^i_j m_j}{\sum_{j \neq i} R^i_j m_j}, \text{ and} \tag{13}$$

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

Proposition 4 provides an analog of proposition 0 when at least one player uses an FJ rule. Without errors, player i's long-run opinion is an average between her own seed x_i and a *composite seed* $\hat{x}_i$ (an average over the others' seeds). The weight p_i defines how player i's own seed influences her long-run opinion, and through the choice of m_i player i has full control over this weight. However, player i has no control over the composite seed $\hat{x}_i$, as the vector of weights $q^i \in \Delta_{n-1}$ is fully determined by A and m_{-i}.

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

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

PROPOSITION 5. If $m_{-i} = 0$ and $m_i > 0$, then $y_i = x_i + [(1 - m_i)/m_i](\xi_i + \Sigma_{j \neq i} R^i_j \xi_j)$, where R is as defined in proposition 4.

Consistent with proposition 3, echo effects rise without bound when m_i gets small. Propositions 4 and 5 imply that if all players but i use DG, all players' opinions will build on x_i only, however small m_i is. Mueller-Frank (2017) makes a similar observation in a model without errors (concluding that learning outcomes are highly sensitive to small departures from DG).

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

$$W_i \equiv \text{var}(\hat{x}_i) + \hat{\varpi}_i, \text{ where } \hat{\varpi}_i = E\hat{\xi}^2_i. \tag{14}$$

Since player i fully controls p_i by adjusting m_i (since $p_i/(1 - p_i) = h_i m_i/(1 - m_i)$), individual i optimally sets p_i so that $p_i/(1 - p_i) = W_i/\sigma^2_i$, and we obtain the following:

PROPOSITION 6. For a given m_{-i}, the optimal choice of m_i and resulting loss L_i satisfy

$$\frac{m_i}{1 - m_i} = \frac{W_i}{h_i \sigma^2_i} \quad \text{and} \quad L_i = \sigma^2_i p_i = \frac{W_i}{1 + W_i/\sigma^2_i}. \tag{15}$$

Since W_i/h_i depends on the network structure and m_{-i} only, proposition 6 allows us to easily characterize equilibrium weights m^*_i, as well as the induced equilibrium losses.

This proposition also implies that the loss L_i is fully determined by W_i. It is instructive to compare L_i with the minimum feasible loss v^* obtained under efficient aggregation of initial opinions—that is, $v^* = \min_q \text{var}(\pi.x)$. This minimum loss satisfies

$$v^* = \sigma^2_i \pi^*_i = \frac{\underline{W}^*_i}{1 + \underline{W}^*_i/\sigma^2_i}, \tag{16}$$

where $\underline{W}^*_i = \min_q \text{var}(q.x_{-i})$.[32] So whenever W_i rises above $\underline{W}^*_i$, the loss L_i rises above v^*. Expression (14) thus highlights the two possible additional sources of losses that player i now faces: (i) the fact that seeds of others may not be efficiently aggregated (i.e., $\text{var}(\hat{x}_i) > \underline{W}^*_i$) and (ii) the presence of the cumulated error term $\hat{\xi}_i$.

[32] This is because $v^* = \min_\pi \text{var}(\pi.x) = \min_{\pi_i} \text{var}(\pi_i x_i + (1 - \pi_i)\underline{W}^*_i)$.

Section IV builds on propositions 4, 5, and 6 to characterize the equilibrium of the rule-choice game. We also see that when errors are small, the cumulated errors $\hat{\xi}_i$ are the preponderant source of inefficiency. We conclude this section with further comments on DG and FJ rules.

F. *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 who never changes opinion: it anchors one's opinion, preventing too much drift. As a result, it also anchors the opinions of one's neighbors and hence the opinions of everyone in the (connected) network.

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

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

$$\frac{p_i}{p_k} = \frac{m_i\rho_i}{m_k\rho_j}. \tag{17}$$

This is an analog to (4) showing that, close to the limit, m_i plays the same role that $1/\gamma_i$ 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 do not put any weight on their own initial signal after the first period, the direct route for that signal to stay relevant is through the weight put on their own previous period's opinion. This source clearly has dwindling importance over time. This gets compensated by the growing weight on the indirect route—each individual i adjusts his or her opinion based on the opinions of their neighbors, and these are in turn influenced by i's past opinions and, through those, by i's initial signal. In DG without transmission errors, the second force at least partly offsets the first one—but this is no longer true when there is any transmission error because of the cumulative effect of noise that comes with the feedback from others.

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

$$Y_i^t = \gamma_i(D_i^{t-1} + \xi_i), \tag{DG}$$

$$Y_i^t = (1 - m_i)A_i Y^{t-1}. \tag{FJ}$$

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

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

IV. Choosing the Rule

A. When There Are No Errors

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

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

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

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

Intuitively, both y_i and the neighbor's composite limit opinion $\hat{y}_i$ are weighted averages between x_i and the composite seed $\hat{x}_i$, with different weights when players do not use DG rules. In equilibrium, i optimally chooses the weighting to reduce variance, so if the equilibrium is not DG, the variance $v(y_i)$ must be *strictly* smaller than the variance $v(\hat{y}_i)$, 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 one. This can happen only if they coincide with the efficient weights π_i^*. Therefore, there is a unique (and efficient) equilibrium outcome.

B. 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 that DG puts zero weight) are dominated from the point of view of the individual decision-maker, as well as being socially suboptimal.

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

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

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

C. *How Large Is the Divergence in Opinions?*

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

Result 3 below shows that because of cumulated errors, the optimal weight put on one's own seed tends to be relatively large—that is, $O(\varpi^{1/3})$ (of the order of $\varpi^{1/3}$).[34] As a result, when ϖ is small the extra weight on one's own seeds becomes the preponderant source of dispersion. These extra weights also determine the equilibrium magnitude of $\hat{\varpi}_i$ and L_i. We have the following:

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

Note that in addition to cumulated errors, there is another source of inefficiency in equilibrium: the fact that seeds are not efficiently weighted. But that inefficiency is $O(\varpi^{2/3})$:[35] a socially optimal choice of weights m_i would trade off more inefficient weighting (larger m) against decreasing the variance of cumulated errors.

The intuition for result 3 runs as follows. The error terms $\hat{\varpi}$ are $O(\varpi/m^2)$. These error terms degrade the quality of information that each i gets (raising W_i above W_i^*), which in turn implies a weighting p_i of i's seed larger than the efficient weighing π_i^*, with $p_i - \pi_i^*$ at least $O(\varpi/m^2)$ (by [15] and [16]). When $m > 0$, players end up weighing seeds differently, but when all m are small, the spread between the weights is also small and $O(m)$. Thus, if p_k represents 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 one, the p_k's must add up to at most $1 + O(m)$, and since the sum $\Sigma_k(p_k - \pi_k^*)$ is

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

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

at least $O(\varpi/m^2)$, m must be at least $O(\varpi/m^2)$ in equilibrium, which gives m at least $O(\varpi^{1/3})$.[36]

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

D. Privately versus Socially Optimal Choices

We already showed that both private and social optima must deviate from DG when there is noise. The next result shows that there is a sense in which, in the presence of noise, the Nash equilibrium is closer to DG than is desirable from the point of view of social welfare maximization.[37]

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

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

$$y_j = (1-\mu_{ji})C^{ji}\tilde{x}_{-i} + \mu_{ji}y_i, \tag{18}$$

where C^{ji} is a probability vector and $\mu_{ji} \in (0,1)$,[38] with μ_{ji} and C^{ji} both independent of m_i.

The expression above highlights that when player i chooses m_i optimally (for him) to minimize the variance of y_i, there is no reason why he would also be minimizing the variance of y_j. Specifically, we use (18) to separate the loss L_j into three terms:

$$L_j = (1-\mu_{ji})^2\text{var}(C^{ji}\tilde{x}_{-i}) + \mu_{ji}L_i + 2(1-\mu_{ji})\mu_{ji}\text{Cov}(C^{ji}\tilde{x}_{-i}, y_i). \tag{19}$$

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

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

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

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

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

E. *Simple Examples*

To conclude this section, we directly compute the equilibrium and socially efficient weights in simple examples to shed further light on the rule choice and information aggregation. We assume that initial opinions are equally informative ($\sigma_i^2 = 1$ for all i) and each player treats all his neighbors symmetrically ($A_{ij} = 1/|N_i|$). We start with the two-player network and next discuss other larger simple networks (directed circle, complete network, and star network).

1. Two-Player Case

Social optimum.—Assuming independent errors, we obtain from (9)

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

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

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

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

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

Welfare is maximized for an m^{**} that optimally trades off these two effects, and the socially efficient weight m^{**} (which minimizes L) can be significantly different from zero 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, $m^{**} \simeq (4\varpi)^{1/4}$.[39]

Nash equilibrium.—We now assume that individuals choose their rules noncooperatively. Applying proposition 5, we obtain $p_i/(1 - p_i) = 1 + \hat{\varpi}_i$, so

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

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

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

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

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

In the absence of noise, $\hat{\varpi}_i = 0$ and player 1 should set m_1 so that $p_1 = 1/2$ (for information aggregation purposes), which requires that $m_1 < m_2$; this explains why there is no equilibrium with positive m (this is the force toward DG). With noise, the variance $\hat{\varpi}_i$ 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 (and hence m_i), which in turn puts a lower bound on equilibrium weights: in equilibrium, $m_1^* = m_2^* = m^*$ and m^* is a solution to

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

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

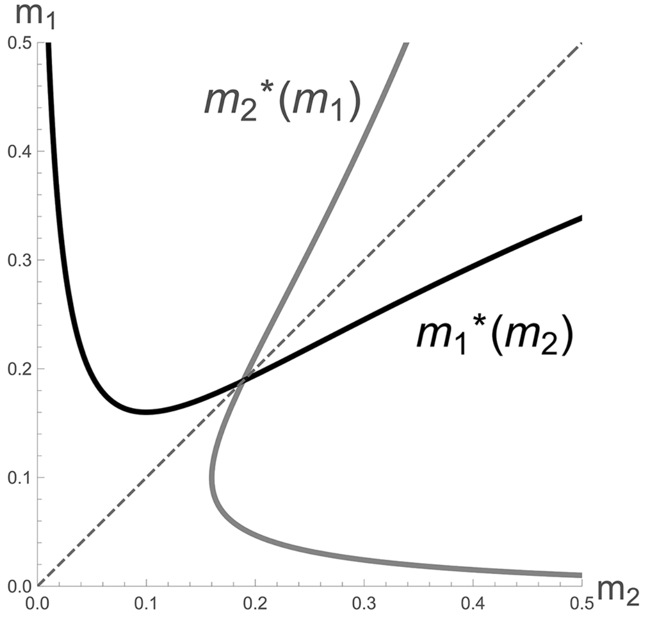

FIG. 1.—Best responses, $\varpi = 0.001$.

2. Larger Networks

Equilibrium weights are obtained using proposition 6: player i's incentive condition yields

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

where $W_i = \mathrm{var}(\hat{x}_i) + \hat{\varpi}_i$. Both h_i and W_i depend only on m_{-i} and the structure of the network, and the equilibrium values m_i^* are obtained by simultaneously solving these equations. For the directed circle and complete network, all players are symmetric, so one easily finds the equilibrium weight m^*. For the star network, we have to examine the incentives of the central player (labeled player 0) and peripheral players separately, and we obtain equilibrium weights m_0^* and m^* for central and peripheral players, respectively. We leave the details of the computation to the appendixes (app. sec. A12 and app. B), focusing on the case where the variance ϖ is small. We report here some notable facts.

In the star network, the central player can have disproportionate influence on the opinions of others, and in equilibrium she chooses m_0^* much below m^* to ensure that this is not the case: for fixed n, $m_0^*/m^* \simeq 1/(n-1)$, and at the large-n limit, $m_0^*/m^* \simeq O(\varpi^{1/3})$. Thus, in effect, at this limit the central player essentially ignores her signal and behaves like a DG player.

It is also interesting to compare the aggregation properties of different networks. Both h_i and W_i depend on the structure of the network, and this eventually affects the performance of the network. For example, for fixed m, the cumulated error terms $\hat{\varpi}_i$ are higher in the star network than in both other networks, because the central player's error term contaminates all other players in a correlated way. The consequence is that the star network performs worse than both the directed circle and the complete network.

Finally, we find that the directed circle performs better than the complete network when n is not too large, because while the terms W_i are similar across these two networks, players have stronger incentives to raise m_i in the directed network.[40] The comparison is reversed for large n—the directed circle yields poorer information aggregation ($\mathrm{var}(\hat{x}_i)$ is higher) and, when persistent errors are independent, poorer averaging of errors.

V. Implications for the Divergence of Opinions and Polarization

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

[40] Technically, this is because h_i is smaller in the directed network, so a given target p_i is achieved with a higher m_i.

aggregates seeds. In a large network, this opinion y^* will essentially coincide with the underlying state θ ($y^* \simeq \theta$), which implies that if we consider two such identical networks, there will be *consensus within* each network and *consensus across* networks.

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

To fix ideas, we consider the case of two large disconnected star networks modeled as above.[41] This description generally fits the maps of social networks in the US population, with the two stars representing Democrats and Republicans (Cox et al. 2020). We assume that in each star network all peripheral players use the same weight m and that central players behave as DG players, simply aggregating peripheral players' opinions.[42] We are interested in the effect of m on the distribution of opinions within the star and across stars. In each star, if the central player is labeled player 0, for any peripheral player i of that star we have

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

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

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

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

$$D = 2v(\bar{y}).$$

The following result establishes a relationship between d and D, where ϖ_0 refers to the variance of ξ_0:

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

Proof. When the central player is DG, $y_0 = \bar{y} + \xi_0$, so for a large network and independent errors this immediately gives $\bar{y} = ((1 - m)\xi_0)/m$

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

[42] In app. sec. A11, we consider the case where central players benevolently choose m_0 to minimize the losses of the peripheral players, given m.

and hence $D \simeq 2\varpi_0/m^2 \simeq 4\varpi_0/d$ for small m. Writing $y_i = (y_i - \bar{y}) + \bar{y}$, we obtain $v(y_i) = (1/2)(d + D)$. Since $D \simeq 4\varpi_0/d$, the loss $v(y_i)$ is minimized for $D \simeq d \simeq 2\varpi_0^{1/2}$ (hence, $m^{**} \simeq \varpi_0^{1/4}$). QED

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

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

VI. Idiosyncratic Errors

We now introduce idiosyncratic errors and assume that

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

where ν_i^t are i.i.d. across individuals and time.[43] We further assume that $E\nu_i^t = 0$ and let $\varpi^0 = \mathrm{var}(\nu_i^t)$. We wish to characterize the (additional) loss generated by these idiosyncratic errors and examine the consequence regarding incentives.

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

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

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

PROPOSITION 7. If $m_i = 0$ for all i, V^t increases without bound. If $m_{i_0} > 0$ for some i_0, $\bar{y}^t$ and V^t both have well-defined limits $\bar{y}$ and V. Besides, $\bar{y} = y^0$, V is independent of x and ξ, and $L_i = L_i^0 + V$. Furthermore, if $m_i \leq \bar{m}$ and $\gamma_i \geq \underline{\gamma}$ for all i, $V_i \geq (\varpi^0/2n)[\underline{\gamma}^2(1 - \bar{m})^2/\bar{m}]$.

For given $m, \gamma > 0$, expected long-run opinions eventually coincide with y_0, but long-run opinions are subject to temporary changes resulting from idiosyncratic communication errors. Proposition 7 shows that, for given γ, these temporary changes are significant and costly and when all m are small.

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

RESULT 6. Fix $\underline{m}$. We have the following:

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 very small γ_i, a player can get rid of the additional variance induced by the idiosyncratic noise.

Note that the incentive to set γ_i arbitrarily small obviously depends on the assumption that players care only 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).

VII. Discussion

In the working paper version of this paper, we discuss various extensions of and possible variations on our base model. We examine the case of biased persistent errors, showing that this provides additional incentives to raise m_i and increases the losses L_i. We show that FJ rules are robust to variations in the communication protocol. We also discuss how nonstationary rules might create further difficulties. Here we focus on one extension, the case of coarse communication, which enables us to discuss the relationship between our work and recent papers on information aggregation in networks when agents' priors are misspecified (Frick, Iijima, and Ishii 2020; Bohren and Hauser 2021). The connection with that literature is that our persistent errors play the same role as a misspecification. One key difference is that we allow players to correct—to some extent (i.e., through the weight m_i)—for the misspecifications that players are subject to.

A. *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. Below we explain 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, Iijima, and Ishii (2020), as well as Bohren and Hauser (2021).

Assume heterogeneous preferences with $\theta_i = \theta + b_i$ characterizing i's value from choosing one over zero, so the optimal action a_i^* is one when $\theta_i > 0$ and zero otherwise.[44] 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 that the b_i's are drawn from identical distribution g (and cumulative denoted G) with full support on $\mathcal{R}$.

As before, we define 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 that $a_i^t = 1$ if $y_i^t + b_i > 0$ and $a_i^t = 0$ otherwise. Each agent i observes the fraction f_i^t of neighbors who choose action zero, which she can use to make an inference $\psi_i(f_i^t)$ about θ and update her opinion using an FJ-like rule:

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

Long-run opinions clearly depend on the inference rule assumed, but there is a natural candidate for ψ_i, the function $\phi \equiv h^{-1}$, where $h(y) \equiv G(-y) = \Pr(y + b_i < 0)$ is the fraction of agents who choose $a = 0$ when their opinions are all equal to y. If others have opinions that are correct and equal to θ, a fraction $f \simeq h(\theta)$ chooses $a = 0$ 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 us assume that

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

where ξ_i represents a persistent error in interpreting f.[45] To fix ideas, we assume correlated errors ($\xi_i = \xi$ for all i) with variance ϖ.

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

[45] As in Frick, Iijima, and Ishii (2020), ξ_i could arise from an erroneous prior $g_i \neq g$, with agents using the inference function $\psi_i = h_i^{-1}$ where $h_i(\theta) = G_i(-\theta)$. The difference $\xi_i(f) \equiv \psi(f) - \phi(f)$ is an error in making inferences. With preferences centered on $\bar{b}$ and agents having an erroneously translated prior centered on $\bar{b}_i$, the error is independent of f and equal to $\xi_i \equiv \bar{b}_i - \bar{b}$. In comparison, Ellison and Fudenberg (1993, sec. 1) examine

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 app. B). 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 drift away from θ and converge to $\theta + [(1 - m)\xi/m]$. 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 $(1 - m)\xi/m$. And agents continue to disagree even in the long run. The consequence regarding social incentives and private incentives is as before, with m^* and m^{**} respectively comparable to $\varpi^{1/3}$ and $\varpi^{1/4}$: agents do not incorporate the damaging echo effect that an m_i set too low produces in their choice of m_i.

B. A Connection to Misspecified Bayesian Models

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

We now introduce, as in Bohren and Hauser (2021), a fraction q of autarkic players who base their choice only on their private signal x_i (thus ignoring the social information). Define $G^0(\theta)$ as the fraction of autarkic

social learning assuming that $b_i = 0$ for all and $\psi_i(f) = f - 1/2$; choices are tilted in favor of the more popular one. Ellison and Fudenberg (1993) find that small enough m's 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 Ellison and Fudenberg [1993], agents receive many signals x_i about the state, but given their assumptions, their model is equivalent to the one proposed here where agents simply receive one signal at the start.)

types who choose $a = 0$ when the state is θ, and to fix ideas, further assume that nonautarkic types have correct priors about G^0. When beliefs of nonautarkic types are concentrated on θ and the true state is θ_0, the fraction f becomes

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

while a fraction

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

would be expected. The observed f will meet expectations when

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

which implies a discrepancy $\Delta = \theta_0 - \theta$ comparable to ξ/q, which thus blows up when q is small.

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

Autarkic types thus play a role similar to our weights m_i, helping to anchor the beliefs of social types.[46] In our setup, the analog of social and autarkic types would be to assume that agents either are DG ($m_i = 0$) or use $m_i = 1$. In contrast, we have assumed that some intermediate m_i is feasible for each agent.

Another difference is that we focus on the optimal choices of m_i from the social or private points of view. In looking for a Nash equilibrium, we decentralize the choice of m_i and endogenize the weight each puts on social versus private information.[47]

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

[46] Note that, unlike Bohren and Hauser (2021), here we find that the fraction q needs to be large enough. This is because, unlike Bohren and Hauser (2021), who assume few states and correct priors over states, here we assumed that subjective priors on θ have full support.

[47] A similar decentralization exercise (endogenizing q) could be done in the Bohren and Hauser (2021) environment with agents choosing ex ante whether to be autarkic or social, with the consequence that in equilibrium, they would have to be indifferent between the two roles and hence incur a significant loss (equal to that of the autarkic type).

that they allow for sufficiently diverse opinions or beliefs, end up being more robust.

VIII. Concluding Remarks

We end the paper with a discussion of issues that we have not dealt with and that may provide fruitful directions for future research. One premise of our model is that everyone has a well-defined initial signal. However, the analysis here would be essentially unchanged if some players did not have an initial opinion to feed the network and were thus setting $m_i = 0$ for the entire process. FJ would aggregate the initial opinions of those who have one.

In real life, many of our opinions come from others and in ways that we are not necessarily aware of, and the existence of a well-defined "initial opinion" could be legitimately challenged. In other words, people may have a choice over the particular opinion they want to hold on to and refer back to (i.e., 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 are soon overwhelmed by noise in transmission. The other (non-DG) players could provide the system with the necessary memory, using the initial communication phase to gradually build up an "initial opinion" based on the reports of their more knowledgeable DG neighbors, and then perpetually seed in that initial opinion into the network. In other words, in an environment where information is heterogeneous and weights m_i are set suboptimally by some, there could be a value for some agent in adopting a more sophisticated strategy in which the initial opinion is updated for some period of time before it becomes anchored. In other words, it may be optimal for some of the less informed to listen and not speak for a while as they build up their own initial opinions before joining the public conversation.

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

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

Appendix A

A1. 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 $\tilde{x}_i = x_i + \xi_i(1 - m_i)/m_i$ denote the modified initial opinion and let $\tilde{x} = (\tilde{x}_i)_i$ denote the vector. Next, define the matrix $B = I - \Gamma + \Gamma(I - M)A$.

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

A2. Proof of Proposition 3

In the main text, we show that when $m_i > 0$ for all i, long-run opinions are weighted averages of modified opinions $\tilde{x}_i$. Lemma A1 below (proved in app. B) generalizes this observation. Define $N^0 \nsubseteq N$ as the set of n_0 agents following DG ($m_i = 0$). Denote by ξ^0 the vector of errors of these players. We have the following:

Lemma A1. Assume that $n_0 < n$. Then $y = P\tilde{x} + Q\xi^0$, where P is an $(n, n - n_0)$-probability matrix.

From lemma A1, $L_i = v(y) \geq (1/n) \min v(\tilde{x}_i)$, which proves proposition 3. QED

A3. Proof of Proposition 4

Assume that $m \gg 0$ so $\tilde{x}_j$ is well defined for all j.[48] For $j \neq i$, let $X_j = m_j\tilde{x}_j + (1 - m_j)A_{ji}y_i$ and let $c_j^i = m_j + (1 - m_j)A_{ji}$. Equation (12) can be written

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

in matrix form to obtain, by definition of Q^i, $y_{-i} = Q^i X$. Note that if $\tilde{x}_j = 1$ for all j and $y_i = 1$, then $y_k = 1$ for all k, so $\Sigma_{j \neq i} Q^i_{kj} c^i_j = 1$ for all k, which implies that

$$\sum_{j \neq i} Q^i_{kj}(1 - m_j)A_{ji} = 1 - \sum_{j \neq i} Q^i_{kj} m_j, \tag{20}$$

and, since Q^i is a positive matrix,[49] $\Sigma_{j \neq i} Q^i_{kj} m_j \leq 1$, so $\Sigma_{j \neq i} R_j{}^i m_j \leq 1$. Equation (20) further implies that

$$y_k = \sum_{j \neq i} Q^i_{kj} m_j \tilde{x}_j + (1 - \sum_{j \neq i} Q^i_{kj} m_j) y_i, \tag{21}$$

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

$$y_i = m_i \tilde{x}_i + (1 - m_i)(\sum_{j \neq i} R^i_j m_j \tilde{x}_j + y_i(1 - \sum_{j \neq i} R^i_j m_j), \tag{22}$$

which, since $m_j \tilde{x}_j = m_j x_j + (1 - m_j)\xi_j$ and $h_i = 1/\Sigma_{j \neq i} R^i_j m_j$, gives the desired expressions (13) for y_i, $\hat{x}_i$, p_i, and $\hat{\xi}_i$. QED

A4. Proof of Proposition 5

Assume that $m_i > 0$ and apply proposition 4, taking the limit where all m_j tend to zero. For A given, Q^i and R^i are uniformly bounded (with a well-defined limit when all m_j tend to zero) and $(1 - p_i)\hat{x}_i$ tends to zero, which concludes the proof. QED

A5. Proof of Proposition 6

Player i optimally sets p_i such that $p_i/(1 - p_i) = v(\hat{x}_i + \hat{\xi}_i)/v(x_i) = W_i/\sigma_i^2$. Substituting p_i, we get the desired expression for L_i. QED

A6. Proof of Result 1

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

Assume by contradiction that $m_j > 0$. Then (13) implies that $m_i > 0$ for all i, so $m \gg 0$. Next, from (22), and letting $r_i = 1/h_i$, we obtain $\hat{y}_i = r_i \hat{x}_i + (1 - r_i) y_i$, hence substituting y_i,

$$\hat{y}_i = (1 - r_i) p_i x_i + (1 - (1 - r_i) p_i) \hat{x}_i \,. \tag{23}$$

[49] $Q^i = \Sigma_{n \geq 0}((I - M^i)(I - \alpha^i)\tilde{A}^i)^n$, so Q^i is nonnegative. If in addition $m_{-i} \ll 1$, and since A is connected, then $Q^i \gg 0$.

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

$$v(y_i) < v(\hat{y}_i) \le \max_{j \neq i} v(y_j), \tag{24}$$

where the second inequality follows from $\hat{y}_i$ being an average of the y_j's. Since (24) 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, $\hat{\pi}^i$ the relative weights on $k \neq i$, and $\hat{\pi}_i = \hat{\pi}^i.x_{-i}$. We have $y_i = \pi_i x_i + (1 - \pi_i)\hat{x}_i$, and modifying γ_i allows the agent to modify π_i without affecting $\hat{x}_i$ (player i increases π_i by decreasing γ_i). Therefore, the optimal choice π_i satisfies

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

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

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

which implies that

$$\pi_i = \pi_i^* + \frac{(1 - \pi_i)(1 - \pi_i^*)}{\sigma_i^2}(v(\hat{x}_i) - W_i^*). \tag{25}$$

The weights π_i must thus be above the efficient weights π_i^*. Since all π_i (and π_i^*) add up to one, one must have $v(\hat{x}_i) = W_i^*$; hence, information aggregation is perfect. QED

Before showing result 2, we start with two lemmas that we also use to prove result 3:

LEMMA A2. For each $j \neq i$, there exists μ_{ji} and a probability vector $C^{ji} \in \Delta_{n-1}$, each independent of m_i, such that

$$y_j = (1 - \mu_{ji})C^{ji}\tilde{x}_{-i} + \mu_{ji}y_i. \tag{26}$$

Proof. This immediately follows from expression (20) in the proof of proposition 4. QED

LEMMA A3. If $\partial L_i/\partial m_i \le 0$, then $\partial L_j/\partial m_i < 0$ for all j.

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

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

We substitute $y_i = p_i x_i + (1 - p_i)(\hat{x}_i + \hat{\xi}_i)$ (see [13]). Since $\tilde{x}_k$ and x_i are independent, and since $\hat{x}_i$, $\tilde{x}_k$, and $\hat{\xi}_i$ do not depend on m_i, we get

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

The terms $\partial p_i/\partial m_i$ and $\text{Cov}(x_k \hat{x}_i)$ are positive, and so are the terms $\text{Cov}(\tilde{x}_k \hat{\xi}_i)$ 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 A3. QED

A7. Proof of Result 2

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

Assume first that all other players use DG. Then by proposition 5, L_i decreases strictly with m_i. Now assume that at least one player j chooses $m_j > 0$. Then $L_i = p_i^2 + (1 - p_i)^2 v(\hat{x}_i + \hat{\xi}_i)$. Whether persistent errors are independent or fully correlated, the variance of $\hat{\xi}_i$ is at least equal to $h_i^2\varpi$, which implies that L_i strictly decreases with p_i when $p_i/(1 - p_i) < h_i^2\varpi$ and hence also with m_i when $m_i/(1 - m_i) < h_i\varpi$, and from lemma A2 we conclude that L_j decreases as well (on this range of m_i). QED

A8. Proof of Result 3

1. Step 1: Lower Bounds on $\bar{m}^i \equiv \max_{j\neq i} m_j$

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

$$\frac{p_i}{1 - p_i} = \frac{v(\hat{x}_i) + v(\hat{\xi}_i)}{\sigma_i^2}, \tag{27}$$

and (25) becomes

$$p_i = \pi_i^* + \frac{(1 - p_i)(1 - \pi_i^*)}{\sigma_i^2}(v(\hat{x}_i) - v_i^* + v(\hat{\xi}_i)). \tag{28}$$

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

While (28) implies a lower bound on p_i, as (25) 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 one. Nevertheless, when all m's are small, $\Sigma_i p_i = 1 + O(m)$ is close to one, and this allows us to bound $v(\hat{\xi}_i)$ (and the difference $v(\hat{x}_i) - v_i^*$), as we now explain.

From proposition 4, each opinion y_i may be written as $y_i = P^i x + (1 - P_i^i)\hat{\xi}_i$, where P^i is a weighting vector (such that $P_i^i = p_i$). Equation (21) 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 driven mostly by the terms $\hat{\xi}_i$. Specifically, let $\bar{m}^i = \max_{j\neq i} m_j$. Equation (21) implies that for all $k \neq i$,

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

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

$$1 - p_i = \sum_{k \neq i} P_k^i \geq \sum_{k \neq i} p_k - Kc\bar{m}^i \geq 1 - \pi_i^* - Kc\bar{m}^i, \tag{29}$$

which, combined with (28), yields for some constant d,

$$\bar{m}^i \geq d\left(v(\hat{x}_i) - v_i^* + \frac{\varpi}{(\bar{m}^i)^2}\right). \tag{30}$$

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

2. Step 2: Upper Bounds on $\bar{m}^i$

Let $r_i = \Sigma_{j \neq i} R_j m_j$ and $\hat{y}_i = \Sigma_{k \neq i} A_{ik} y_k$. With transmission errors, we obtain

$$\hat{y}_i = (1 - r_i)p_i x_i + (1 - (1 - r_i)p_i)(\hat{x}_i + \hat{\xi}_i) + \bar{\xi}_i,$$

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

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

for some constant c and d (independent of ϖ and m). Since $v(\hat{y}_i) \leq \max v(y_k)$, 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 $\bar{m}^{i_0}$ and all m_j with $j \neq i_0$ are $O(\varpi^{1/3})$ and that m_{i_0} is thus *at least* $O(\varpi^{1/3})$.

It only remains to check that m_{i_0} cannot be large. From (29), $p_{i_0} \leq \pi_{i_0}^* + O(\varpi^{1/3})$, and since $p_{i_0} \geq 1/(1 + r_{i_0}/m_{i_0})$, we conclude that all m_i (and thus $\bar{m}^i$) are $O(\varpi^{1/3})$, which further implies that all variances $v(\hat{\xi}_i)$ are $O(\varpi^{1/3})$.

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

A9. Proof of Result 4

This follows from lemma A3 since at equilibrium $\partial L_i/\partial m_i = 0$. QED

A10. Proof of Expression (17)

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

A11. Proof of (Generalized) Result 5

Rather than assuming that the central player is DG, here we consider a central player who uses her seed x_0 optimally to minimize the loss $v(\bar{y})$, given m. We have $\bar{y} = (1 - m)y_0$ and $y_0 = m_0 x_0 + (1 - m_0)(\bar{y} + \xi_0)$. This gives $\bar{y} = (1 - m)$

$(p_0 x_0 + (1 - p_0)(\xi_0/m))$, where the central player controls p_0. The variance $v(\bar{y})$ is minimized for $p_0/(1 - p_0) = \varpi_0/m^2$, and we get $v(\bar{y}) = (1 - m)^2[(\varpi_0/m^2)/(1 + \varpi_0/m^2)]$. So long as $m \gg (\varpi_0)^{1/2}$, we obtain $D \simeq 4\varpi_0/d$ as for the DG case. Note that when $m \leq O(\varpi_0)^{1/2}$, cumulated errors are potentially huge and the (benevolent) central player mitigates them by choosing a large m_0; since she is benevolent, the loss cannot exceed one (the variance of her own seed). QED

A12. Network Comparisons

We conclude this appendix with network comparisons. We derive proposition A0 (see below), on which the discussion in the main text is based and that we prove in appendix B.

To facilitate network comparisons, we assume initial signals of identical precision ($\sigma_i^2 = 1$), so that the efficient weighting of signals is $\pi_i^* = 1/n$ and $\underline{W}_i^* \equiv \text{var}(\pi_{-i}^*.x_{-i}) = 1/(n-1)$. All players are subject to a processing error ξ_i, with same variance ϖ. From proposition 6, player i's incentives yield

$$\frac{m_i}{1 - m_i} = W_i/h_i, \tag{31}$$

where $W_i = \text{var}(\hat{x}_i) + \text{var}(\hat{\xi}_i)$. Both h_i and W_i depend only on m_{-i} and the structure of the network, and the equilibrium values m_i^* are obtained by simultaneously solving these equations. Given these equilibrium values, we can then compute W_i^* and hence (by proposition 6) the equilibrium loss L_i^*. To measure how losses L_i depart from the minimum loss $\underline{L}_i^*$, we define

$$\hat{\Delta}_i \equiv W_i - \underline{W}_i^*,$$

which characterizes the size of the inefficiency resulting from the inefficient aggregation of others' signals and cumulated errors. Defining

$$\rho_i \equiv \frac{p_i}{1 - p_i} - \frac{1}{n-1} = \frac{m_i h_i}{1 - m_i} - \frac{1}{n-1},$$

the equilibrium condition can thus be written

$$\rho_i = \hat{\Delta}_i,$$

which has the following economic interpretation: the relative weight on x_i (relative to other signals) should exceed the efficient weighting by $\hat{\Delta}_i$.

We compare three n-player networks: the *complete network*, where each player is connected to all others; the *directed circle*, where information transmission is directed and one-sided (player i communicates to player $i - 1$, who communicates to $i - 2$, etc.—player 0 is player n); and the *star network*, which consists of $n - 1$ peripheral players labeled $k = 1, \ldots, n - 1$ and a central player, labeled 0, who aggregates the opinions of the peripheral players.

For each network, we characterize h_i, $\hat{x}_i$, and $\hat{\xi}_i$ (and hence ρ_i and $\hat{\Delta}_i$), indicating a superscript c for the complete network, d for the directed circle, and s for the star network. We next solve for equilibrium, focusing on the limit cases where ϖ is small (for a fixed n) and where for a fixed small ϖ, n gets large. For the complete network and the directed circle, we solve for a symmetric equilibrium. For the star network, we solve for an equilibrium where all peripheral players use the same weight m and the central player, labeled player 0, uses m_0. We obtain the following:

Proposition A0. For fixed $n \geq 3$ and small ϖ, $\hat{\Delta}_d^* < \hat{\Delta}_c^* < \hat{\Delta}_s^*$. For fixed small ϖ, at the large-n limit, $\hat{\Delta}_c^* < \hat{\Delta}_d^* < \hat{\Delta}_s^*$. These comparisons hold whether errors are independent or correlated. Furthermore, for the star network, $m_0^*/m^* \simeq 1/(n-1)$ for fixed n and small ϖ, and $m_0^*/m^* \leq (2\varpi)^{1/3}$ at the large-n limit.

References

Acemoglu, D., M. A. Dahleh, I. Lobel, and A. Ozdaglar. 2011. "Bayesian Learning in Social Networks." *Rev. Econ. Studies* 78 (4): 1201–36.

Alatas, V., A. Banerjee, A. G. Chandrasekhar, R. Hanna, and B. A. Olken. 2016. "Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia." *A.E.R.* 106 (7): 1663–704.

Aumann, R. J. 1976. "Agreeing to Disagree." *Ann. Statis.* 4 (6): 1236–39.

Axelrod, R. 1984. *The Evolution of Cooperation.* New York: Basic Books.

Bakshy, E., S. Messing, and L. A. Adamic. 2015. "Exposure to Ideologically Diverse News and Opinion on Facebook." *Science* 348 (6239): 1130–32.

Banerjee, A. 1992. "A Simple Model of Herd Behavior." *Q.J.E.* 107 (3): 797–817.

Banerjee, A., and O. Compte. 2023. "Consensus and Disagreement: Information Aggregation under (Not So) Naive Learning." Working Paper no. 29897, NBER, Cambridge, MA.

Bertrand, M., and E. Kamenica. 2023. "Coming Apart? Cultural Distances in the United States over Time." *American Econ. J. Appl. Econ.* 15 (4): 100–141.

Bikhchandani, S., D. Hirshleifer, and I. Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *J.P.E.* 100 (5): 992–1026.

Bohren, J. A., and D. N. Hauser. 2021. "Learning with Heterogeneous Misspecified Models: Characterization and Robustness." *Econometrica* 89 (6): 3025–77.

Boxell, L., J. Conway, J. N. Druckman, and M. Gentzkow. 2022. "Affective Polarization Did Not Increase during the COVID-19 Pandemic." *Q. J. Polit. Sci.* 17 (4): 491–512.

Compte, O., and A. Postlewaite. 2018. *Ignorance and Uncertainty.* Econometric Society Monographs. Cambridge: Cambridge Univ. Press.

Cox, D. A., R. Streeter, S. J. Abrams, and J. Clemence. 2020. "Socially Distant: How Our Divided Social Networks Explain Our Politics." Report, Survey Center American Life.

Dasaratha, K., B. Golub, and N. Hak. 2023. "Learning from Neighbours about a Changing State." *Rev. Econ. Studies* 90 (5): 2326–69.

DeGroot, M. H. 1974. "Reaching a Consensus." *J. American Statis. Assoc.* 69 (345): 118–21.

DeMarzo, P. M., D. Vayanos, and J. Zwiebel. 2003. "Persuasion Bias, Social Influence, and Unidimensional Opinions." *Q.J.E.* 118 (3): 909–68.

Ellison, G., and D. Fudenberg. 1993. "Rules of Thumb for Social Learning." *J.P.E.* 101 (4): 612–43.

———. 1995. "Word-of-Mouth Communication and Social Learning." *Q.J.E.* 110 (1): 93–125.

Eyster, E., and M. Rabin. 2010. "Naive Herding in Rich-Information Settings." *American Econ. J. Microeconomics* 2 (4): 221–43.

Frick, M., R. Iijima, and Y. Ishii. 2020. "Misinterpreting Others and the Fragility of Social Learning." *Econometrica* 88 (6): 2281–328.

Friedkin, N. E., and E. C. Johnsen. 1990. "Social Influence and Opinions." *J. Math. Soc.* 15 (3/4): 193–206.
———. 1999. "Social Influence Networks and Opinion Change." *Advances Group Processes* 16:1–29.
Fudenberg, D. 1998. *The Theory of Learning in Games.* Cambridge, MA: MIT Press.
Genest, C., and J. V. Zidek. 1986. "Combining Probability Distributions: A Critique and an Annotated Bibliography." *Statis. Sci.* 1 (1): 114–35.
Gentzkow, M. 2016. "Polarization in 2016." Working paper.
Gentzkow, M., and J. M. Shapiro. 2011. "Ideological Segregation Online and Offline." *Q.J.E.* 126 (4): 1799–839.
Gentzkow, M., M. Wong, and A. T. Zhang. 2021. "Ideological Bias and Trust in Information Sources." Working paper.
Golub, B., and M. O. Jackson. 2010. "Naive Learning in Social Networks and the Wisdom of Crowds." *American Econ. J. Microeconomics* 2 (1): 112–49.
Golub, B., and E. Sadler. 2017. "Learning in Social Networks." Working paper.
Guess, A. M. 2021. "(Almost) Everything in Moderation: New Evidence on Americans' Online Media Diets." *American J. Polit. Sci.* 65 (4): 1007–22.
Jackson, M. O. 2008. *Social and Economic Networks.* Princeton, NJ: Princeton Univ. Press.
Jackson, M. O., S. Malladi, and D. McAdams. 2019. "Learning through the Grapevine and the Impact of the Breadth and Depth of Social Networks." Working paper.
Jadbabaie, A., P. Molavi, A. Sandroni, and A. Tahbaz-Salehi. 2012. "Non-Bayesian Social Learning." *Games and Econ. Behavior* 76 (1): 210–25.
Levy, G., and R. Razin. 2015. "Correlation Neglect, Voting Behavior, and Information Aggregation." *A.E.R.* 105 (4): 1634–45.
Levy, R. 2021. "Social Media, News Consumption, and Polarization: Evidence from a Field Experiment." *A.E.R.* 111 (3): 831–70.
Molavi, P., A. Tahbaz-Salehi, and A. Jadbabaie. 2018. "A Theory of Non-Bayesian Social Learning." *Econometrica* 86 (2): 445–90.
Mossel, E., A. Sly, and O. Tamuz. 2015. "Strategic Learning and the Topology of Social Networks." *Econometrica* 83 (5): 1755–94.
Mueller-Frank, M. 2017. "Robust Non-Bayesian Learning." Working paper.
Mueller-Frank, M., and C. Neri. 2021. "A General Analysis of Boundedly Rational Learning in Social Networks." *Theoretical Econ.* 16 (1): 317–57.
Rosenberg, D., E. Solan, and N. Vieille. 2009. "Informational Externalities and Emergence of Consensus." *Games and Econ. Behavior* 66 (2): 979–94.
Sethi, R., and M. Yildiz. 2012. "Public Disagreement." *American Econ. J. Microeconomics* 4 (3): 57–95.
———. 2016. "Communication with Unknown Perspectives." *Econometrica* 84 (6): 2029–69.
———. 2019. "Culture and Communication." Working paper.
Sobel, J. 2014. "On the Relationship between Individual and Group Decisions." *Theoretical Econ.* 9 (1): 163–85.
Sunstein, C. R. 2001. *Echo Chambers: Bush v. Gore, Impeachment, and Beyond.* Princeton, NJ: Princeton Univ. Press.
Vives, X. 1993. "How Fast Do Rational Agents Learn?" *Rev. Econ. Studies* 60 (2): 329–47.
———. 1997. "Learning from Others: A Welfare Analysis." *Games and Econ. Behavior* 20 (2): 177–200.