Homophily and Influence

Arnold Polanski
University of East Anglia

Fernando Vega-Redondo
Bocconi University & IGIER

April 26, 2018

Abstract

We study how learning and influence co-evolve in a social network, eventually determining *both* the pattern of social influence across individuals and the distribution of their opinions or beliefs. Our model generalizes the classical model of DeGroot (1974) to multidimensional opinions and a finite learning span. This generalization allows us to rely on an intuitive notion of agent similarity – widely used in practice – to define a homophily-based condition that endogenizes in a stark manner the pattern of inter-agent influence that can prevail at equilibrium.

Our analysis of the model starts by establishing the existence of equilibrium, which is followed by its characterization in some simple contexts. Next, we show that, at equilibrium, the strength of the link between any two agents is given by its "support" – roughly, the amount of third-part influence impinging on both agents. This result leads to the key insight that disconnected groups may fail to integrate even if numerous links are created across them. It also leads to the identification of a set of sufficient conditions for which social segmentation remains a robust state of affairs – in particular, dynamically stable under a natural adjustment process.

Keywords: social learning; homophily; influence; echo chambers; integration.
JEL classif. codes: D83, D85.

1 Introduction

In this paper, we develop the idea that the network channeling social influence is shaped in conjunction with the learning process that unfolds on it – that is, we propose a model where the two dimensions, influence and learning, co-evolve. This novel perspective allows us to shed light on the important issue of how social learning (e.g. the formation of individual opinions in a social context) is affected by the relaxation of the common, but unrealistic, assumption that the pattern of social influence stays fixed. Indeed, one expects that, in the real world, "influence weights" are usually affected by the ongoing learning process of agents, and this may end up having a substantial and lasting effect on the final outcome.

In particular, it may exacerbate the tendency of an initially cohesive society to become disegregated into separate groups that hardly interact. This, in fact, has been highlighted as one of the distinct, and also worrying, features of modern societies. For, today, despite the unquestionable abundance of easily reachable information, individuals with dissenting views often tend to be scarcely exposed to each other. Paradoxically, this means that, de facto, the population may end up living in a relatively poor information environment. The detrimental social consequences of the "echo chambers" thus induced can hardly be exaggerated.[1]

To shed light on the problem, we propose a learning model that builds upon the classical setup proposed by DeGroot (1974) – studied in the economic literature by, among others, DeMarzo, Vayanos and Zwiebel (2003) and Golub and Jackson (2010, 2012). In the DeGroot framework, the opinion (or belief[2]) held by an agent adjusts over time by combining linearly her immediately preceding opinions and those of others. The vectors of weights specifying how each agent impinges on the learning of any other agent in the population define what is called the influence network/matrix, which fully governs the overall social-learning process. Under the twin assumptions that the number of learning rounds is *unbounded* and the influence network is connected, it can be shown that mild regularity conditions imply that the population converges to consensus, i.e. all agents eventually hold the same opinion.

The situation, however, is interestingly different if, as in our case, learning involves *finite* learning spells. Under these circumstances, the pattern of social influence must play a key role in modulating how much convergence of opinions is attained through social learning. That is, in general, the network positions of different agents will influence the extent to which their initial differences in opinions may persist (or perhaps even widen) through social interaction. Naturally, one of the reasons for postulating a finite learning process is that it adds a welcome dose of realism: in the real world, where new issues often come and go quite fast, the time span during which social attention is focused on any one of them is relatively short. Our main motivation here, however, is of a modeling nature. For, as we explain next, finite learning spells play a crucial role in generating the heterogeneity of agent behavior on which we build our network-formation theory, i.e. the theory that endogenously determines the prevailing pattern of social influence.

[1] A good illustration of the problem is documented by Adamic and Glance (2005), who study the deep divide between conservative and liberal blogs in the period preceding the U.S. Presidential Election of 2004. With a different focus, Bishop (2008) discusses the geographical basis for this phenomenon and its negative effect on social cohesion. For a more general approach to the problem, we refer to Boutyline and Willer (2017), who show that the interaction bias is more prevalent the more extreme are the political views held by individuals. Finally, another interesting illustration is provided by Golub and Jackson (2012), who describe how political prior alignment segmented the information (and hence ended up segmenting the opinion as well) of the American public when the question arose in 2003 as to whether or not Iraq had weapons of mass destruction.

[2] In principle, we could apply this rule to any continuous behavioral trait with a compact range.

In our model, the key postulate we invoke for endogenizing influence is that of *homophily* – that is, the tendency of individuals to associate (or link) with others who share a similar trait, behavior, or (as in our context) opinions and beliefs. The importance of homophily in determining how individuals construct their social network has been widely documented in the sociological literature. For a good early account of its pervasiveness in social environments, the reader is referred to the seminal work by Lazarsfeld and Merton (1954), while a modern survey of recent evidence can be found in the review by McPearson, Smith-Lovin and Cook (2001).[3] Concerning the specific problem that concerns us here, namely, how homophily impinges on opinion and belief formation, Golman, Loewenstein, Moene, and Zarri (2016) carry out an insightful discussion of modern literature, covering a rich range of empirical evidence on what they call the innate human preference for "belief consonance." As they explain, in many cases this urge for consonance leads to belief clustering, i.e. the choice "to associate with – that is, become friends with, work with, and even have romantic relationships with – others who share their beliefs" (see op. cit., p. 177).

In line with the previous discussion, it is worth emphasizing that, in the present paper, homophily is taken to be a *descriptive postulate*, not a normative one. For, in general, if agents were to choose optimally with whom to connect, it is plausible that, in many cases (e.g. when information gathering is the main objective), they would prefer to do so with others who are different from them (hence likely to hold complementary information). For a good illustration of the potentially negative effects of homophily on social learning, the reader is referred to Golub and Jackson (2012) and Lobel and Sadler (2013). In different contexts (DeGroot learning framework in the first case, Bayesian in the second) they show that homophily can be detrimental to welfare by slowing down, or blocking altogether, the ability of a population to learn some relevant state of the world through interactive learning.

In a nutshell, our homophily-based approach to endogenizing the network of social influence involves tailoring the weight of each link to the corresponding similarity of opinions/beliefs of the connected agents. In this respect, an important property that will be assumed on the underlying environment is that it is rich enough to span opinions in multiple dimensions. Thus, for example, individuals may hold opinions on a number of different topics: economic, political, religious, etc. Or, even if they restrict to just one such category, say the economic one, their concern covers a wide range of different issues such as growth, unemployment, inflation, or income distribution. Such richness of the "topic space," however, cannot be accommodated by the traditional DeGroot model, which is inherently one-dimensional.[4] This leads us to extending the theoretical framework proposed by DeG-

[3]For a fuller perspective on the phenomenon, see also Cohen (1977), Kandel (1978) Marsden (1987, 1988), Alexander et al. (2001), Moody (2001), and Knecht et al. (2010).

[4]Even though multidimensional versions of the model have been formulated – see e.g. DeMarzo, Vayanos and Zwiebel (2003) – they deal with each dimension independently and thus lack one of the features that is important in the real world as well as our model, namely, that correlations across the different dimensions (in part, brought about by social interaction itself) carry relevant information and have interesting implications.

root, endowing it with a mathematical structure (that of random variables) that not only accounts for the required multidimensionality but provides as well a natural way to measure opinion similarity by focusing on correlations.

Formally, our approach is grounded on the equilibrium notion we label *Equilibrium Influence Matrix* (EIM). Given some arbitrary attention/observation network (a binary network specifying exogenous constraints) the idea underlying the EIM concept is simple:

At equilibrium, each individual should "apportion the influence" among the agents she observes in proportion to their relative similarity (i.e. correlation) of opinions.

With this equilibrium concept in hand, our main objective is to explore how the patterns of influence induced by an EIM depend on the following three features of the environment:

(a) the structure of the initial signals received by the agents (which define their initial opinions);

(b) the depth of the learning process (how many rounds of opinion adjustment are conducted within each learning spell);

(c) the attention network (since, naturally, an agent can exert influence only on those who pay attention to her).

Our analysis starts by establishing a *basic existence result*: an EIM exists for any set of initial signals, learning depth, and attention network. This equilibrium matrix is then characterized for some useful scenarios, e.g. the *benchmark cases* where the number of learning rounds is unbounded or the attention network is complete. Our main result, however, applies generally to any attention network and learning depth. It involves the identification of a topological measure for the influence network that characterizes the equilibrium influence strengths prevailing at any equilibrium. Specifically, the requirement is that the equilibrium influence that any agent i exerts on another one j must be proportional to the accumulated normalized influence that all other agents exert on *both* i and j. We call this magnitude the *support* of the relationship between i and j.[5]

The important insight to be gathered from this result is that segmentation of the population into groups may be hard to overcome by the mere establishment of "bridging links" across them. Indeed, we show that even if many such links exist, they will fail to integrate the population unless they are properly "supported." In other words, the conclusion is that, at equilibrium, only bridging links that enjoy high support may channel significant "influence" across groups and hence promote integration. This stands in interesting contrast with

[5]This notion is conceptually related to what Easley and Kleinberg (2010) call "neighborhood overlap" and Jackson, Rodríguez-Barraquer and Tan (2012) label"link support." One important difference, however, is that our measure of support reflects the accumulated third-party influence – a real variable – not just the number of agents exerting that influence.

the celebrated dictum in the social network literature put forward by Granovetter (1983) – "the strength of weak links," which applies to information rather than influence. In our context, weak links are *not* strong, in the sense of being effective channels of novel information.[6] They are strong because they have substantial support (or what, in an informational vein, could be understood as some measure of redundancy).

Finally, to understand better the issue of social segmentation, the last part of the paper focuses on a particularly stark version of the problem and studies it dynamically. We start the analysis by considering two groups that are disconnected *across* (i.e. pay no attention to each other) and are fully connected *within* (hence provide the best local support structure). Then we suppose that some links are added to the underlying attention network with the aim of *bridging influence* between the two groups. The question we pose is under what conditions is segmentation a robust phenomenon. Or, more formally, when is it a locally stable equilibrium for the implicit dynamics underlying the EIM concept. Our answer to this question is in line with, and complements, the aforementioned analysis. Besides confirming that segmentation *dynamically* persists if the bridging links are few and weakly supported, it also shows that a similar conclusion obtains if the learning spells are short, the groups are large, and/or asymmetric in size. This enriches our understanding of the issue of social segmentation, by highlighting additional factors that play a relevant role and casting the problem in a dynamic manner.

The rest of the paper is divided into three more sections. Section 2 presents the basic framework, Section 3 carries out the analysis, Section 4 concludes with a summary. For the sake of a smooth presentation, all formal proofs are gathered in the Appendix.

2 The model

2.1 The pattern of communication and influence

Consider a given population of agents, $N = \{1, 2, ..., n\}$, who are connected as specified by an *exogenous attention network* on the set N. The links are binary (i.e. of the same type and intensity) and directional (connect one agent to another in a particular direction). Formally, it is convenient to represent this network through an adjacency matrix $L \equiv (l_{ij})_{i,j \in N}$ with $l_{ij} = 1$ if player i is connected to agent $j \neq i$ and $l_{ij} = 0$ otherwise. When such a connection exists, the interpretation is that agent i pays attention to, and hence *can* be influenced by (or learn from) agent j.

For the case $i = j$, we shall posit that $l_{ii} = \eta \geq 0$, i.e. every agent pays attention to herself as parametrized by η. Intuitively, in a dynamic setup, this parameter can be

[6] A similar point has been made by Uzzi (1997), Reagans and McEvily (2003), or Aral and Van Alstyne (2011), when they stress the importance of what is sometimes labeled "bandwidth" of links.

viewed as capturing the extent to which agents display persistence (or "stubbornness) in their opinions. Large values for η reflect high persistence, meaning that the revised opinions of each individual are largely based on their own preceding opinions. In contrast, $\eta = 0$ corresponds to the extreme polar case where the revision of every agent is based only on the input received from the other agents, while she attributes zero weight to her own previous opinions. For most of the analysis in this paper the parameter η plays no significant role. Thus, for notational simplicity, we shall usually focus on the case $\eta = 1$, in which each agent pays as much attention to herself as to any other. The case $\eta = 0$ will be used occasionally for expositional purposes, when it is useful to explain some of the main ideas effectively.

The attention network introduced above represents the underlying structure that is assumed to channel/restrict any bilateral interaction across agent pairs. It can be interpreted as reflecting the *exogenous* features of the situation (say, geographical, ethnic, linguistic, or hierarchical) that constrain whether or not an agent can directly influence another agent. Inter-agent influence, therefore, must respect the attention network, in the sense that it can flow from agent i to j only if the latter pays attention to the former (i.e. only if $l_{ji} = 1$). In fact, the pattern of effective inter-agent influence can itself also be formalized as a network – what we shall call the *influence network* – and represented through a corresponding matrix $A \equiv (a_{ij})_{i,j \in N}$. However, in contrast to the attention network, the influence network is *weighted* and *endogenous*. We explain each of these two important characteristics in turn.

(i) A typical entry a_{ij} of the matrix A is interpreted as the (relative) intensity with which agent j influences i – i.e. the weight that i attributes to j's opinions in shaping her own. For convenience, every such measure of bilateral influence is taken to lie in the interval $[0, 1]$ and every row vector of individual influences, $\mathbf{a}_i \equiv (a_{ij})_{j \in N}$, is assumed to add up to unity. Thus, A is a *row-stochastic matrix*.

(ii) In contrast with the attention matrix, the key feature that characterizes the influence matrix is that it is *endogenous*. Specifically, the pattern of inter-agent influence is endogenized through an equilibrium concept that requires "homophily-based consistency" between the influence weights and the behavioral correlations induced by the social adjustment process.

To formulate precisely the equilibrium concept mentioned above, the following core components of the model need to be formally introduced: (a) the behavioral adjustment dynamics; (b) the notion of homophily; (c) the required consistency between the former two. To do so is the objective of the following three subsections.

2.2 Learning dynamics

The starting point of our model is the learning framework proposed by DeGroot (1974), which can be summarized as follows. Time proceeds discretely, $t = 0, 1, 2, ...$, and at every

t each agent $i \in N$ holds some opinion $x_i(t)$, identified with a point in a common and pre-specified compact interval, say $[0, 1]$. Thus, given some initial signals, $\beta_i \in [0, 1]$ for every $i \in N$, the opinions of individuals evolve over time as follows:

$$x_i(t) = \sum_{j=1}^{n} a_{ij} x_j(t-1) \qquad (i = 1, 2, ..., n;\ t = 1, 2, ...)$$

with $x_i(0) = \beta_i$ for every agent i. Or, in compact matrix form, we may write

$$\mathbf{x}(t) = A\mathbf{x}(t-1) \qquad (t = 1, 2, ...) \tag{1}$$

where A is the prevailing (fixed) influence matrix specified as above, and $\mathbf{x}(t) = (x_1(t), ..., x_n(t))'$ is the column vector of opinions at t.[7]

In this paper we extend DeGroot's framework in two complementary directions: (a) the learning span is given by some K, a parameter of the model (typically finite); (b) opinions are multidimensional. As explained, the motivation for (a)-(b) is two-fold. First, these two features allow for a more realistic description of social learning in real-world environments that are both fast-changing and complex. Second, they yield a theoretical setup that is tractable and rich enough to address our primary modeling objective, namely, the formulation of a homophily-based framework where the pattern of influence can be suitably endogenized.

To model multidimensional opinions, suppose that there is a (finite and common) set of relevant issues/dimensions, $\Omega = \{\omega_1, \omega_2, ..., \omega_D\}$, over which every agent i holds a particular opinion $x_i(\omega_d) \in [0, 1]$ for each $\omega_d \in \Omega$. For example, $x_i(\omega_d)$ could be interpreted as the fraction of time or budget that, according to i, should be devoted to issue ω_d. In general, not all issues are taken to be equally prominent, or frequent, in social discourse or in agents' minds. Rather, there are some specific weights associated to each of them, as given by a discrete measure on Ω. For simplicity, we assume that this measure is common across individuals and it is normalized so that it can be formally treated as a probability. We denote it by $\mathbb{P}$, with $\mathbb{P}(\omega_d)$ reflecting the prominence of the issue $\omega_d \in \Omega$.

Thus, mathematically, a (multidimensional) opinion displayed by any given agent i can be conceived as a random variable[8] $\widetilde{x}_i : \Omega \to \mathbb{R}$, where the set of issues Ω plays the role of the underlying sample space on which the random variable is defined and $\mathbb{P}(\cdot)$ that of the corresponding probability density. A useful feature of this formalism is that it allows one to rely on the statistical notion of correlation (here we shall rely on the standard Pearson correlation coefficient) to assess the opinion similarity of any two agents. Indeed, this is

[7] As usual, the notation $(\cdot)'$ applied to vectors and matrices stands for the corresponding transpose.

[8] Throughout we shall use the tilde to denote constructs, such as agents' opinions, that display the formal structure of a random variable even though, as discussed, in the present case its interpretation is not stochastic.

the typical approach implemented by the recommender systems that are widely used with considerable success on the Internet (e.g. by Amazon.com, Booking.com, or Netflix.com).[9]

To fix ideas, consider the simple example where there are only three agents (1, 2, and 3) and also just three issues that are of agents' concern – say, climate change (ω_c), education (ω_e), and Internet quality (ω_i), each with the same prominence: $\mathbb{P}(\omega_c) = \mathbb{P}(\omega_e) = \mathbb{P}(\omega_i) = 1/3$. Let us suppose that there are just three "opinion levels" that individuals attach to each issue, which could be interpreted as the levels of government funding that are available for each issue. Denote these levels by ℓ (low), m (medium), and h (high) – all three numbers in the interval $[0, 1]$ – and assume for simplicity that they satisfy $m = (\ell + h)/2$.

Next suppose that the opinions $\widetilde{x}_k$ displayed by each individual $k = 1, 2, .3$ are as follows:

$$\begin{aligned}&\widetilde{x}_1(\omega_c) = h; \quad \widetilde{x}_1(\omega_e) = \ell; \quad \widetilde{x}_1(\omega_i) = m\\&\widetilde{x}_2(\omega_c) = m; \quad \widetilde{x}_2(\omega_e) = h; \quad \widetilde{x}_2(\omega_i) = \ell\\&\widetilde{x}_3(\omega_c) = m; \quad \widetilde{x}_3(\omega_e) = \ell; \quad \widetilde{x}_3(\omega_i) = h.\end{aligned}$$

Then, it is straightforward to compute that the Pearson correlation coefficient among the opinions of the three agents (cf. (4) below) are:

$$corr(\widetilde{x}_1, \widetilde{x}_2) = -\frac{1}{2}; \quad corr(\widetilde{x}_1, \widetilde{x}_3) = \frac{1}{2}; \quad corr(\widetilde{x}_2, \widetilde{x}_2) = -1. \tag{2}$$

Thus, even though the opinions of agents 2 and 3 are perfectly anti-correlated, individual 1's opinion is mildly correlated in opposite directions with the opinion of those two agents. This provides a simple illustration of the diversity of situations that our similarity measure can deliver when opinions are multidimensional.

In the extended DeGroot framework, the social-learning dynamics is defined as a direct counterpart of that posited for the classical setup. That is, we simply replace (1) by the following generalization:

$$\tilde{\mathbf{x}}(t) = A\tilde{\mathbf{x}}(t-1) = A^t\tilde{\mathbf{x}}(0), \qquad (t = 1, 2, ..., K). \tag{3}$$

The above formulation simply considers multidimensional opinions (thus, formally, random variables) and the learning spell lasts K periods (typically finite). Of course, just as for the standard framework, the learning process defined by (3) requires the specification of some initial conditions. These are given by a set of initial signals $\tilde{\boldsymbol{\beta}} \equiv (\tilde{\beta}_i)_{i \in N}$. For simplicity, these signals are taken to display the same format as the opinions themselves and hence directly define initial opinions: $\tilde{\mathbf{x}}(0) = \tilde{\boldsymbol{\beta}}$. In principle, we shall allow such initial opinions to be correlated, possibly reflecting some overlapping sources of information across agents.

[9] See, for example, Jannach *et al.* (2010) and Ricci *et al.* (2011), who discuss the so-called collaborative-filtering systems. As in our case, the objective of these recommender systems is to identify individuals with similar characteristics, as identified by the correlation of their previously observed decisions across a number of relevant situations. For additional references on this topic, the reader may refer to Bobadilla *et al.* (2012), Rajaraman *et al.* (2012), or Resnick *et al.* (1994).

Thus the vector of random variables $\tilde{\boldsymbol{\beta}}$ may display an arbitrary covariance matrix, which we shall denote by $\Sigma_{\tilde{\beta}}$. Note that for the degenerate case where $\Sigma_{\tilde{\beta}} = \mathbf{0}$ (i.e. a matrix with all zero entries) we are led, if $K \to \infty$, to a context that is equivalent to the one posited by DeGroot. For, in this case, opinions are constant over all dimensions (hence, essentially, one-dimensional) and the learning spells are unbounded, just as in the classical setup.

2.3 Homophily and equilibrium

Homophily is defined as the tendency of individuals to interact preferentially with those who display similar characteristics. In our case, as explained above, these characteristics are agents' multidimensional opinions, formalized as random variables over a common state space given by the set of relevant issues. As advanced, an advantage of this formalization is that it provides a natural way to measure the similarity between the opinions held by any given pair of individuals i and j, $\widetilde{x}_i$ and $\widetilde{x}_j$, through their corresponding Pearson correlation coefficient:

$$\rho_{ij}(A) \equiv corr(\widetilde{x}_i(A), \widetilde{x}_j(A)) = \frac{c_{ij}}{\sqrt{c_{ii} c_{jj}}}, \quad \text{where} \quad (c_{ij})_{i,j \in N} = A^K \Sigma_{\tilde{\beta}} (A^K)'. \tag{4}$$

Hence, heuristically, we declare two agents as similar if they *adjust* their behavior along convergent directions. As explained in the previous subsection, this same criterion can be, and indeed is, widely used nowadays by many firms and various organizations. Given the extensive information available on the previous choices of individuals across Internet platforms, such information is found to be very useful to target both products and/or messages.[10]

Having defined our notion of inter-agent similarity, we turn now to introducing our homophily-based notion of equilibrium. Conceptually, it is quite simple: an influence matrix defines an equilibrium if the pattern of opinion similarity it induces reflects its own pattern of influence – that is each agent attributes to each of the agents to whom she pays attention an influence that is proportional to the extent at which their opinions are correlated at the end of a learning spell. To proceed formally, it is convenient to make the simplifying assumption (which shall hold throughout) that the covariance matrix $\Sigma_{\tilde{\beta}}$ of the initial signals is finite and has only non-negative entries.[11] Then, all opinion correlations are non-negative

[10] Clearly, there are other reasonable approaches one can use to measure the similarity of random variables. For example, a simple option would be to gauge the similarity of two agents i and j by the distance between their mean opinions, $|E[\widetilde{x}_i] - E[\widetilde{x}_k]|$. If we think of Ω as containing mostly political issues, the "expected" or "average" opinion of an individual could be identified with her *ideology* and then small values of the aforementioned distance would capture ideological proximity. One of the important advantages of our correlation-based measure is its invariance with respect to linear affine transformations. Thus, in particular, it can be measured independently of the separate scales used to define individual opinions.

[11] If some signals were negatively correlated, Definition 5 could be adapted by replacing the correlations $\rho_{ij}(\cdot)$ by the non-negative truncation $\rho^{+}_{ij}(\cdot) \equiv \max\{\rho_{ij}(\cdot)\,, 0\}$. Nothing essential in our analysis would be affected by this adaptation.

and the equilibrium concept that endogenizes the influence matrix can be suitably defined as follows.

Definition 1 *Fix the learning time span K as well as the vectorial random variable $\tilde{\beta} \equiv \tilde{\mathbf{x}}(0)$ determining the initial signals. Let $\rho_{ij}(A)$ stand for the Pearson's correlation coefficient between the post-communication opinions $\widetilde{x}_i(A)$ and $\widetilde{x}_j(A)$ induced by social learning under an influence matrix A. Then, matrix $A^* \equiv \left(a^*_{ij}\right)_{i,j\in N}$ is called an* equilibrium influence matrix *(EIM) if*

$$a^*_{ij} \propto l_{ij}\rho_{ij}(A^*) \qquad (i, j = 1, 2, ..., n). \tag{5}$$

The matrix A^* will be also called homophily-consistent.

Note that, in view of the twin restriction that

(a) an agent i can be influenced only by those agents j whom she pays attention to (i.e. whose $l_{ij} \neq 0$);

(b) only relative influence matters and hence the total influence exerted on every agent i is normalized to unity (i.e. $\sum_{j\in N} a_{ij} = 1$),

the equilibrium condition (5) simply embodies the aforementioned requirement of proportionality between similarity and influence.

Intuitively, one can interpret the homophily-consistency condition (5) as a fixed (or stationary) point of an adjustment process in which, after every learning spell, agents reconsider their influence weights so that they match (i.e. are proportional to) the observed correlations of opinions. By way of a simple illustration, one could think of Twitter users, who reconsider weekly the time and intensity they are going to devote to follow each of their contacts. During the working days, they are too busy to do that, so they simply read the tweets of their chosen contacts according to the plan. Then, on the week-end, they assess the relative (dis)agreement they have experienced with each of the individuals they follow and redraw the plan for the ensuing week accordingly.

We end this section with an example and a remark that bear on interesting aspects of the model and should help understand it better. The former contains a simple two-agents, two-issues example; the latter makes a methodological point concerning the role of opinion multidimensionality in the analysis.

Example 1 (Two agents, two issues) *Building on the Twitter analogy outlined above, let us consider the simple framework where there are two agents and two issues that are equally prominent in the social discourse, say, global warming and political polarization. At the beginning of each week, both agents acquire (e.g., from the press or from the news channels) some new information on these issues and assign a value to each of them in light of this new information.*

Each value reflects the urgency of the corresponding issue in agent's mind. Suppose that the respective values for agent 1 and agent 2 are,

$$(\beta_1(\omega_1), \beta_1(\omega_2)) = (0.6, 0.3), \quad (\beta_2(\omega_1), \beta_2(\omega_2)) = (0.51, 0.49).$$

Hence, based on the current week's news, agent 1 considers global warming to be "twice as urgent" as political polarization, while agent 2 assigns almost equal importance to both issues. Combining this with the equal relevance of both issues, the latter vectors define the agents' initial opinions $\tilde{\mathbf{x}}(0) = \tilde{\beta} = (\tilde{\beta}_1, \tilde{\beta}_2)$. Viewing them as random variables, we compute the corresponding covariance matrix and cross-agent correlation:

$$\Sigma_{\tilde{\beta}} = \begin{pmatrix} .0001 & .001 \\ .001 & .0225 \end{pmatrix}, \quad \rho_{12}(\tilde{\beta}) = \frac{0.001}{\sqrt{0.0001 \cdot 0.0225}} = 2/3.$$

In particular, note that the covariance between $\tilde{\beta}_1$ and $\tilde{\beta}_2$ is calculated as the issue-weighted sum of the cross-agent products of the deviations from the respective agent averages,

$$Cov(\tilde{\beta}_1, \tilde{\beta}_2) = \textstyle\sum_{d=1}^{2} \mathbb{P}(\omega_d) \left\{ \beta_1(\omega_d) - E[\tilde{\beta}_1] \right\} \left\{ \beta_2(\omega_d) - E[\tilde{\beta}_2] \right\} =$$
$$0.5\,\{0.6 - 0.45\}\,\{0.51 - 0.5\} + 0.5\,\{0.3 - 0.45\}\,\{0.49 - 0.5\} = 0.001.$$

Thus, in the example, the process starts with a correlation of 2/3 between agents' opinions. Then, during the rest of the week, the agents exchange tweets while they keep adjusting their opinions upon reading others' tweets. Given the uniform adjacency matrix L, $l_{ij} = 1$ for $i, j = 1, 2$, assume, for simplicity, that there is only one round of communication within the week, and that the matrix A applied in this round has each agent assigning the weight 2/3 to her own opinion and 1/3 to that of the other. This leads to a pair of opinions $\tilde{\mathbf{x}}(1) = (\tilde{x}_1(1), \tilde{x}_2(1))$ at the end of the week that can be computed by (3):

$$\begin{pmatrix} x_1(\omega_1) & x_1(\omega_2) \\ x_2(\omega_1) & x_2(\omega_2) \end{pmatrix} = \begin{pmatrix} 0.57 & 0.363 \\ 0.54 & 0.427 \end{pmatrix} = \begin{pmatrix} 2/3 & 1/3 \\ 1/3 & 2/3 \end{pmatrix} \begin{pmatrix} .6 & .3 \\ .51 & .49 \end{pmatrix}.$$

Then, we can compute the corresponding covariance matrix and correlation coefficient as follows:

$$\Sigma_{\tilde{\mathbf{x}}(1)} = \begin{pmatrix} .01071 & .00585 \\ .00585 & .00319 \end{pmatrix}, \quad \rho_{12}(\tilde{\mathbf{x}}(1)) = 1.$$

If we assume that the influence matrix to be used the following week is adjusted so that the revised influence is proportional to correlation, the new influence matrix A' can be readily found to be:

$$A' = \begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix}$$

It is also easy to verify that A' is an EIM in this example. Actually, we show in Proposition 3 in the next section that A' is the only connected EIM. Finally, let us advance that the dynamic adjustment process of the influence matrix illustrated in this example will play an important role in Section 3.4 when we study the phenomena of bridging and segmentation from a dynamic perspective.

Remark 1 (Opinion multidimensionality) *In our model, opinions are defined simultaneously on different issues (or dimensions). The question we address here is how much of such a multidimensionality is necessary to render our approach meaningful. To make our point in the starkest manner, let us suppose that, when regarded as random variables, the agents' initial opinions are independent*[12] *and their variances are identical, i.e.* $\Sigma_{\tilde{\beta}} = \sigma \mathbf{I}$, *where* $\mathbf{I}$ *is the identity matrix and* $\sigma > 0$. *The common variance* σ *can be interpreted as a measure of how much the alternative issues can be viewed as genuinely different in agents' minds – if they were not, we would have* $\sigma = 0$ *and all the opinion dimensions would be redundant. In this simplified context, we may extend previous notation and rewrite the homophily-consistency condition (5) as follows:*

$$a^*_{ij} = \frac{l_{ij}\rho_{ij}(A^*, \sigma)}{\sum_{k \in N} l_{ik}\rho_{ik}(A^*, \sigma)} \qquad (i, j = 1, 2, ..., n), \tag{6}$$

where $\rho_{ij}(A; \sigma)$ *stands for the correlation between agents* i *and* j *when the influence matrix is* A *and the opinion variance is* σ. *From (6), it can be easily verified that,* given any attention matrix L, *the corresponding correlation coefficients* $\rho_{ij}(A, \sigma)$ *do* not *depend on the magnitude of* $\sigma > 0$. *It follows, therefore, that the same EIM* A^* *satisfies (6) for any fixed value of* σ, *and hence also in the limit* $\sigma \searrow 0$. *This allows one to view* A^* *as an EIM for* any *opinion multidimensionality,* no matter how small. [13]

3 Analysis

In this section, we carry out the analysis of the model. First, Subsection 3.1 starts our analysis of endogenous EIMs by addressing the basic issue of equilibrium existence. Subsection 3.2 characterizes EIMs in some benchmark cases. Subsection 3.3 discusses the relationship between homophily-consistency (i.e. equilibrium) and a suitable notion of link support or overlap. Finally, Subsection 3.4 relies on the insights obtained in Subsection 3.3 to study the relationship between communication/attention bridges and integration in some illustrative cases.

3.1 Equilibrium existence

Our primary interest in this paper is to study the characteristics of endogenous patterns of influence as captured by the equilibrium notion in Definition 1. Thus a first basic question that must be addressed is the following: Does there always exist an equilibrium influence matrix that satisfies the required consistency between the learning outcome and the influence weights? A positive answer to this question is provided by the next result.

[12] Such independence is not essential for the argument and it is done just for the sake of formal simplicity.

[13] Methodologically, the approach is reminiscent of that pursued in other strands of literature (e.g. evolutionary game theory – cf. Kandori *et al.* (1993) or Young (1993)), where a similar approach has been used as a powerful selection device.

Proposition 1 *Consider an arbitrary adjacency matrix $L \equiv (l_{ij})_{i,j \in N}$ with $l_{ij} \in \{0, 1\}$, any learning depth $K \geq 1$, and any vectorial random variable $\tilde{\beta}$ governing agents' initial opinions. An influence matrix A^* that satisfies (5) always exists.*

Proof. See Appendix.

The proof of Proposition 1 follows from a standard fixed-point argument applied to the vector field $F : (\Delta^{n-1})^n \rightarrow (\Delta^{n-1})^n$ given by:[14]

$$F_{ij}(A; K, L, \Sigma_{\tilde{\beta}}) \equiv F_{ij}(A) \equiv \frac{l_{ij}\, \rho_{ij}(A)}{\sum_{k \in N} l_{ik}\, \rho_{ik}(A)} \qquad (i, j = 1, 2, ..., n), \tag{7}$$

which maps row-stochastic matrices A of size $n \times n$ satisfying

$$\forall i, j \in N, \quad l_{ij} = 0 \Rightarrow a_{ij} = 0 \tag{8}$$

into other such matrices $\tilde{A} = F(A)$, also satisfying (8). In Subsection 3.4, we apply $F(\cdot)$ as an homophily-based adjustment rule that, operating across learning spells revises the influence weights between attention-connected agents so as to match their corresponding opinion correlation.

Figures 1 and 2 illustrate the EIMs for a randomly generated attention network with $n = 25$ nodes. They are obtained by iterating the vector field F defined by (7) from a situation where every agent assigns the same influence weight to each of her neighbors in the attention network.

In Figure 1, we consider the case where the learning span is the shortest possible, $K = 1$. We observe that the resulting pattern of equilibrium influence displays a marked structure with small subsets of agents maintaining high-influence links among themselves while the links with all other agents carry much lower influence. This implies that the end opinions in the population will tend to be quite heterogeneous, with agents fragmented into opinion clusters that display only light cross-interactions.

In contrast, Figure 2 displays an EIM that is spanned on the same attention network as in Figure 1 but with communication proceeding for $K = 100$ rounds. In this alternative case, the learning process eventually leads to a situation where agents hold the same opinions. Such a perfect convergence among individual opinions induces in turn an influence matrix where every agent attributes the same influence weight to each of her neighbors (and herself). The induced EIM is therefore very different from the one resulting for $K = 1$ as it is determined by the exogenous attention structure alone.

[14] As usual, $\Delta^{n-1} \subset \mathbb{R}^n_+$ denotes the $(n-1)-$dimensional simplex.

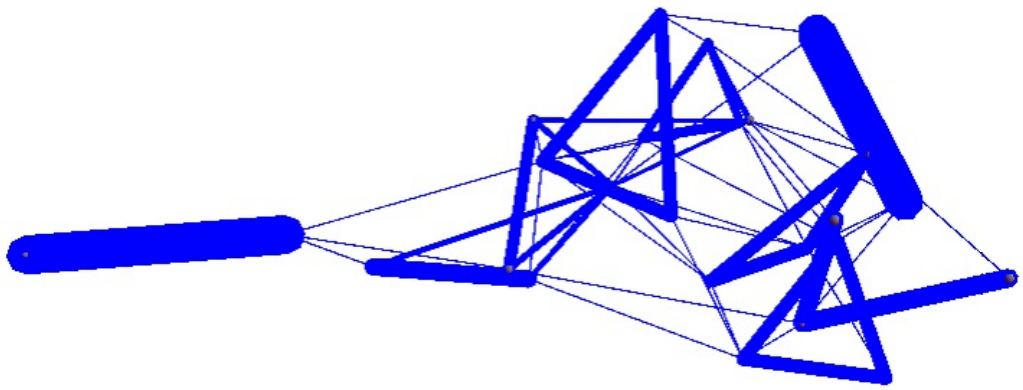

Figure 1: Graphical representation of an EIM $A^* \equiv (a^*_{ij})_{i,j\in N}$ for an underlying attention network with $n = 25$ nodes that is generated as a realization of an Erdös-Rényi random network. Initial signals are determined by identical random variables (with unit variances) that are uncorrelated across nodes. The learning depth is set to $K = 1$. The EIM is obtained through the iteration of (7) from initial conditions in which every agent attributes a uniform weight to all her connections in the attention network. The thickness of each edge ij is proportional to $\frac{a^*_{ij}+a^*_{ji}}{2}$. Self links are not shown.

3.2 Some benchmark cases

In this subsection we consider two benchmark scenarios that are useful in different respects. First, we show that the outcome in Figure 2 can be generalized to any connected[15] attention matrix as long as the learning span K is unbounded. This is the content of the following result.

PROPOSITION 2 *Consider a connected attention network with adjacency matrix $L \equiv (l_{ij})_{i,j\in N}$, $l_{ij} \in \{0, 1\}$, and any vectorial random variable $\tilde{\beta}$ governing agents' initial opinions. Then, if the learning depth $K \to \infty$, the unique connected influence matrix $A^* \equiv (a^*_{ij})_{i,j\in N}$ is given by*

$$a^*_{ij} = \frac{l_{ij}}{\sum_{k\in N} l_{ik}} \qquad (i, j = 1, 2, ..., n),$$

and the induced bilateral correlations among end opinions satisfy $\rho_{ij}(A^*) = 1$ for all $i, j \in N$.

Proof. See Appendix.

Thus, when the learning depth is unbounded, the only connected influence matrix that meets the equilibrium requirement is the one that matches the attention network and assigns uniform weights across neighbors. This result makes the important point that a connected

[15] Naturally, for attention networks (whose directed links are discrete) the notion of connectedness simply requires that there exist a directed path between any two nodes/agents.

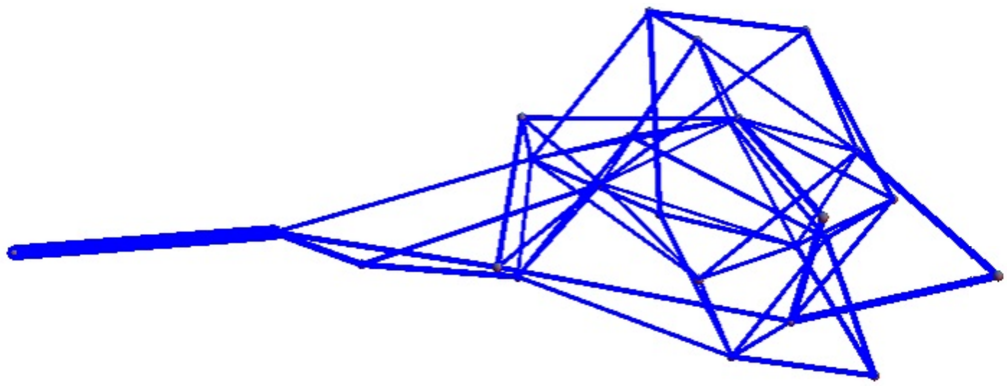

Figure 2: Graphical representation of an EIM $A^* \equiv (a^*_{ij})_{i,j\in N}$ under the same conditions as for Figure 1 except that the learning depth is $K = 100$.

influence pattern with non-trivial local structure can arise only if the learning process falls short of delivering social consensus – for, if such consensus obtains, the attention network fully imposes its structure on an EIM. As an unbounded K leads to fully correlated opinions, no feedback from the learning process will have any impact on the equilibrium pattern of influence. However, the final (consensus) opinions will still typically depend on the architecture of the communication (attention) network.[16] It follows from this discussion that in order to understand the rise of non-trivial influence structures it is essential to contemplate a finite time span in learning.

Another interesting benchmark case concerns situations where the population is segmented (endogenously) into separate influence components and, for each of these components, the underlying attention structure is complete. Admittedly, this is a stark setup but, as we shall explain, it is suggestive of some general forces that shape communication among groups. We shall find it particularly useful in studying "bridging and segmentation" in Subsection 3.4. Our analysis of the situation described above builds upon the following result.

PROPOSITION 3 *Let the attention network be* completely *connected (i.e. the adjacency matrix* $L \equiv (l_{ij})_{i,j\in N}$ *satisfies* $l_{ij} = 1$ *for all* $i, j \in N$*). Then, for any given learning span* $K \geq 1$*, the unique connected influence network has the adjacency matrix* $A = Q(n)$ *where,* $Q(n) \equiv (q_{ij}(n))_{i,j\in N}$ *stands for the row (and also column) stochastic square matrix with* $q_{ij}(n) = \frac{1}{n}$ *for all* $i, j = 1, 2, .., n$.

Proof. See Appendix.

[16] As shown by De Marzo *el al.* (2003) in the context of non-probabilistic opinions, the end common opinion $x^*(A)$ will be a convex combination of the initial opinions β_i of each agent i, $x^* = \sum_{i\in N} \omega_i \beta_i$, where each $\omega_i = \frac{|\{j: l_{ji}=1\}|}{\sum_{k\in N} |\{j: l_{jk}=1\}|}$, i.e. it is the unit-normalized in-degree of agent i in the attention network.

This result asserts that if the attention matrix is such that every agent is connected to every other agent (and signals are non-negatively correlated), then the only connected influence matrix that meets the requirement of homophily consistency is the matrix $Q(n)$ where every agent influences *directly* each of the other agents with the same weight. The fact that this matrix is homophily-consistent is quite clear: under $Q(n)$, every agent is in effect exposed to the same convex combination of the initial signals, which in turn supports the uniform-influence matrix as an equilibrium. On the other hand, the fact that this matrix is the unique EIM among those that define a connected influence network is less apparent. It follows from a cumulative effect, in which indirect influence rises correlation, which in turn increases direct influence (and thus correlation) up to the point where only a fully symmetric pattern of influence arises at equilibrium.

Building upon Proposition 3, the following two straightforward corollaries follow.

COROLLARY 1 *Let $\{N_s\}_{s=1}^{r}$ represent an r-element partition of the agent population N and denote $n_s \equiv |N_s|$ for each $s = 1, 2, ..., r$, with $n_1 + ... + n_r = n$. Assume a complete attention network for each N_s. Suppose further that $\Sigma_{\tilde{\beta}} \equiv (\sigma_{ij})_{i,j\in N}$ satisfies $\sigma_{k,\ell} = 0$ if $k \in N_s$ and $\ell \in N_{s'}$ with $s \neq s'$. Then, if $Q(n_r) \equiv (q_{ij}(n_r))_{i,j=1}^{n_r}$ stands for the uniform square matrix of dimension n_r with $q_{ij}(n_r) = \frac{1}{n_r}$ for all $i, j = 1, 2, .., n_r$ and $0_{n_r \times n_s}$ stands for a matrix of dimension $n_r \times n_s$ consisting of all zero, the block diagonal matrix*

$$A = \begin{pmatrix} Q(n_1) & 0_{n_1\times n_2} & \cdots & 0_{n_1\times n_r} \\ 0_{n_2\times n_1} & Q(n_2) & \cdots & 0_{n_2\times n_r} \\ \vdots & \vdots & \ddots & \vdots \\ 0_{n_r\times n_1} & 0_{n_r\times n_2} & \cdots & Q(n_r) \end{pmatrix}$$

is an EIM, upon suitable index labeling.

COROLLARY 2 *Under the conditions contemplated in Proposition 3, if the prevailing EIM is connected then, for every $K \geq 1$, all individuals hold identical (multidimensional) opinions after the first round of learning, i.e. $\tilde{x}_i(A\,;\,t, \tilde{\beta}) = \tilde{x}_j(A\,;\,t, \tilde{\beta})$ for all $i, j \in N$ and $t \geq 1$.*

Corollary 1 simply reflects the comments preceding the statement of Proposition 3 where we explained that, in the absence of signal correlation across some subsets of the population, these subsets can be treated as independent sub-populations. On the other hand, Corollary 2 readily follows from the observation that if the influence exerted on, and received from, others is uniformly spread, then from the first period of interaction onward all agents combine the same initial signals in the same manner.

So far we have ignored what is a major issue in our framework, namely, the problem of equilibrium multiplicity. To bring home the point sharply, suppose that individual signals are uncorrelated and the population is partitioned into two subsets between which no influence flows (i.e. there is no agent in any of these subsets who has any influence on some agent of the other). Then, if the homophily consistency holds *separately* for each of the subsets, it also holds *globally* and the situation defines an equilibrium. This holds, in fact, even if the attention network is fully connected. The simple reason is that, if there is no interaction/learning across those two subsets of agents, their opinions at the end of a learning spell must be uncorrelated just as their initial signals are. Then, our equilibrium notion requires that no cross-group links with positive weights exist in the influence network. In contrast, Proposition 3 above shows that (provided the attention network is complete) a completely connected influence network with uniform weights that spans the two populations also satisfies our equilibrium condition. Moreover, such an influence network induces identical end opinions of all agents for all values to the learning time span K. This example illustrates vividly that equilibrium multiplicity is unavoidable in our model for some parameter configurations. In general, equilibrium multiplicity is a deep and difficult issue, a detailed study of which we leave for future research.

3.3 Influence and support

In this subsection, we identify a topological measure of the influence network that characterizes the equilibrium strengths of interpersonal influences at equilibrium. We start by defining, for some given $n \times n$ binary adjacency matrix $L \equiv (l_{ij})_{i,j \in N} \in \{0,1\}^{n \times n}$ and any pair of nodes ij, the following coefficient:

$$\varphi_{ik}(L) \equiv \frac{\sum_{s=1}^{n} l_{is} l_{ks}}{(\sum_{s=1}^{n} l_{is}^2)^{\frac{1}{2}} (\sum_{s=1}^{n} l_{ks}^2)^{\frac{1}{2}}}. \tag{9}$$

Thus, $\varphi_{ik}(L)$ is a ratio of the number of nodes connected in L to i *and* k by their incoming links over the geometric mean of i's and k's out-degrees. In other words, it is a normalized measure of i's and j's shared neighborhood in L.

We note that (9) is closely related to the notion of *link support* defined for undirected networks by Jackson, Rodríguez-Barraquer and Tan (2012). They declare that a link ik is supported in some underlying network g if there exists an agent s, different from i and k, such that the links is and ks also belong to g. Thus, in contrast with (9), their measure is undirected, a link being (fully) supported when it is part of at least one triad. As these authors explain, their notion of support is very different from the classical one of clustering, and indeed the same applies to (9) as well.[17] They find that support is one of the key features of social networks required for the existence of robust cooperation equilibria.

There is also a close similarity of (9) with the notion of *neighborhood overlap*, defined

[17] The notion of support is link-based whereas that of clustering is node-based – it specifies the fraction of neighbors of a *node* who are themselves connected. For example, if a pair of connected agents are both part of various *common* " clusters" (groups) but these clusters are not connected, the support of their link can be very high but the clustering of each of them very low.

in Easley and Kleinberg (2010) as the ratio,

$$\frac{\text{number of nodes who are neighbors of both } i \text{ and } k}{\text{number of nodes who are neighbors of at least one of } i \text{ or } k} \tag{10}$$

where in the denominator we do not count i or k themselves (even though i and k are taken to be neighbors). For example, if i and k have no common acquaintances the ratio (10) is equal to zero and the link ik represents a potential bridge between two different communities. The neighborhood overlap (10) is also known as *link clustering coefficient* in the network literature (e.g., Pajevic and Plenz, 2012).

The notion of support that we shall use here expands on that defined in (9) by allowing for an arbitrary influence (hence weighted) matrix A as the argument of $\varphi_{ik}(\cdot)$. In this manner, we obtain a topological measure of *support* (we shall also refer to it as *neighborhood overlap*) that can be applied to influence matrices $A \equiv (a_{ij})_{i,j \in N} \in [0,1]^{n \times n}$. That is, we define:

$$\varphi_{ik}(A) \equiv \frac{\sum_{s=1}^{n} a_{is} a_{ks}}{(\sum_{s=1}^{n} a_{is}^2)^{\frac{1}{2}} (\sum_{s=1}^{n} a_{ks}^2)^{\frac{1}{2}}}. \tag{11}$$

This measure is symmetric (i.e. $\varphi_{ik}(A) = \varphi_{ki}(A)$) and normalized ($\varphi_{ik}(A) \in [0,1]$) whenever $\varphi_{ik}(A)$ is well-defined, i.e. whenever the ith and the kth row of A have at least one non-zero entry each. In our context, the latter condition always holds as we require that each node is connected (i.e. pays attention) to itself. As an immediate generalization, we shall also be interested in using, instead of the matrix of direct (immediate) influence A, the matrix $A^K \equiv B \equiv (b_{ij})_{ij \in N}$ capturing the indirect influence b_{ij} that an agent j has on agent i after K rounds of learning. Since B inherits from A all its essential properties, the K-step support $\varphi_{ik}(A^K)$ displays the same properties as the (one-step) support $\varphi_{ik}(A)$.

Based on the notion of K-step support, we obtain the following topological characterization of equilibrium influence matrices.

Proposition 4 *Given the attention network with adjacency matrix* $L \equiv (l_{ij})_{i,j \in N}$ *with* $l_{ij} \in \{0,1\}$, *the learning depth* $K \geq 1$ *and the covariance matrix of initial opinions* $\Sigma_{\tilde{\beta}} \equiv \sigma \cdot \mathbf{I}$ *for some* $\sigma > 0$, *the EIM* $A^* \equiv (a_{ij}^*)_{i,j \in N}$ *verifies:*

$$a_{ik}^* \propto l_{ik} \varphi_{ik}((A^*)^K) \qquad (i,k = 1,2,...,n). \tag{12}$$

Proof. See Appendix.

Heuristically, this result reflects the idea that common K-order partners help "support" the relationship between i and j, and therefore the more such partners i and j share, the stronger their relationship. More specifically, in our context the support comes from the fact that, by being subject to common influence, agents i and k will tend to strengthen the correlation of their behavior and hence, by homophily, their own link as well.

For the particular case of $K = 1$, Proposition 4 leads to the stark implication that link strength between two nodes is *directly proportional* to their neighborhood overlap, i.e. to the (weighted) number of neighbors they have in common. Intuitively, this property suggests that, in endogenous networks, strong links will tend to be arranged in "triangles." The literature – see e.g. Kossinets and Watts, (2006, 2009) and Kumpula *et al.* (2007) – often rationalizes such configurations through the claim that, among strong links, the principle of triadic/transitive closure applies (i.e. the friend of a friend tends to become a friend). Our model, however, provides a quite different explanation for this phenomenon: it is not the strength of the links that brings about the transitivity in connections; rather, it is that only those links that are well supported (and hence form part of some triangles) can be strong at equilibrium.

The aforementioned considerations bear on the important issue of social integration. Specifically, they address the question whether an originally fragmented society – divided into disjoint groups that neither exchange information nor exert any influence across them – may persist in such a segmented situation even after cross-group links are created to avert it. To illustrate starkly that such a state of affairs is indeed possible, we close this section with the discussion of a simple example that, relying on Proposition 4, conveys some of the main insights. These insights are then developed formally in the next subsection, where the equilibrium (hence static) analysis carried out here is extended to a dynamic setting.

Consider a situation where the initial population N is partitioned into two groups, $G_1 = \{1, 2, ..., n_1\}$ and $G_2 = \{n_1+1, n_1+2, ..., n_1+n_2\}$, with equal cardinalities, $n_1 = n_2 = n/2$. Suppose that, *within* each of these groups, the agents are placed linearly along a ring and each of them pays attention (i.e. is connected in the attention network) to the two other agents adjacent to her on either side. To be concrete, let us suppose that each agent $i \in G_1$ is connected with agents $[i \pm 1]_{n/2}$ and $[i \pm 2]_{n/2}$ (where $[\cdot]_{n/2}$ stands for the module $n/2$ operator) and each agent $j \in G_2$ is two-way connected to agents $(\frac{n}{2} + [j \pm 1]_{n/2})$ and $(\frac{n}{2} + [j \pm 2]_{n/2})$. Of course, if those attention links were the only ones in place, the population would be separated into two disjoint and completely unrelated sub-populations, no influence whatsoever flowing across them. Thus, we shall suppose instead that each agent $i \in G_1$ is also connected to agents $(\frac{n}{2} + [i \pm 2]_{n/2})$ from G_2. (See Figure 3 for a schematic description of the network considered.)

The network considered displays a total of n bilateral attention links established across the two groups. However, despite the existence of so many cross-links, it is easy to see that with the vanishing individual persistence (parametrized by η), the equilibrium weight of any link between G_1 and G_2 also vanishes. More precisely, this weight is of order η as this parameter approaches zero. To understand this conclusion, let us simplify the analysis and focus on the limit case where $\eta = 0$, which implies that the learning process displays no individual persistence. Then, it follows from Proposition 4 that, in order for any given

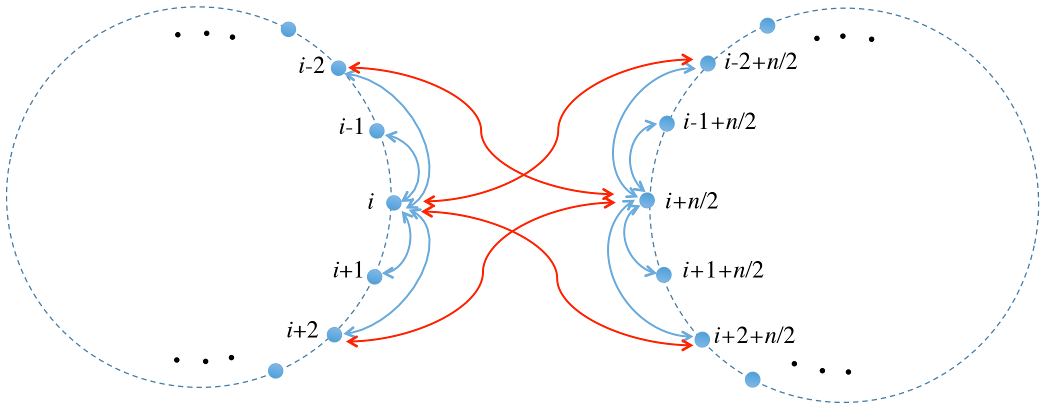

Figure 3: Diagrammatic illustration of the attention network considered in the example proposed here. The nodes in each group, G_1 and G_2, are linearly arranged along the corresponding ring. The neighborhoods of typical nodes i and $i+n/2$ are depicted. Solid links connect agents within a group, whereas dotted link connect agents of different groups.

cross link ij ($i \in G_1$, $j \in G_2$) to display some positive weight at equilibrium, it must be that both i and j are influenced by some common agent k, different from both i and j. The intuition here is that, when there is no opinion persistence, it is only through such third-party influence that some correlation between i's and j's opinions can arise at equilibrium. However, in the present example, none of the cross-group links has any third-party support (cf. Figure 3), which implies that their weight at equilibrium must be zero if $\eta = 0$. [18]

In fact, it turns out that an analogous line of reasoning applies for configurations of the attention network that are much denser than the one considered in Figure 3. For example, there is still no third-party support (and hence zero influence weight at equilibrium) for the cross-group links of an attention network when, in addition to links already indicated, each agent i is connected to agents $i \pm (2 + 3r) + n/2$ for $r = 1, 2, ...$ until the moves along the ring in the two opposite directions eventually meet. This amounts to a number of links of order n^2, comparable to that of the complete network. However, even such very high density of connections (indeed, much denser than the one prevailing within each group) is unable to induce any cross-group influence and the two groups remain isolated as far as genuine communication is concerned.

What is the obstacle to integration in our example? There is a twin problem. On the one hand, neither group has a dense local structure that can strongly support cross-group links. On the other hand, the cross-group links themselves do not take advantage of whatever local structure there is in each group – which is indeed weak but could also be exploited more effectively (e.g. if two adjacent agents in one group were connected to the same individual

[18] It can be readily checked that the only symmetric EIM A^* where both G_1 and G_2 display some within-group influence has, for all $i \in G_1$, $a^*_{ij} = 1/4$ for $j = [i \pm 1]_{n/2}$ and $[i \pm 2]_{n/2}$, while $a^*_{ij} = 0$ otherwise. An analogous statement applies to group G_2.

in the other group). In the following subsection, we study in detail the second problem and hence choose to abstract from the first one. Specifically, we study a context that is analogous to the present one but where the two groups are *completely connected within.* There is, hence, no shortage of local structure where cross-group links might be embedded. Then, it turns out that other aspects such as the relative size of each group, the depth K of the learning process, or the topology of the cross-group links become important. As we shall explain, another significant difference with the analysis that has been pursued here is that our approach will not be static but dynamic. That is, it will not focus on whether or not integrative (non-segmented) equilibria exist but whether segmented equilibria are robust, i.e. dynamically stable. After all, as we know from Corollary 1, a segmented equilibrium for two completely connected groups always exists, so a key question is whether these equilibria are robust or not. For, even if integrative equilibria do exist, the segmented one may be robust and hence difficult to escape from.

3.4 Bridging

We consider the problem of bridging and segmentation in the following stylized context. The overall population N is partitioned into two groups, $G_1 = \{1, 2, ..., n_1\}$ and $G_2 = \{n_1 + 1, n_1 + 2, ..., n_1 + n_2\}$, with $n_1 + n_2 = n$. Each of these two groups is completely connected in the attention network, while the initial opinions/signals of all agents are uncorrelated with variances, common within each group G_q, respectively denoted by ϖ_q.

Initially, the prevailing influence matrix is of the form:

$$A^0 = \begin{pmatrix} Q(n_1) & 0_{n_1 \times n_2} \\ 0_{n_2 \times n_1} & Q(n_2) \end{pmatrix}, \tag{13}$$

where $0_{n_r \times n_s}$ stands for a matrix of dimension $n_r \times n_s$ consisting of all zeros and $Q(n_r)$ represents the uniform matrix with all entries equal to $1/n_r$. As stated in Corollary 1, regardlessly of whether there exist attention links *across* the two groups or how many there are, the influence matrix A^0 defines an EIM. Thus, to fix ideas, let us suppose that no such cross-group links initially exist. That is, the adjacency matrix $L^0 \equiv (l^0_{ij})_{i,j \in N}$ defining the attention network satisfies:

$$l^0_{ij} = l^0_{ji} = 1 \Leftrightarrow i, j \in G_q \qquad (i, j = 1, 2, ..., n;\ q = 1, 2). \tag{14}$$

Then, starting from this situation, suppose that some inter-group attention links can be created, which leads to the following question: How should these two groups be connected so as to destabilize what would otherwise remain a segmented equilibrium configuration? Or reciprocally, we can ask: Under what conditions are those cross-group connections ineffective in breaking the stability of the segmented configuration?

In order to cast this question in mathematical form, we need to make explicit a dynamical adjustment process for influence networks. In line with our discussion in Subsection 2.3 (see our Twitter example there) and Subsection 3.1 (see the discussion following Proposition 1), we posit the following discrete and synchronous *adjustment process*.

Let the stages of the process (i.e. separate learning spells) be indexed by $\tau = 0, 1, 2, ...$ and denote by $A(\tau) \equiv (a_{ij}(\tau))_{i,j \in N}$ the influence matrix prevailing at a typical τ. Then, we postulate the following law of motion:

$$a_{ij}(\tau + 1) = \frac{l_{ij}\, \rho_{ij}(A(\tau))}{\sum_{k \in N} l_{ik}\, \rho_{ik}(A(\tau))} \qquad (i, j = 1, 2, ..., n;\ \tau = 0, 1, 2, ...), \tag{15}$$

or, using the compact vector-field notation in (7), we may simply write:

$$A(\tau + 1) = F(A(\tau)) \qquad (\tau = 0, 1, 2, ...). \tag{16}$$

Before stating formally our main result in this subsection, we introduce some notation. Let the cross-group attention links to be added to the original ones be collected in the adjacency (binary) matrix $V \equiv (v_{ij})_{i,j \in N}$ such that

$$\forall i, j \in N,\ \forall q \in \{1, 2\}, \quad [i, j \in G_q \Rightarrow v_{ij} = 0]. \tag{17}$$

The new attention network is given, therefore, by the adjacency matrix $L = L^0 + V$, which simply adds the new links across groups to the complete array of those pre-existing within groups. Finally, for any two subsets of agents, M and M' (possibly singletons), we define

$$V_M^{M'} \equiv \sum_{i \in M, j \in M'} v_{ij}, \tag{18}$$

as the number of inter-group attention links from the set M to the set M'.

Our main result in this subsection specifies *sufficient* conditions for the segmented configuration (13) to be asymptotically stable for the dynamical system (16) when the prevailing attention matrix is L. Intuitively, this means that even if a "small" amount of communication started to flow across the new cross-group links, the adjustment process would eventually lead to the original configuration where neither group has any influence on the other. In this sense, such a segmented configuration can be viewed as dynamically robust. As standard in the theory of dynamical systems, such local stability is assessed in terms of the eigenvalues of the Jacobian matrix $\mathcal{J}(A^0) \equiv \left(\frac{\partial F_{ij}}{\partial a_{k\ell}}(A^0)\right)_{i,j,k,\ell \in N}$ of the vector field (7) evaluated at $A = A^0$.

Proposition 5 *Consider the attention matrix $L = L^0 + V$, where L^0 is given by (14) and V satisfies (17). Assume that the initial opinions are independent across agents with common variances ϖ_q within each group G_q. Then, the largest eigenvalue in absolute value of the*

Jacobian matrix $\mathcal{J}(A^0)$ is a convex combination of the elements from the following set:

$$Z \equiv \left\{ \frac{n^i \varpi^j V_i^{G^j} + n^j \varpi^i V_j^{G^i} + (K-1)\left(\varpi^i V_{G^j}^{G^i} + \varpi^j V_{G^i}^{G^j}\right)}{n^i \sqrt{\varpi^i \varpi^j n^i n^j}} \right\}_{\{ij: v_{ij}=1\}}, \tag{19}$$

where $K \geq 1$ is the depth of the learning process, $V_M^{M'}$ is defined in (18), and $G^i \equiv G_q$, $n^i \equiv n_q$, and $\varpi^i \equiv \varpi_q$ if $i \in G_q$.

Proof. See Appendix.

Corollary 3 *Under the conditions of Proposition 5, the EIM A^0 is asymptotically stable for the adjustment dynamics given by (15) if* $\max\{z \in Z\} < 1$.

The intuition for Proposition 5 and Corollary 3 is as follows. Given that in the segmented configuration both groups, G_1 and G_2, display an internal pattern of influence that is uniform, the correlation of behavior within them is almost perfect, even as some (small) cross-group influence starts to flow. Then, the adjustment to the internal uniformity after such (small) perturbation can be traced to the changes experienced by the weight associated to cross-groups links, i.e. those in the set $\{ij : v_{ij} = 1\}$. These are, in a sense, a suitable "base" to determine the overall effect of those perturbations. It follows, therefore, that the change along the main eigenvector component, as measured by its corresponding eigenvalue, can be obtained as the composition (i.e. convex combination) of the effects associated to the different cross-group links. Then, given such a composed effect, the conclusion of Corollary 3 readily follows. In particular, if the impact along each of the contemplated dimensions is lower than unity, so happens for their convex combination and thus a standard sufficient condition for local stability obtains: the largest-modulus eigenvalue (which is real) is less than unity.

In view of (19), it is worth noting that the stability/robustness of the segmented configuration is favored, *ceteris paribus*, by the following:

1. low learning depth;
2. large and/or similar group sizes;
3. few cross-component links;
4. low number of triangles in the attention network that cover cross-component links.

The features considered in $1-4$ accord well with intuition as to what should be the factors hindering the establishment of influence across the two groups. Heuristically, the considerations involved in each case can be understood as follows.

1' If the learning span is short (low K), the process does not allow for enough rounds of interaction for the initial opinions to draw close together. In turn, this limits inter-group correlation and, by homophily consistence, leads to low influence weights across different groups.

2' If the groups are large (high n^i and/or n^j), the *internal* (uniform) influence will tend to offset the external influence channeled through the cross-group links. This effect will be stronger the more similar the group sizes are.

3' Each cross-group link is a potential channel of reciprocal influence between the two groups. Increasing the number of such channels (i.e., increasing $V_{G^j}^{G^i}$ and/or $V_{G^i}^{G^j}$) reinforces the cross-group communication (opinion correlation) and, hence, supports the weight of each separate links as a channel of influence.

4' As we already know, the pattern of cross-group connections matters for equilibrium influence. In particular, as neighborhood overlap enhances opinion correlation, a cross-group link ij will gain equilibrium weight as i or/and j have more cross-group links themselves (which form, then, cross-group triangles). With just one round of learning, this number increases with $V_i^{G^j}$ and $V_j^{G^i}$; for longer learning spans, it is also $V_{G_j}^{G^i}$ that matters.

We close our discussion in this section by presenting three remarks that highlight some interesting insights derived from Proposition 5.

REMARK 2 (BRIDGING REVISITED) *At the end of Subsection 3.3, we discussed a simple example that illustrated the potential robustness of social segmentation in the face of (a large number of) bridging links established between the disconnected groups of the population. That example may have been useful but had a number of limitations. One of them was the static (equilibrium) approach of the example, which contrasts with the fact that the robustness of segmented equilibria is best addressed dynamically. A second limitation was that the sharpness of our conclusions relied on a vanishing value of η, the parameter that modulates the opinion persistence of individuals. Both of these aspects can be dealt with by relying on Proposition 5 – for not only the analysis is explicitly dynamic but the result is established for non-vanishing opinion persistence ($\eta = 1$).*

To concentrate our attention on the issue of bridging, let us make the same simplifying assumptions as in the indicated example: both groups display the same size and variance of initial opinions ($n_1 = n_2 = n/2$; $\varpi_1 = \varpi_2 = \varpi$), and every learning spell involves just one round ($K = 1$). Then, the set (19) boils down to,

$$Z \equiv \left\{ \frac{V_i^{G^j} + V_j^{G^i}}{n/2} \right\}_{\{ij : v_{ij}=1\}}. \tag{20}$$

Consider now any pattern of cross-group links V that satisfies:

$$\forall i \in N, \quad \sum_{j \in N} v_{ij} < n/4.$$

That is, every agent in each group G_q is connected to less than half of the individuals in the other group $G_{q'}$, $q' \neq q$. It readily follows from (20) that for every element $z_{ij} \in Z$, associated to the cross-group link ij, we have $z_{ij} < 1$. Proposition 5 implies, then, that the segmented configuration is asymptotically stable and that this can happen even when the total number of cross-group links is of order n^2.

In contrast, there is the possibility of destabilizing the segmented configuration with much fewer cross-group links if these are arranged differently. What matters is not just the number of such links but the specific pattern in which they are arranged. As a simple illustration, let us select two agents, $i_1 \in G_1$ and $i_2 \in G_2$, and suppose each of them is connected to all agents in the other group (hence, there is a total of $n-1$ cross-group links). Then, again referring to (20), we find that $z_{ij} \geq 1 + 2/n$ for each cross-group link ij, which means that these links suffice to destabilize the segmented equilibrium. And this, of course, is done with a total of $n-1$ cross-group links, a much smaller number than in the previous construction that allowed a total number of those links to grow at the order of n^2!

REMARK 3 (UNRESTRICTED CROSS-GROUP COMMUNICATION) *Suppose that the population N is partitioned into two equally numerous groups, $G_1 = \{1, ..., \frac{n}{2}\}$ and $G_2 = \{\frac{n}{2}+1, ..., n\}$, for some even $n \geq 2$. All initial agent opinions are independent and display the same variance ϖ ($= \varpi_1 = \varpi_2$). Furthermore, the underlying attention network $L \equiv (l_{ij})_{i,j \in N}$ is complete, so that $l_{ij} = 1$ for all $i, j \in N$.*

Under the described conditions, the set Z in (19) contains the common element

$$\frac{\varpi n^2/4 + \varpi n^2/4 + (K-1)\left(\varpi n^2/4 + \varpi n^2/4\right)}{\varpi n^2/4} = 2K > 1. \tag{21}$$

It then follows from Proposition 5 that the segmented EIM A^0 is asymptotically unstable for any learning depth $K \geq 1$. The segmentation it induces is therefore bound to be destabilized unless there exist some barriers to inter-group communication in the form of missing links in the attention network L.

It should be emphasized, however, that (21) is a local condition that guarantees only that small perturbations from the equilibrium influence matrix will typically bring the system away from it. It does not necessary imply that, once the system is destabilized in this manner, the adjustment dynamics (15) will lead the population to an integrated state involving cross-group influence.[19] An analogous word of warning is applicable to the interpretation of (21) below, which identifies conditions for dynamic instability but in a quite different context.

[19] Interestingly, we find that such a transition towards integration does arise in practice for the specific instances studied through numerical simulations.

REMARK 4 *[The role of a leader] Suppose that the population* N *consists of two groups: on the one hand, the singleton group* $G_1 \equiv \{1\}$ *consisting of the leader alone; on the other, the group* $G_2 \equiv \{2, 3, ..., n\}$ *that includes the remaining* $n-1$ *agents. Suppose that the leader is influenced by no one – i.e. the underlying adjacency matrix* $L \equiv (l_{ij})_{i,j \in N}$ *has* $l_{1i} = 0$ *for all* $i \in G_2$*, while in the group* G_2*, there are* $m(\leq n-1)$ *agents who pay attention to the leader. For simplicity, let this set of followers be given by* $S = \{2, 3, ..., m+1\}$*. Then, the whole matrix* L *can be characterized as follows:*

$$\forall i, j \in N, \quad l_{ij} = 1 \quad \Longleftrightarrow \quad \{[i, j \in G_2] \ \vee \ [(i \in \{1, ..., m+1\} \wedge j = 1)]\}$$

Consider now an original configuration where the leader has no influence on the group G_2 *but this group (which is completely connected in the attention network) displays a uniform pattern of internal influence. This equilibrium situation is formalized by the EIM* A^0 *given in (13) for* $n_1 = 1$ *and* $n_2 = n-1$*. Then, if we denote by* $\lambda \equiv \frac{m}{n}$ *the fraction of followers in group* G_2*, a necessary condition for homophily-based adjustment to destabilize such an equilibrium is that the common element of the set (19) exceeds unity. That is, we must have:*

$$\frac{(n-1)\varpi_1 + (K-1)\varpi_1\lambda(n-1)}{(n-1)\sqrt{\varpi_1\varpi_2(n-1)}} = \frac{1 + (K-1)\lambda}{\sqrt{n-1}}\sqrt{\frac{\varpi_1}{\varpi_2}} > 1. \tag{22}$$

The above inequality holds if, ceteris paribus, the learning time span is long enough (high K*), the variance (i.e. informativeness) of the signals received by the leader – the agent exerting all the influence – is sufficiently high (large* ϖ_1*), the variance of the signals received by followers is sufficiently low (small* ϖ_2*), the size of the group* G_2 *on which the leader exerts (directly or indirectly) some influence is small (low* $n-1$*), and the fraction of those directly influenced by the leader is large (high* λ*).*

4 Conclusion

We have started this paper by studying a model of social learning on a *given* social network that extends the classical one proposed by DeGroot (1974) in two relevant dimensions:

- opinions are multidimensional (and modeled as non-degenerate random variables);
- the length of learning spells (i.e. the number of learning rounds) is typically finite.

The above features lead to a context where, even under customary regularity conditions, the learning outcome generally falls short of full consensus. This allows one to identify the extent to which agents' final positions correlate, depending on the nature of the initial signals/opinions, the architecture of the network, and the length of the learning process.

In a such a richer learning environment, we addressed the main objective of the paper: an endogeneization of the influence network that

(a) respects the communication restrictions imposed by an exogenous attention network;

(b) is consistent with a notion of homophily that requires the strength of a relationship (influence) to be proportional to the correlation of behavior (opinions).

An important consequence induced by the postulated homophily-based consistency is that, in equilibrium, the strength of every link must be proportional to its support in a suitably defined neighborhood. This has been then brought to bear on the important issue of social integration, a social-network phenomenon with far-reaching consequences. In particular, we have shed light on what features of the environment can make social segmentation a robust phenomenon and, correspondingly, the approaches (e.g. the creation of effective "bridges") that would allow to overcome it.

Naturally, a proper welfare and policy analysis of the problem of social cohesion requires a meaningful assessment of benefits and costs. This, however, was absent from our theoretical framework and represents one of the important avenues for future research. Indeed, depending on how cohesion is defined, it may be the case that not always more of it is better – for example, along certain dimensions too much integration can be detrimental to the preservation of valuable diversity.

Finally, another important and somewhat related objective for future research is the introduction of some extent of payoff-guided behavior into the learning environment. The model we have proposed is purely behavioral and, in this sense, quite mechanical. Other paradigms of learning that could be studied from an analogous perspective include, e.g., observational learning (Bala and Goyal, 1998), Bayesian learning (Gale and Kariv, 2003), or a mixture of boundedly-rational and Bayesian learning (Mueller-Frank, 2014). It is conceivable that when, as we have done for the modified DeGroot model, such alternative forms of learning are combined with a homophily-based influence, the results will be quite similar to those presented here. This overarching approach, however, raises the following central issue to start with: what is the rationale for homophily? Is it, as many social scientists suggest, a purely innate tendency of human beings that has been shaped by natural selection? Or is it, instead, a behavior that has some payoff basis, at least in some contexts and under certain circumstances? Addressing these and related questions must obviously be a central concern in a properly normative analysis of the problem.

5 Appendix

Proof of Proposition 1: For a given (finite) learning depth K, adjacency matrix L and the signal covariance matrix $\Sigma_{\tilde{\beta}}$, the vector field $F(A; K, L, \Sigma_{\tilde{\beta}}) : (\Delta^{n-1})^n \to (\Delta^{n-1})^n$, defined in (7), maps an $n-$dimensional stochastic matrix A into another $n-$dimensional stochastic matrix (Δ^{n-1} is an $n-$dimensional simplex). $F(\cdot)$ is continuous as it involves only a finite number of continuous matrix operations when $K < \infty$. As $(\Delta^{n-1})^n$ is compact and convex, Brouwer fixed-point theorem implies that F has a fixed point A^*. For $K \to \infty$, we construct an EIM in the proof of Proposition 2 below. ■

Proof of Proposition 2: If A^* is an EIM for a fixed adjacency matrix L and signal covariance matrix $\Sigma_{\tilde{\beta}}$, then $a^*_{ii} > 0$ because $l_{ii} = 1$ and $\rho_{ii}(A^*) = 1$ for each $i \in N$. Hence, A^* must be aperiodic (e.g., Golub & Jackson, 2010). It is well known that for a connected (i.e., irreducible), aperiodic and stochastic matrix A^* each row of $\lim_{K\to\infty}(A^*)^K$ is equal to the left eigenvector of A^* associated to the eigenvalue 1. Then, all correlations $\rho_{ik}(A^*)$, computed by (4) for $i, k \in N$, are equal to one. Substituting unit correlations into the definition (5) of EIM yields the claim, $a^*_{ik} = l_{ik}/\sum_{s=1}^{N} l_{is}$ for all $i, k \in N$. ■

For the proof of Proposition 3, we need the following Lemma.

LEMMA 1 *For row-stochastic and strictly positive* $n \times n$ *matrices* $A \equiv (a_{ik})_{i,k\in N}$ *and* $S \equiv (s_{ik})_{i,k\in N}$ *such that* $A \neq Q(n)$, $a_{ii} \geq a_{ik}$ *and* $a_{ik}/a_{ii} = a_{ki}/a_{kk}$ *for all* $i, k = 1, ..., n$,

$$\delta(A) < \delta(AS'), \quad \textit{where} \quad \delta(X) \equiv \min_{i,k} \frac{x_{ik}x_{ki}}{x_{ii}x_{kk}}.$$

Proof: As A, S are strictly positive and row-stochastic, it holds for all $i, k = 1, ..., n$,

$$\begin{aligned} &\min_s a_{is} \leq \textstyle\sum_{s=1}^{n} a_{is}s_{ks} = (AS')_{ik} \leq \max_s a_{is} = a_{ii} \Rightarrow \\ &\min_s a_{is} \leq \min_s (AS')_{is} \leq (AS')_{ii} \leq \max_s (AS')_{is} \leq a_{ii} \Rightarrow \\ &\min_s \frac{a_{is}}{a_{ii}} \leq \min_s \frac{(AS')_{is}}{(AS')_{ii}}, \quad \forall i \Rightarrow \min_{i,s} \frac{a_{is}}{a_{ii}} \leq \min_{i,s} \frac{(AS')_{is}}{(AS')_{ii}}. \end{aligned} \tag{23}$$

Note that $a_{ii} = \max_s a_{is} > \min_s a_{is}$ for at least one i. Otherwise, $a_{ii} = a_{is}$ for all i, s which implies $a_{is} = 1/n$ as A is row-stochastic. This, however, would contradict $A \neq Q(n)$. Therefore, the inequalities in (23) are strict for at least one i. Then, the last equality in (23) is also strict and it follows from it and from the symmetry condition $a_{ik}/a_{ii} = a_{ki}/a_{kk}$,

$$\begin{aligned} \delta(A) &= \min_{i,k} \frac{a_{ik}a_{ki}}{a_{ii}a_{kk}} = \min_{i,k} (\frac{a_{ik}}{a_{ii}})^2 = (\min_{i,k} \frac{a_{ik}}{a_{ii}})^2 \\ &< (\min_{i,k} \frac{(AS')_{ik}}{(AS')_{ii}})^2 \leq \min_{i,k} \frac{(AS')_{ik}(AS')_{ki}}{(AS')_{ii}(AS')_{kk}} = \delta(AS'). \end{aligned}$$

■

Proof of Proposition 3: The proof that $Q(n)$ is an EIM follows from the fact that $Q(n) = Q(n)' = Q(n)^K$ and $Q(n)XQ(n) = const * Q(n)$ for any matrix X. From (4) follows then that $\rho_{ik}(Q(n)) = 1$ for all $i, k \in N$. The substitution of unit correlations into the definition (5) shows that $Q(n)$ is an EIM for the completely connected attention network L.

The difficult task is to prove that $Q(n)$ is the unique connected EIM. First, we show that any connected EIM $A \equiv (a_{ik})_{i,k\in N}$ on completely connected L must be strictly positive. To see this, consider nodes i, j and k such that $a_{ij}a_{jk} > 0$,

$$a_{ij}a_{jk} > 0 \Rightarrow a_{ij} > 0 \Rightarrow \rho_{ij}(A) = \rho_{ji}(A) > 0 \Rightarrow a_{ji} > 0,$$
$$a_{ij}a_{jk} > 0 \Rightarrow a_{jk} > 0 \Rightarrow \rho_{jk}(A) = \rho_{kj}(A) > 0 \Rightarrow a_{kj} > 0.$$

Then, even if signals are uncorrelated, $\tilde{x}_i(A)$ is positively correlated with $\tilde{x}_k(A)$ via $\tilde{x}_j(A)$,

$$\rho_{ij}(A)\rho_{jk}(A) > 0 \Rightarrow \rho_{ik}(A) > 0.$$

We obtain, therefore,

$$a_{ij}a_{jk} > 0 \Rightarrow a_{ik}a_{ki} > 0.$$

As A is connected, this argument propagates to all links in L. We conclude, therefore, that all elements in the connected EIM A are strictly positive.

On the other hand, the condition (5) implies for the EIM A and the covariance matrix,

$$C \equiv (c_{ik})_{i,k\in N} \equiv A^K \Sigma_{\tilde{\beta}} (A^K)',$$

the following equalities,

$$\frac{a_{ik}}{a_{ii}} = \frac{a_{ki}}{a_{kk}} = \rho_{ik}(A) = \frac{c_{ik}}{(c_{ii}c_{kk})^{1/2}} \Rightarrow \frac{a_{ik}a_{ki}}{a_{ii}a_{kk}} = \frac{c_{ik}c_{ki}}{c_{ii}c_{kk}}, \quad \forall i,k \in N. \tag{24}$$

The necessary condition for (24) can be expressed with the matrix operator $\delta(\cdot)$,

$$\delta(A) = \delta(C), \quad where \quad \delta(X) \equiv \min_{i,k} \frac{x_{ik}x_{ki}}{x_{ii}x_{kk}} \quad for \quad X \equiv (x_{ik})_{i,k\in N}. \tag{25}$$

The following properties of $\delta(\cdot)$ are easily verified,

$$\delta(X') = \delta(X), \quad \delta(cX) = \delta(X), \quad \delta(D_1 X D_2) = \delta(X),$$

where X is an arbitrary matrix with positive entries, D_1 and D_2 are diagonal matrices and c is a constant. By the last property, we obtain from (25),

$$\delta(A) = \delta(C) = \delta(CD_2) \equiv \delta(AS'), \tag{26}$$

where $S' \equiv A^{K-1}\Sigma_{\tilde{\beta}}(A^K)'D_2$ and D_2 is a diagonal matrix that normalizes the sum of each column in $A^{K-1}\Sigma_{\tilde{\beta}}(A^K)'$. Hence, S' is a column-stochastic strictly positive matrix for any $K \geq 1$. In Lemma 1, we prove for strictly positive row-stochastic matrices S and $A \neq Q(n)$ such that $a_{ii} \geq a_{ik}$ and $a_{ik}/a_{ii} = a_{ki}/a_{kk}$ for all i, k that $\delta(A) < \delta(AS')$. As this contradicts (26), we conclude that only $A = Q(n)$ can be an EIM for the completely connected L. ■

Proof of Proposition 4: For the learning depth $K \geq 1$, signal covariance matrix $\Sigma_{\tilde{\beta}} = \sigma\mathbf{I}$ (where $\sigma > 0$ and $\mathbf{I}$ is the identity matrix) and an $n \times n$ matrix A, the correlation (4) is readily verified to be identical with the link support (11),

$$\rho_{ik}(A) = \frac{\sum_{s=1}^n b_{is}b_{ks}}{(\sum_{s=1}^n b_{is}^2)^{\frac{1}{2}} (\sum_{s=1}^n b_{ks}^2)^{\frac{1}{2}}} = \varphi_{ik}(B), \quad \forall i,k \in N, \quad where \quad B \equiv A^K.$$

Then, the condition (5) for an EIM A^* can be written as,

$$a^*_{ij} = l_{ij} a^*_{ii} \rho_{ij}(A^*) = l_{ij} a^*_{ii} \varphi_{ik}((A^*)^K).$$

■

For the proof of Proposition 5 we need the following notation and lemmas:

Let $\sigma = (\sigma_1, ..., \sigma_n)$ and define $I(\sigma)$ as an $n \times n$ diagonal covariance matrix of n IID signals with the diagonal elements $(I(\sigma))_{ii} = \sigma_i > 0$. For the learning depth $K \geq 1$, the $n \times n$ influence matrix $A \equiv (a_{ik})_{i,k \in N} \geq 0$ and the binary adjacency matrix $L \equiv (l_{ik})_{i,k \in N} \in \{0,1\}^{n \times n}$, we define

$$\begin{aligned} cov(A) &\equiv A^K I(\sigma)(A^K)^T, \quad \rho_{ik}(A) \equiv \frac{cov_{ik}(A)}{\sqrt{cov_{ii}(A) cov_{kk}(A)}}, \\ \widetilde{F}_{ik}(A; L) &\equiv \frac{l_{ik}\rho_{ik}(L \cdot A)}{\sum_{s=1}^n l_{is}\rho_{is}(L \cdot A)}, \qquad i,k = 1, ..., n, \end{aligned} \tag{27}$$

where AB and $A^K = A...A$ are products of (compatible) matrices A and B, while $A \cdot B$ is the Hadamard product, i.e., $(A \cdot B)_{ik} = a_{ik} b_{ik}$. Note that $\widetilde{F}_{ik}(A;L)$ is identical to $F_{ik}(A;L)$, as defined in (7), when $A \leq L$.

Further, let $G_1 = \{1, ..., n_1\}$ and $G_2 = \{n_1 + 1, ..., n_1 + n_2\}$ be two groups with n_1 and n_2 nodes, $n_1 + n_2 = n$. For any $i = 1, ..., n$, let $G^i = G_q$, $n^i = n_q$ and $\varpi^i = \varpi_q$ if $i \in G_q$ $(q = 1, 2)$. Let

$$\begin{aligned} A^0 &\equiv \begin{pmatrix} Q(n_1) & 0_{n_1 \times n_2} \\ 0_{n_2 \times n_1} & Q(n_2) \end{pmatrix}, \quad L^0 \equiv sign(A^0), \quad I(\varpi) \equiv \begin{pmatrix} \varpi_1 I_{n_1 \times n_1} & 0_{n_1 \times n_2} \\ 0_{n_2 \times n_1} & \varpi_2 I_{n_1 \times n_2} \end{pmatrix}, \\ where, \quad & Q(n_r) \equiv (q_{ik}(n_r))_{i,k=1}^{n_r}, \quad q_{ik}(n_r) \equiv \frac{1}{n_r}, \quad \varpi \equiv (\varpi_1, ..., \varpi_1, \varpi_2, ..., \varpi_2). \end{aligned}$$

Lemma 2 *For an* $n \times n$ *matrix* $U \equiv \{u_{ij}\}_{i,j \in N}$ *and* $ik \in n \times n$*,*

$$\frac{d\rho_{ik}(A^0 + \omega U)}{d\omega}|_{\omega=0} = \begin{cases} \frac{n^i \varpi^k U_i^{G^k} + n^k \varpi^i U_k^{G^i} + (K-1)(\varpi^k U_{G^i}^{G^k} + \varpi^i U_{G^k}^{G^i})}{\sqrt{\varpi^i \varpi^k n^i n^k}}, & i,k : G^i \neq G^k, \\ 0, \quad i,k : G^i = G^k, \end{cases}, \tag{28}$$

where $U_M^{M'} \equiv \sum_{i \in M, j \in M'} u_{ij}$*.*

Proof: We define the real matrix $A(\omega) \equiv A^0 + \omega U$ and the matrix-valued function

$$f^K(\omega) \equiv cov(A(\omega)) = A(\omega)^K I(\varpi)(A(\omega)^K)', \tag{29}$$

with a recursive structure,

$$\begin{aligned} f^K(\omega) &= A(\omega)^K I(\varpi)(A(\omega)^K)' = A(\omega)A(\omega)^{K-1} I(\varpi)(A(\omega)A(\omega)^{K-1})' \\ &= A(\omega)A(\omega)^{K-1} I(\varpi)(A(\omega)^{K-1})' A(\omega)' = A(\omega) f^{K-1}(\omega) A(\omega)', \\ f^0(\omega) &= I(\varpi). \end{aligned} \tag{30}$$

We use the product rule $(h \cdot g)' = h'g + hg'$ to compute the derivative of (30),

$$\begin{aligned} \frac{df^K(\omega)}{d\omega} &= Uf^{K-1}(\omega)A(\omega)' + A(\omega)(\frac{df^{K-1}(\omega)}{d\omega}A(\omega)' + f^{K-1}(\omega)U') \\ &= A(\omega)\frac{df^{K-1}(\omega)}{d\omega}A(\omega)' + Uf^{K-1}(\omega)A(\omega)' + A(\omega)f^{K-1}(\omega)U', \\ \frac{df^0(\omega)}{d\omega} &= 0. \end{aligned} \tag{31}$$

We solve (31) by successive substitution and evaluate $\frac{df^K(\omega)}{d\omega}$at $\omega = 0$,

$$\frac{df^K(0)}{d\omega} = A^0\frac{df^{K-1}(0)}{d\omega}A^0 + \Psi + \Psi^T = ... = (K-1)A^0(\Psi + \Psi^T)A^0 + \Psi + \Psi^T, \tag{32}$$

where $A^k(0) = A^0$ for any $k = 1, ..., K$, $\Psi \equiv Uf^K(0)$ and

$$f^K(0) = A^0 I(\varpi) A^0 \Rightarrow f^K_{ik}(0) \equiv \begin{cases} \varpi^i/n^i, & G^i = G^k, \\ 0, & G^i \neq G^k. \end{cases} \tag{33}$$

From (32) and (33), it can be verified directly that

$$\frac{df^K_{ik}(0)}{d\omega} = \frac{n^i\varpi^k U_i^{G^k} + n^k\varpi^i U_k^{G^i} + (K-1)(\varpi^k U_{G^i}^{G^k} + \varpi^i U_{G^k}^{G^i})}{n^i n^k}. \tag{34}$$

By applying the quotient rule $(h/g)' = (h'g - hg')/g^2$, we obtain the derivative of the correlation function,

$$\begin{aligned} &\frac{d\rho_{ik}(A(\omega))}{d\omega}|_{\omega=0} = \frac{d(f^K_{ik}(\omega)/\sqrt{f^K_{ii}(\omega)f^K_{kk}(\omega)})}{d\omega}|_{\omega=0} \\ &= \begin{cases} \sqrt{\frac{n^i n^k}{\varpi^i\varpi^k}}\frac{df^K_{ik}(0)}{d\omega}, & i,k : G^i \neq G^k, \\ \frac{n^i}{\varpi^i}(\frac{df^K_{ik}(0)}{d\omega} - \frac{1}{2}(\frac{df^K_{ii}(0)}{d\omega} + \frac{df^K_{kk}(0)}{d\omega})), & i,k : G^i = G^k, \end{cases} \end{aligned} \tag{35}$$

where we used (33) to substitute for $f^K_{..}(0)$. The formula (28) obtains then by substituting $\frac{df^K_{..}(0)}{d\omega}$ from (34). In particular, $G^i = G^k$ implies $\varpi^i = \varpi^k$ and $n^i = n^k$ and all terms in (34) cancel out in this case. ■

Lemma 3 *For binary* $n \times n$ *matrices* $V \equiv (v_{\tau\omega})_{\tau,\omega\in N}$ *and* $U^{st} \equiv (u^{st}_{\tau\omega})_{\tau,\omega\in N}$ *such that* $v_{\tau\omega} = 0$ *if* $G^\tau = G^\omega$ *and* $u^{st}_{st} = 1$, $u^{st}_{\tau\omega} = 0$ *for* $\tau\omega \neq st$,

$$\begin{aligned} &\frac{d}{d\omega}\widetilde{F}_{ik}(A^0 + \omega U^{st}; L^0 + V)|_{\omega=0} = \\ &\begin{cases} \frac{n^i\varpi^k(\Upsilon^{st})_i^{G^k} + n^k\varpi^i(\Upsilon^{st})_k^{G^i} + (K-1)(\varpi^k(\Upsilon^{st})_{G^i}^{G^k} + \varpi^i(\Upsilon^{st})_{G^k}^{G^i})}{n^i\sqrt{\varpi^i\varpi^k n^i n^k}}, & if \quad v_{ik} = 1, \\ 0, \quad if \quad v_{st} = 0, \end{cases} \end{aligned} \tag{36}$$

where $\Upsilon^{st} \equiv (L^0 + V) \cdot U^{st}$.

Note that the two cases in (36) are neither mutually exclusive nor collectively exhaustive (but they are the only relevant ones for the proof of Proposition 5).

Proof: We define $\Lambda \equiv (\lambda_{ik})_{i,k\in N} \equiv L^0 + V$, $A^{st}(\omega) \equiv \Lambda \cdot (A^0 + \omega U^{st}) = A^0 + \omega \Upsilon^{st}$ and the normalization factor

$$\eta_i(\omega, \Lambda) \equiv \textstyle\sum_{s=1}^n \lambda_{is}\rho_{is}(A^{st}(\omega)), \quad i = 1, .., n.$$

We note that $\eta_i(0, \Lambda) = \sum_{s\in G^i} l_{is}^0 \cdot 1 = n^i$ due to the fact that $v_{is} = 0$ when $s \in G^i$ (i.e. $G^s = G^i$) and,

$$\rho(A^0) = L^0 = \begin{pmatrix} 1_{n_1\times n_1} & 0_{n_1\times n_2} \\ 0_{n_2\times n_1} & 1_{n_2\times n_2} \end{pmatrix}. \tag{37}$$

Then, we calculate the following derivative by applying the quotient rule,

$$\frac{d\widetilde{F}_{ik}(A^0 + \omega U^{st}; \Lambda)}{d\omega}|_{\omega=0} = \frac{d}{d\omega}\frac{\lambda_{ik}\rho_{ik}(A^{st}(\omega))}{\eta_i(\omega, \Lambda)}|_{\omega=0} = \tag{38}$$
$$\frac{\lambda_{ik}}{(n^i)^2}(\frac{d\rho_{ik}(A^{st}(\omega))}{d\omega}n^i - \rho_{ik}(A^0)\frac{d\eta_i(\omega, \Lambda)}{d\omega})|_{\omega=0},$$

where,

$$\frac{d\eta_i(\omega, \Lambda)}{d\omega} = \textstyle\sum_{s=1}^n \lambda_{is}\frac{d\rho_{is}(A^{st}(\omega))}{d\omega}. \tag{39}$$

For $v_{ik} = 1$ the definition of V implies $G^i \neq G^k$ and, then, we have $\rho_{ik}(A^0) = l_{ik}^0 = 0$ and $\lambda_{ik} = v_{ik}$ by (37). Then, (38) takes the form,

$$\frac{d\widetilde{F}_{ik}(A^0 + \omega U^{st}; \Lambda)}{d\omega}|_{\omega=0} = \frac{v_{ik}}{n^i}(\frac{d\rho_{ik}(A^{st}(\omega))}{d\omega})|_{\omega=0}, \tag{40}$$

which after substitution from (28) specializes to the expression in (36). In order to prove (36) for $v_{st} = 0$, we consider three mutually exclusive and collectively exhaustive cases.

1) $v_{st} = 0$ and $G^s = G^t$ and $G^i \neq G^k$: This is a special case of the expression in (36) with $\Upsilon^{st} = U^{st}$ and $(\Upsilon^{st})_i^{G^k} = (\Upsilon^{st})_k^{G^i} = (\Upsilon^{st})_{G^i}^{G^k} = (\Upsilon^{st})_{G^k}^{G^i} = 0$.

2) $v_{st} = 0$ and $G^s = G^t$ and $G^i = G^k$: Then, $\frac{d\rho_{ik}(.)}{d\omega}|_{\omega=0} = 0$ by (28) and $\rho_{ik}(A^0, \varpi, K) = 1$ by (37). Hence, we obtain from (38),

$$\frac{d\widetilde{F}_{ik}(A^0 + \omega U^{st}; \Lambda)}{d\omega}|_{\omega=0} = -\frac{\lambda_{ik}}{(n^i)^2}\frac{d\eta_i(\omega, \Lambda)}{d\omega})|_{\omega=0},$$

which vanishes after substitution from (39),

$$\frac{d\eta_i(\omega, \Lambda)}{d\omega} = \textstyle\sum_{s=1}^n \lambda_{is}\frac{d\rho_{is}(.)}{d\omega} = \sum_{s:G^s=G^i} 1 \cdot 0 + \sum_{s:G^s\neq G^i} 0 \cdot \frac{d\rho_{is}(.)}{d\omega} = 0.$$

3) $v_{st} = 0$ and $G^s \neq G^t$: In this case, $\Upsilon^{st} = 0$ and

$$\frac{d}{d\omega}\widetilde{F}_{ik}(A^0 + \omega U^{st}; \Lambda)|_{\omega=0} = \frac{d}{d\omega}\frac{\lambda_{ik}\rho_{ik}(\Lambda \cdot (A^0 + \omega U^{st}))}{\sum_{s=1}^n \lambda_{is}\rho_{is}(\Lambda \cdot (A^0 + \omega U^{st}))} =$$
$$\frac{d}{d\omega}\frac{\lambda_{ik}\rho_{ik}(A^0 + \omega\Upsilon^{st})}{\sum_{s=1}^n \lambda_{is}\rho_{is}(A^0 + \omega\Upsilon^{st})} = \frac{d}{d\omega}\widetilde{F}_{ik}(A^0 + \omega\Upsilon^{st}; \Lambda)|_{\omega=0} = 0.$$

$\blacksquare$

Proof of Proposition 5: For a vector valued function f, column vectors $\mathbf{x}$, $\mathbf{u}$ and a real number ω, the first order approximation of f at $\mathbf{x}$ is computed as,

$$f(\mathbf{x}+\omega\mathbf{u}) \approx f(\mathbf{x}) + \omega\frac{\partial f(\mathbf{x})}{\partial \mathbf{x}}\mathbf{u} \Rightarrow \frac{d}{d\omega}f(\mathbf{x}+\omega\mathbf{u})|_{\omega=0} = \frac{\partial f(\mathbf{x})}{\partial \mathbf{x}}\mathbf{u}, \tag{41}$$

where $\partial f(\mathbf{x})/\partial\mathbf{x}$ is the Jacobian of f at $\mathbf{x}$. In particular, for the vector $\mathbf{u}^t$ such that $\mathbf{u}^t_t = 1$ and $\mathbf{u}^t_\nu = 0$ for all $v \neq t$,

$$\frac{d}{d\omega}f_i(\mathbf{x}+\omega\mathbf{u}^t)|_{\omega=0} = (\frac{\partial f(\mathbf{x})}{\partial \mathbf{x}}\mathbf{u}^t)_i = \frac{\partial f_i(\mathbf{x})}{\partial x_t},$$

where $\frac{\partial f}{\partial \mathbf{x}}\mathbf{u}^t$ is the t^{th} column of the Jacobian $\partial f(\mathbf{x})/\partial\mathbf{x}$. By the same token, in our context we obtain for the $n \times n$ matrix $U^{st} \equiv (u^{st}_{\tau\omega})_{\tau,\omega\in N}$ such that $u^{st}_{st} = 1$ and $u^{st}_{ik} = 0$ for all $ik \neq st$,

$$\frac{d}{d\omega}\widetilde{F}_{ik}(A^0+\omega U^{st}; L^0+V)|_{\omega=0} = \frac{\partial \widetilde{F}_{ik}(A; L^0+V)}{\partial a_{st}}|_{A=A^0} \equiv \mathcal{J}_{ik,st}, \tag{42}$$

where $ik \in n\times n$ indexes the row and $st \in n \times n$ indexes the column in the $n^2 \times n^2$ Jacobian matrix $\mathcal{J} \equiv \left(\frac{\partial \widetilde{F}_{ij}}{\partial a_{st}}\right)_{i,j,s,t\in N}$. From (42) and Lemma 3, we obtain the relevant entries in $\mathcal{J}$ and its transposed $\mathcal{J}'$ as illustrated in the tables below,

$$\mathcal{J}_{ik,st} = \left\{\begin{array}{lll} & v_{st}=0 & v_{st}=1 \\ v_{ik}=0 & 0 & ? \\ v_{ik}=1 & 0 & \geq 0 \end{array}\right\} \Rightarrow \mathcal{J}'_{ik,st} = \left\{\begin{array}{lll} & v_{st}=0 & v_{st}=1 \\ v_{ik}=0 & 0 & 0 \\ v_{ik}=1 & ? & \geq 0 \end{array}\right\} \tag{43}$$

Then, from the system of eigenvalue equations $\mathcal{J}'\mathbf{e} = \lambda\mathbf{e}$ for $\lambda \neq 0$ it follows that $e_{ik} = 0$ when $v_{ik} = 0$, where $\mathbf{e} = \{e_{ik}\}_{i,k\in N}$ is an $n^2 \times 1$ eigenvector of $\mathcal{J}'$ with entries indexed by $ik \in n \times n$. Hence, in light of (43), only elements of $\mathcal{J}'_{ik,st}$ with $v_{ik} = 1$ and $v_{st} = 1$ appear in the eigen equations for $\mathcal{J}'$. Therefore, for the computation of the eigenvalues and eigenvectors of $\mathcal{J}'$, we can think of all entries $\mathcal{J}'_{ik,st}$ as equal to zero except when $v_{ik} = 1$ and $v_{st} = 1$, in which case they are non-negative.

By the Perron-Frobenius Theorem, the largest eigenvalue of a nonnegative square matrix is real and positive and has an associated nonnegative eigenvector. Hence, for the system of eigen equations $\mathcal{J}'\mathbf{e} = \lambda\mathbf{e}$, we have that $\mathbf{e} > 0$ if $\lambda > 0$ is the Perron-Frobenius eigenvalue of $\mathcal{J}'$ (and, hence, of $\mathcal{J}$). Then, we can write the sum of the eigen equations as,

$$\textstyle\sum_{ik:v_{ik}=1} e_{ik} \sum_{st:v_{st}=1} \mathcal{J}'_{st,ik} = \lambda \sum_{ik:v_{ik}=1} e_{ik} > 0. \tag{44}$$

Dividing (44) by $\sum_{ik:v_{ik}=1} e_{ik}$ shows that λ is a convex combination of the values in the set

$$\{\textstyle\sum_{st:v_{st}=1}\mathcal{J}'_{st,ik}\}_{ik:v_{ik}=1} = \{\sum_{st:v_{st}=1}\mathcal{J}_{ik,st}\}_{ik:v_{ik}=1}.$$

For each ik such that $v_{ik} = 1$, we compute $\sum_{st:v_{st}=1} \mathcal{J}_{ik,st}$ by substituting for $\mathcal{J}_{ik,st}$ from (36) with $\Upsilon^{st} \equiv (L^0 + V) \cdot U^{st} = U^{st}$ as $v_{st} = 1$ (and, hence, $G^s \neq G^t$),

$$\sum_{st:v_{st}=1} \frac{n^i \varpi^k (U^{st})_i^{G^k} + n^k \varpi^i (U^{st})_k^{G^i} + (K-1)(\varpi^k (U^{st})_{G^i}^{G^k} + \varpi^i (U^{st})_{G^k}^{G^i})}{n^i \sqrt{\varpi^i \varpi^k n^i n^k}}$$
$$= \frac{n^i \varpi^k V_i^{G^k} + n^j \varpi^i V_k^{G^i} + (K-1)(\varpi^k V_{G^i}^{G^k} + \varpi^i V_{G^k}^{G^i})}{n^i \sqrt{\varpi^i \varpi^k n^i n^k}}.$$

■

References

[1] Adamic, L. and N. Glance (2005): "The Political Blogosphere and the 2004 U.S. Election: Divided They Blog," *LinkKDD'05, Proceedings of the 3rd International Workshop on Link Discovery*, 36-43.

[2] Alexander C, Piazza M, Mekos D, Valente T. (2001): "Peers, schools, and adolescent cigarette smoking," *Journal of Adolescent Health* **29**(1):22-30.

[3] Aral, S. and M.Van Alstyne (2011): "The Diversity-Bandwidth Trade-off," *American Journal of Sociology* **117**, 90-171.

[4] Bala, V. and S. Goyal (1998): "Learning from Neighbors," *Review of Economic Studies* **65**, 595-621.

[5] Bishop, B. (2008): *The Big Sort: Why the Clustering of Like-Minded America is Tearing Us Apart*, New York: Mariner Books.

[6] Bobadilla, J., F. Ortega, A. Hernando, and J. Bernal (2012): " A Collaborative filtering approach to mitigate the new user cold start problem", *Knowledge Based Systems* **26**, 225-238.

[7] Boutyline A. and R. Willer (2017): "The Social Structure of Political Echo Chambers: Ideology and Political Homophily in Online Communication Networks", *Political Psychology* **38**, 551-569.

[8] Cohen JM. (1977): "Sources of peer group homogeneity", *Sociology of Education* **50**(4):227-241.

[9] DeGroot, M. H. (1974): "Reaching a Consensus", *Journal of the American Statistical Association* **69**, 118-121.

[10] DeMarzo, P.M., D. Vayanos and J. Zwiebel (2003): "Persuasion Bias, Social Influence, and Unidimensional Opinions", *The Quarterly Journal of Economics*, **118**, 909-968.

[11] Easley, D. and J. Kleinberg (2010): “Networks, Crowds, and Markets: Reasoning About a Highly Connected World”, Cambridge: Cambridge University Press.

[12] Gale, D. and S. Kariv (2003): “Bayesian learning in social networks,” *Games and Economic Behavior* **45**, 329-346.

[13] Golman, R., G. Loewenstein, K. O. Moene, and L. Zarri (2016): “The Preference for Belief Consonance,” *Journal of Economic Perspectives* **30**, 165-188.

[14] Golub, B. and M. O. Jackson (2010): “Naïve Learning in Social Networks: Convergence, Influence, and the Wisdom of Crowds”, *American Economic Journal: Microeconomics* **2**, 112-149.

[15] Golub, B. and M. O. Jackson (2012): “How Homophily Affects the Speed of Learning and Best Response Dynamics,” *Quarterly Journal of Economics* **127**, 1287-1338.

[16] Granovetter, M. (1973): “The Strength of Weak Ties”, *American Journal of Sociology* **78**, 1360-1380.

[17] Granovetter, M. (1983): “The Strength of Weak Ties: A Network Theory Revisited”, *Sociological Theory*, **1**, 201-233.

[18] Jackson, M.O., T. Rodríguez-Barraquer and X. Tan (2012): “Social Capital and Social Quilts: Network Patterns of Favor Exchange”, *American Economic Review* **102**, 1857-1897.

[19] Jannach, D., M. Zanker, A. Felfernig, and G. Friedrich (2010): *Recommender Systems: An Introduction*, Cambridge: Cambridge University Press.

[20] Kandori, M., G. Mailath, and R. Rob (1993): “Learning, Mutation, and Long-run Equilibria in Games”, *Econometrica* **61**, 29-56.

[21] Kandel DB (1978): “Homophily, selection, and socialization in adolescent friendships”, *American Journal of Sociology* **84**(2):427-436.

[22] Knecht A, Snijders TAB, Baerveldt C, Steglich CEG, Raub W. (2010), “Friendship and Delinquency: Selection and Influence Processes in Early Adolescence”, *Social Development* **19**(3):494-514.

[23] Kossinets, G., Watts, D.J., (2006): “Empirical Analysis of an Evolving Social Network”, *Science* **311** (5757): 88-90.

[24] Kossinets, G., Watts, D.J., (2009): “Origins of Homophily in an Evolving Social Network,” *American Journal of Sociology* **115**(2), 405-450

[25] Kumpula, J.M., Onnela, J.P., Saramaki, J., Kaski, K., Kertesz, J. (2007): “Emergence of communities in weighted networks”, *Physical Review Letters* **99**, 228701.

[26] Lazarsfeld, P.F. and R.K. Merton (1954): "Friendship as a social process: a substanive and methodological analysis", in M. Berger (ed.), Freedom and Control in Modern Society, New York: Van Nostrand.

[27] Lobel I. and E. Sadler (2013): "Preferences, Homophily, and Social Learning", mimeo.

[28] Marsden, P.V. (1987): "Core Discussion Networks of Americans", *American Sociological Review* **52**, 122-313.

[29] Marsden, P.V. (1988): "Homogeneity in Confiding Relations", *Social Networks* **10**, 57-76.

[30] McPherson, M., L. Smith-Lovin and J. M. Cook (2001): "Birds of a Feather: Homophily in Social Networks", *Annual Review Sociology* **27**, 415-44.

[31] Moody, J. (2001): "Race, School Integration, and Friendship Segregation in America", *The American Journal of Sociology* **107**(3), 679-716.

[32] Mueller-Frank, M. (2014): "Does one Bayesian make a difference?," *Journal of Economic Theory* **154**, 423-452.

[33] Rajaraman, A., J. Leskovec, and J.D. Ullman (2012): *Mining of Massive Datasets*, Cambridge: Cambridge University Press.

[34] Reagans, R., and B. McEvily (2003): "Network Structure and Knowledge Transfer: The Effects of Cohesion and Range", *Administrative Science Quarterly* **48**, 240-67.

[35] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, J. Riedl (1994): "GroupLens: an open architecture for collaborative filtering of netnews," in *Proceeding of the ACM Conference on Computer Supported Cooperative Work*, pp. 175-186.

[36] Ricci, F., L. Rokach, and B. Shapira (2011): "Introduction to Recommender Systems Handbook", *Recommender Systems Handbook*, Springer, pp. 1-35.

[37] Uzzi, B. (1997): "Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness", *Administrative Science Quarterly* **42**, 35-67.

[38] Young, P. (1993): "The Evolution of Conventions", *Econometrica* **61**, 29-56.