

Information diffusion in networks with the Bayesian Peer Influence heuristic

Gilat Levy, LSE and Ronny Razin, LSE¹

Abstract: Repeated communication in networks is often considered to impose large information requirements on individuals, and for that reason, the literature has resorted to use heuristics, such as DeGroot's, to compute how individuals update beliefs. In this paper we propose a new heuristic which we term the Bayesian Peer Influence (BPI) heuristic. The BPI accords with Bayesian updating for all (conditionally) independent information structures. More generally, the BPI can be used to analyze the effects of correlation neglect on communication in networks. We analyze the evolution of beliefs and show that the limit is a simple extension of the BPI and parameters of the network structure. We also show that consensus in society might change dynamically, and that beliefs might become polarized. These results contrast with those obtained in papers that have used the DeGroot heuristic.

1 Introduction

Repeated communication in groups and more generally in networks is often considered to impose large informational requirements on individuals. Individuals may be unaware of the structure of the network, so that while they know who they communicate with, they might not know their neighbors' neighbors. This implies that it may be very difficult to trace the path that a piece of information takes in an environment with repeated communication.

The network literature has typically taken one of two avenues. One avenue is the fully rational approach whereby individuals are fully aware of the network and the equilibrium and update using Bayes rule (see Acemoglu et al 2014). The second avenue is to assume that individuals follow a particular heuristic when updating. A leading example is the DeGroot heuristic, where individuals average their's and others' beliefs, as in Golub and Jackson (2010) and De Marzo et al (2003). These are two polar ways to model information diffusion, one based on full rationality and the other based on an adhoc heuristic.

¹Levy: Department of Economics, LSE, Houghton St. WC2A 2AE, UK; g.levy1@lse.ac.uk. Razin: Department of Economics, LSE, Houghton St. WC2A 2AE, UK; r.raizin@lse.ac.uk. We thank Erik Eyster, Francesco Nava, Tristan Gagnon-Bartsch, Matthew Rabin and Alireza Tahbaz-Salehi for helpful discussions. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 681579.

In this paper we analyze information diffusion in networks by using a new heuristic which is based on rational foundations for all information structures which are conditionally independent. Specifically, we assume a simple model of communication in which individuals sincerely transmit their beliefs to each other.² Sobel (2014) and Levy and Razin (2016) show that if individuals believe that their marginal information sources generate (conditionally) independent signals, then upon communication, Bayesian updating yields a belief that is proportional to a simple multiplication of their posterior beliefs. Thus, if $p(\omega)$ is a common prior about a state of the world $\omega \in \Omega$ (finite), and $q^i(\omega)$ is the posterior belief of individual i that the state is ω following some signal realisation, the resulting belief following communication of posteriors in a group of n such individuals is

$$\frac{\frac{1}{p(\omega)^{n-1}} \prod_{i \in N} q^i(\omega)}{\sum_{v \in \Omega} \frac{1}{p(v)^{n-1}} \prod_{i \in N} q^i(v)},$$

where the formula can be easily extended to non-common priors, as well as to more general perceptions of correlation (Levy and Razin 2016). In this paper we adopt this formula as an updating heuristic, which we term the Bayesian Peer Influence (BPI) heuristic. We analyse information diffusion in networks using the BPI.³

There are several advantages to using the BPI. First, it is a heuristic which is rational for all environments in which the information sources of individuals are truly (conditionally) independent. A benefit of using the BPI in these environments is that it is “information structure free”: Both the modeler and individuals in the network do not need to know the exact information structures of others in order to compute the properties of beliefs in the network. We therefore do not need to make any specific assumptions about information structures. Second, in other, more complicated, environments, the BPI is simple to compute. The BPI is an order-free heuristic, which -as we show- lends itself easily to the computation of limit beliefs.

The BPI also allows us to isolate the implications of correlation neglect arising from repeated communication from other types of incorrect information processing biases. Correlation neglect has recently attracted attention in the literature and is a natural bias to arise in network communication.⁴ As information flows in a network, individuals may be

²The assumption of sincerity is quite reasonable in the context of information diffusion in networks; as is the case in most of the literature, for such environment it is common to assume that individuals are not strategic (see the survey in Jackson 2011). The assumption that people communicate their beliefs is motivated by the difficulty to remember and communicate the exact details of information structures.

³The BPI is used for the case of binary states and particular information structures by Duffie and Manso (2010) and Eyster and Rabin (2010, 2014).

⁴Ortoleva and Snowberg (2015) analyze the effect of correlation neglect on the polarisation of beliefs. De Marzo et al (2003) and Gagnon-Bartsch and Rabin (2015) study how it affects the diffusion of

unaware that they are exposed to the same information they have been exposed to in the past, and thus if they treat information as independent, correlation neglect is likely to arise. When using the BPI to analyze repeated communication in a network, we are in effect assuming that correlation neglect is the only departure from rationality.

Our main results are the following. First, we characterise the limit beliefs of repeated communication in networks and show that they are easily computed. Similarly to Golub and Jackson (2010) and De Marzo et al (2003) the limit beliefs depend both on the initial information in the network and the structure of the network. Beliefs converge to the mode of a belief that extends the BPI to account for the initial beliefs and the eigenvector centrality of the individuals in the network. In particular, beliefs converge to the mode of $\Pi_{j \in N}(q_0^j(\omega))^{\alpha^j}$ where $q_0^j(\omega)$ is the initial belief of individual j and α^j is the j 'th element of the vector α that is parallel to the Perron–Frobenius eigenvector.

The result about the mode allows us to provide sharp predictions about the limit beliefs, which qualitatively differ from those obtained under the DeGroot heuristic. Specifically, a recent theoretical and experimental literature has focused on the question of whether polarised beliefs in society arise because of correlation neglect (see for example Glaeser and Sunstein 2009, Schkade et al 2000, and Sobel 2014 for an alternative view). In our model polarisation arises as beliefs become degenerate in the limit. Using the BPI we show how polarisation depends both on the network configuration as well as on the nature of the initial belief. Relatedly, the BPI also implies that consensus might change dynamically; that is, even when all have the same beliefs -as long as these are not degenerate- individuals will continue to update from each other. These two results cannot arise within the DeGroot framework where limit beliefs are always in the convex hull of initial group beliefs and in which when consensus is reached there is no further updating. This implies that the BPI can account for phenomena such as “Groupthink” whereby even group homogeneity implies polarisation.

We show that with the BPI an individual’s influence on the group depends both on his centrality but also on the quality of his information. In particular, the variance of an individual’s belief is important. An individual holding beliefs with high variance has little effect on others’ beliefs. Again, this contrasts with the DeGroot heuristic under which only the expectation of an individual’s belief (or some exogenous parameters) can determine his influence. For example, a uniform belief will imply that an individual has no influence on others in our model, no matter his centrality in the network. In contrast, in the DeGroot model, such an individual will be influential as long as his expectation is

information in social networks. Glaeser and Sunstein (2009) and Levy and Razin (2015a, 2015b) explore the implications for group decision making in political applications. Recent experimental evidence is in Eyster and Weizsacker (2011), Kallir and Sonsino (2009) and Enke and Zimmermann (2013).

different than others' and the higher is his centrality in the network.

Our paper relates to several strands of the literature. We contribute to the literature on information diffusion in networks (Golub and Jackson 2010, 2012) by suggesting a new heuristic, which, as in machine learning, is motivated by a “Naive-Bayes” approach. Our analysis complements a recent paper by Molavi, Tahbaz-Salehi and Jadbabaie (2015). Whereas we assume rationality up to correlation neglect, Molavi et al (2015) take an axiomatic approach, focusing on imperfect recall. They characterize a family of heuristics which embed the BPI as a special case. Although our focus is on communication, our paper is also related to the social learning literature which studies how individuals learn from others' actions (e.g., Bala and Goyal 1998). We show that society can converge to be fully confident in one state of the world (the mode of the adjusted BPI), which is not necessarily the true state of the world.⁵

Some recent experimental literature examines the predictions of the theory of information diffusion in networks, mainly comparing the fully Bayesian updating with the DeGroot heuristic. Choi et al (2012) find that the Bayesian model fits the data well in three-player networks, although such networks are perceived to lack statistical power to distinguish between the DeGroot and the fully Bayesian model. Grimm and Mengel (2014) show that the DeGroot model is better than the Bayesian at predicting correct guesses, while the Bayesian model is better at predicting whether consensus arises. They conclude that agents are semi-Bayesian in the sense that they do take into consideration the network structure when updating, albeit in a rudimentary way. Mueller-Frank and Neri (2013) show that consensus is hard to achieve when agents are heterogenous, even when the network is strongly connected.⁶ Our analysis also falls between the Bayesian model and the naive one, and our results illustrate how consensus can be dynamic which is an interesting avenue to explore in future experiments. Finally, Philippos et al (2017) consider how individuals update when they consider different issues (our analysis focuses on a unidimensional state of the world) and show how the type of unidimensional disagreement that arises depends on the communication channels.

⁵See also Eyster and Rabin (2010) and Gagnon-Bartch and Rabin (2015). Other papers such as Guarino and Jehiel (2013) show that society can still learn the truth, even in the presence of information processing biases.

⁶Both Grimm and Mengel (2014) and Mueller-Frank and Neri (2013) suggest alternative updating approaches, as we do here. We differ as our heuristic is always Bayesian in the first period, and departs from rationality only because of correlation neglect.

2 The Model

Information and initial conditions: There is a finite set N of individuals who have a common uniform prior on the state $\omega \in \Omega$, where Ω is finite. The analysis can be easily extended to non-uniform prior (as shown in the introduction) as well as to known non-common priors. Each individual $j \in N$ holds an initial posterior $q_0^j(\omega)$. We assume that $q_0^j(\omega) > 0$ for any $\omega \in \Omega$ and any j , and that these initial posterior beliefs were derived by the individuals observing signals coming from some (conditionally) independent information sources. Let $\mathbf{q}_0 = (q_0^j(\cdot))_{j \in N}$ denote the vector of original beliefs.

The network and communication: The individuals in N are organized in a network. Let \mathbf{T} denote the matrix of links where $T_{ij} = 1$ if there is a link between i and j and 0 otherwise. For simplicity we assume that the links are non-directed. Note also that $T_{ii} = 1$. At any period, individuals i, j communicate their beliefs to each other if $T_{ij} = 1$. After each period $k \geq 1$, individual i updates her belief using the BPI heuristic (see below) to a new posterior $q_k^i(\omega)$. When communicating at period $k + 1$, individuals truthfully reveal their beliefs, $q_k^i(\omega)$.

The Bayesian Peer Influence (BPI) heuristic: Let $\mathbf{q}_k^i = (q_k^j(\cdot))_{j \in N | T_{ij}=1}$ denote the vector of beliefs observed by individual i at period $k + 1$. If the individual thinks that these beliefs have been derived from independent information sources, then Bayesian updating implies that after observing q_k^i , individual i updates to the belief:⁷

$$q_{k+1}^i(\omega) = \frac{\prod_{j \in N | T_{ij}=1} q_k^j(\omega)}{\sum_{v \in \Omega} \prod_{j \in N | T_{ij}=1} q_k^j(v)}.$$

Note that the BPI heuristic is order-free, information-structure free, and rational whenever all beliefs stem from conditionally independent information sources. While in the first stage of communication information is truly independent across individuals, this will not be the case in the second period as information in the network will become repeated. We assume that individuals use the BPI at any stage. Thus, while belief updating is fully rational in the first stage of communication, correlation neglect arises at later stages and updating is then not fully Bayesian.

3 Limit beliefs

We now find conditions on the environment under which the limit beliefs exist and characterise them. Formally, the primitives of the environment are the network structure and

⁷See Levy and Razin (2016) Proposition 1 and Sobel (2014) Proposition 5.

the set of original beliefs, $(\mathbf{T}, \mathbf{q}_0)$. Note that at period k , the individual will have been exposed to the posterior of j a number of times equal to all the possible paths in the network from j to i that involve up to k steps. If individuals use the BPI, the order in which information was heard does not matter. In matrix notation, the posterior after k rounds can be written as,

$$q_{k+1}^i(\omega) = \frac{\prod_{i \in N} (q_0^i(\omega))^{\mathbf{T}_{ij}^k}}{\sum_{v \in \Omega} \prod_{i \in N} (q_0^i(v))^{\mathbf{T}_{ij}^k}}, \quad (1)$$

where $\mathbf{T}^k = \underbrace{\mathbf{T} \times \mathbf{T} \times \dots \times \mathbf{T}}_{k \text{ times}}$.

Therefore, to study the convergence of this process we need to study the convergence of \mathbf{T}^k . If \mathbf{T} is irreducible and aperiodic, then it is primitive and we can use the Perron–Frobenius Theorem about the convergence of primitive matrices.

In particular let $\boldsymbol{\alpha} = (\alpha^1, \dots, \alpha^n) \in \mathbb{R}_+^n$ be the Perron-Frobenius eigenvector of \mathbf{T} . The following is an assumption about \mathbf{q}_0 and $\boldsymbol{\alpha}$.

Assumption 1 $\arg \max_{\omega \in \Omega} \prod_{j \in N} (q_0^j(\omega))^{\alpha^j}$ is unique.

Note that Assumption 1 will be satisfied generically when randomly choosing \mathbf{q}_0 and \mathbf{T} .

Before moving on to our main result, we illustrate with examples what happens when Assumption 1 is violated.

Let $\Omega = \{0, 1\}$. Consider a network with four agents $N = \{1, 2, 3, 4\}$ where the agents are on a circle and each agent is connected to her two neighbours, i.e.,

$$\mathbf{T}' = \begin{pmatrix} 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \end{pmatrix}.$$

This network is connected and aperiodic and its Perron-Frobenius eigenvector is $(1, 1, 1, 1)$.

Example 1 (beliefs never change and remain at \mathbf{q}_0) : Let the initial beliefs be $\mathbf{q}_0(1) = (\frac{3}{4}, \frac{3}{4}, \frac{1}{4}, \frac{1}{4})$. Note that with these beliefs and network matrix \mathbf{T}' , $\prod_{j \in N} (q_0^j(1))^{\alpha^j} = \prod_{j \in N} (q_0^j(0))^{\alpha^j} = (\frac{3}{4})^2 (\frac{1}{4})^2$ and so Assumption 1 is not satisfied. For this starting point, \mathbf{q}_0 , updating according to the BPI heuristic leads to a second period vector of posterior beliefs which is the same as \mathbf{q}_0 . To see this take agent 2 for example. At the end of period 1 his belief will be

$$\begin{aligned}
q_1^2(1) &= \frac{q_0^1(1)q_0^2(1)q_0^3(1)}{q_0^1(1)q_0^2(1)q_0^3(1) + q_0^1(0)q_0^2(0)q_0^3(0)} = \frac{\frac{3}{4} \frac{3}{4} \frac{1}{4}}{\frac{3}{4} \frac{3}{4} \frac{1}{4} + \frac{1}{4} \frac{1}{4} \frac{3}{4}} \\
&= \frac{3}{4} = q_0^2(1)
\end{aligned}$$

Similarly this will be the case for the others, as the neighbours' information always cancels out. So while each agent's beliefs converges, these beliefs do not converge to the same one.

Example 2 (beliefs alternate and never converge): Assume the same state space and network as above but now consider the initial beliefs $\mathbf{q}_0(1) = (\frac{3}{4}, \frac{1}{4}, \frac{3}{4}, \frac{1}{4})$. Note again that with these beliefs and network matrix \mathbf{T}' , $\prod_{j \in N} (q_0^j(1))^{\alpha^j} = \prod_{j \in N} (q_0^j(0))^{\alpha^j} = (\frac{3}{4})^2 (\frac{1}{4})^2$ and so Assumption 1 is not satisfied. Note that for this starting point, \mathbf{q}_0 , updating according to the BPI heuristic leads to a second period vector of posterior beliefs which is a permutation of \mathbf{q}_0 . To see this take agent 2 for example. At the end of period 1 his belief will be

$$\begin{aligned}
q_1^2(1) &= \frac{q_0^1(1)q_0^2(1)q_0^3(1)}{q_0^1(1)q_0^2(1)q_0^3(1) + q_0^1(0)q_0^2(0)q_0^3(0)} = \frac{\frac{3}{4} \frac{1}{4} \frac{3}{4}}{\frac{3}{4} \frac{1}{4} \frac{3}{4} + \frac{1}{4} \frac{1}{4} \frac{3}{4}} \\
&= \frac{3}{4} > \frac{1}{4} = q_0^2(1)
\end{aligned}$$

The opposite will happen for agent 3:

$$\begin{aligned}
q_1^3(1) &= \frac{q_0^4(1)q_0^3(1)q_0^4(1)}{q_0^4(1)q_0^3(1)q_0^4(1) + q_0^4(0)q_0^3(0)q_0^4(0)} = \frac{\frac{1}{4} \frac{3}{4} \frac{1}{4}}{\frac{1}{4} \frac{3}{4} \frac{1}{4} + \frac{3}{4} \frac{1}{4} \frac{3}{4}} \\
&= \frac{1}{4} < \frac{3}{4} = q_0^3(1)
\end{aligned}$$

Therefore, at each period beliefs permute, with each agent acquiring the beliefs of his (left) neighbour. So in this case again beliefs do not converge.

We are now ready to state our main result.

Proposition 1 *Assume that \mathbf{T} is connected (irreducible) and aperiodic and that q_0 and \mathbf{T} satisfy Assumption 1. Then there exists a vector $\alpha = (\alpha^1, \dots, \alpha^n) \in \mathbb{R}_+^n$ such that: (i) the limit posterior beliefs of all players converge to a degenerate belief on the maximiser of $\prod_{j \in N} (q_0^j(\omega))^{\alpha^j}$. (ii) the vector α is parallel to the Perron-Frobenius eigenvector of \mathbf{T} .*

To see the intuition of the result, consider the complete network. Note that after the first period of communication, all would have the same beliefs, at $q_1(\omega) \equiv \frac{\prod_{i \in N} q_0^i(\omega)}{\sum_{\nu \in \Omega} \prod_{i \in N} q_0^i(\nu)}$.

After the second period, beliefs would be at $q_2(\omega) \equiv \frac{(q_1(\omega))^n}{\sum_{\nu \in \Omega} (q_1(\nu))^n}$ and more generally at the k 'th period, beliefs would be at $q_k(\omega) \equiv \frac{(q_1(\omega))^{n^{k-1}}}{\sum_{\nu \in \Omega} (q_1(\nu))^{n^{k-1}}}$. Thus, when $k \rightarrow \infty$, we would have that for any two states ω and ω' ,

$$\frac{q_k(\omega)}{q_k(\omega')} = \left[\frac{q_1(\omega)}{q_1(\omega')} \right]^{n^{k-1}} \rightarrow_{k \rightarrow \infty} \infty \text{ if } q_1(\omega) > q_1(\omega') \text{ and otherwise to } 0.$$

Therefore, the mode of the α -weighted BPI is important in determining the limit beliefs. The following example illustrates the implications of Proposition 1.

Example 3: Let us consider the following simple example. Consider three players. Player 1 is connected to both 2 and 3, while each of them is only connected to player 1 (a ‘‘star’’ network). Thus, the \mathbf{T}^* matrix is (recall that every player is connected to herself as well):

$$\mathbf{T}^* = \begin{matrix} & 1 & 2 & 3 \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & 1 & 1 & 0 \\ & 1 & 0 & 1 \end{matrix}.$$

When communication is repeated infinitely, then beliefs in society converge to a distribution that puts equal weights on a subset of the modes of $q^*(\omega) = \frac{(q_0^1(\omega))^{\sqrt{2}} \bullet q_0^2(\omega) \bullet q_0^3(\omega)}{\sum_{v \in \Omega} (q_0^1(v))^{\sqrt{2}} \bullet q_0^2(v) \bullet q_0^3(v)}$. This arises from calculating the eigenvector of \mathbf{T}^* which is $(\sqrt{2}, 1, 1)$. Thus, the procedure is simple, and the nature of the consensus depends on the matrix \mathbf{T}^* . As mentioned above, generically, $q^*(\omega)$ will have a unique mode. For example, for any $(\mathbf{T}, \mathbf{q}_0)$, any open set of perturbations in \mathbf{q}_0 will result in a subset of these inducing a unique mode, while only some specific perturbations will induce a set of modes.

4 Implications

We now use Proposition 1 to derive testable predictions about repeated communication in networks. The predictions will also allow us to differentiate our model from the DeGroot (1974) heuristic, according to which individuals average others’ and their’s beliefs (see De Marzo et al 2003 and Golub and Jackson 2010, 2012). Note that one crucial difference is methodological. That is, the DeGroot heuristic does not generally correspond to Bayesian updating.⁸ The BPI heuristic on the other hand coincides with Bayesian updating whatever the information structures are, as long as they are conditionally independent. Using the BPI therefore allows us to understand what are the biases in information processing

⁸This is so only in a limited set of environments. One example is when each information structure is a normal distribution.

that we are modelling. In particular the BPI departs from Bayesian updating due to individuals neglecting the correlation across their information sources (which arises with repeated communication). This bias has been recently identified by experiments in finance and political economy.⁹

We henceforth assume that the conditions of Proposition 1 are satisfied so that we can discuss the implications of the limit beliefs and focus on the generically unique mode of $q^*(\omega)$, denoted by ω^* .

4.1 Polarisation of beliefs and consensus dynamics

A recent literature, e.g., Glaeser and Sunstein (2009), attempts to explain the phenomena of group polarisation.¹⁰ We can say that group polarisation arises whenever the limit beliefs are not in the convex hull of q_0 . Note that when considering the DeGroot heuristic, the limit consensus beliefs cannot become polarised as they must be contained in the convex hull of individuals' beliefs due to the averaging procedure of DeGroot. On the other hand, with the BPI, given that we start with beliefs that are full support, and beliefs generically become degenerate on ω^* , the mode of $q^*(\omega)$, then we have polarisation.

The literature also discusses the relation between correlation neglect and polarisation of beliefs in groups, as for example in Glaeser and Sunstein (2009) and Schkade et al (2010). Sobel (2014) illustrates that correlation neglect is not a necessary condition for polarisation as it could arise also with Bayesian rational decision makers. Similarly, as the BPI corresponds with Bayesian updating for some environments, also in our model polarisation can arise for rational individuals. Consider for example the fully connected network and assume that there is only one round of communication. For a large enough n , for many vectors of beliefs \mathbf{q}_0 , beliefs will become degenerate. However, as we show below, whether the polarised beliefs accord with the truth or not does depend on whether correlation neglect arises or not. After one round of communication beliefs are always rational with the BPI, whereas repeated communication will entail convergence to the truth only in some environments.

The polarisation property of the BPI implies that small biases towards one state in the initial set of beliefs can loom large, as the example below illustrates:

Example 4 (Small biases loom large): Suppose that a share $1 - \varepsilon$ of the network has uniform beliefs while a share ε has beliefs which put a weight $\frac{1}{|\Omega|} + \nu$ on some state ω^* (and naturally a weight of $\frac{1}{|\Omega|} - \nu$ on some other state(s)). In this case, $q^*(\omega)$ has a unique

⁹See Ortoleva and Snowberg (2015), Eyster and Weizsacker (2011), Kallir and Sonsino (2009) and Enke and Zimmermann (2013).

¹⁰Experiments on group polarisation were initiated by Stoner (1968).

mode on ω^* . Thus, even for a small ν and ε , as well as for a low network centrality of the share ε of individuals with the biased beliefs, society converges to degenerate beliefs on ω^* . Thus, correlation neglect implies that small biases towards one state will accentuate in the long term.

Another simple observation is how consensus can be achieved but is not static. In the DeGroot model, once a consensus is achieved, it will never change, as beliefs are combined by averaging others. In contrast, with the BPI, correlation neglect implies that individuals will keep on learning from others and substantially change their beliefs even when their posterior beliefs are similar -as long as they are not degenerate. Therefore, with the BPI heuristic consensus can change dynamically.

Example 5 (Changing Consensus): To illustrate, consider an example with binary states, 1 or 0, and a complete network. The posterior belief of each individual is $q_0^i \in (0, 1)$ which denotes the probability that the state is 1. If all individuals have the same beliefs, $q^i = q^j \equiv q$, with the DeGroot heuristic these would remain unchanged. In contrast, with the BPI, beliefs would become $\frac{q^n}{q^n + (1-q)^n}$ after the first period. While consensus in society would remain, beliefs change and individuals would converge to have degenerate beliefs on state 1 if $\frac{q^n}{q^n + (1-q)^n} > \frac{1}{2}$ and on state 0 if $\frac{q^n}{q^n + (1-q)^n} < \frac{1}{2}$.

4.2 A Measure of Influence

How can we measure the influence of an individual? As the result in Proposition 1 shows, what is important is the effect of individuals on the mode of $q^*(\cdot)$. The mode ω^* satisfies

$$\frac{q^*(\omega^*)}{q^*(v)} = \frac{\prod_{j \in N} (q_0^j(\omega^*))^{\alpha^j}}{\prod_{j \in N} (q_0^j(v))^{\alpha^j}} > 1 \text{ for all } v \neq \omega^*,$$

and thus individual i 's effect on the mode depends on both his centrality in the network, α^i , as well as on the properties of his beliefs. We now explore these two ways of measuring the influence of an individual.

Definition 2: *Suppose the limit beliefs change from a degenerate belief on state ω^* to a degenerate belief on state $\hat{\omega}^*$. The new belief becomes closer to j 's beliefs compared with the old belief if $q_0^j(\hat{\omega}^*) > q_0^j(\omega^*)$.*

Proposition 2: *Suppose that α^j increases to $\hat{\alpha}^j$. Then the new limit beliefs $\hat{\omega}^*$ either does not change or if it does, it becomes closer to j 's beliefs.*

Proof: Note that

$$\frac{\hat{q}^*(\omega)}{\hat{q}^*(v)} = \frac{q^*(\omega) q_0^j(\omega)^{\hat{\alpha}^j - \alpha^j}}{q^*(v) q_0^j(v)^{\hat{\alpha}^j - \alpha^j}}.$$

As a result, $\frac{\hat{q}^*(\omega)}{\hat{q}^*(v)} > (<) \frac{q^*(\omega)}{q^*(v)}$ for all $q_0^j(\omega) > (<) q_0^j(v)$ implying that either the mode remains the same or it changes to another mode $\hat{\omega}^* \neq \omega^*$ but in this case $\frac{\hat{q}^*(\omega^*)}{\hat{q}^*(\hat{\omega}^*)} < \frac{q^*(\omega^*)}{q^*(\hat{\omega}^*)}$ implying $q_0^j(\omega^*) < q_0^j(\hat{\omega}^*)$ as desired. ■

Remark: Analogous to network centrality, one can be more influential if he puts more weight on his posterior as compared to his peers'. One way to model this with the BPI is to assume that an overconfident individual i , at period $k + 1$, updates his beliefs to $\frac{q_k^i(\omega)^\Upsilon \prod_{j|T_{ij}=1, j \neq i} q_k^j(\omega)}{\sum_{v \in \Omega} q_k^i(v)^\Upsilon \prod_{j|T_{ij}=1, j \neq i} q_k^j(v)}$, where $\Upsilon > 1$. It is easy to see then that the consensus would shift towards individuals i with $\Upsilon > 1$ (as others incorporate the beliefs of individual i) and as Υ grows large, i would become more and more influential. Overconfidence is therefore a substitute to network centrality.

We now focus on how changes in initial beliefs $q_0^j(\cdot)$ affect the limit beliefs in the group. Suppose one can order the set of states in some natural way. We say that ω is closer to ω^{*j} than v if ‘‘travelling’’ along this order from v to ω^{*j} you pass through ω . Consider then beliefs for individual j that are unimodal (single-peaked), and let ω^{*j} be the mode of the beliefs of individual j .

Definition 3: An individual j 's beliefs $\hat{q}_0^j(\cdot)$ become more confident in state ω^{*j} compared to $q_0^j(\cdot)$ if

$$\frac{\hat{q}_0^j(\omega)}{\hat{q}_0^j(v)} > \frac{q_0^j(\omega)}{q_0^j(v)} \text{ for any } \omega \text{ that is closer to } \omega^{*j} \text{ than } v.$$

What happens to the mode of $\hat{q}^*(\cdot)$ which follows individual j becoming more confident in her own mode? Note that

$$\frac{\hat{q}^*(\omega)}{\hat{q}^*(v)} = \frac{q^*(\omega)}{q^*(v)} \frac{q_0^j(v)}{q_0^j(\omega)} \frac{\hat{q}_0^j(\omega)}{\hat{q}_0^j(v)}.$$

Thus for any ω that is closer to ω^{*j} than v , we have $\frac{\hat{q}^*(\omega)}{\hat{q}^*(v)} > \frac{q^*(\omega)}{q^*(v)}$, and specifically $\frac{\hat{q}^*(\omega^{*j})}{\hat{q}^*(v)} > \frac{q^*(\omega^{*j})}{q^*(v)}$ for all $v \neq \omega^{*j}$. However, for all ω closer to ω^{*j} than ω^* , we have that $\frac{\hat{q}^*(\omega)}{\hat{q}^*(\omega^*)} > \frac{q^*(\omega)}{q^*(\omega^*)}$. Thus, it must be that either the new mode remains at ω^* , or that it moves to $\hat{\omega}^*$ that is closer to ω^{*j} . So then again we have:

Proposition 3: Consider single peaked beliefs for individual j . If individual j 's beliefs become more confident in his mode then the limit beliefs become closer to j 's beliefs. Moreover, for any full support vectors of beliefs $(q_0^1, q_0^2, \dots, q_0^n)$, there exists a K such that if $\frac{q_0^j(\omega^{*j})}{q_0^j(v)} > K$ for all $v \neq \omega^{*j}$, then $\omega^* = \omega^{*j}$.

At the extreme, when individual j becomes almost fully confident in state ω^{*j} , and in the limit of such sequences of beliefs we would have $\frac{q_0^j(\omega^{*j})}{q_0^j(v)} \rightarrow \infty$ for all $v \neq \omega^{*j}$, we would

also have $\frac{q^*(\omega^{*j})}{q^*(v)} \rightarrow \infty$. Thus, by continuity, an individual with a sufficiently informative original signal would induce society to converge to his beliefs.

4.3 Convergence to the truth in large networks

Do individuals reach the correct beliefs in the limit? In the first period of communication, the BPI heuristic is consistent with rationality (when information is initially independent). Thus, with the BPI, in a one-off interaction where each individual has independent information, information is aggregated in the sense that updating is rational. But in a repeated interaction, as individuals neglect the correlation in their information, this is not necessarily the case. Individuals' learning might actually unravel, as illustrated in Example 5 above.¹¹

Still, as in Golub and Jackson (2010), in large societies organized in sufficiently balanced networks, information would be fully aggregated in the limit. Suppose that the true state is $\bar{\omega}$ and individuals receive independent (and informative) signals. Golub and Jackson (2010) define a sequence of networks $(T_n)_{n=1}^\infty$ to be “wise” if beliefs converge (in probability) to be degenerate on $\bar{\omega}$. Let $q_{\infty,n}^i(\omega)$ be the limit beliefs of individual i for a network T_n as defined in (1), when $k \rightarrow \infty$. A network is wise then if $\text{plim}_{n \rightarrow \infty} \max_{i \leq n} |q_{\infty,n}^i(\bar{\omega}) - 1| = 0$.

Proposition 4 *Suppose that $(T_n)_{n=1}^\infty$ is a sequence of irreducible and aperiodic matrices such that $\alpha_n^i \rightarrow_{n \rightarrow \infty} \alpha_n^j$ for all i, j . Then the network is wise.*

To see the intuition, consider a complete network. Once society is large enough, then already after the first round of communication, beliefs would be fairly concentrated around $\bar{\omega}$, and repeated communication will only accentuate this as it leads the beliefs to the mode. For more general networks, if $q^*(\omega)$ mimics the correct rational beliefs, which is the case when $\alpha^i \rightarrow_{n \rightarrow \infty} \alpha^j$ for all i, j , then learning will arise. That is, correlation neglect that arises from repeated communication is mitigated by the rational learning in the initial stages.

5 Conclusion

In this paper we advocate the use of a “rational” heuristic to explore information diffusion in repeated exchanges in networks. The BPI heuristic is rational in the sense that it follows Bayesian updating for *all* information structures that satisfy conditional independence across individuals. We have shown that the limit beliefs converge to the mode of a simple extension of the BPI, which implies that beliefs become polarised in society, and that

¹¹This is related to the “unlearning” result of Gagnon-Bartsch and Rabin (2015).

consensus beliefs will still change. These results differ from those derived in previous literature using the DeGroot heuristic.

One possible avenue for further research is to consider the co-evolution of information transmission and the network itself. For example, as in the homophily literature (e.g., Golub and Jackson 2012), individuals may not communicate with those that have beliefs which differ substantially from their own. Thus while this fosters polarisation, it also makes belief in society more immune to those with strong beliefs. Individuals with strong beliefs are more likely to be different from others, and may not be as influential as in our model. This may affect the speed at which society polarises.

6 Appendix

Proof of Proposition 1: Note that the posterior of individual i at stage $k + 1$ is given by

$$\frac{\Pi_j(q_0^j(\omega))^{\mathbf{T}_{ij}^k}}{\sum_{\omega \in \Omega} \Pi_j(q_0^j)^{\mathbf{T}_{ij}^k}}$$

Where $\mathbf{T}^k = \underbrace{\mathbf{T} \times \mathbf{T} \times \dots \times \mathbf{T}}_{k \text{ times}}$.

If \mathbf{T} is strongly connected (irreducible) and aperiodic, then it is primitive and we can use the Perron–Frobenius Theorem for primitive matrices. This theorem states:

1. There is a positive real number $r > 1$, which is an eigenvalue of \mathbf{T} and any other eigenvalue λ satisfies $|\lambda| < r$.
2. There exists a right eigenvector v of \mathbf{T} with eigenvalue r such that all components of v are positive (respectively, there exists a positive left eigenvector w).
3. $\lim_{k \rightarrow \infty} \mathbf{T}^k / r^k = \mathbf{P}$ where $\mathbf{P} = vv^T$, where the left and right eigenvectors for \mathbf{T} are normalized so that $w^T v = 1$. Moreover, the convergence is exponential.

Therefore we have that $\frac{\mathbf{T}^k}{r^k}$ converging to a matrix of rank 1. Let $\boldsymbol{\alpha}^T = \mathbf{w}^T$.

We now show that this implies that all individuals' beliefs converge to a degenerate distribution on a subset of the modes of $\Pi_{j \in N}(q_0^j(\omega))^{\alpha^j}$. For simplicity assume that there is only one mode of $\Pi_{j \in N}(q_0^j(\omega))^{\alpha^j}$ and denote it by ω^* .

Consider some linear order on Ω and for any stage k let $F_k(\omega) = \frac{\sum_{\omega' \leq \omega} (\Pi_{j \in N}(q_0^j(\omega'))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}})^{v_i r^k}}{\sum_{\nu \in \Omega} (\Pi_{j \in N}(q_0^j(\nu))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}})^{v_i r^k}}$

be the cumulative probability function of individual i . v_i is the relevant coordinate for i in the eigenvector v .

Let N_k be the cardinality of the support of the beliefs at period k , which is bounded by $N_1 = \cap_j \text{Supp}(q_0^j(\omega))$ and is therefore finite for any k . If, as we assume, the beliefs are

full support than the cardinality is just that of Ω .

Claim 1: Let $\omega < \omega^*$, then $\lim_{k \rightarrow \infty} F_k(\omega) = 0$.

Proof: Let $\omega < \omega^*$ and let $\varepsilon = |\Pi_{j \in N}(q_0^j(\omega^*))^{\alpha^j} - \max_{\omega' \leq \omega} \Pi_{j \in N}(q_0^j(\omega'))^{\alpha^j}|$. Also, let \bar{k} be such that for all $k > \bar{k}$, $|\Pi_{j \in N}(q_0^j(\omega^*))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}} - \max_{\omega' \leq \omega} \Pi_{j \in N}(q_0^j(\omega'))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}}| > \frac{\varepsilon}{2}$.

$$\begin{aligned} \text{Remember that } F_k(\omega) &= \frac{\sum_{\omega' \leq \omega} (\Pi_{j \in N}((q_0^j(\omega'))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}})_{v_i r^k})}{\sum_{v \in \Omega} (\Pi_{j \in N}((q_0^j(v))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}})_{v_i r^k})} = \frac{\sum_{\omega' \leq \omega} \left(\frac{\Pi_{j \in N}((q_0^j(\omega'))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}})_{v_i r^k}}{\max_{v \leq \omega} \Pi_{j \in N}(q_0^j(v))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}}} \right)_{v_i r^k}}{\sum_{\omega \in \Omega} \left(\frac{\Pi_{j \in N}((q_0^j(\omega))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}})_{v_i r^k}}{\max_{v \leq \omega} \Pi_{j \in N}(q_0^j(v))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}}} \right)_{v_i r^k}} \\ &\leq \frac{N_1}{\sum_{\omega \in \Omega} \left(\frac{\Pi_{j \in N}((q_0^j(\omega))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}})_{v_i r^k}}{\max_{v \leq \omega} \Pi_{j \in N}(q_0^j(v))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}}} \right)_{v_i r^k}}, \text{ as in the nominator we have a sum of expressions which} \\ &\text{are all less than 1 or 1.} \end{aligned}$$

Therefore,

$$F_k(\omega) \leq \frac{N_1}{\sum_{\omega \in \Omega} \left(\frac{\Pi_{j \in N}((q_0^j(\omega))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}})_{v_i r^k}}{\max_{v \leq \omega} \Pi_{j \in N}(q_0^j(v))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}}} \right)_{v_i r^k}} < \frac{N_1}{((1+\zeta))^{v_i r^k}} \rightarrow \lim_{k \rightarrow \infty} 0, \text{ for } \zeta = \frac{\varepsilon/2}{\max_{v \leq \omega} \Pi_{j \in N}(q_0^j(v))^{\frac{\mathbf{T}_{ij}^k}{v_i r^k}}}. \square$$

Claim 2: Let $\omega > \omega^*$, then $\lim_{k \rightarrow \infty} F_k(\omega) = 1$.

Proof: The proof follows the proof of Claim 1 by focusing on $1 - F(\omega)$. $\square \blacksquare$

Proof of Proposition 4: Note that if we have a set of n independent posteriors, by Proposition 1, $\frac{\Pi_{j \in N}(q_0^j(\omega))}{\sum_{v \in \Omega} \Pi_{j \in N}(q_0^j(v))}$ will converge to the true distribution with $n \rightarrow \infty$. Recall that the true parameter is $\bar{\omega}$, and thus this distribution will have $\frac{\Pi_{j \in N}(q_0^j(\bar{\omega}))}{\sum_{\omega \in \Omega} \Pi_{j \in N}(q_0^j(\omega))} \rightarrow 1$. Note also that $\frac{\Pi_{j \in N}(q_0^j(\omega))^{\alpha^j}}{\sum_{\omega \in \Omega} \Pi_{j \in N}(q_0^j(\omega))^{\alpha^j}}$, when $\alpha^j = c$ for all j , would replicate exactly the same distribution for all n .

Note now that for any $\varepsilon > 0$, there exists n' , such that for all $n > n'$, $|\alpha_n^i - \alpha| < \varepsilon$ for any i . Note that as the true state is $\bar{\omega}$, the belief $\frac{\Pi_{j \in N}(q_0^j(\omega))^{\alpha^j}}{\sum_{\omega \in \Omega} \Pi_{j \in N}(q_0^j(\omega))^{\alpha^j}}$ will have one mode, on $\bar{\omega}$.

We also know that for any n , there exist k' , such that for any $k > k'$, then $\lim_{k \rightarrow \infty} F_k^n(\omega)$ is a degenerate distribution on the modes of $\frac{\Pi_{j \in N}(q_0^j(\omega))^{\alpha^j}}{\sum_{\omega \in \Omega} \Pi_{j \in N}(q_0^j(\omega))^{\alpha^j}}$.

Therefore, there exists an n' , such that for any $n > n'$, $\lim_{k \rightarrow \infty} F_k^n(\omega)$ converges to a degenerate distribution on the modes of $\frac{\prod_{j \in N} (q_0^j(\omega))^{\alpha_j}}{\sum_{\omega \in \Omega} \prod_{j \in N} (q_0^j(\omega))^{\alpha_j}}$. For this part we can essentially repeat the proof in Proposition 2 with $\omega^* = \bar{\omega}$. That is, for any $n > n'$, there exists $k > k'$, such that $|\prod_{j \in N} (q_0^j(\bar{\omega}))^{\frac{\alpha_j}{v_i r^k}} - \max_{\omega' \leq \omega} \prod_{j \in N} (q_0^j(\omega'))^{\frac{\alpha_j}{v_i r^k}}| > \frac{\varepsilon}{2}$, and the rest follows. ■

References

- [1] Acemoglu D, Bimpikis K, Ozdaglar A. (2014). Dynamics of information exchange in endogenous social networks. *Theor. Econ.* 9:41–97.
- [2] Bala, V. and S. Goyal (1998). Learning from neighbours. *Review of Economic Studies* 65, 595-621.
- [3] Choi, S. and D. Gale and S. Kariv (2012), Social Learning in Networks: A Quantal Response Equilibrium Analysis of Experimental Data, *Review of Economic Design* 16(2), 135-157.
- [4] DeGroot, M. (1974). Reaching a Consensus. *Journal of the American Statistical Association*, Volume 69, Issue 345. 118- 121.
- [5] De Marzo PM, Vayanos D, Zwiebel J. (2003). Persuasion bias, social influence and unidimensional opinions. *Q. J. Econ.* 118:909–68
- [6] Duffie, D. and G. Manso (2010), Information percolation, *American Economic Journal: Microeconomics*, Vol. 2, No. 1, pp. 100-111.
- [7] Enke, B. and F. Zimmerman. (2013). Correlation neglect in belief formation.
- [8] Eyster, E. and M. Rabin. (2014). Extensive imitation is irrational and harmful. *The Quarterly Journal of Economics* [Internet]. 2014;qju021.
- [9] Eyster, E. and M. Rabin. (2010). Naïve Herding in Rich-Information Settings. *American Economic Journal: Microeconomics*, 2(4): 221-43.
- [10] Eyster, E. and G. Weizsäcker. (2011). Correlation Neglect in Financial Decision-Making. *Discussion Papers of DIW Berlin* 1104.
- [11] Gagnon-Bartsch, T. and M. Rabin (2015), Naive Social Learning, Mislearning, and Unlearning, mimeo.
- [12] Glaeser, E. and Cass R. Sunstein (2009), Extremism and social learning. *Journal of Legal Analysis*, Volume 1, Number 1.

- [13] Golub B, Jackson MO. (2010). Naïve learning in social networks and the wisdom of crowds. *Am. Econ. J. Microecon.* 2(1):112–49.
- [14] Golub B, Jackson MO. (2012). How Homophily Affects the Speed of Learning and Best-Response Dynamics. *Quarterly Journal of Economics*, pp. 1287–1338.
- [15] Grimm, V. and F. Mengel (2014), Experiments on Belief Formation in Networks, mimeo.
- [16] Guarino, A. and P. Jehiel. (2013). Social Learning with Coarse Inference. *American Economic Journal: Microeconomics*, 5(1): 147-74.
- [17] Jackson, M. (2011). An Overview of Social Networks and Economic Applications, in the *The Handbook of Social Economics*, edited by J. Benhabib, A. Bisin, and M.O. Jackson, North Holland Press.
- [18] Jadbabaie, A., A., P. Molavi, and A. Tahbaz-Salehi (2015), Foundations of Non-Bayesian Social Learning, mimeo.
- [19] Kallir, I. and Sonsino, D. (2009). The Perception of Correlation in Investment Decisions. *Southern Economic Journal* 75 (4): 1045-66.
- [20] Levy, G. and R. Razin (2016). Combining Forecasts: Why Decision Makers Neglect Correlation. mimeo.
- [21] Levy, G. and R. Razin. (2015a). Correlation neglect, voting behaviour and information aggregation. *American Economic Review* 105: 1634-1645.
- [22] Levy, G. and R. Razin. (2015b). Does polarisation of opinions lead to polarisation of platforms? The case of correlation neglect, forthcoming, *Quarterly Journal of Political Science*, Volume 10, Issue 3.
- [23] Mueller–Frank, M. and C. Neri, Social learning in networks: Theory and experiments, mimeo.
- [24] Ortoleva, P. and E. Snowberg. (2015). Overconfidence in political economy. *American Economic Review* 105: 504-535.
- [25] Philippos L., O. Troumpounis, and N. Tsakas (2017), Communication and the emergence of a unidimensional world, mimeo.
- [26] Schkade , D., Sunstein C. and D. Kahneman, (2000), Deliberating about Dollars: The Severity Shift, (John M. Olin Program in Law and Economics Working Paper No. 95.

- [27] Sobel, J. (2014). On the relationship between individual and group decisions, *Theoretical Economics* 9, 163–185.
- [28] Stoner, J. A. F. (1968). Risky and cautious shifts in group decisions: the influence of widely held values. *Journal of Experimental Social Psychology*: 442-459.