

HOW HOMOPHILY AFFECTS THE SPEED OF LEARNING AND BEST-RESPONSE DYNAMICS

BENJAMIN GOLUB
MATTHEW O. JACKSON

We examine how the speed of learning and best-response processes depends on homophily: the tendency of agents to associate disproportionately with those having similar traits. When agents' beliefs or behaviors are developed by averaging what they see among their neighbors, then convergence to a consensus is slowed by the presence of homophily but is not influenced by network density (in contrast to other network processes that depend on shortest paths). In deriving these results, we propose a new, general measure of homophily based on the relative frequencies of interactions among different groups. An application to communication in a society before a vote shows how the time it takes for the vote to correctly aggregate information depends on the homophily and the initial information distribution. *JEL* Codes: D83, D85, I21, J15, Z13

I. INTRODUCTION

There are pronounced disagreements in society about factual issues. In October 2004, 47% of Republican poll respondents believed that Iraq had weapons of mass destruction just before the 2003 invasion of that country, as opposed to only 9% of Democrats. These disagreements can be highly persistent. Sixteen months later, in March 2006, the percentages had changed to 41% for Republicans and 7% for Democrats.¹ This kind of disagreement occurs on many other important factual questions—for instance, whether temperatures on Earth are increasing over time.²

Earlier versions of this article were titled “How Homophily Affects the Speed of Contagion, Best Response and Learning Dynamics.” Jackson gratefully acknowledges financial support from the NSF under grants SES–0647867 and SES–0961481. Golub gratefully acknowledges financial support from an NSF Graduate Research Fellowship, as well as the Clum, Holt, and Jaedicke fellowships at the Stanford Graduate School of Business. We thank S. Nageeb Ali, Arun Chandrasekhar, Christoph Kuzmics, Carlos Lever, Yair Livne, Norma Olaiola Ortega, Amin Saberi, Federico Valenciano, as well as Elhanan Helpman, Larry Katz, and the anonymous referees for comments and suggestions. Finally, we thank Cyrus Aghamolla, Andres Drenik, Lev Piarsky, and Yiqing Xing for careful readings of earlier drafts and suggestions on the exposition.

1. The polling data appear in “Iraq: The Separate Realities of Republicans and Democrats” World Public Opinion (2006).

2. This is documented in “Little Consensus on Global Warming: Partisanship Drives Opinion” Pew Research Center (2006). We emphasize that these

© The Author(s) 2012. Published by Oxford University Press, on behalf of President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2012), 1287–1338. doi:10.1093/qje/qjs021.

Advance Access publication on May 7, 2012.

How long does it take for beliefs to reach a consensus about issues of broad interest? What determines the extent of disagreement? Why might consensus be reached among certain subgroups of a population much more quickly than among a population as a whole? Understanding convergence times can help us understand if and when we should expect a consensus to be reached, and whether a society's beliefs should settle down quickly or continue to shift for substantial amounts of time.

The answers to these questions lie partly in the networks of relationships that are critical determinants of how people update their beliefs and how they choose their behaviors.³ In this article we examine how the speed of convergence of agents' behaviors and beliefs depends on network structure in a model that is rich enough to capture the segregation patterns that are pervasive in social networks.

Although social networks are naturally complex, they nonetheless exhibit fundamental patterns and regularities. We focus on the impact of two of the most fundamental aspects of network architecture: homophily and link density. Link density refers to a measure of the number of relationships per capita in a society. Homophily, a term coined by Lazarsfeld and Merton (1954), refers to the tendency of individuals to associate disproportionately with others who are similar to themselves. Indeed, homophily is one of the most pervasive and robust tendencies of the way people relate to each other (see McPherson, Smith-Lovin, and Cook 2001 for a survey).

Although homophily has been documented across a wide array of different characteristics, including ethnicity, age, profession, and religion, there is little modeling of how homophily affects behavior. Intuitively, segregation patterns in a network are very important for processes of behavior updating, learning, and diffusion, so it is essential to develop models of homophily's effects. One example of this is from Rosenblat and Möbius

disagreements concern purely factual questions, not policy matters. Differences in initial information, which could be related to attempts to justify policy positions, can lead to long-standing disagreements about facts, as we shall see in our analysis.

3. In our analysis, we abstract away from preference differences over policies and focus directly on updating over information without any incentives to distort or manipulate the spread of information. Of course, it is difficult to separate any factual question from its policy implications, and policy preferences might play some role in the answer to these questions. Nonetheless, network structure also plays a key role, so we focus on isolating that issue.

(2004), who show that if agents' preferences are based on a (one-time) weighted average of taste parameters of their neighbors, then the resulting preferences can depend on group separation.⁴ However, although there is a literature on how agents responding to neighbors' behaviors converge to equilibrium, and it is known that network structure matters (e.g., see the survey by Jackson and Yariv 2011), there has been no work relating homophily to the speed of convergence. We address this gap by: (1) working with a model of networks that captures both homophily and link density; and (2) studying how these properties affect simple but important benchmark updating processes that are relevant in economic settings. There turns out to be a clean relationship between the convergence speeds of updating processes and the structure of the social networks on which they operate, and we characterize that dependence.

The model of networks that we study, which we refer to as the *multi-type random network* model, allows there to be an arbitrary number of groups making up arbitrary fractions of society. The probability of a link between two nodes depends on which groups they are in. Thus, the model is an extension of classical random graph models that allows for arbitrary heterogeneity in types and allows us to tune two fundamental network characteristics: link density and homophily.⁵

Using the multi-type random network model as a base, we focus on a simple updating process, called *average-based updating*, in which agents set their next-period behaviors or beliefs based on the average choices of their peers—as in standard peer effects models, with a social network structure defining the peer relationships (for example, see Calvó-Armengol, Patacchini, and Zenou 2009 and Bramoullé, Djebbari, and Fortin 2009). This type of process is relevant in a variety of

4. See their Section III.C. Other previous work in economics focuses on models of homophily's origins (Currarini, Jackson, and Pin 2009, 2010; Bramoullé et al. 2012) and rigorous foundations for measuring the extent of segregation (Echenique and Fryer 2007). DeMarzo, Vayanos, and Zwiebel (2003) discuss some aspects of how segregation can affect the updating process they study, but do not formally model linking behavior as being affected by group membership.

5. Variations of such random graph models appear in the stochastic block modeling literature (e.g., Fienberg, Meyer, and Wasserman 1985 and Holland, Laskey, and Leinhardt 1983) and the community detection literature (e.g., Copic, Jackson and Kirman 2009). The focus in those literatures is on fitting and the estimation of the underlying (latent) types, not dynamic processes occurring on networks.

applications. A simple one is a coordination setting in which an agent finds it optimal to match the average behavior of his or her social neighbors and best responds myopically to their last-period choices. Another updating process in this class is the classic DeGroot (1974) model of learning and consensus in which one's opinion tomorrow about some unknown quantity is an average of the opinions of one's friends today. The end of Section II.D includes some evidence, both experimental and theoretical, on why average-based updating is a useful benchmark.⁶

Our focus is on how long it takes a society to reach a consensus or equilibrium via an average-based updating process—and, in particular, how this depends on homophily. Although we provide results in the general model, the way network structure affects convergence speed is most easily seen within the context of a special kind of random network that we define in Section II.C, called the *equal-sized islands model*. In that model, agents come in several different types, with an equal number of agents of each type. Agents link to other agents of the same type with a probability that is different (typically higher) than the probability with which they link to agents of other types. Homophily is defined as the difference of these two probabilities, normalized by a measure of the overall linking probability (see equation [2] for a formula). In this setting, the time it takes for average-based updating processes to converge is increasing and convex in homophily and proportional to the logarithm of population size, but it is essentially unaffected by any change in overall link density (holding homophily and population size fixed) as long as the density exceeds a low threshold.

This special case captures the essence of what the article aims to do: operationally define homophily and show how it affects convergence time. To extend this exercise beyond the simple structure of the islands model, we introduce a general measure of homophily that works in any multi-type random network. This measure is called *spectral homophily*. It is defined by considering

6. One could model the updating of beliefs regarding, say, weapons of mass destruction in Iraq using a standard Bayesian learning framework. But because of the extreme complexity of the relevant Bayesian calculations, and the fact that the rational models tend to predict fast convergence to consensus (Parikh and Krasucki 1990; DeMarzo, Vayanos and Zwiebel 2003; Mossel and Tamuz 2010), we do not believe this is a particularly appropriate model either in its predictions or its mechanics for the types of applications that concern us here.

a matrix that captures relative densities of links between various pairs of groups and taking its second-largest eigenvalue.

This article makes two main contributions. The first is to introduce spectral homophily and argue that this quantity actually measures homophily—that it coincides with intuitive notions of homophily and segregation. We do this by showing that spectral homophily specializes to obvious “hands-on” (i.e., eigenvalue-free) measures of homophily in important classes of networks and by proving an explicit interpretation of spectral homophily in terms of segregation patterns or “fault lines” in the group structure. The second contribution is to show how spectral homophily affects the convergence time of average-based updating processes. Indeed, our main theorem generalizes the islands result: for average-based updating processes, the time to converge is increasing and convex in spectral homophily; it is directly proportional to the logarithm of population size; but it is essentially unaffected by link density as long as the density exceeds a low threshold.

An intuition for the results is that in the average-based updating processes that we examine, agents are influenced by their acquaintances, and the relative weights on different types of agents affects an agent’s interim belief. If one doubles connectivity without changing the relative weights of interactions across groups, the influence on a typical agent’s beliefs from different groups is unaltered, and so the speed to convergence is unaltered. In contrast, rearranging the relative weights between groups affects the speed of convergence. Effectively, under nontrivial homophily, within-group convergence is relatively fast, and the main impediment to convergence is the process of reaching consensus across groups. The more integrated groups become, the faster overall convergence to a consensus becomes.

One of our main innovations is to relate the rate of convergence to imbalances of interactions at the level of *groups* and to show that the specific interactions among agents can be ignored. This is important, as it allows one to work with a very reduced form rather than considering many details of the interaction structure, which can be especially difficult to measure or even keep track of in large societies. This also distinguishes our approach from previous work on the convergence of various dynamic processes, in which it is known that the second-largest eigenvalue of networks is important for the convergence rate (see, e.g., Montenegro and Tetali 2006 for results on Markov

processes). By showing that the specifics of the full network can be neglected, we characterize convergence rates in terms of group-level homophily. Indeed, this perspective reduces the study of certain processes in large random networks to the study of much simpler aggregate statistics that measure interactions between groups. As demonstrated, such a simplification allows for clean characterizations of convergence times that are not available otherwise (see Section III.B for a concrete illustration). This approach also has implications for statistical work. Because macroscopic, group-level statistics are much more easily obtained than data on the full network structure (exactly who links with whom), our results can be used to simplify the work of an econometrician who is using network data and justify this simplification rigorously.

To provide some context for the results regarding the convergence time of average-based updating, we briefly compare it to the analogous convergence time for a simple contagion process. In the simplest contagion process, a node becomes infected or activated immediately after any of its neighbors are infected. Because one neighbor suffices to transmit the contagion, the relative weights on different types are no longer critical, and instead overall link density becomes the main determinant of convergence speed. This contrast provides more insight into both types of processes and makes it clear that the role of network structure in determining convergence speed depends in intuitive ways on the type of dynamic process.

As an application of the results and intuitions, we examine a simple model of voting on a policy in a homophilous society whose members see signals that are correlated with a true state of nature. To keep the setting simple, we abstract away from different individual preferences and focus on heterogeneity in information and interaction patterns. If agents knew the true state of the world—for example, whether there is a threat to their national security—all of them would vote the same way. The agents, however, are uncertain about the true state—they start with different signals about it—and they have a chance to communicate and update their beliefs before voting. The two distortions in the society come from homophily between the groups and an initial bias in the distribution of information throughout the society—so that there is correlation between an agent's type (say, demographic group) and the signal seen by that agent. The problematic case is one in which the signals that are in favor of the truly better policy

are held mainly by the minority group, while the majority group happens to have an initial bias of signals in favor of the worse policy. We study how long it takes for such a society to overcome the homophily and initial bias in the distribution of information and vote “correctly.” If the society votes before any communication takes place, then all agents vote their initial information and the majority vote is correct (because there are more correct signals than incorrect ones in the society). After communication starts, each group aggregates the information within it and, for a while, the majority group is inclined to vote for the wrong policy. It then takes time for communication across groups to overcome the homophily and lead to a correct overall vote. We show that the time to a correct vote is proportional to an intuitive quantity that is increasing in the homophily in the society, corresponding precisely to our general measure of homophily specialized to this context. Indeed, the time to a correct vote is proportional to our “consensus” time, incorporating our measure of homophily weighted by a factor that captures the bias in the initial information.

The article is organized as follows. Section II lays out the model of networks, our measure of homophily, and the updating process we focus on. Section III presents the main results on how homophily affects the speed of average-based updating processes. Section IV strengthens the basic results to characterize large-scale dynamics of updating and analyze the sources of long-run disagreement. Section V applies the model to a simple voting setting. Section VI relates spectral homophily to more basic, “hands-on” network segregation measures. Section VII contrasts an average-based process with a contagion that travels along shortest paths in terms of how each is affected by homophily. Section VIII concludes. All proofs are in the Online Appendix.

II. THE MODEL: NETWORKS, HOMOPHILY, AND LEARNING

We work with a simple model of network structure that generalizes many random network models and is well-suited for identifying the relative roles of network density and homophily.

II.A. Multi-type Random Networks

Given a set of n nodes $N = \{1, \dots, n\}$, a network is represented via its adjacency matrix: a symmetric n -by- n matrix \mathbf{A} with

entries in $\{0, 1\}$. The interpretation is that $A_{ij} = A_{ji} = 1$ indicates that nodes i and j are linked, and the symmetry restricts attention to undirected networks.⁷ Let $d_i(\mathbf{A}) = \sum_{j=1}^n A_{ij}$ denote the degree (number of links) of node i , the basic measure of how connected a node is. Let $d(\mathbf{A})$ denote average degree. Finally, let

$$D(\mathbf{A}) = \sum_i d_i(\mathbf{A})$$

be the sum of all degrees in the network, which is twice the total number of links.

Agents or nodes have “types,” which are the distinguishing features that affect their propensities to connect to each other. Types might be based on any characteristics that influence agents’ probabilities of linking to each other, including age, race, gender, profession, education level, and even behaviors.⁸ For instance, a type might consist of the 18-year-old female African Americans who have completed high school, live in a particular neighborhood, and do not smoke. The model is quite general in that a type can embody arbitrary lists of characteristics; which characteristics are included will depend on the application. There are m different types in the society. Let $N_k \subset N$ denote the nodes of type k , so the society is partitioned into the m groups, (N_1, \dots, N_m) . Let $n_k = |N_k|$ denote the size of group k , $\mathbf{n} = (n_1, \dots, n_m)$ be the corresponding vector of cardinalities, and n denote the total number of agents.

A *multi-type random network* is defined by the cardinality vector \mathbf{n} together with a symmetric m -by- m matrix \mathbf{P} , whose entries (in $[0, 1]$) describe the probabilities of links between various types.⁹ The resulting random network is captured via its adjacency matrix, which is denoted by $\mathbf{A}(\mathbf{P}, \mathbf{n})$.¹⁰ In particular,

7. Although we conjecture that our results can be extended to directed networks without much change in the statements, some of our proof techniques take advantage of the symmetry of the adjacency matrix, so we are not sure of the modifications that might be needed in examining directed networks.

8. However, we do not allow types to depend on behaviors or beliefs that are endogenous to the model, leaving this interesting potential extension for future work.

9. We assume a numbering of agents such that $N_1 = \{1, 2, \dots, n_1\}$, $N_2 = \{n_1 + 1, n_1 + 2, \dots, n_1 + n_2\}$, and so on. Given this convention, it is possible to recover the partition from just the vector \mathbf{n} .

10. The modeling of the network structure as random with certain probabilities amounts to assuming that the large-scale structure of the network is exogenous to the updating process being studied. This kind of stochastic network with

$\mathbf{A}(\mathbf{P}, \mathbf{n})$ is built by letting the entries A_{ij} with $i > j$ be independent Bernoulli random variables that take on a value of 1 with probability $P_{k\ell}$ if $i \in N_k$ and $j \in N_\ell$. That is, the entry $P_{k\ell}$ captures the probability that an agent of type k links to an agent of type ℓ . We fill in the remaining entries of $\mathbf{A}(\mathbf{P}, \mathbf{n})$ by symmetry: $A_{ij} = A_{ji}$. We set $A_{ii} = 0$ for each i .¹¹ Unless otherwise noted, we use $\mathbf{A}(\mathbf{P}, \mathbf{n})$ to denote a random matrix, and \mathbf{A} without an argument to refer to a given deterministic matrix.¹²

The multi-type random model subsumes many other random networks. The seminal random network model of Erdős and Rényi is a special case, as are many cases of the model based on degree distributions of Chung and Lu (2002).¹³ One can also view the probabilities in the matrix \mathbf{P} as arising from distances between some underlying locations, either physical or abstract. One can even include different sociabilities, so that some groups consist of individuals who, for whatever reasons, form more relationships on average than others. Thus, it need not be that all nodes have the same expected number of connections; the network can have a nontrivial degree distribution.¹⁴

type-dependent heterogeneity can arise from a strategic friendship formation process with some randomness in the order of meetings, as in Currarini, Jackson, and Pin (2009).

11. Under the assumptions of our main results (see Definition 3), the specification of the diagonal does not make a difference: one can forbid self-links, require them, or do anything in between. The proofs in the Online Appendix go through regardless: a single link for each agent does not significantly affect the aggregate dynamics of the processes we study.

12. For individual entries, we drop the arguments (\mathbf{P}, \mathbf{n}) , but the matrix in question will be clear from context.

13. This can also be seen as a cousin of some of the statistical models that have been used to capture homophily in networks, such as various \mathbf{p}^* and exponential random graph models (e.g., see the references and discussion in Jackson 2008b). There are variations on it in the computer science literature called planted multi-section models, for example McSherry (2001). An early version of this type of model was introduced by Diaconis and Freedman (1981) in a study on the psychology of vision, independently of its introduction in the stochastic block modeling literature mentioned earlier.

14. There are, of course, networks which this approach is not well suited for modeling: strict hierarchies, regular lattices, and so on, even though homophily can and does occur in these networks.

II.B. A General Measure of Homophily

We now provide a general definition of homophily based on the probabilities of interaction between various types, and then show how it works in an important special case.

Let $Q_{k\ell}(\mathbf{P}, \mathbf{n}) = n_k n_\ell P_{k\ell}$ be the expected total contribution to the degrees of agents of type k from links with agents of type ℓ ; when $k \neq \ell$, this is simply the expected number of links between k and ℓ . Also, let $d_k[\mathbf{Q}(\mathbf{P}, \mathbf{n})] = \sum_\ell Q_{k\ell}(\mathbf{P}, \mathbf{n})$ be the expected total degree of nodes of type k .

Let $\mathbf{F}(\mathbf{P}, \mathbf{n})$ be a matrix of the same dimensions as \mathbf{P} with entries (throughout we take $0/0 = 0$)

$$F_{k\ell}(\mathbf{P}, \mathbf{n}) = \frac{Q_{k\ell}(\mathbf{P}, \mathbf{n})}{d_k[\mathbf{Q}(\mathbf{P}, \mathbf{n})]}.$$

Thus, $F_{k\ell}$ is the expected fraction of their links that nodes of type k will have with nodes of type ℓ . This simplifies things in two respects relative to the realized random network. First, it works with groups (or representative agents of each type) rather than individual nodes; second, it works with ratios of expected numbers of links rather than realized numbers of links. With this matrix defined, we can formulate a general homophily measure.

DEFINITION 1. The *spectral homophily* of a multi-type random network (\mathbf{P}, \mathbf{n}) is the second-largest¹⁵ eigenvalue of $\mathbf{F}(\mathbf{P}, \mathbf{n})$. We denote it by $h^{\text{spec}}(\mathbf{P}, \mathbf{n})$.

The spectral homophily measure is based on first simplifying the overall interaction matrix to that of the expected interaction across groups, and then looking at a particular part of the spectrum of that matrix: the second-largest eigenvalue. On an intuitive level, a second-largest eigenvalue captures the extent to which a matrix can be broken into two blocks with relatively little interaction across the blocks. Indeed, in Section VI, we present a formal result showing that spectral homophily picks

15. It is easily checked, analogously to Fact 1 in Online Appendix 1, that the matrix $\mathbf{F}(\mathbf{P}, \mathbf{n})$ is similar to the symmetric matrix with entries $\frac{Q_{k\ell}(\mathbf{P}, \mathbf{n})}{\sqrt{d_k[\mathbf{Q}(\mathbf{P}, \mathbf{n})]d_\ell[\mathbf{Q}(\mathbf{P}, \mathbf{n})]}}$, and so all the eigenvalues of \mathbf{F} are real. To define the second-largest eigenvalue, list the eigenvalues of this matrix ordered from greatest to least by absolute value, with positive eigenvalues listed first if there are ties. An eigenvalue is listed the same number of times as the number of linearly independent eigenvectors it has. The second eigenvalue in the list is called the second-largest eigenvalue.




FIGURE 1

Islands networks with low and high homophily are shown in (a) and (b), respectively. Nodes that are shaded differently are of distinct types.

up “fault lines” created by segregation in the network. Here, we illustrate this in the context of a special case.

II.C. A Special Case: The Islands Model

For an illustration of the general definitions, it is useful to consider a situation in which groups are equal-sized and all biased in the same way. In particular, links within a type are more probable than links across types, and the probability of those across types does not depend on the specifics of the types in question. We call this case the “islands” model.

More precisely, the *islands model* is the special case of the multi-type random networks model such that (1) each type (group) has the same number of agents; and (2) an agent only distinguishes between agents of one’s own type and agents of a different type. Moreover, all agents are symmetric in how they do this. Formally, in the multi-type random network notation, we say the multi-type random network (\mathbf{P}, \mathbf{n}) is an islands network with parameters (m, p_s, p_d) if:

- there are m islands and their sizes, n_k , are equal for all k ;
- $P_{kk} = p_s$ for all k ; and
- $P_{k\ell} = p_d$ for all $k \neq \ell$, where $p_d \leq p_s$ and $p_s > 0$.

The idea that agents only distinguish between “same” and “different” agents in terms of the linking probabilities is surprisingly accurate as a description of some friendship patterns (e.g., see Marsden 1987 and note 7 in McPherson, Smith-Lovin, and Cook 2001).

Figure I depicts two different random networks generated by the islands model, with different linking probabilities.

In the context of the islands model, it is easy to define homophily. Let

$$(1) \quad p = \frac{p_s + (m - 1)p_d}{m}$$

be the average linking probability¹⁶ in the islands model (i.e., the probability that two agents drawn uniformly at random are linked).

One natural measure of homophily then compares the difference between same and different linking probabilities to the average linking probability, with a normalization of dividing by the number of islands, m :

$$(2) \quad h^{\text{islands}}(m, p_s, p_d) = \frac{p_s - p_d}{mp}.$$

Note that this is equivalent to Coleman’s (1958) homophily index specialized to the islands model:

$$\frac{\frac{p_s}{mp} - \frac{1}{m}}{1 - \frac{1}{m}}.$$

This is a measure of how much a group’s fraction of same-type links ($\frac{p_s}{mp}$) exceeds its population share ($\frac{1}{m}$), compared to how big this difference could be ($1 - \frac{1}{m}$).¹⁷

The measure h^{islands} captures how much more probable a link to a node of one’s own type is than a link to a node of any other type, and varies between 0 and 1, presuming that $p_s \geq p_d$. If a node only links to same-type nodes (so that $p_d = 0$), then the average linking probability p becomes $\frac{p_s}{m}$ and so $h^{\text{islands}} = 1$, while if nodes do not pay attention to type when linking, then $p_s = p_d$ and $h^{\text{islands}} = 0$. Indeed, the purpose of the m in the denominator

16. This calculation is approximate: A typical agent is viewed as having $(m - 1)$ times as many potential intergroup links as potential intragroup links. In reality, because an agent cannot have a link to himself (but *can* have a link to anyone outside his group), this ratio is slightly different. The error of the approximation vanishes as the islands grow large.

17. To see this, note that Coleman’s index can be rewritten as $\frac{p_s - p_d}{(m - 1)p}$, and then from (1) it follows that $p_s - p_d = \frac{m - 1}{m}(p_s - p_d)$; substituting verifies the equivalence.

is to restrict the variation of the measure exactly to the interval $[0, 1]$ (under the assumption that $p_d \leq p_s$).

This simple measure of homophily introduced is equal to the spectral homophily.

PROPOSITION 1. If (\mathbf{P}, \mathbf{n}) is an islands network with parameters (m, p_s, p_d) , then

$$h^{\text{islands}}(m, p_s, p_d) = h^{\text{spec}}(\mathbf{P}, \mathbf{n}).$$

II.D. The Average-Based Updating Processes

The processes we focus on are ones where agents' behaviors or beliefs depend on some average of their neighbors' behaviors or beliefs. Applications include those where agents dynamically and myopically best respond, trying to match the average behavior of their neighbors, as in common specifications of peer effects models. This also includes belief updating rules based on a model of updating and consensus-reaching that was first discussed by French (1956) and Harary (1959), and later articulated in its general form by DeGroot (1974).

Definition Average-based updating processes are described as follows. Given a network \mathbf{A} , let $\mathbf{T}(\mathbf{A})$ be defined by $T_{ij}(\mathbf{A}) = \frac{A_{ij}}{d_i(\mathbf{A})}$. Beginning with an initial vector of behaviors or beliefs $\mathbf{b}(0) \in [0, 1]^n$, agent i 's choice at date t is simply

$$b_i(t) = \sum_j T_{ij}(\mathbf{A})b_j(t - 1).$$

That is, the agent matches the average of his or her neighbors' last-period choices. In matrix form, this is written as:

$$\mathbf{b}(t) = \mathbf{T}(\mathbf{A})\mathbf{b}(t - 1)$$

for $t \geq 1$. It follows that

$$\mathbf{b}(t) = \mathbf{T}(\mathbf{A})^t \mathbf{b}(0).$$

The process is illustrated in Figure II.

In Online Appendix 2, we examine a variation of the model in which agents always put some weight on their initial beliefs:

$$\mathbf{b}(t) = (1 - \alpha)\mathbf{b}(0) + \alpha\mathbf{T}(\mathbf{A})\mathbf{b}(t - 1).$$




FIGURE II

An illustration of the average-based updating process in which beliefs (depicted by the shading of the boxes) range from 0 (white) to 1 (black). Some nodes begin with beliefs of 0 and others begin with beliefs of 1 (Panel A). Over time, as nodes average the information of their neighbors (Panels B and C), the process converges to a situation where nodes' beliefs are shades of gray (Panel D).

Although in such a model a consensus is not reached, our results about speed of convergence have direct analogs there. Therefore, we focus on the case where $\alpha = 1$ in the main text and refer the interested reader to Online Appendix 2 for the analogous statements in the other context.

Interpretations One simple interpretation of this process is as a myopic best-response updating in a pure coordination game. For example, suppose the agents have utility functions

$$u_i(\mathbf{b}) = - \sum_j \frac{A_{ij}}{d_i(\mathbf{A})} (b_j - b_i)^2,$$

with one interpretation being that agents receive disutility from interacting with neighbors making different choices (e.g., about language or information technology). The Nash equilibria clearly consist of strategy profiles such that all agents in a component

choose the same behavior. As we will discuss, under mild conditions, myopic best responses lead to equilibrium and we analyze the speed of convergence.

In a different interpretation based on the updating of beliefs, each agent begins with some belief $b_i(0)$ at time 0. If the initial beliefs $b_i(0)$, as i ranges over N , are independent and identically distributed draws from normal distributions around a common mean, then the linear updating rule at $t=1$ corresponds to Bayesian updating for estimating the true mean, as discussed by DeMarzo, Vayanos, and Zwiebel (2003).¹⁸ The behavioral aspect of the model concerns times after the first round of updating. After the first period, a Bayesian agent would adjust the updating rule to account for the network structure and the differences in the precision of information that other agents might have learned. However, due to the complexity of the Bayesian calculation, the DeGroot process assumes that agents continue using the simple averaging rule in later periods as well.¹⁹ DeMarzo, Vayanos, and Zwiebel (2003) argue that continuing to update according to the same rule can be seen as a boundedly rational heuristic that is consistent with psychological theories of persuasion bias. It is important to note that agents do learn new things by continuing to update, as information diffuses through the network as agents talk to neighbors, who talk to other neighbors, and so forth.

Some further remarks on the averaging model and its empirical relevance, as well as theoretical properties, can be found shortly.

Convergence As long as the network is connected and satisfies a mild technical condition, the process will converge to a limit.²⁰

18. Under the coordination game interpretation, the model involves no self-weight ($A_{ii} = 0$), while in the belief-updating interpretation, $A_{ii} = 1$ is natural. It turns out that for our asymptotic results, there is no change that arises from introducing either assumption to the multi-type random graph framework. This is because (as can be seen via calculations very similar to those in the proof of Theorem 2 in Online Appendix 1) the spectral norm of the difference between the updating matrices under these different assumptions tends to 0 as n grows large.

19. Although Bayesian updating can be complicated in general, there are settings where results can be deduced about the convergence of Bayesian posteriors in social networks, such as those studied by Acemoglu et al. (2011) and by Mueller-Frank (2012).

20. If the communication network is directed, then convergence requires some aperiodicity in the cycles of the network and works with a different segmentation

In the results about random networks, the assumptions ensure that the networks satisfy the assumptions of the lemma with a probability tending to one as n grows.²¹

LEMMA 1. If \mathbf{A} is connected and has at least one cycle of odd length (this is trivially satisfied if some node has a self-link), then $\mathbf{T}(\mathbf{A})^t$ converges to a limit $\mathbf{T}(\mathbf{A})^\infty$ such that $(\mathbf{T}(\mathbf{A})^\infty)_{ij} = \frac{d_j(\mathbf{A})}{D(\mathbf{A})}$.

Lemma 1 follows from standard results on Markov chains²² and implies that for any given initial vector of beliefs $\mathbf{b}(0)$, all agents' behaviors or beliefs converge to an equilibrium in which consensus obtains. That is:

$$\lim_{t \rightarrow \infty} \mathbf{b}(t) = \mathbf{T}(\mathbf{A})^\infty \mathbf{b}(0) = (b, b, \dots, b) \quad \text{where} \quad b = \sum_j b_j(0) \cdot \frac{d_j(\mathbf{A})}{D(\mathbf{A})}. \quad (3)$$

Therefore, the relative influence that an agent has over the final behavior or belief is his or her relative degree. The rough intuition behind convergence is fairly straightforward. With a connected network, some of the agents who hold the most extreme views must be interacting with some who are more moderate, and so the updating reduces the extremity of the most extreme views over time. The linearity of the updating process ensures that the moderation is sufficiently rapid that all beliefs converge to the same limit. The precise limit depends on the relative influence of various agents in terms of how many others they interact with.

Measuring Speed: Consensus Time Consider some network \mathbf{A} and the linear updating process with the matrix $\mathbf{T}(\mathbf{A})$. To measure how fast average-based updating processes converge, we simply examine how many periods are needed for the vector describing all agents' behaviors or beliefs to get within some distance ε of its limit. The measure of deviation from consensus we use has a simple interpretation. At each moment in time, there

into components, but still holds quite generally, as discussed in Golub and Jackson (2010).

21. For example, Theorem 1 shows that under these assumptions, agents converge to within a specified distance of consensus beliefs in the worst case, which could not happen if the network were not connected.

22. For example, see Golub and Jackson (2010) and Chapter 8 of Jackson (2008b) for details, background, and a proof.

are twice as many “messages” sent as there are links in the network (for the information/learning interpretation of the updating model)—two messages across each link. Let $\mathbf{m}(t)$ be this vector of messages for some ordering of the directed links. We define the distance from consensus at time t to be the root mean square distance of $\mathbf{m}(t)$ from its limit $\mathbf{m}(\infty)$. For the network \mathbf{A} and starting beliefs \mathbf{b} , we denote this distance (the *consensus distance*) at time t by $CD(t; \mathbf{A}, \mathbf{b})$.

This measurement of distance between messages corresponds to a weighted version of the usual (l_2) norm. More specifically, given two vectors of beliefs \mathbf{v} and \mathbf{u} , define $\|\mathbf{v} - \mathbf{u}\|_{\mathbf{w}} = [\sum_i w_i (v_i - u_i)^2]^{\frac{1}{2}}$. The distance of beliefs at time t from consensus is then $CD(t; \mathbf{A}, \mathbf{b}) = \|\mathbf{T}(\mathbf{A})^t \mathbf{b} - \mathbf{T}(\mathbf{A})^\infty \mathbf{b}\|_{\mathbf{s}(\mathbf{A})}$, where we use the weights $\mathbf{s}(\mathbf{A})$ defined by $\mathbf{s}(\mathbf{A}) = (\frac{d_1(\mathbf{A})}{D(\mathbf{A})}, \dots, \frac{d_n(\mathbf{A})}{D(\mathbf{A})})$. In other words, $\|\mathbf{T}(\mathbf{A})^t \mathbf{b} - \mathbf{T}(\mathbf{A})^\infty \mathbf{b}\|_{\mathbf{s}(\mathbf{A})}^2$ is a weighted sum of differences between current beliefs and eventual beliefs, with agent i 's term weighted by his or her relative degree. This is equivalent to the “messages” interpretation because an agent with degree $d_i(\mathbf{A})$ sends a share of the messages in the network given exactly by $s_i(\mathbf{A}) = \frac{d_i(\mathbf{A})}{D(\mathbf{A})}$.

Then the consensus time is the time it takes for this distance to get below ε :

DEFINITION 2. The *consensus time* to $\varepsilon > 0$ of a connected network \mathbf{A} is

$$CT(\varepsilon; \mathbf{A}) = \sup_{\mathbf{b} \in [0,1]^n} \min\{t : CD(t; \mathbf{A}, \mathbf{b}) < \varepsilon\}.$$

The need to consider different potential starting behavior or belief vectors \mathbf{b} is clear, because, if one starts with $b_i(0) = b_j(0)$ for all i and j , then equilibrium or consensus is reached instantly. Thus, the “worst case” \mathbf{b} will generally have behaviors or beliefs that differ across types and is useful as a benchmark measure of how homophily matters; taking the supremum in this way is standard in defining convergence times. This is closely related to mixing time, a standard concept of convergence for analyzing Markov chains.²³

23. For a discussion of various measures of convergence speed, see Montenegro and Tetali (2006).

Why is consensus time a good measure to study? Fundamentally, average-based updating is about reaching an equilibrium or consensus through repeated interaction. In many applications, agents may never fully reach an equilibrium or consensus, and whether they get close depends on the speed of convergence. Therefore, an important measure of the speed of convergence is obtained by asking how many rounds of communication it takes for beliefs to get within a prespecified distance of their limit.²⁴

Why Average-Based Updating? Despite the simplicity of average-based updating, it has a number of appealing properties. In the coordination game application, it obviously leads agents to converge to an efficient equilibrium in a very simple and decentralized way.

In the learning application, too, it turns out that the process often leads to efficient or approximately efficient outcomes in the long run. In particular, Golub and Jackson (2010) analyze conditions under which this “naive” updating process converges to a fully rational limiting belief (that is, the Bayesian posterior conditional on all agents’ information) in a large society. The conditions require²⁵ no agent to be too popular or influential. These conditions will be satisfied with a probability going to one in the settings that we study here—for example, under the regularity conditions of Definition 3—and so the naive beliefs will eventually approach a fully rational limit.

To obtain the same learning outcome by behaving in a fully Bayesian manner when the network is not common knowledge, agents would have to do a crushing amount of computation (see Mueller-Frank 2012 for some discussion of this issue, as well as explicit Bayesian procedures). Thus, if computation has even a tiny cost per arithmetic operation, the averaging heuristic can have an enormous advantage, both from an individual and group perspective. Indeed, this argument appears to be borne

24. The particular measure of distance from consensus that we have defined, CD, can be small even if a few agents deviate substantially from the eventual consensus. One could focus instead on the *maximum* deviation from consensus across agents. We believe all our results would be similar under this alternative specification. Note also that consensus time to ε as we define it is a lower bound on the time it would take for such an alternative distance measure to fall below ε .

25. See Section II.C of that paper for the simple result relevant to the present framework.

out by some recent empirical evidence. In an experiment seeking to distinguish different models of learning, Chandrasekhar, Larreguy, and Xandri (2010) placed experimental subjects in the learning setting described above.²⁶ They find that the subjects' updating behavior is better described by repeated averaging models than by more sophisticated rules.²⁷

III. HOW HOMOPHILY AFFECTS THE SPEED OF CONVERGENCE

III.A. The Main Result

This section presents the main result. A few definitions and notations are needed first. Throughout the section, we consider sequences of multi-type random networks, with all quantities (e.g., the matrix of intergroup linking probabilities \mathbf{P} and the vector of group sizes \mathbf{n}) indexed by the overall population size, n . We generally omit the explicit indexing by n to avoid clutter.

The next definition catalogs several regularity conditions on a sequence of multi-type random networks that will be assumed in the theorem.

DEFINITION 3.

- (1) A sequence of multi-type random networks is *sufficiently dense* if the ratio of the minimum expected degree to $\log^2 n$ tends to infinity. That is:

$$\frac{\min_k d_k[\mathbf{Q}(\mathbf{P}, \mathbf{n})]}{\log^2 n} \rightarrow \infty.$$

26. See also Choi, Gale, and Kariv (2005) for earlier experiments on learning in simple networks.

27. Corazzini et al. (2011) also report results that favor a behavioral DeGroot-style updating model over a Bayesian one. Nevertheless, there are some nuances in what the most appropriate model of boundedly rational updating might be, and how it depends on circumstances. Möbius, Phan, and Szeidl (2010) find some experimental evidence that in situations where information is “tagged” (so that agents not only communicate their information but also where a piece of information came from), the overweighting of some information that may take place under the DeGroot process can be avoided. The straight DeGroot model seems more appropriate when such tagging is difficult, which can depend on the nature of the information being transmitted, the size of the society, and other details of the process.

- (2) A sequence of multi-type random networks has *no vanishing groups* if

$$\liminf_n \frac{n_k}{n} > 0.$$

- (3) A sequence of multi-type random networks has *interior homophily* if

$$0 < \liminf_n h^{\text{spec}}(\mathbf{P}, \mathbf{n}) \leq \limsup_n h^{\text{spec}}(\mathbf{P}, \mathbf{n}) < 1.$$

- (4) Let \underline{P} denote the smallest nonzero entry of \mathbf{P} and \bar{P} denote the largest nonzero entry. A sequence of multi-type random networks has *comparable densities* if:

$$0 < \liminf_n \frac{\bar{P}}{\underline{P}} \leq \limsup_n \frac{\bar{P}}{\underline{P}} < \infty.$$

The sufficient density condition ensures that with a probability tending to 1, all nodes will be path-connected to each other.²⁸ The condition of no vanishing groups ensures that all groups have a positive share of the population. The condition of interior homophily requires that homophily not grow arbitrarily large or approach 0.²⁹ Finally, the comparable densities condition ensures that positive interaction probabilities do not diverge arbitrarily: they may be quite different, but their ratios must remain bounded.

The next definition is simply used to state the theorem's conclusion compactly.

DEFINITION 4. Given two sequences of random variables $x(n)$ and $y(n)$, we write $x(n) \approx y(n)$ to denote that for any $\varepsilon > 0$, if n is large enough, then the probability that

$$\frac{(1 - \varepsilon)y(n)}{2} \leq x(n) \leq 2(1 + \varepsilon)y(n)$$

is at least $1 - \varepsilon$.

28. The minimum needed for the network to be connected asymptotically almost surely is for the degrees to grow faster than $\log n$. The condition here is a little stronger than this and turns out to be what is needed to prove the tight asymptotic characterizations of convergence time we are about to present.

29. The case of no homophily is dealt with in detail in Chung, Lu, and Vu (2004), and the case of homophily approaching 1 may lead the network to become disconnected, in which case there can be no convergence at all. We leave the study of that more delicate situation (where the rate of homophily's convergence to 1 will be important) to future work.

In other words, $x(n) \approx y(n)$ indicates that the two (random) expressions $x(n)$ and $y(n)$ are within a factor of 2 (with a vanishingly small amount of slack) for large enough n with a probability going to 1.

With these preliminaries out of the way, we can state the main result.

THEOREM 1. Consider a sequence of multi-type random networks satisfying the conditions in Definition 3. Then, for any $\gamma > 0$:

$$CT\left(\frac{\gamma}{n}; \mathbf{A}(\mathbf{P}, \mathbf{n})\right) \approx \frac{\log(n)}{\log\left(\frac{1}{|\lambda_{\text{spec}}(\mathbf{P}, \mathbf{n})|}\right)}.$$

This result says that the speed of convergence of an average-based updating process is approximately proportional to $\log\left(\frac{1}{|\lambda_{\text{spec}}(\mathbf{P}, \mathbf{n})|}\right)$. Moreover, the speed of the process essentially depends only on population size and homophily. The approximation for consensus time on the right-hand side is always within a factor of 2 of the true consensus time. Properties of the network other than spectral homophily can change the consensus time by at most a factor of 2 relative to the prediction made based on spectral homophily alone.

Note that the matrix $\mathbf{F}(\mathbf{P}, \mathbf{n})$ introduced in Section II.B is invariant to multiplying all linking probabilities by the same number, and so the estimate above is invariant to homogeneous density shifts. Indeed, in Proposition 2, we prove something stronger than this.

The intuition behind why degree does not enter the expression in Theorem 1 is as follows. If one doubles each agent's number of links, but holds fixed the proportion of links that an agent has to various groups, then the amount of weight that agents place on various groups is unaffected. In the DeGroot process, each group quickly converges to a meta-stable internal belief, and then the differences in beliefs across groups are primarily responsible for slowing down convergence to a global consensus (see Section IV for more on this). It is the relative weights that agents put on their own groups versus other groups that determine the speed of this convergence. This is exactly what is captured by the homophily measure. Because these relative weights do not change under uniform density shifts, neither does convergence speed.

Finally, the parameter $\varepsilon = \frac{\gamma}{n}$ in the consensus time $\text{CT}(\frac{\gamma}{n}; \mathbf{A}(\mathbf{P}, \mathbf{n}))$ deserves some explanation. This choice is not essential. Indeed, Proposition A.5 in Online Appendix 1 shows that for any choice of ε , the consensus time $\text{CT}(\varepsilon; \mathbf{A}(\mathbf{P}, \mathbf{n}))$ can be characterized to within a fixed additive constant, and the inverse proportionality to $\log(\frac{1}{h^{\text{spec}}(\mathbf{P}, \mathbf{n})})$ remains unchanged. The intuition for choosing $\varepsilon = \frac{\gamma}{n}$ can be described as follows. Under the assumption of comparable densities and no vanishing groups, all agents have an influence of order $\frac{1}{n}$ on the final belief. That is, the final belief is a weighted average of initial beliefs, and each agent's weight is of order $\frac{1}{n}$.³⁰ Therefore, if we begin with an initial condition where one agent has belief 1 and all others have belief 0, then the limiting consensus beliefs will be of order $\frac{1}{n}$. Suppose we want a measure of consensus time that is sensitive to whether the updating process has equilibrated in this example. Then, to consider consensus to have been reached, the distance from consensus, as measured by the distance $\text{CD}(t; \mathbf{A}(\mathbf{P}, \mathbf{n}), \mathbf{b})$ of Section II.D, should be $\frac{\gamma}{n}$ for some small constant $\gamma > 0$. This amounts to requiring that agents should be within a small percentage of their final beliefs. Therefore, setting $\varepsilon = \frac{\gamma}{n}$ results in a consensus time measure that is sensitive to whether a single agent's influence has diffused throughout the society.

III.B. Applications to Specific Classes of Networks

The Islands Model We can immediately give two applications of this result. First, recall the islands model of Section II.C. There, we showed that if (\mathbf{P}, \mathbf{n}) is an islands network, then

$$h^{\text{spec}}(\mathbf{P}, \mathbf{n}) = h^{\text{islands}}(m, p_s, p_d) = \frac{p_s - p_d}{mp}.$$

This is a simple and hands-on version of the spectral homophily measure. Theorem 1 then immediately implies the following concrete characterization of consensus time.

COROLLARY 1. Consider a sequence of islands networks with parameters (m, p_s, p_d) satisfying the conditions in Definition 3.

30. The formal proof relies on the fact that each agent's influence is proportional to his degree, as stated in Section II.D, and the fact that each agent's degree is very close to his expected degree (see Lemma A.4 in Online Appendix 1).

Then, for any $\gamma > 0$:

$$CT\left(\frac{\gamma}{n}; \mathbf{A}(\mathbf{P}, \mathbf{n})\right) \approx \frac{\log(n)}{\log\left(\frac{1}{|h^{\text{islands}}(m, p_s, p_d)|}\right)} = \frac{\log(n)}{\log\left|\frac{mp}{p_s - p_d}\right|}.$$

Note that if p_s and p_d are scaled by the same factor, then p scales by that factor, too, so the estimate above is unaffected.

This example illustrates why a group-level perspective is useful, and how it goes beyond what we knew before. It is fairly straightforward from standard spectral techniques to deduce that $CT\left(\frac{\gamma}{n}; \mathbf{A}(\mathbf{P}, \mathbf{n})\right)$ is approximately

$$\frac{\log(n)}{\log\left|\frac{1}{\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n})))}\right|},$$

where $\lambda_2(\mathbf{T})$ is the second-largest eigenvalue of \mathbf{T} . Because $h^{\text{spec}}(\mathbf{P}, \mathbf{n})$ is the second-largest eigenvalue of a closely related matrix (recall Section II.B), one might ask whether Theorem 1 really yields much new insight.

We argue that it does. Recall that $\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n}))$ has dimensions n -by- n , which typically makes it a large matrix, and that it is a random object, with zeros and positive entries scattered throughout. It is not at all obvious, a priori, what its second eigenvalue is, or how it relates to the large-scale linking biases. Theorem 1 allows us to reduce this hairy question about $\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n}))$ to a question about the much smaller deterministic matrix \mathbf{P} (whose dimensionality is the number of *groups*, not agents), and obtain the formula of Corollary 1. We are not aware of other methods that can yield such a formula. This demonstrates the power of the group-level approach.

Two Groups The analysis of the islands model is special in that all groups have the same size. To obtain simple expressions that allow for heterogeneity in group size, we restrict attention to two groups, that is, $m = 2$. This echoes the intuitions of the islands model and again illustrates the main points cleanly.

For the two-group model, the vector \mathbf{n} has two entries (the two group sizes) and we focus on a case such that $P_{11} = P_{22} = p_s$ while $P_{12} = P_{21} = p_d$, and $p_s > p_d$.

In contrast to the islands model, there is no longer a homogeneous link density, since the two groups can differ in

size. Thus, the average link probability (allowing self-links) for group k is

$$p_k = \frac{n_k p_s + n_{-k} p_d}{n_1 + n_2},$$

where $-k$ denotes the group different from k .

Coleman's (1958) homophily index specialized to a group k is

$$h_k = \frac{\frac{n_k p_s}{n p_k} - \frac{n_k}{n}}{1 - \frac{n_k}{n}} = \frac{n_k}{n - n_k} \frac{p_s - p_k}{p_k}.$$

Recall that this is a measure of how much a group's fraction of same-type links ($\frac{n_k p_s}{n p_k}$) exceeds its population share ($\frac{n_k}{n}$), compared to how big this difference could be ($1 - \frac{n_k}{n}$).

We define the two-group homophily measure as the weighted average of the groups' Coleman homophily indices:

$$h^{\text{two}}(p_s, p_d, \mathbf{n}) = \frac{n_2}{n} h_1 + \frac{n_1}{n} h_2.$$

Here, weighting each homophily by the other group's size accounts for the relative impact of each group's normalized homophily index, which is proportional to the size of the group with which a given group interacts.

Again, note that the homophily measure, $h^{\text{two}}(p_s, p_d, \mathbf{n})$, is insensitive to uniform rescalings of the link density and depends only on relative fractions of links going from one group to another. It is 0 in a case where the link probabilities within groups are the same as across groups; it is 1 when all links are within groups.

With the definitions in hand, we state the characterization of consensus time in the case of two groups.

COROLLARY 2. Consider a sequence of two-group random networks (as already described) satisfying the conditions in Definition 3. Then, for any $\gamma > 0$:

$$\text{CT}\left(\frac{\gamma}{n}; \mathbf{A}(\mathbf{P}, \mathbf{n})\right) \approx \frac{\log(n)}{\log\left(\frac{1}{|h^{\text{two}}(p_s, p_d, \mathbf{n})|}\right)}.$$

Thus, consensus time depends only on the size of the network and on the weighted average of the groups' Coleman homophily indices!

III.C. The Speed of Average-Based Updating Is Invariant to Uniform Density Shifts

Theorem 1 and the special cases we have just discussed suggest that consensus times do not depend on density but only on ratios of linking densities. That is, if linking probabilities are adjusted uniformly, then the estimate of consensus time in Theorem 1 is essentially unaffected. The following result strengthens this conclusion.

PROPOSITION 2. Consider a sequence of multi-type random networks (\mathbf{P}, \mathbf{n}) and another $(\mathbf{P}', \mathbf{n})$, where $\mathbf{P}' = c\mathbf{P}$ for some $c > 0$. Under the conditions of Theorem 1, the ratio of consensus times

$$\frac{\text{CT}(\frac{\gamma}{n}; \mathbf{A}(\mathbf{P}, \mathbf{n}))}{\text{CT}(\frac{\gamma}{n}; \mathbf{A}(\mathbf{P}', \mathbf{n}))}$$

converges in probability to 1.³¹

III.D. How the Main Result Is Obtained

In this section, we give an outline of the main ideas behind Theorem 1. There are two pieces to this. One is the role of the second eigenvalue as a measure of speed, which follows from known results in Markov theory. The other, which is the major technical innovation in our article, is to show that the interactions that need to be considered are only at the group level,

31. For any $\delta > 0$, we can find large enough n such that the ratio of the two Online Appendix 1small enough δ , consensus times is in the interval $[1 - \delta, 1 + \delta]$ with probability at least $1 - \delta$.

The reason that this proposition is not an immediate corollary of Theorem 1 is as follows. According to Theorem 1, both consensus times $\text{CT}(\frac{\gamma}{n}; \mathbf{A}(\mathbf{P}, \mathbf{n}))$ and $\text{CT}(\frac{\gamma}{n}; \mathbf{A}(\mathbf{P}', \mathbf{n}))$ are approximately

$$\frac{\log(n)}{\log(\frac{1}{h^{\text{spec}}(\mathbf{P}, \mathbf{n})})} \approx \frac{\log(n)}{\log(\frac{1}{h^{\text{spec}}(\mathbf{P}', \mathbf{n})})}$$

with the equality holding since h^{spec} is invariant to degree shifts. But the \approx of Theorem 1 allows each consensus time to deviate by a factor of 2 from the estimate, so that, a priori, the two consensus times might differ by a factor of as much as 4. The proposition shows that this is not the case.

Density adjustments that are not uniform will change the ratios of interaction across groups. By Theorem 1, the consensus time would then change in a way that depends on how the second-largest eigenvalue of $\mathbf{F}(\mathbf{P}, \mathbf{n})$ is affected.

and only need to be considered in terms of expectations and not actual realizations.

Consensus Time and Second Eigenvalues The following lemma relates consensus time to the second-largest eigenvalue (in magnitude) of the realized updating matrix.

LEMMA 2. Let \mathbf{A} be connected, $\lambda_2(\mathbf{T}(\mathbf{A}))$ be the second-largest eigenvalue in magnitude of $\mathbf{T}(\mathbf{A})$, and $\underline{s} := \frac{\min_i d_i(\mathbf{A})}{D(\mathbf{A})}$ be the minimum relative degree. If $\lambda_2(\mathbf{T}(\mathbf{A})) \neq 0$, then for any $0 < \varepsilon \leq 1$:

$$\left[\frac{\log\left(\frac{1}{(2\varepsilon)}\right) - \log\left(\frac{1}{\underline{s}^{1/2}}\right)}{\log\left(\frac{1}{|\lambda_2(\mathbf{T}(\mathbf{A}))|\right)} \right] \leq \text{CT}(\varepsilon; \mathbf{A}) \leq \left[\frac{\log\left(\frac{1}{\varepsilon}\right)}{\log\left(\frac{1}{|\lambda_2(\mathbf{T}(\mathbf{A}))|\right)} \right].$$

If $\lambda_2(\mathbf{T}) = 0$, then for every $0 < \varepsilon < 1$ we have $\text{CT}(\varepsilon; \mathbf{A}) = 1$.

If ε is fairly small, then the bounds in the lemma are close to each other and we have a quite precise characterization in terms of the spectrum of the underlying social network. However, the lower bound in Lemma 2 includes a term $\log\left(\frac{1}{\underline{s}^{1/2}}\right)$, which can grow as n grows. In Online Appendix 1, Proposition A.5 shows that this can be dispensed with under the assumptions of Definition 3.

The proof of this result follows standard techniques from the literature on Markov processes and their relatives (Montenegro and Tetali 2006).

Relating Second Eigenvalues to Large-Scale Network Structure As mentioned above in Section III.B, a result like Lemma 2 has the limitation that the second eigenvalue of a large random matrix does not immediately yield intuitions about how group structure affects convergence rates; for large populations, even computing this eigenvalue precisely can be a challenge. Thus, our goal is to reduce the study of this object to something simpler.

To this end, we present the main technical result: a “representative-agent” theorem that allows us to analyze the convergence of a multi-type random network by studying a much smaller network in which there is only one node for each *type* of agent. We show that under some regularity conditions, the

second eigenvalue of a realized multi-type random graph converges in probability to the second eigenvalue of this representative agent matrix—namely, the matrix $\mathbf{F}(\mathbf{P}, \mathbf{n})$ introduced in Section II.B. That eigenvalue is precisely the spectral homophily, $h^{\text{spec}}(\mathbf{P}, \mathbf{n})$.

This result is useful for dramatically simplifying computations of approximate consensus times both in theoretical results and in empirical settings, because now the random second eigenvalue of the updating matrix can be accurately predicted knowing only the relative probabilities of connections across different types, as opposed to anything about the precise realization of the random network. Indeed, this result is the workhorse used in Online Appendix 1 to prove all the propositions about the islands and two-group cases already discussed.

THEOREM 2. Consider a sequence of multi-type random networks described by (\mathbf{P}, \mathbf{n}) that satisfies the conditions of Definition 3 (i.e., is sufficiently dense; and has no vanishing groups, interior homophily, and comparable densities). Then for any $\delta > 0$ if n is sufficiently large,

$$|\lambda_2(\mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n}))) - \lambda_2(\mathbf{F}(\mathbf{P}, \mathbf{n}))| \leq \delta,$$

with probability at least $1 - \delta$.

Theorem 2 is a law of large numbers for spectra of multi-type random graphs. Large-number techniques are a central tool in the random graphs literature; they show that various important properties of random graphs converge to their expectations, which shows that these locally haphazard objects have very predictable global structure. The closest antecedent to this particular theorem is by Chung, Lu, and Vu (2004) for networks *without* homophily. Their result shows that expectations rather than realizations are important in some particular limiting properties of a class of random graph models. Our theorem is the first of its kind to apply to a model that allows homophily and the associated heterogeneities in linking probabilities, which eliminates the sort of symmetry present in many random graph models. The proof employs a strategy similar to that of Chung, Lu, and Vu (2004). That strategy relies on decomposing the random matrix representing our graph into two pieces: an “orderly” piece whose entries are given by linking




FIGURE III

Panel (A) is shaded according to the type-based probabilities of linking between nodes, with darker shades indicating higher probabilities, and Panel (b) is a single realized matrix drawn from the multi-type random network model with Panel (B) probabilities, shaded according to the actual realizations of links.

probabilities between nodes of various types, and a noisy piece due to the randomness of the actual links. By bounding the spectral norm of the noise, we show that, asymptotically, the second eigenvalue of the orderly part is, with high probability, very close to the second eigenvalue of the random matrix of interest. Then we note that computing the second eigenvalue of the orderly part requires dealing only with a representative-agent matrix.

Figure III provides an idea of why Theorem 2 holds. Figure IIIa presents the type-based probabilities of linking, in a case with 300 nodes and three groups (each of 100 nodes) with varying probabilities of linking within and across groups represented by the shading of the diagram. In Figure IIIb, the picture is broken into 300×300 pixels, where a pixel is shaded black if there is a link between the corresponding nodes and is white if there is no link. This is a picture of one randomly drawn network where each link is independently formed with the probability governed by the multi-type random network model with the probabilities in Figure IIIa. One sees clearly the law of large numbers at work as the relative shadings of the expected and realized matrices coincide rather closely. Though this is harder to see in a picture, the same will be true of the important parts of the spectra of the two matrices.

IV. INTERIM DYNAMICS AND SEPARATING THE SOURCES OF DISAGREEMENT

In this section, we focus on the dynamics of behaviors to tease apart the effects of homophily and initial disagreement and thus clarify the implications of the model.

IV.A. Reducing to the Dynamics of Representative Agents

Suppose that initial beliefs $b_i(0)$ of agents of type k (so that $i \in N_k$) are independent random variables distributed on $[0, 1]$ with mean μ_k . Let $\boldsymbol{\mu} \in \mathbb{R}^m$ be the vector of these means with as many entries as types.

Recall the definition of $\mathbf{F}(\mathbf{P}, \mathbf{n})$ from Section II.B. The (k, ℓ) entry of this matrix captures the relative weight of type k on type ℓ . Fixing $\mathbf{F}(\mathbf{P}, \mathbf{n})$, define the vector $\bar{\mathbf{b}}(t) \in \mathbb{R}^m$ by

$$\bar{\mathbf{b}}(t) = \mathbf{F}(\mathbf{P}, \mathbf{n})^t \boldsymbol{\mu}.$$

This is an updating process in which there is one representative agent for each type, that starts with that type’s average belief, and then the representative agents update according to the group updating matrix $\mathbf{F}(\mathbf{P}, \mathbf{n})$. We call this the representative-agent updating process. We can then define a vector $\hat{\mathbf{b}}(t) \in \mathbb{R}^n$ by the property that if agent i is of type k , then $\hat{b}_i(t) = \bar{b}_k(t)$. That is, $\hat{\mathbf{b}}(t)$ gives to each agent a belief equal to the belief of the representative agent of his type.

Then we have the following result, which states that the real process is arbitrarily well-approximated by the representative agent updating process for large enough networks.

PROPOSITION 3. Fix a sequence of multi-type random networks described by (\mathbf{P}, \mathbf{n}) that satisfies the conditions in Definition 3. Consider initial beliefs drawn as described above, and let³² $\mathbf{b}(t) = \mathbf{T}(\mathbf{A}(\mathbf{P}, \mathbf{n}))^t \mathbf{b}(0)$. Then, given a $\delta > 0$, for any sufficiently large n and all $t \geq 1$, the inequality³³ $\|\mathbf{b}(t) - \hat{\mathbf{b}}(t)\|_{\mathbf{e}/n} \leq \delta$ holds with probability at least $1 - \delta$.

This proposition shows that in large enough (connected) random networks, the convergence of beliefs within type occurs

32. Thus, $\mathbf{b}(t)$ is a random variable, which is determined by the realization of the random matrix $\mathbf{A}(\mathbf{P}, \mathbf{n})$.

33. Here \mathbf{e} is the vector of all ones, and so we are using an equal-weight ℓ_2 norm: $\|\mathbf{v} - \mathbf{u}\|_{\mathbf{e}} = \left[\frac{1}{n} \sum_i (v_i - u_i)^2 \right]^{1/2}$.

quickly—essentially in one period—and that all of the difference is across types afterward. To understand why this is the case, note that under the connectivity assumption in Definition 3, each agent is communicating with many other agents as n becomes large, and thus the idiosyncratic noise of any single agent is washed away as n grows—even with just one period of communication. Moreover, each agent of a given type has a similar composition of neighbors, in terms of percentages of various types. Therefore, it is only the difference across types that remains after one period of communication. This proposition then allows us to focus on the representative-agent updating process with only an arbitrarily small loss in accuracy, even after just one period.

IV.B. Separating and Estimating the Sources of Disagreement

To clarify the roles of initial disagreement and homophily, it is enough (and clearest) to examine the case of two groups, as described in Section III.B. Let $D_i = n_i(p_s n_i + p_d n_{-i})$ be the expected total degree of group i and D be the total degree $D = D_1 + D_2$. Writing $h = h^{\text{two}}(p_s, p_d, \mathbf{n})$, we can compute³⁴

$$\bar{b}_i(t) = \frac{D_i + D_{-i}h^t}{D} \cdot \mu_i + \frac{D_{-i}(1 - h^t)}{D} \cdot \mu_{-i}.$$

Beliefs converge to a weighted average of initial beliefs, with each group’s mean getting a weight proportional to its total degree, D_k . The difference in beliefs is then

$$\bar{b}_1(t) - \bar{b}_2(t) = h^t(\mu_1 - \mu_2).$$

Thus, disagreement at a given time is always proportional to initial disagreement, but its impact decreases by a factor that decays exponentially in time based on the level of homophily.

We can also write³⁵

$$\log(\bar{b}_1(t) - \bar{b}_2(t)) = (\log h) \cdot t + \log(\mu_1 - \mu_2).$$

Consequently, given data on the average disagreement between types, $\bar{b}_1(t) - \bar{b}_2(t)$, at several different times, running a regression of $\log(\bar{b}_1(t) - \bar{b}_2(t))$ on t would estimate both the logarithm of

34. The formulas can be verified inductively using the law of motion $\bar{\mathbf{b}}(t + 1) = \mathbf{F}(\mathbf{P}, \mathbf{n})\bar{\mathbf{b}}(t)$, recalling that $\bar{\mathbf{b}}(t) = \mu$.

35. We assume $p_s > p_d$, so $h > 0$.

homophily (as the coefficient on t) and the logarithm of initial disagreement (as the intercept). Note that the validity of this procedure does not depend on the sizes of the different groups or the values of p_s and p_d , nor does it require any adjustments for these (typically unknown) quantities. Therefore, when the model holds, it provides a simple way to distinguish what part of disagreement is coming from differences in initial information or inclinations, and what part is coming from homophily in the network. Extending these results to more groups, as well as richer distributions of initial beliefs (allowing for some correlation across agents' beliefs), presents interesting directions for future study.

IV.C. Consequences for Interpreting the Main Results

A central finding of this article is that homophily slows convergence. The mechanism by which this occurs in our model is as follows. Homophily, through the type-dependent network-formation process, causes “fault lines” in the topology of a network when agents have linking biases toward their own type. That is, there are relatively more links among agents of the same type, and fewer links between agents of different types. This creates the potential for slow convergence if agents on different sides of the fault lines start with different beliefs. Since consensus time is a worst-case measure (recall Definition 2), it equals the time to converge starting from such initial beliefs; this is a natural initial point to the extent that types not only correspond to network structure but also to differences in initial information.

In the model, agents' types play a direct role *only* in determining the network topology.³⁶ We view this as a plausible and important channel through which homophily affects communication processes. In some situations, however, it may be desirable to also think of types as separately determining agents' initial beliefs, rather than focusing on a worst-case initial condition.

36. In particular, agents' types have no formal role in the definition of the updating process or of consensus time, although, of course, they affect both things through the structure of the network. Holding fixed a given random network generated by the multi-type random graph model, we could “scramble” the type labels—that is, reassign them at random—and neither the updating process nor the consensus time would change, though some interpretations might. Thus, our main results about homophily should be interpreted as referring to the types that were relevant in network formation in the multi-type random network setting.

The results given in this section show how to do this. More generally, one may think of agents having two different kinds of type: a “network type”—the traits that determine the probabilities of linking—and a “belief type”—the traits that govern initial beliefs $b_i(0)$. Then the calculation in Section IV.B can be interpreted as follows: if network type and belief type are highly correlated (so that μ_1 is significantly different from μ_2), then we should see a high persistence of disagreement (assuming there is homophily based on the network type). But if they are uncorrelated, so that $\mu_1 = \mu_2$, then there should be no persistent disagreement, regardless of linking biases.³⁷

V. AN APPLICATION: VOTING IN A SOCIETY WITH HOMOPHILY

We now provide an application that illustrates some of the concepts and results. The application is one in which a society that exhibits homophily sees some signals that are correlated with a true state of nature, and then the agents communicate to update their beliefs. After communicating for some time, the agents vote on a policy. The question is: do they vote correctly? The answer depends on homophily. Even when the society has more correct signals than wrong ones, and it is guaranteed eventually to converge to a situation where a majority holds the correct view, in the medium run homophily can cause incorrect majority rule decisions.³⁸

To keep the setting simple, we abstract away from individual preferences and focus on heterogeneity in information and in interaction patterns. All agents want their votes to match the true state.

The model is as follows: a society of n agents consists of two groups; one forms a fraction $M \in (\frac{1}{2}, 1)$ of the society and is referred to as the “majority group”; the other is referred to as the “minority group.”

We work in the setting of Section IV. There is a true state of nature, $\omega \in \{0, 1\}$. Agents see signals that depend not only on the

37. We thank an anonymous referee for suggesting this discussion.

38. Neilson and Winter (2008) study deliberation through linear updating followed by voting. They point out that voting before convergence occurs can yield results different from the eventual consensus of the society; we extend their framework to study how these deliberative processes are affected by large-scale homophily in a connected network.

state but also on the group they are in. In particular, agents in the majority have a probability μ of seeing a signal that is equal to the state ω , and probability $1 - \mu$ of seeing a signal $1 - \omega$, which is the opposite of the true state. As for the minority group, their probability of seeing a correct signal is ν , and otherwise they see an incorrect signal. Conditional on ω , all these signals are independent.

The two probabilities μ and ν are chosen so that the overall expected fraction of agents in the society with a correct signal is some given $p > \frac{1}{2}$. Thus, $M\mu + (1 - M)\nu = p$ and so

$$\nu = \frac{p - M\mu}{1 - M}.$$

We assume that each agent, irrespective of type, has the same expected number of links and, consequently, approximately the same influence on the final belief. Thus, regardless of the initial distribution of who sees which signal, the weighted averaging of beliefs will converge in a large society to $b_i(\infty) = p$ if the state is $\omega = 1$ and $b_i(\infty) = 1 - p$ if the state is $\omega = 0$. We consider a voting rule so that, at time t , an individual votes “1” if $b_i(t) > \frac{1}{2}$ and “0” if $b_i(t) \leq \frac{1}{2}$. This will eventually lead society to a correct decision. Moreover, note that if agents vote without any communication (that is, based on initial beliefs $b_i(0)$), then there will be a correct majority vote, since a fraction $p > \frac{1}{2}$ of the agents will vote for the correct state. It is in intermediate stages—such that agents have had some communication, but not yet converged—where incorrect votes may occur.

The groups exhibit homophily. The group-level matrix $\mathbf{F}(\mathbf{P}, \mathbf{n})$ of relative linking densities is³⁹:

$$\mathbf{F}(\mathbf{P}, \mathbf{n}) = \begin{pmatrix} 1 - f & f \\ fQ & 1 - fQ \end{pmatrix}.$$

A majority agent has a fraction f of his or her links to the minority group, whereas a minority agent has a fraction fQ of his or her

39. Note that this multi-type random network departs from our two-group setting introduced in Section III.B, in that a majority agent’s probability of linking to a majority agent may differ from the probability that a minority agent links to a minority agent (whereas, in the basic two-type model, both probabilities would be equal to the same number p_s). Nevertheless, it fits into the general multi-type framework. Throughout this section, we assume that the regularity conditions of Definition 3 are satisfied by the sequence of multi-type random networks.

links to the majority group, where $Q = \frac{M}{1-M}$ (as required for a case of reciprocal communication). Suppose that $1 - M \geq f > 0$, so that there is a bias toward linking to one's own type (if there were no bias, then f should equal $1 - M$).⁴⁰ In a situation in which there is no bias in how signals are distributed across the population, the communication will quickly aggregate information and lead to correct voting. The interesting case is when there is some bias in how signals are distributed across the groups.

Without loss of generality, suppose the true state turns out to be $\omega = 1$. If there is no bias in how signals are distributed across groups, then a fraction pM of these signals are observed by the majority group, and a fraction $p(1 - M)$ of them are observed by the minority group. If there is a bias, then the "correct" signals will be more concentrated among either the majority group or the minority group. It is easy to see that if they are concentrated among the majority group, voting will tend to be correct from the initial period onward, and so is not led astray by communication. However, if the correct signals turn out to be more concentrated among the minority, then it is possible for short-term communication to lead to incorrect voting outcomes. Recall that μ is the fraction of majority agents who observe the correct signal of 1. From now on we focus on the case $\mu < \frac{1}{2}$. Under the assumption that p is the overall expected fraction of correct signals in the population, the minimum value that μ can take is $\frac{p(1-M)}{M}$, which corresponds to the case where every minority agent sees a correct signal and then the remaining correct signals are observed by the majority group.

In the initial period $t = 0$, before any communication, all agents vote based on their signals, so there is a correct vote, with a fraction p agents voting "1" and $(1 - p)$ of the agents voting "0." Now let us consider what happens with updating.

We use Proposition 3, which allows us to reduce the large-population dynamics to the representative agents with a vanishing amount of error. We can then use the matrix $\mathbf{F}(\mathbf{P}, \mathbf{n})$ to deduce that, after one period of updating, the beliefs of the majority group will be

$$b_{\text{Maj}}(1) = p - \left(1 - \frac{f}{1 - M}\right)(p - \mu)$$

40. The number $f = 1 - M$ is the correct unbiased link fraction for majority agents if agents are allowed self-links, but otherwise it would be $\frac{(1-M)n}{(n-1)}$.

and the beliefs of the minority will be

$$b_{\text{Min}}(1) = p + Q \left(1 - \frac{f}{1 - M} \right) (p - \mu).$$

Thus, in a case where $f < 1 - M$ and $\mu < p$, the majority will have a lower belief than the average signal. This corresponds to homophily (f below its uniformly mixed level of $1 - M$) and a bias toward error in the initial signal distribution of the majority (μ below the probability p that a randomly chosen agent has a correct signal).

This presents an interesting dynamic. If agents can vote before any communication has taken place, then they will vote correctly. After an initial round of communication, the majority of beliefs can be biased towards the wrong state, but then again in the long run the society will reach a correct consensus. So, how long will it take for the voting behavior to converge to being correct again after some communication? This will depend both on the homophily and the bias in the signal distribution. In particular, the general expression for beliefs after t periods of updating is⁴¹

$$b_{\text{Maj}}(t) = p - \left(1 - \frac{f}{1 - M} \right)^t (p - \mu)$$

and

$$b_{\text{Min}}(t) = p + Q \left(1 - \frac{f}{1 - M} \right)^t (p - \mu).$$

Here we see the dynamics of homophily explicitly. The deviation of beliefs after t periods from their eventual consensus is proportional to the initial bias in signal distribution times a factor of

$$\left(1 - \frac{f}{1 - M} \right)^t,$$

which captures how homophilous the relationships are. Recall that f is the fraction of the majority group's links to the minority

41. This is seen as follows. If $b_{\text{Maj}}(t - 1) = p - a$ and $b_{\text{Min}}(t - 1) = p + aQ$, then $b_{\text{Maj}}(t) = p - (1 - f)a + fQa$, which is then rewritten as $b_{\text{Maj}}(t) = p - a(1 - f(1 + Q))$. Noting that $1 + Q = \frac{1}{1 - M}$ leads to the claimed expression. The averaging of overall beliefs to p provides the corresponding expression for b_{Min} .

group, which in a world without homophily would be $1 - M$, and with homophily is below $1 - M$. The impact of homophily decays exponentially in time.

The first period in which voting will return to being correct is the first t such that

$$\left(1 - \frac{f}{1 - M}\right)^t < \frac{p - \frac{1}{2}}{p - \mu},$$

which depends on how biased the initial signal distribution is, and on an exponentially decaying function of homophily. Treating the representative dynamics approximation as exact for the moment and ignoring integer constraints⁴², the time to the correct vote is

$$(4) \quad \frac{\log \frac{p - \frac{1}{2}}{p - \mu}}{\log \left(1 - \frac{f}{1 - M}\right)}.$$

The expression resembles our earlier results, with $(1 - \frac{f}{1 - M})$ corresponding to the homophily. In fact, the matrix $\mathbf{F}(\mathbf{P}, \mathbf{n})$ has a second-largest eigenvalue

$$h^{\text{spec}}(\mathbf{P}, \mathbf{n}) = 1 - f - fQ = 1 - \frac{f}{1 - M}$$

and so (4) is exactly

$$\frac{\log \left(\frac{p - \frac{1}{2}}{p - \mu}\right)}{\log |h^{\text{spec}}(\mathbf{P}, \mathbf{n})|},$$

which resembles the formulas in our earlier results.

From the formula, one can immediately deduce that the time to a correct vote is decreasing in the fraction of majority agents having the correct initial signal, decreasing in the overall fraction of correct signals p , and increasing and convex in homophily (becoming arbitrarily large as homophily becomes extreme).

Thus, when deliberation occurs in the setting of our model before a vote, the efficiency of electoral outcomes (measured by the time it takes to be able to get a correct vote) depends not only on the overall quality of information distributed throughout society

42. It can be seen from Proposition 3 that for fixed values of the parameters in this section, taking n large enough will result in the true t being off by at most 1 relative to this estimate.

at the beginning but also on how it is distributed (what fraction of majority agents get correct information) and on the segregation patterns in communication, as measured by homophily.

VI. WHAT SPECTRAL HOMOPHILY MEASURES

Homophily has consequences for updating processes because it creates “fault lines” in the patterns of interactions among groups. Homophily makes it possible to draw a boundary in the group structure, separating it into two pieces so that there are relatively few links across the boundary and relatively many links *not* crossing the boundary. Therefore, an appropriate global measure of homophily should find a boundary where that disparity is strongest and quantify it.⁴³

In this section, we show that the spectral homophily measure accomplishes this. We do this by proving an estimate on spectral homophily in terms of a more “hands-on” quantity that we call *degree-weighted homophily*.

Let $M = \{1, \dots, m\}$ be the set of groups. First, we define a notion of the weight between two collections of groups.

DEFINITION 5. Let $\mathbf{F}(\mathbf{P}, \mathbf{n})$ be as defined in Section II.B. For two subsets of groups, $B, C \subseteq M$, let

$$W_{B,C} = \frac{1}{|B||C|} \sum_{(k,\ell) \in B \times C} F_{k\ell} F_{\ell k}.$$

The quantity $W_{B,C}$ keeps track of the relative weight between two (possibly overlapping) collections of groups B and C , and is a measure that ranges between 0 and 1. The numerator measures the total intensity of interaction between groups in the collection B and groups in the collection C . The denominator is the product of the sizes of the two sets B and C . With this definition in hand, we define a notion of degree-weighted homophily.

DEFINITION 6. Given any subset of groups $\emptyset \subsetneq B \subsetneq M$, let the *degree-weighted homophily* of (\mathbf{P}, \mathbf{n}) relative to B be

43. Previous work taking a different approach in the same spirit is discussed in Diaconis and Stroock (1991).

defined by

$$\begin{aligned} \text{DWH}(B; \mathbf{P}, \mathbf{n}) &= \frac{W_{B,B} + W_{B^c,B^c} - 2W_{B,B^c}}{|B|^{-2} \sum_{k \in B} d_k(\mathbf{Q}(\mathbf{P}, \mathbf{n}))^{-2} + |B^c|^{-2} \sum_{k \in B^c} d_k(\mathbf{Q}(\mathbf{P}, \mathbf{n}))^{-2}}, \end{aligned} \tag{5}$$

where the W 's are computed relative to $\mathbf{F}(\mathbf{P}, \mathbf{n})$.

The numerator keeps track of how much of the weight in \mathbf{F} falls within B and within B^c , as well as (with the opposite sign) how much weight goes between these sets of nodes. Indeed, links within B or its complement B^c increase the degree-weighted homophily, whereas links between the two subsets decrease it. The denominator is a normalizing value, which guarantees that the quantity defined in equation (5) is no greater than 1 in absolute value.⁴⁴

To see that the degree-weighted homophily has an intuitive interpretation, consider a very simple special case. Suppose $|B| = \frac{m}{2}$ and that every group has the same expected number of links. Then,

$$\text{DWH}(B; \mathbf{P}, \mathbf{n}) = \frac{\#(\text{links within } B \text{ or } B^c) - \#(\text{links from } B \text{ to } B^c)}{\#(\text{total links})}, \tag{6}$$

where all quantities are expectations.

Let

$$\text{DWH}(\mathbf{P}, \mathbf{n}) = \max_{\emptyset \subset B \subset M} |\text{DWH}(B; \mathbf{P}, \mathbf{n})|.$$

Thus, the degree-weighted homophily (DWH) of a given network is the maximum level of degree-weighted homophily across different possible splits of the network.⁴⁵

The point of this section is the following lemma, which shows that DWH provides a lower bound on the spectral homophily.

LEMMA 3. If $\mathbf{Q}(\mathbf{P}, \mathbf{n})$ is connected (viewed as a weighted network), then

$$|h^{\text{spec}}(\mathbf{P}, \mathbf{n})| \geq |\text{DWH}(\mathbf{P}, \mathbf{n})|.$$

44. See Section E in Online Appendix 1 for details.

45. This is related, intuitively speaking, to a weighted version of a minimum cut, although this degree-weighted homophily measure turns out to be the right one for our purposes.

The key implication is that the second eigenvalue can be related to a “hands-on” ratio of weights in and out of groups. This provides intuition as to why it measures homophily and relates to the slowdown of averaging processes.

This DWH lower bound on the spectral homophily is tight in the islands and two-group models, as can be verified by simple calculations. Thus, by the remark made in Section II.C, DWH coincides with the Coleman homophily index for the islands model; by the formula in Section III.B it also has a simple relationship with Coleman homophily in the two-group case.

Under some additional assumptions, a general complementary upper bound can be established, which is not quite tight, but reasonably good when the number of groups is not too large (as explored in an extension, Golub and Jackson 2011).

VII. COMPARING AVERAGE-BASED UPDATING WITH DIRECT CONTAGION

It is useful to compare average-based updating with a different sort of transmission process. It turns out that the two processes are affected in very different ways by homophily and density. This shows that the averaging aspect of the average-based updating process (though probably not the exact linear functional form) is essential for producing the results discussed above. It also shows, more generally, that different natural processes have starkly different dependencies on the network structure.

VII.A. *Direct Contagion Processes*

Loosely, let us say that a dynamic process on a network is a “direct transmission process” if it is characterized by travel along shortest paths⁴⁶ in the network. More formally, we simply consider a direct contagion process to be any process such that the time to converge is proportional to the average shortest path between nodes in a network.⁴⁷

46. Standard network definitions, such as that of a shortest path, are omitted. They can be found in Jackson (2008b).

47. As will become clear, one could replace “average shortest path” with “maximum length shortest path” (network diameter) and the conclusions below would still hold.




FIGURE IV

Illustration of the simple diffusion process where every node passes information to all of its neighbors at each date, $t = 1, 2, \dots$. Panels A through F show how an initial piece of information gradually spreads through the whole network.

This covers a variety of processes. For example, consider a game where an agent is willing to choose action 1 (e.g., buy a new product) as soon as at least one of his or her neighbors does. Let us examine the (myopic) best response dynamics of such a process. If we begin with some random agent taking action 1, what is the time that it will take for the action to spread to others in the society, on average? That time will be determined by the average network distance between the initial agent who takes action 1 and any other agent in the society. This class of pure contagion processes also models the spread of some diseases, ideas or rumors—where an agent is either “infected” or not, “informed” or not, and so forth—and where the time it takes for something to diffuse from one agent to another is proportional to the length of the shortest path between them. The class also includes broadcast processes, where nodes communicate to all neighboring nodes in each period, as well as processes where the network is explicitly navigated by a traveler using some sort of addressing system. Such contagion processes serve as a useful point of comparison to the average-based updating that we have considered. An example of a direct contagion process operating is depicted in Figure IV.

Direct contagion processes have an obvious measure of speed, which is simply the average shortest path length between pairs of nodes in the network. We denote this random variable by

$\text{AvgDist}(\mathbf{A}(\mathbf{P}, \mathbf{n}))$. If one is worried about the longest time it could take to pass from some node to some other node, then the diameter of the network is the right measure; we write this as $\text{Diam}(\mathbf{A}(\mathbf{P}, \mathbf{n}))$.⁴⁸

Such direct contagion processes are obviously idealized, but they can easily extend to analyze more realistic phenomena. For example, it may be more plausible to posit that each node sends the news to each neighbor with probability $\pi < 1$, and the decisions are independent. It turns out that this process can be analyzed using the simpler contagion process outlined above in which the transmission is certain. Given a network on which the “noisy” process is supposed to happen, one merely considers a subnetwork in which edges of the original network are included with probability π and excluded with probability $1 - \pi$, independently of each other. The deterministic broadcasting process operating on this sparser subnetwork is equivalent to the noisy process on the original network.

VII.B. The Speed of Contagion

Before discussing the speed of a direct contagion process, we provide one more definition. Let $\tilde{d}(\mathbf{P}, \mathbf{n}) = \frac{\sum_k (d_k(\mathbf{Q}(\mathbf{P}, \mathbf{n})))^2 n_k}{D(\mathbf{Q}(\mathbf{P}, \mathbf{n}))}$ be the *second-order average degree*, which is a useful quantity in analyzing asymptotic properties of multi-type random networks.⁴⁹ If the average degree $d_k(\mathbf{Q}(\mathbf{P}, \mathbf{n}))$ is the same across groups, then this is just the average degree, but more generally it weights degrees quadratically across groups.

We need to restrict attention to settings satisfying certain regularity conditions. In particular, we suppose that

- (i) there exists $M < \infty$ such that $\max_{k,k'} \frac{d_k(\mathbf{Q}(\mathbf{P}, \mathbf{n}))}{d_{k'}(\mathbf{Q}(\mathbf{P}, \mathbf{n}))} < M$,
- (ii) $\tilde{d}(\mathbf{P}, \mathbf{n}) \geq (1 + \varepsilon) \log n$ for some $\varepsilon > 0$,
- (iii) $\frac{\log \tilde{d}(\mathbf{P}, \mathbf{n})}{\log n} \rightarrow 0$, and

48. For the multi-type random networks that we examine, it turns out that the average distance and diameter are effectively the same. This is because, as long as clustering remains bounded away from 1, the majority of pairs of nodes in a large network are at the maximum possible distance from each other.

49. There is a subtlety in the definition. It resembles the second moment of the degree distribution divided by the first moment. But it is important to note that in computing the numerator, one first takes average degrees within groups and then squares those averages. This makes it somewhat different from a second moment taken across nodes.

(iv) there exists $\varepsilon > 0$ such that $\frac{\min_{k,k'} P_{kk'}}{\max_{k,k'} P_{kk'}} > \varepsilon$.

These conditions admit many cases of interest and can be understood as follows: condition (i) requires that there is not a divergence in the expected degree across groups so that no group completely dominates the network; (ii) ensures that the average degree grows with n fast enough so that the network becomes connected with a probability going to 1, making communication possible; (iii) implies that the average degree grows more slowly than n , as otherwise the shortest path degenerates to being of length 1 or 2 (which is not of much empirical interest); and (iv) ensures that there is some lower bound on the probability of a link between groups relative to the overall probability of links in the network. This last condition ensures that groups do not become so homophilous that the network becomes disconnected. Nevertheless, the conditions can accommodate any arbitrarily high fixed level of homophily, since M and ε are arbitrary parameters.

By adapting a theorem of Jackson (2008a) to our setting, we derive the following characterizations of the average (and maximum) distance between nodes in a multi-type random network. We say a statement holds *asymptotically almost surely* if, for every $\delta > 0$, it holds with probability at least $1 - \delta$ in large enough societies.

PROPOSITION 4. If the random network process (\mathbf{P}, \mathbf{n}) satisfies (i)–(iv), then, asymptotically almost surely, the network is connected. Moreover, the average distance between nodes and the network diameter are asymptotically proportional to $\frac{\log n}{\log(\tilde{d}(\mathbf{P}, \mathbf{n}))}$: the average distance between nodes satisfies⁵⁰

$$\text{AvgDist}(\mathbf{A}(\mathbf{P}, \mathbf{n})) \in (1 + o(1)) \frac{\log n}{\log(\tilde{d}(\mathbf{P}, \mathbf{n}))};$$

and the diameter of the largest component satisfies

$$\text{Diam}(\mathbf{A}(\mathbf{P}, \mathbf{n})) \in \Theta\left(\frac{\log n}{\log(\tilde{d}(\mathbf{P}, \mathbf{n}))}\right).$$

50. The notation $o(1)$ indicates a factor going to 0 as n goes to infinity and Θ indicates proportionality up to a fixed finite factor.

Proposition 4 tells us that although homophily can change the basic structure of a network, it does not affect the average shortest path length between nodes in the network. Moreover, we have a precise expression for that average distance which is the same as for an Erdős-Rényi random network where links are formed uniformly at random with the same average degree. Effectively, as we increase the homophily, we increase the density of links within a group but decrease the number of links between groups. The result is perhaps somewhat surprising in showing that these two effects balance each other to keep average path length unchanged; more precisely, any deviations from the formula $\frac{\log n}{\log(d(\mathbf{P}, \mathbf{n}))}$ that are introduced by homophily only affect the result by adjusting a multiplicative constant, and not in the asymptotic rates.

The intuition behind the proposition can be understood in the following manner. Suppose that every node had a degree of d and that the network was a tree. Then the k -step neighborhood of a node would capture roughly d^k nodes. Setting this equal to n leads to a distance of $k = \frac{\log n}{\log d}$ to reach all nodes, and given the exponential expansion, this would also be the average distance. The proposition shows that this is exactly how the average distance behaves even when the network is not a tree and exhibits substantial clustering, even when we introduce noise to the network so that nodes do not all have the same degree, and even when we add substantial homophily to the network.

COROLLARY 3. Consider a process that has an expected convergence time proportional to the average distance between nodes. If it is run on two different random network formation sequences satisfying (i)–(iv) that have the same second-order average degree as a function of n , then the ratio of the expected convergence times of the two different random network sequences tends to 1, asymptotically almost surely.

The foregoing results tell us that the average distance is asymptotically not affected by homophily, and that the diameter is affected only up to a fixed finite factor, provided there is some minimal level of intergroup connectivity. Thus, direct contagion processes are not affected by homophily, but are affected by the link density in a society.

This provides an interesting contrast with the average-based updating processes, and helps clarify when and why homophily

matters. These results show that homophily is not affecting average path lengths in a network, and so not changing the distances that information has to travel. The density of the network plays that role.⁵¹ In the average-based updating processes that we have considered, it is *relative* connectivity from group to group that matters, not absolute distances. Homophily is critical in determining such connectivity ratios and thus consensus time in processes that are dependent on relative interaction across groups, while density is inconsequential. In contrast, if all that matters is overall distance, then homophily does not make a difference and instead only density is important.

VII.C. A Remark on Asymptotics

In Theorem 1, as well as in Proposition 4, a convergence time has a network structure statistic in the denominator and a $\log n$ in the numerator. In particular,

$$\text{CT}\left(\frac{\gamma}{n}; \mathbf{A}(\mathbf{P}, \mathbf{n})\right) \approx \frac{\log(n)}{\log\left(\frac{1}{h^{\text{spec}}(\mathbf{P}, \mathbf{n})}\right)},$$

while the average distance satisfies (with $a \sim b$ meaning that $\frac{a}{b}$ tends to 1 in probability)

$$\text{AvgDist}(\mathbf{A}(\mathbf{P}, \mathbf{n})) \sim \frac{\log n}{\log(\tilde{d}(\mathbf{P}, \mathbf{n}))}.$$

Thus it may appear that the asymptotic behavior of these two quantities in n is similar. This need not be the case. In particular, assumptions (ii) and (iii) allow for a wide range of variation in the second-order average degree $\tilde{d}(\mathbf{P}, \mathbf{n})$: it may grow as slowly as $\log^2 n$ or as quickly as $n^{\frac{5}{\log n}}$. On the other hand, under the assumptions of Definition 3, the denominator in the consensus time estimate cannot diverge to infinity. Thus, the growth rates of average distance and consensus time may behave very

51. One other thing is worth pointing out. The multi-type random network model allows for many different degree distributions, simply by allowing different groups to have different expected degrees. Nonetheless, the average distance between nodes depends only on the (second-order) average degree and not on any other moments, provided (i–iv) are satisfied. Other aspects of the degree distribution can matter in determining average distance if the conditions are violated. This happens, for example, in other classes of random graph models that have extreme variation in the degree distribution (unbounded variance as the number of nodes increases), as in scale-free networks.

differently under the assumptions for which our results are valid.

Moreover, the two formulas above, taken together, imply some more precise quantitative information about the comparison between consensus time and the time it takes for an infection to spread. In particular,

$$\frac{\text{CT}(\frac{\lambda}{n}; \mathbf{A}(\mathbf{P}, \mathbf{n}))}{\text{AvgDist}(\mathbf{A}(\mathbf{P}, \mathbf{n}))} \approx \frac{\log(\tilde{d}(\mathbf{P}, \mathbf{n}))}{\log\left(\frac{1}{|h^{\text{spec}}(\mathbf{P}, \mathbf{n})|}\right)}.$$

Therefore, when the asymptotics have “kicked in” and the propositions of this article give close approximations (which, in numerical experiments, occurs for values of n on the order of 1,000), we can make an explicit prediction, with an error bound, about which process is faster and by how much, knowing only the quantities in the above ratio.

In particular, suppose there are two large multi-type networks with the same second-order average degree $\tilde{d}(\mathbf{P}, \mathbf{n})$. One network has a spectral homophily of 1.2, while the other has a spectral homophily of 5. Then the above ratio of convergence times will be at least 4 times as large for the more homophilous network (because $\frac{\log(\frac{1}{1.2})}{\log(\frac{1}{5})}$ is about 8.8, and there is an approximation error of a little more than a factor of 2). By Proposition 4, this change in ratio is not due to the average distance, which stayed essentially the same, but purely due to the change in consensus time. This highlights the sense in which homophily matters for average-based updating but not for direct contagion.

Because the approximations \sim and \approx in the expressions work well even for networks of about 1,000 nodes, there is fairly precise quantitative information in these results beyond the asymptotics in n .

VIII. CONCLUDING REMARKS

Homophily has long been studied as a statistical regularity in the structure of social interactions. In this article, we propose a general measure of it—spectral homophily—and show that it is useful for characterizing the speed of convergence in natural average-based updating processes.

Homophily slows down the convergence of average-based updating according to the simple formula of Theorem 1, and

network density does not matter for the asymptotics of consensus time. In stark contrast, only density, and not homophily, matters for the speed of contagion processes.

The different relationships that we have uncovered between network characteristics and the speed of information flow have strong intuitions behind them. Under average-based updating processes, what is critical to convergence is the weight that nodes put on nodes of other types relative to those of their own type. Increasing the overall number of links while maintaining the homophily will not speed up convergence. With direct contagion processes only the average distance between pairs of agents matters. Adding homophily changes who is close to whom, but it does not change the average lengths of the shortest paths branching out from each node. Although the benchmarks we have analyzed turn out to represent extreme points, they offer some insight into the key elements of network structure that matter for learning and diffusion processes.

One can also compare our results to those of rational learning models that have been studied elsewhere, with the goal of understanding which models are more appropriate in different settings. Mueller-Frank's (2012) study of dynamic Bayesian updating in an arbitrary network entails that adding links to a network can only decrease the heterogeneity of observed outcomes. In average-based updating, adding links can strictly increase the observed heterogeneity of choices (by decreasing the rate of convergence) if those links are added in a way that increases homophily. Thus, these two models make different predictions about comparative statics, and these predictions are testable.⁵² While there have been some studies of how segregation based on exogenous characteristics, such as race, affects certain observed outcomes, such as school performance and happiness (e.g., Echenique, Fryer, and Kaufman 2006), we are not aware of any empirical work that has explored in detail whether such segregation leads to slower convergence to consensus or greater cross-sectional variation in beliefs or behaviors.⁵³ Field

52. Of course, some additional theoretical work would have to be done to make the predictions directly comparable; the comparison of the results in this discussion is somewhat heuristic.

53. The work of Bisin, Moro, and Topa (2011) asks some related questions in the course of a structural empirical exercise focused on smoking behavior in schools.

experiments may be a valuable tool to approach these important questions (see the discussion in Section II.D).

Beyond the conclusions about dynamics on networks, our results relating the second-largest eigenvalue of the interaction matrix to homophily, as well as the representative-agent theorem, could be useful more broadly. These provide general theoretical tools for a parsimonious mean-field approach to studying networks with homophily. The representative-agent updating matrix (and thus the spectral homophily) can be estimated consistently as long as we can obtain a consistent estimate of the *relative* interaction probabilities an agent has with various types. For example, in the network data on high school friendships in the National Longitudinal Study of Adolescent Health (see Golub and Jackson 2012), the relative frequencies of friendships among various races may be reasonably estimated based on the relative frequencies of nominations in surveys. On the other hand, the *absolute* densities of links between various races are much more difficult to infer, due to subjects' imperfect recall and an upper limit on the number of friends that can be named in the survey. This suggests econometric questions—which, to the best of our knowledge, are open—about how to estimate relative interaction frequencies most efficiently and how the resulting estimators of spectral homophily behave. The approach also naturally raises the issue of what other global properties of large networks can be estimated accurately using convenient summary statistics that avoid collecting too much local information; this is a potential avenue of further research.

Our results highlight the importance of understanding homophily to understand the functioning of a society. This is, of course, a first step and suggests many avenues for further research, of which we mention only some obvious ones. One clear direction for future work is considering processes that are not in either of the benchmark classes that we have examined here. The techniques used in deriving our conclusions are tailored to these two classes, so making headway on other kinds of processes will most likely involve developing some substantial new theoretical approaches.⁵⁴ For example, an interesting area to explore and

54. There are also questions regarding what sort of updating people employ. An average-based updating rule makes the most sense when people find it difficult to communicate precisely what signals they have observed and also what others have observed and so forth, but instead can only communicate aggregated beliefs.

compare results with would be the study of coordination games on networks (different from the ones we have studied), where it has also been found that network structure can affect both the strategic choices (Young 1998; Morris 2000; Jackson 2008b) and the speed of convergence (Ellison 2000; Montanari and Saberi 2010). Another interesting question is how belief dynamics look in a model of Bayesian updating with heterogeneous priors (depending on an agent's type, which may correlate with network position), and whether that sort of process can be readily distinguished in the data from myopic updating under homophily.

The study of how homophily affects communication naturally raises the issue of how advances in technology and changes in the media affect the structure of agents' information networks and the outcomes of communication. As pointed out by Rosenblat and Möbius (2004), technology that makes interactions easier may lead to greater network density (and lower average path length) even as it increases homophily among groups: agents may interact more but use the better technology to seek out agents more like themselves to interact with. Our results tie this back to the speed of convergence to a consensus in different processes and show that average-based updating may become slower after the introduction of a new communication technology.

Mass media outlets play a key role in the flow of information. Thus, it is important to incorporate their effects into models of communication. In our setting, this issue raises three broad questions—a conceptual one, an empirical one, and a theoretical one—which we believe to be fruitful avenues for future work. The conceptual one is how to model a media outlet in the type of framework studied in this article. It can be modeled as a “forceful agent” with its own agenda that influences others more than it is influenced (see, e.g., Acemoglu, Ozdaglar, and ParandehGheibi 2010), or as a very widely observed agent that simply aggregates and rebroadcasts along others' information.

However the media may be modeled, there is an empirical question of how agents are influenced by the various media content that they consume. Gentzkow and Shapiro (2012) show that the consumption of online media content, as measured by website visits, exhibits much less segregation than typical offline

In cases where it is easy to communicate signals directly, the weights that agents use in updating may adjust more over time. There is some preliminary evidence in this direction in a paper by Möbius, Phan, and Szeidl (2010).

face-to-face interaction (somewhat surprisingly, in view of the potential for greater choice-driven segregation online). Since disagreement nevertheless persists, it is reasonable to consider the possibility, as Gentzkow and Shapiro do, that agents weight and process the different information they receive very differently depending on its source. A liberal may consult conservative media but not change his or her views, or even move to the left as a result. Thus, it is important to be able to estimate the weights that describe agents' updating and to understand how those weights depend on the media outlet and the issue under discussion.

There is then a theoretical question of how to extend our results when there is heterogeneity in the weights agents place on different neighbors, which can include media outlets. The baseline model considered in this article has every neighbor who is listened to being weighted equally. When heterogeneity in weights is present, the specification of the group-level updating matrix should change to account for the weights. We conjecture that, once this is done, analogs of our results—ones that reduce the study of updating in large networks to the study of homophily in small deterministic representative-agent matrices—should be available. However, such results do not seem easy to obtain using the proof techniques we have used, and would constitute a substantial technical advance as well as a valuable tool for applied modeling and estimation.

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at QJE online (qje.oxfordjournals.org).

HARVARD UNIVERSITY AND MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
STANFORD UNIVERSITY,
SANTA FE INSTITUTE, AND CANADIAN INSTITUTE FOR
ADVANCED RESEARCH

REFERENCES

- Acemoglu, Daron, Asuman Ozdaglar, and Ali ParandehGheibi, "Spread of (Mis)information in Social Networks," *Games and Economic Behavior*, 70 (2010), 194–227.

- Acemoglu, Daron, Munther Dahleh, Ilan Lobel, and Asuman Ozdaglar, "Bayesian Learning in Social Networks," *Review of Economic Studies*, 78 (2011), 1201–1236.
- Bisin, Alberto, Andrea Moro, and Giorgio Topa *The Empirical Content of Models with Multiple Equilibria in Economies with Social Interactions*. (Federal Reserve Bank of New York Staff Report No. 504, 2011).
- Bramoullé, Yann, Sergio Currarini, Matthew O. Jackson, Paolo Pin, and Brian Rogers, "Homophily and Long-Run Integration in Social Networks," *Journal of Economic Theory* (forthcoming 2012).
- Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin, "Identification of Peer Effects through Social Networks," *Journal of Econometrics*, 150 (2009), 41–55.
- Calvó-Armengol, Antoni, Eleonora Patacchini, and Yves Zenou, "Peer Effects and Social Networks in Education," *Review of Economic Studies*, 76 (2009), 1239–1267.
- Chandrasekhar, Arun, Horacio Larreguy, and Juan Pablo Xandri *Testing Models of Social Learning on Networks: Evidence from a Lab Experiment in the Field*. (Consortium on Financial Systems and Poverty Working Paper, 2010).
- Choi, Syngjoo, Douglas Gale, and Shachar Kariv "Behavioral Aspects of Learning in Social Networks: An Experimental Study," in *Advances in Applied Microeconomics*, ed. John Morgan (Berkeley, CA: BE Press, 2005).
- Chung, Fan, and Linyuan Lu, "The Average Distance in Random Graphs with Given Expected Degrees," *Proceedings of the National Academy of Sciences*, 99 (2002), 15879–15882.
- Chung, Fan, Linyuan Lu, and Van Vu, "The Spectra of Random Graphs with Given Expected Degrees," *Internet Mathematics*, 1 (2004), 257–275.
- Coleman, James S., "Relational Analysis: The Study of Social Organizations with Survey Methods," *Human Organization*, 17 (1958), 28–36.
- Copic, Jernej, Matthew O. Jackson, and Alan Kirman, "Identifying Community Structures from Network Data," *Journal of Theoretical Economics (Contributions)*, 9 (2009), Article 30.
- Corazzini, Luca, Filippo Pavesi, Beatrice Petrovich, and Luca Stanca *Influential Listeners: An Experiment on Persuasion Bias in Social Networks*. (CISEPS Research Paper No. 1/2011, 2011).
- Currarini, Sergio, Matthew O. Jackson, and Paolo Pin, "An Economic Model of Friendship: Homophily, Minorities and Segregation," *Econometrica*, 77 (2009), 1003–1045.
- , "Identifying the Roles of Race-Based Choice and Chance in High School Friendship Network Formation," *Proceedings of the National Academy of Sciences*, 107 (2010), 4857–4861.
- DeGroot, Morris H., "Reaching a Consensus," *Journal of the American Statistical Association*, 69 (1974), 118–121.
- DeMarzo, Peter, Dimitri Vayanos, and Jeffrey Zwiebel, "Persuasion Bias, Social Influence, and Uni-Dimensional Opinions," *Quarterly Journal of Economics*, 118 (2003), 909–968.
- Diaconis, Persi, and David Freedman, "On the Statistics of Vision: The Julez Conjecture," *Journal of Mathematical Psychology*, 24 (1981), 112–138.
- Diaconis, Persi, and Daniel Stroock, "Geometric Bounds for Eigenvalues of Markov Chains," *The Annals of Applied Probability*, 1 (1991), 36–61.
- Echenique, Federico, and Roland Fryer, "A Measure of Segregation Based on Social Interaction," *Quarterly Journal of Economics*, 122 (2007), 441–485.
- Echenique, Federico, Roland Fryer, and Alex Kaufman, "Is School Segregation Good or Bad?" *American Economic Review*, 96 (2006), 265–269.
- Ellison, Glenn, "Basins of Attraction, Long Run Stochastic Stability and the Speed of Step-by-Step Evolution," *Review of Economic Studies*, 67 (2000), 17–45.
- Fienberg, Stephen E., Michael M. Meyer, and Stanley S. Wasserman, "Statistical Analysis of Multiple Sociometric Relations," *Journal of the American Statistical Association*, 80 (1985), 51–67.
- French, John, "A Formal Theory of Social Power," *Psychological Review*, 63 (1956), 181–194.

- Gentzkow, Matthew, and Jesse M. Shapiro, "Ideological Segregation Online and Offline," *Quarterly Journal of Economics*.
- Golub, Benjamin, and Matthew O. Jackson, "Does Homophily Predict Consensus Times? Testing a Model of Network Structure via a Dynamic Process," *Review of Network Economics* (forthcoming).
- , "Naïve Learning in Social Networks: Convergence, Influence, and the Wisdom of Crowds," *American Economic Journal: Microeconomics*, 2 (2010), 112–149.
- , "Network Structure and the Speed of Learning: Measuring Homophily Based on its Consequences," *Annals of Economics and Statistics* (forthcoming 2011), available at ssrn.com/abstract=1784542 (accessed June 27, 2012).
- Harary, Frank "A Criterion for Unanimity in French's Theory of Social Power," in *Studies in Social Power*, ed. Dorwin Cartwright, (Ann Arbor: University of Michigan Press, 1959).
- Holland, Paul W., Kathryn B. Laskey, and Samuel Leinhardt, "Stochastic Blockmodels: First Steps," *Social Networks*, 5 (1983), 109–137.
- Jackson, Matthew O. "Average Distance, Diameter, and Clustering in Social Networks with Homophily," in *Proceedings of the Workshop on Internet and Network Economics (WINE 2008)*, ed. Christos Papadimitriou, and Shuzhong Zhang, (New York: Springer-Verlag, 2008).
- *Social and Economic Networks*. (Princeton, NJ: Princeton University Press, 2008).
- Jackson, Matthew O., and Leeat Yariv "Diffusion, Strategic Interaction, and Social Structure," in *Handbook of Social Economics*, ed. Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, (San Diego, CA: North Holland, 2011), 645–678.
- Lazarsfeld, Paul F., and Robert K. Merton "Friendship as a Social Process: A Substantive and Methodological Analysis," in *Freedom and Control in Modern Society*, ed. M. Berger, (New York: Van Nostrand, 1954).
- Marsden, Peter V., "Core Discussion Networks of Americans," *American Sociological Review*, 52 (1987), 122–131.
- McPherson, Miller, Lynn Smith-Lovin, and James M. Cook, "Birds of a Feather: Homophily in Social Networks," *Annual Review of Sociology*, 27 (2001), 415–444.
- McSherry, Frank, "Spectral Partitioning of Random Graphs," *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC 2001)* (2001).
- Meyer, Carl D. *Matrix Analysis and Applied Linear Algebra*. (Philadelphia, PA: SIAM, 2000).
- Möbius, Markus, Tuan Phan, and Adam Szeidl *Treasure Hunt*. (Mimeo, 2010).
- Montanari, Andrea, and Amin Saberi, "The Spread of Innovations in Social Networks," *Proceedings of the National Academy of Sciences*, 107 (2010), 20196–20201.
- Montenegro, Ravi, and Prasad Tetali, "Mathematical Aspects of Mixing Times in Markov Chains," *Foundations and Trends in Theoretical Computer Science*, 1 (2006), 237–354.
- Morris, Stephen, "Contagion," *Review of Economic Studies*, 67 (2000), 57–78.
- Mossel, Elchanan, and Omer Tamuz *Efficient Bayesian Learning in Social Networks with Gaussian Estimators*. arXiv:1002.0747 (Working Paper, 2010).
- Mueller-Frank, Manuel, "A General Framework for Rational Learning in Social Networks," *Theoretical Economics* (2012), forthcoming.
- Neilson, William S., and Harold Winter, "Votes Based on Protracted Deliberations," *Journal of Economic Behavior & Organization*, 67 (2008), 308–321.
- Parikh, Rohit, and Paul Krasucki, "Communication, Consensus, and Knowledge," *Journal of Economic Theory*, 52 (1990), 178–189.
- Pew Research Center, "Little Consensus on Global Warming: Partisanship Drives Opinion," Report, Washington, DC, 2006. Available at www.people-press.org/report/280/ (accessed June 27, 2012).
- Rosenblat, Tanya S., and Markus M. Möbius, "Getting Closer or Drifting Apart," *Quarterly Journal of Economics*, 119 (2004), 971–1009.

- World Public Opinion, "Iraq: The Separate Realities of Republicans and Democrats," Report, Washington, DC, 2006. Available at www.worldpublicopinion.org/pipa/articles/brunitedstatescanadara/186.php (accessed June 27, 2012).
- Young, H. Peyton. *Individual Strategy and Social Structure*. (Princeton NJ: Princeton University Press, 1998).